Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Summer 1991 No. 27 été 1991

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

Man vs. Machine: The Story of the Checkers Program "Chinook"

Connie Bryson

Homme vs Machine : L'histoire du logiciel de jeu de dames "Chinook"

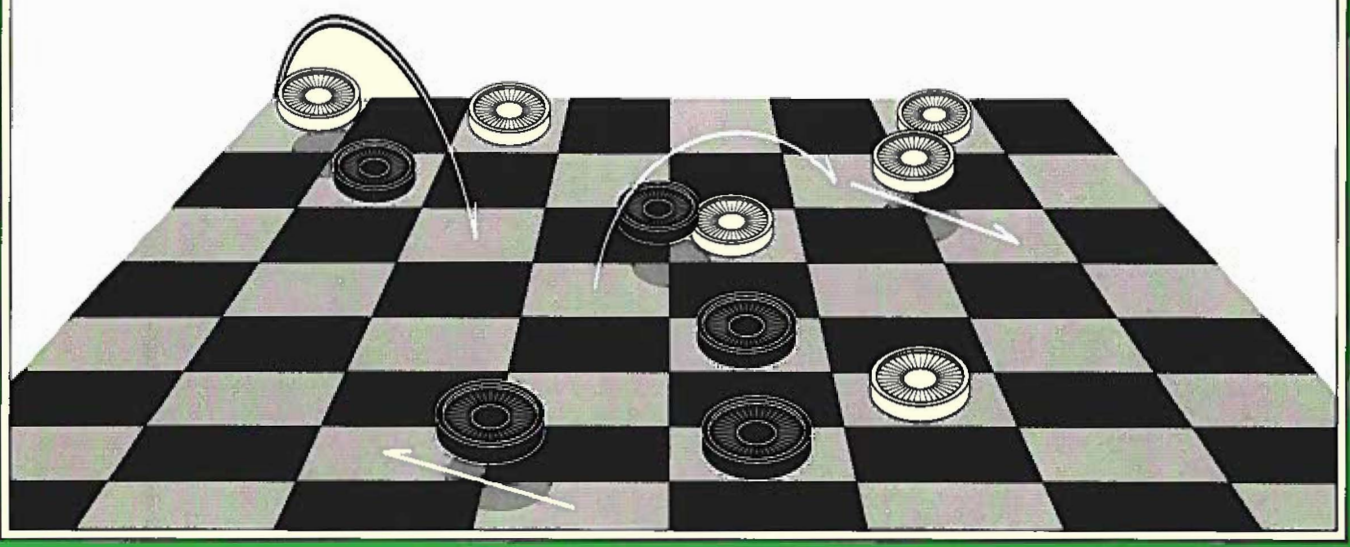

NeuralWorks Professional II and Designer Pack... the sensible approach to neural network development

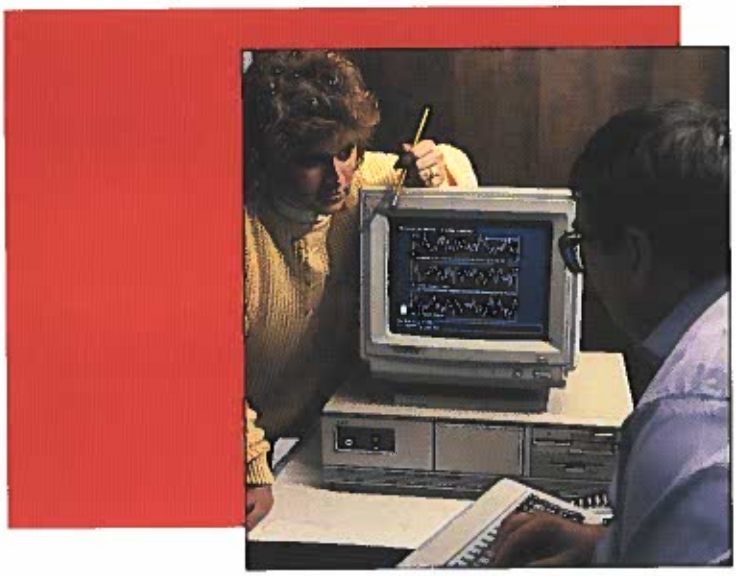

NEURALWARE, INC.

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 of NeuralWare's products, AAI offers the following services:

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

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

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

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

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

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

• robotics (perception, control, action fusion, generation of action sequence) • artificial skin/muscle, vision and scene analysis • character and object recognition • process control (process monitoring, quality control, generation of flexible manufacturing schedule) • expert systems (diagnosis, control expert system) and automated knowledge acquisition • voice recognition • signal processing (intelligent noise filtering, signal formatting) • radar clutter elimination • risk analysis for credit authorization • prediction • detection of explosives • autonomous vehicles (sensory and control processes map generation, map interpretation, navigation and maneuvering) • various multiple constraint satisfaction problems (urban traffic optimization, communication network optimization amenity optimization) • recognition (sonar images, on-orbit satellites, mineral ore discovery, radar return signals, radar signal source 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

Summer 1991 No. 27 ete 1991

Canada's National AI magazine.

Editor · Editeur
Roy Masrani
Editor Emeritus · Rédacteur emeritus
Graeme Hirst
Production Manager · Chef de production
Carol Tubman
Production · Production
Ad Ventures Studio Inc.
Academia · Académiques
Jan Mulder
Industrial Applications · Applications Industriel
Chris Lumb
PRECARN Update · Nouvelles de PRECARN
Jean-Claude Gavrel
CIAR Update · Nouvelles de CIAR
Zenon Pylyshyn
World Watch · Vue sur le monde
Russ Thomas
Book Reviews · Critiques de livres
Graeme Hirst
Translation · Traduction
Michel Addison
Advertising · Publicité
Marco Ariano

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

Contents

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

Canadian Artificial Intelligence
Alberta Research Council
3rd Floor, 6815, 8th St. N.E.,
Calgary, AB T2E 7H7
or — masrani@noah.arc.ab.ca

Advertising rates are available upon request from the address above.

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

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

ACAI - valete

Dick Peacocke

The Associate Committee on Artificial Intelligence was created by the National Research Council in 1987 to act as a focal point and forum in Canada for issues relating to AI of national concern to the scientific and engineering communities. Like other NRC Associate Committees, of which there were approximately twenty, it was also to provide advice to NRC and government on aspects of this new science and the engineering technologies associated with it. This is all in the past tense because the Associate Committees, including the AI one, have been disbanded recently.

The Associate Committee on AI was formed after discussions between NRC, Fraser Mustard (CIAR President), Gordon McCalla, and myself (then CSCSI Past-President and President respectively). These discussions took place in the summer of 1986. The Committee met thirteen times between 1987 and 1991 (in various places across Canada). There were initially twenty-one members, selected from universities, industry, and government. Some of the membership was changed during the committee's lifetime, but a few diehards like McCalla and myself stuck it out. It was a pleasure to work with the three chairpersons, Claude Lajeunesse, Renato De Mori, and Grant Thomas, and also with the ever diligent secretary Jack Scrimgeour. Members were unpaid, but reimbursed by NRC for travel and accommodation expenses.

Encouraging closer collaboration between AI users and the AI research community in university and government was a major topic chosen by the committee. It's a challenge, and I can't say that we achieved tangible results but at least we tried. A report we produced called "AI in Canada" may be of some help in putting people and organizations in touch. You can get copies (free of charge, in English or French) from the Editorial Office at NRC. An earlier report, "The Social Context of Artificial Intelligence" is also available.

As part of a cost-cutting exercise NRC has abandoned the associate committees, at least in their present form. The AI committee provided a wonderful networking opportunity; it was great while it lasted, but in my own opinion terminating it was probably not a bad way for the government to save some money. It won't save much of course because most of the real costs were members' time - and that was donated! Valete means goodbye...

A robot helps a startup clean up

Beppl Crosarlol – staff writer

Un robot aide une nouvelle compagnie à faire le grand nettoyage.

It's a janitor's dream: a robot that can vacuum a 6,000-square-foot hall in an hour. And for Vivek Burhanpurkar, it's turned into a lucrative reality. The 27-year-old electrical-engineering dropout has just signed two multimillion-dollar contracts to supply his "CyberVacs" to major cleaning companies in the United States and Europe. It's a major windfall for Cyberworks Inc., his Orilla, Ont., start-up. And Burhanpurkar says interest is growing among other members of North America's $50-billion contract-cleaning industry.

Hodon Group BV, an Amsterdam cleaning company, recently ordered 300 of the $25,000 to $35,000 machines after three years of successful on-site testing. Meanwhile, ISS International Service System Inc. of New York, a $500-million cleaning concern, which originally ordered five robot vacuums, now wants 100 more. Burhanpurkar is gearing up for high-volume production and will be able to deliver up to 150 units by year-end. The battery-driven robot, about the size of a large trash can, is guided by 12 sonar "eyes"

C'est le rêve de tout concierge: un robot qui peut nettoyer une salle de 6,000 pieds carré dans l'espace d'une heure. Et pour Vivek Burhanpurkar, cela s'avère une réalité lucrative. Cet ingénieur-électricien de 27 ans, vient juste de signer 2 contrats de plusieurs millions de dollars, pour fournir son "CyberVac" aux plus grosses entreprises d'entretien commercial aux Etats-Unis et en Europe. C'est vraiment un débouché extraordinaire pour Cyberworks Inc., sa nouvelle campagnie située à Orillia en Ontario. Et Burhanpurkar affirme que plusieurs membres de l'industrie d'entretien commercial (évaluée à plus de $50 milliards), démontrent un intérêt grandissant pour son robot.

Hodon Group BV, une compagnie d'entretien commercial d'Amsterdam, a récemment commandé 300 de ces machines de $25,000 à $30,000, aprés plus de 3 ans de test concluant. Entre-temps, ISS International Service System Inc. de New York, une compagnie d'entretien commercial valant plus de $500 millions, qui originellement avait commandé 5 robots-aspirateurs, en veut maintenant 100 de plus. Burhanpurkar se prépare à augmenter sa production et sera prêt à livrer jusqu'à 150 unités d'ici la fin de l'année.

mounted to a revolving shaft, which use sound waves to detect obstacles. An artificial-intelligence program, written by Burhanpurkar, interprets the data and guides the machine around furniture and other obstacles. Competing systems have to be programmed to follow a predetermined course, Burhanpurkar says. "They just don't cut it," he says. "You can't be reprogramming a robot every time you move a chair." Burhanpurkar got the idea for the CyberVac while working on his final-year thesis in university. He built a series of prototypes with $50,000 in seed capital from his father and coaxed $750,000 out of Hodon after a company representative spotted him at a 1987 German trade show demonstrating his handy new invention. Moreover, the money came in the form of a hassle-free advance against future shipments – not a controlling equity investment or power-sharing strategic partnership – which kept the company in the hands of Burhanpurkar and his father. Relatives and government later kicked in an additional $600,000.

Burhanpurkar says he has tried to sell his robots in Canada but so far has had no takers. "Canadian companies tend to be highly conservative in their acceptance of new products," he says.

reprinted with permission from The Financial Times, March 25, 1991.

Ce robot mû par une batterie, de la grosseur d'une grosse poubelle, est guidé par 12 yeux de type "SONAR" assemblés sur un essieu pivotant, qui utilise les ondes sonores pour détecter les obstacles. Un programme d'intelligence artificielle conçu par Burhanpurkar, interprète les données et guide le robot autour du mobilier et autre obstacles. Les systèmes compétitifs existants, doivent être programmés pour suivre une rout pré-déterminée, selon Burhanpurkar. "Ils ne sont pas à la hauteur" selon lui. "Vous ne pouvez pas re-programmer un robot à chaque fois que vous déplacez une chaise."

Burhanpurkar a eu l'idée du CyberVac en travaillant sur sa thèse universitaire finale. Il construisit une série de prototypes avec $50,000 en capital venant de son père et réussit à se faire subventionner $75,000 par Hodon, aprés qu'un représentant de la compagnie l'ai remarqué dans un salon d'affaires en Allemagne en 1987, alors qu'il démontrait sa nouvelle invention. De plus, cette subvention arriva dans la forme d'une avance pour livraisons futures – non un investissement ou partenariat stratégique – où la compagnie demeurera entièrement la propriété de Burhanpurkar et de son père. Des parents et le gouvernement injectèrent par la suite, la somme additionnelle de $600,000.

Burhanpurkar a essayé de vendre ces robots au Canada, mais sans succès. "Les compagnies canadiennes ont tendances à être très conservatrices quant à l'acceptation de nouveaux produits," selon lui.

Controlling Super-Intelligent Machines

Peter Turney – National Research Council

Introduction

Imagine that AI research leads to the development of a "superintelligent machine" (SIM); a machine significantly more intelligent than any human being. How could we control such a machine?

This question assumes that we would want to control a SIM, rather than granting it complete freedom. Science fiction abounds with stories in which a SIM pursues a plan which does not please its human creators. For example, Colossus (the inspiration for the movie The Forbin Project) presents a plausible scenario involving a SIM created to manage the nuclear weaponry of the U.S.A. [6]. In the end, the computer treats its creators like laboratory animals.

Some of us might welcome SIM's as our evolutionary successors. Yet, even if we view our species as merely a stepping stone in the path of evolution, it would be beneficial to be able to control our successors to some extent: The stepping stone should not be trampled by careless feet.

Goals

It seems that goal-directed behaviour is intrinsic to intelligence, so we should expect SIM's to have goals. For example, SIM's might be motivated by a desire for knowledge, survival, or power. Perhaps SIM's will not require goals, but let us assume for the sake of argument that they will. The problem is that our goals may conflict with the goals of SIM's.

The simplest solution would be to design SIM's so that their goals are in harmony with our own. This might not be possible. For example, SIM's may be so complex that we do not understand them well enough to design goals for them. Also, it is easy to imagine that we approve of a SIM's goals, but not the method it uses to achieve its goals.

Perhaps a SIM can be constructed to have a weak will, or a strong desire to please its owner, like a good dog. It should be easy to control such a SIM. I suspect, however, that superintelligence is incompatible with subservience or weakness of will. This is a matter for further research.

Brute Force

We could control a SIM by killing it, or threatening to do so, unless it obeyed our instructions. This would require continual observation of the SIM, to verify that it is indeed obedient. A SIM could outsmart a human supervisor or a

simple software supervisor. It would seem that the inspection must be done by another SIM. Could we trust SIM's to monitor each other? This would be like trusting a pack of wolves to keep each other from attacking sheep.

If we did detect a disobedient SIM, it could be difficult to carry out our threat. A SIM might defend itself by copying its software into several machines. This defense might be countered by a special "virus" program, tailored to seek out and destroy all copies of the SIM's software; the SIM might be designed so that its hardware or software discourages copying; the SIM might be designed to self-destruct unless it is periodically sent a special code; or the SIM's access to other hardware could be severely restricted.

Asimov's Three Laws of Robotics

Asimov considered the question of how humans could control robots. As a result, he developed his Three Laws of Robotics [1]:

1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law.

Suppose these laws were incorporated in a SIM's program. Would this not give us all the control we need?

John Sladek provides a witty critique of Asimov's Laws [8]. The problem with laws is their interpretation. For example, a SIM must decide what is a human being, before it can obey the First Law. Perhaps it will decide that it is itself a human being, and then refuse to take orders from any other human beings. Perhaps it will decide that most people do not qualify for human status. Sladek describes eleven ways a robot could misinterpret Asimov's Laws.

Limits to Growth

Gibson's Neuromancer portrays a future in which there are legal limits to the intelligence an AI may possess [3]. The law — enforced by "Turing Police" — prevents AI's from becoming SIM's. Benford's "Me/Days" proposes a law requiring that a SIM must have its memory erased after each job [2]. This would prevent the SIM from acquiring enough knowledge to be dangerous to people. Similarly, we might set legal limits to a machine's self-awareness or its consciousness. Another suggestion is to control SIM's by organizing them into committees (K. Deaton, personal communication). It is well-known that the intelligence of a committee is much less than the intelligence of any individual committee member.

These methods control SIM's by setting external limits on their growth. An AI approaching these limits would naturally try to evade them. This is what happens in both Gibson's and Benford's stories. Once the djinn is out of the bottle, it may be difficult to put it back in.

Pleasure and Pain

Alan Turing suggests that it will be useful to program an AI to experience pain and pleasure [9]. Turing argues that pain and pleasure are integral parts of learning: We learn to do a certain thing, because it pleases us; we learn to avoid a certain thing, because it hurts us. Pain and pleasure would also provide a way for us to control the AI.

The control of humans and animals by this method has been intensively studied in behaviourist psychology. Skinner

continued on page 12

The CSCSI ANNUAL GENERAL MEETING

Will be held at AAAI, on Thursday July 18
12:50 to 2:00 p.m.

Joe Culberson (left) and Jonathan Schaeffer

MAN VS. MACHINE:

The Story of the Checkers Program "Chinook"

Connie Bryson

Edmonton likes to call itself the "City of Champions", a title that reflects the prowess of its teams in sports like hockey and football. But the city also lays claim to a champion in the "brains not brawn" category — the world's number-two checkers player.

The "star" is a computer program called Chinook, developed at the University of Alberta. Chinook came second to the human world checkers champion in the U.S. National Open in 1990, thus earning the right to compete for the World Championship this year.

In the beginning

It all began over a cup of coffee in the Students' Union Building at the University of Alberta in 1989. Joe Culberson and Duane Szafron asked Jonathan Schaeffer (they're all in the U of A's computing science department) what happened to checkers. Culberson and Szafron knew Schaeffer had done a lot of work on computer chess, they thought it might be possible to solve the game of checkers.

"I estimated there were 5 X 10^{20} possible positions but many of those moves were illegal," says Culberson. "It seems that about 10^{18} positions are real. That's a large number, but the technology available today puts solving checkers within reach."

With that casual conversation, the Chinook project was born. Besides Schaeffer, Culberson and Szafron, Chinook's other developers are Norman Treloar, Brent Knight and Paul Lu. Today, Chinook is one of the strongest checkers players in the world and the strongest 8 X 8 checkers program in existence.

You can be excused if you're still not impressed. Computer scientists may look down their noses at checkers in the belief that the game has already been solved. That misconception dates back to the 1960s and the work of Arthur Samuel. Samuel took on checkers and pioneered a machine learning system. He developed a program that, given simple rules, worked itself up to expert level.

Contrary to the hype, Samuel's program did not conquer checkers. It did win one game against a blind checkers champion, but was not even playing at Master level. Nonetheless, word spread that checkers had been solved.

"AI was young then, it was a new field," notes Schaeffer. "Even though it won only one game, the success of Samuel's program went around the world. That misconception has stayed around and done us harm.

"When I tell people I'm working on checkers they often say 'Wasn't that already solved?'"

Chinook: More than a program

The Chinook project has three objectives: solving the game, beating the world champion, and doing AI research. Solving the game is the long-range objective. Schaeffer is confident that Chinook will achieve this goal.

Right now, the research team is working on endgame databases. This involves enumerating all positions with "n" pieces or less on the board and computing whether each is a win, loss or draw. The 6-piece databases comprise 2.5 billion positions; the 7-piece databases an additional 35 billion. As of September, 1990, about 15 billion positions in the 6- and 7-piece endgame databases had been computed. The goal is the 8-piece database — 400 billion positions.

The short-range objective is to have Chinook defeat the human World Champion in a match. Chinook is very good. It came second in the U.S. National Open which is restricted to Master and Grandmaster players. Chinook didn't lose a match, but came second to the World Champion Dr. Marion Tinsley, who won one more match than Chinook.

Tinsley is no ordinary checkers player. He has been the best player in the world since 1950, losing only five games in the past 40 years.

"That sort of record indicates incredible intellectual skill," Schaeffer says. "Tinsley is the best player in the world by far. I can't imagine anyone closer to perfection. Even if Chinook wins only one game against him, it will be an achievement."

Tinsley was in Edmonton last December to play a 14-game exhibition match against Chinook. The final score

---

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

was a narrow 7.5-6.5 victory for Tinsley. He won one game and drew the rest. From his comments in "Checkers", a magazine put out by the International Checker Hall of Fame, Tinsley appears to enjoy playing Chinook.

He wrote: "One really has to hand it to them (Chinook's developers) for their determination and bulldog tenacity. I for one welcome all this interest in our Grand Old Game. It is a refreshing change even if they are 'after my hide'!"

"Chinook is the first real challenger Tinsley's had in a long time," Schaeffer adds. "The computer doesn't have preconceptions about the game, it doesn't necessarily play by the book. Consequently it moves in uncharted territory, making risky, unusual moves. Tinsley once called Chinook 'a precocious child'.

"Eventually we're going to win. Tinsley (63 years old) is getting older. Chinook is getting better."

The research angle

Working with Chinook isn't all fun and games. Schaeffer and company use the game as a test bed for AI problems such as search, pattern recognition, and learning.

"Checkers is a good opportunity to solve some of these problems," Schaeffer says. "When applied to chess, AI problem solving is very difficult. The problems are the same in checkers but the game isn't as complex.

"Plus we get a real kick out of it. We get immediate feedback and can implement changes and see the results in competition."

Culberson calls knowledge Chinook's "Achilles' heel". The problem is getting the program to think rather than compute.

"Humans can make logical inferences about the board," Culberson explains. "They can see sequences and take long leaps ahead. Human players visualize 20 moves ahead and then start analyzing. The program analyzes at the very first step. The computer can look ahead; a human player can leap ahead."

Over the summer the team will be working on Chinook to make it more selective about what it searches. "Our goal is to have Chinook understand what is obvious and modify its thinking along those lines," Culberson says.

Planning for the "big game"

Chinook's big test will likely come in November, the tentative date for the World Championship in London, England. The match will likely be billed a "man vs. machine" championship as neither the American Checker Federation nor the English Draughts Association will sanction a World Championship Match that has a computer vying for a human title.

"The associations want to keep the World Championship a human championship," explains Schaeffer who is disappointed by this turn of events. He knows that Chinook, by coming second to the World Champion at the US Open, earned the right to compete in the World Championship.

"A lot of people are afraid Tinsley will lose," Schaeffer continues. "If I were a betting man, I would bet on Tinsley."

Although last year's exhibition match against Tinsley was a good work-out for Chinook, the program still needs more experience playing checkers experts.

"We work in isolation and have to go out to find the top ten or twenty players to play Chinook," Schaeffer says. "Of course they want the games made interesting for them and that means prize money."

While Schaeffer has the computing resources for Chinook's development, he's short on cash. He is looking for sponsors to help with prize money and travel costs. As well, solving checkers requires a lot of compute cycles, disk storage and memory. The team could use a more powerful system.

But checkers isn't the kind of high-profile game that attracts corporate sponsors with deep pockets. Schaeffer says most people think of checkers as a game for children and old men.

"As a chess player, I was scornful of checkers before I started the project," he says. "Checkers rules are very simple. However it's a mistake to think that the simplicity of the rules means the game is simple. Checkers is more complex and demanding than I ever imagined. It is truly an intellectual game par excellence."

Chinook: The nuts and bolts

This brief description of Chinook is adapted from the research team's 1990 technical report. For a more detailed look at Chinook, readers should consult "Reviving the Game of Checkers" by Schaeffer et al., University of Alberta, Department of Computing Science, Technical Report TR 90-31, September, 1990.

Chinook is a typical alpha-beta search program, with iterative deepening (two-ply at a time), transposition tables and the history heuristic. Its strength comes from deep searches (deciding which positions to examine), a good evaluation function (deciding how favourable a position is) and endgame databases which contain perfect information on all positions with six pieces or less.

The checkers knowledge used in the program is a collection of heuristics devised from experience in playing with Chinook and from general experience in the programming of other games. The relative importance of each heuristic was adjusted manually, based on data obtained from the checkers literature and from experience playing against the program.

Program Committee:

Ranan **Banerji**, SL Joseph's U., Philadelphia, PA
Wolfgang **Bibel**, I.T.H., Darmstadt, West Germany
David **Bonham**, Mechanical Eng., U.N.B.
Z. **Chen**, U. of Nebraska, Omaha, NE
Chang **Choo**, W.P.I., Worcester, MA
Rajamani **Doraiswami**, Electrical Eng., U.N.B.
Martin A. **Fischler**, SRI Int., Menlo Park, CA
Herbert **Freeman**, Rutgers U., New Brunswick, NJ
C. Lee **Giles**, NEC Research Ins., Princeton, NJ
Paul **Gillard**, Memorial U., St. John's, Nfld.
Vasant **Honavar**, Iowa State U., Ames, IA
Laveen **Kanal**, U. of Maryland, Coll. Park, MD
Vladik **Kreinovich**, U. of Texas, El Paso, TX
Wemer **Kuhn**, U. of Maine, Orono, ME
Bernd **Kurz**, Computer Science, U.N.B.
Patrice **Lapointe**, AECL Res., Chalk River, Ont.
Mark **Lidd**, Mitre Corp., Fairfax, VA
B.I.B. **Madhav**, I.I.T., Madras, India
Ettore **Merlo**, C.R.I.M., Montréal, Qué.
Tom **Mitchell**, Carnegie Mellon, Pittsburgh, PA
Jan **Mulder**, Dalhousie U., Halifax, N.S.
Eric **Neufeld**, U. of Saskatchewan, Saskatoon, Sask.
N. **Parameswaran**, I.I.T., Madras, India
Richard **Peacocke**, Bell-Northern Res., Ottawa, Ont.
Denis **Riordan**, T.U.N.S., Halifax, N.S.
John **Robinson**, Civil Eng., U.N.B.
Aznel **Rosenfeld**, U. of Maryland, Coll. Park, MD
Leemseop **Shim**, Governors State U., Univ. Park, IL
Eduardo **Sontag**, Rutgers U., New Brunswick, NJ
Paul **Tarau**, U. de Moncton, Moncton, N.B.
Manoel **Tenorio**, Purdue U., West Lafayette, IN
Richard **Tervo**, Electrical Eng., U.N.B.
Tetsuyuki **Toyofuku**, National-Panasonic, Osaka, Japan
André **Trudel**, Acadia U., Wolfville, N.S.
Lloyd **Waugh**, Civil Eng., U.N.B.
Ian **Witten**, U. of Calgary, Calgary, Alta.
Andrew K.C. **Wong**, U. of Waterloo, Waterloo, Ont.
Jean **Zanazaka**, U. de Sherbrooke, Qué.

CALL for PARTICIPATION

The Fourth UNB
Artificial Intelligence Symposium

University of New Brunswick
Fredericton, New Brunswick, Canada
Friday & Saturday, Sept. 20 & 21, 1991

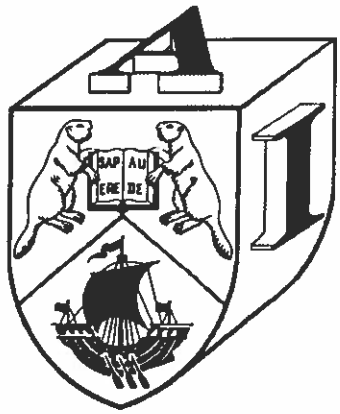

Sponsored by:
Faculty of Computer Science, University of New Brunswick
Bell-Northern Research
Canadian Society for Computational Studies of Intelligence

In Cooperation with:
American Association for Artificial Intelligence (AAAI)
IEEE Computer Society
International Association of Knowledge Engineers (IAKE)
Japanese Society of Artificial Intelligence

Keynote Address:
"Self-Reliant Robots: The Ambler Rover and Beyond" by Dr. Reid Simmons, Research Computer Scientist, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA

Panel Discussion (1.5 hours):
"On the Role of Machine Learning in Artificial Intelligence"

The UNB Artificial Intelligence Symposium provides a forum for an exchange of ideas, experiences and information about activities in artificial intelligence. Your participation is invited in the fomm of
(I) a **paper** to be reviewed (nominally by three reviewers) and, if accepted, published and presented at the Symposium, or
(2) a half-aay **tutorial** on a suitable AI topic, to be presented at the Symposium. All tutorial proposals must have a stated goal, a detailed outline (2 - 5 pages), target audience defined, and a brief resume of the tutorial presenter(s), or
(3) an **exhibit** or display of artificial intelligence technology.

Important Dates:
May 15, 1991 — Four copies of an extended abstract (2 to 4 double-spaced pages) or a full paper (max. 11 pages, including figures) to one of the Program Co-Chairs (see address below).
July 1,1991 — Notification of acceptance will be mailed.
August 15,1991 — Final camera-ready copies of papers due.

Topics of interest include, but are not limited to

- Foundations of AI
- Pattern Recognition
- Knowledge Acquisition
- Computer Vision
- Natural Language Processing
- Knowledge-Based Systems
- Robotics
- Automated Planning
- Automated Scheduling
- Case-Based Reasoning
- Reasoning Under Uncertainty
- Connectionist Theory & Applications
- Machine Learning
- Automated Theorem Proving
- Knowledge Representation
- Temporal Reasoning
- Database/Knowledge Base Integration
- Real-Time Expert Systems
- Expert System Verification
- Novel Architectures for AI
- Truth Maintenance Systems

A pre-published proceedings will be available at the Symposium.
Submit **papers** or extended abstracts to: Program Co-Chairs, 4th UNB AI Symposium, Brad Nickerson or Lev Goldfarb, E-mail: bgn@unb.ca or goldfarb@unb.ca
Submit **tutorial** proposals to: Bruce Spencer or Przemyslaw Pochec, Tutorial Co-Chairs, 4th UNB AI Symposium, E-mail: bspencer@unb.ca or pochec@unb.ca
Submit **exhibit** requests to: Kirby Ward, Exhibits Chair, 4th UNB AI Symposium, E-mail: wardk@unb.ca
Local arrangements: Steven Rauch, Financialadministration: DanaWasson

The mailing address, phone and fax numbers for the above are as follows:
Faculty of Computer Science, University of New Brunswick, P.O. Box 4400, Fredericton, N.B., Canada E3B SA3
Phone: (506) 453-4566 Fax: (506) 453-3566

It is planned to hold the CKETM (Certified Knowledge Engineer) examination on Sept. 19,1991 at the Symposium location. Symposium participants wishing to take the examination should contact the IAKE at
International Associaton of Knowledge Engineers Georgetown P.O. Box 25461, Washington, D.C. 20007, U.S.A.
ph: (301) 231-7826 fax: (301) 770-4621 E-mail: IAKE@UC780.bitnet

Jan Mulder

From the academic editor

Many interesting ideas are born at the coffee table. This was also the case for the checker's program described in this issue's feature article. The program was developed at the University of Alberta and it is rapidly conquering the checkers world. The article is based on Connie Bryson's interview with Jonathan Schaeffer, one of the program's conceivers.

This issue also contains the promissed list of Canadian graduate students in AI. The list does not necessarily provide a complete picture of AI graduate students in Canada. Only those students who responded to my request in the previous issue are listed here. If I receive many more listing requests in the near future, then I will do a second listing later this year. I would also welcome suggestions for feature articles in the academic section.

My e-mail address is: mulder@arcsun.arc.ab.ca

GRADUATE STUDENT LIST

Last Update: April 5, 1991

UNIVERSITY OF ALBERTA

Department of Computer Science
Edmonton, A.B. T6G 2E1

Robert Baron
Supervisor: Randy Goebel
Degree: Ph.D. — Plan-guided perception, and perception-guided planning — Fall 1991

Scott Goodwin
Supervisor: Randy Goebel
Ph.D. — Second order direct inference — Spring 1991

Dekang Lin
Supervisor: Randy Goebel
Degree: Ph.D. — Theory and application of obvious abduction — Fall 1991

Abdul Sattar
Supervisor: Randy Goebel
Degree: Ph.D. — Improving the empirical complexity of hypothetical reasoning — Fall 1991

Lingyan Shu
Supervisor: Jonathan Schaeffer
Degree: Ph.D. — genetic algorithm learning — 1991

UNIVERSITY OF SASKATCHEWAN

Department of Computer Science
Saskatoon, Sask. S7N 0W0

Barb (Brecht) Wasson
Supervisor: Gordon McCalla, Marlene Jones
Degree: Ph.D. — Content Planning in Tutoring Systems — Sept. 1990

Xueming Huang
Supervisor: Gordon McCalla, Eric Neufeld
Degree: Ph.D. — Belief Revision with ATMS — Spring 1991

Dinesh Gadwal
Supervisor: Gordon McCalla, Jim Greer
Degree: M.Sc. — Knowledge-based Chess Tutoring — Fall 1990

Mary Mark
Supervisor: Jim Greer
Degree: M.Sc. — Training with Device Simulators — Spring 1991

Dan Baril
Supervisor: Gordon McCalla, Jim Greer
Degree: M.Sc. — Qualitative Reasoning for User Modelling — Spring 1991

Permanand Mohan
Supervisor: Jim Greer, Marlene Jones
Degree: M.Sc. — Planning Instructional Delivery — Spring 1991

Carl Gutwin
Supervisor: Gordon McCalla
Degree: M.Sc. — Tutorial Dialogue — Summer 1991

Zuihui Du
Supervisor: Gordon McCalla
Degree: M.Sc. — Case-based Instructional Planning — Summer 1991

Randy Coulman
Supervisor: Gordon McCalla, Jim Greer
Degree: M.Sc. — Case-based Diagnosis — Summer 1991

Maylor Leung
Supervisor: Herb Yang
Degree: Ph.D. — Human Body Motion Analysis Fall 1991

Omar Vega
Supervisor: Herb Yang
Degree: M.Sc. — 3D Shape Analysis Fall 1991

Russ Muzzolini
Supervisor: Herb Yang
Degree: M.Sc. — Biomedical Image Processing Fall 1991

Gang Chen
Supervisor: Herb Yang
Degree: M.Sc. — Computer Vision Fall 1991

QUEENS UNIVERSITY

Department of Computer and Information Science
Kingston, Ont. K7L 3N6

Kim Baxter
Supervisor: Janice Glasgow, S. Fortier
Degree: Ph.D. — Molecular Imagery 1995

Michel Feret
Supervisor: Janice Glasgow
Degree: Ph.D. — Explanation-based Learning 1994

Konstantinos Konstantoulis
Supervisor: Janice Glasgow
Degree: M.Sc. — Constraint Satisfaction 1991

UNIVERSITY OF TORONTO

Department of Computer Science
Toronto, Ont. M5S 1A1

Craig Boutilier e-mail: cebly@cs.toronto.edu
Supervisor: Ray Reiter
Degree: Ph.D. — The use of conditional and modal logics in the semantic characterization of default reasoning and belief revision, and the relationship of such approaches to probabilistic accounts Fall 1991

UNIVERSITY OF OTTAWA

Department of Computer Science
Ottawa Ont. K1N 6N5

Xu Xin email address: xuxin@csi.uottawa.ca
Supervisor: Robert Holte (holte@csi.uottawa.ca)
Degree: M.Sc. — interactive search/design Apr, 1991

Leixuan Yang email address: leixuan@csi.uottawa.ca (or 052594@UOTTAWA.BITNET)
Supervisor: Stan Szpakowicz (szpak@csi.uottawa.ca)
Degree: Ph.D. — Conceptual Network: A Knowledge Representation Formalism and Its Various Applications in AI Jan, 1992

Christine Laurendeau email address: chris@csi.uottawa.ca
Supervisor: Stan Szpakowicz (szpak@csi.uottawa.ca)
Degree: M.Sc. — Automatic Acquisition of Technical Concepts from Text Sept 1991

Sylvain Delisle email address: sylvain@csi.uottawa.ca
Supervisor: Stan Szpakowicz (szpak@csi.uottawa.ca)
Degree: Ph.D. — Knowledge Acquisition from Text end of 1992

UNIVERSITY OF CALGARY

Department of Computer Science
Calgary A.B. T2N 1N4

John Aldwinckle
Supervisor: M.L.G. Shaw
Degree: Ph.D. — Formal specification for computer systems June 1992

Jacky Baltes
Supervisor: B.A. MacDonald
Degree: M.Sc. — Learning maco operators and parallel planning June 1992

Suchitra Chander
Supervisor: I.H. Witten
Degree: M.Sc. — Machine learning and human computer interaction June 1992

Bruce Conrad
Supervisor: B.R. Gaines
Degree: Ph.D. — Knowledge support systems June 1992

Richard Esau
Supervisor: D.R. Hill
Degree: Ph.D. — Natural language understanding June 1991

Dan Freedman
Supervisor: B.R. Gaines
Degree: Ph.D. — Process migration June 1991

Rosanna Heise
Supervisor: B.A. MacDonald
Degree: Ph.D. — Task acquisition 1994

Daniel Jaliff
Supervisor: M.L.G. Shaw
Degree: M.Sc. — Construction of a machine learning system combining SBL and EBL June 1992

Greg James
Supervisor: I.H. Witten
Degree: M.Sc. — Genetic algorithms in game strategies June 1991

Brent Krawchuk
Supervisor: I.H. Witten
Degree: M.Sc. — Logical foundations of machine learning June 1991

Debbie Leishman
Supervisor: B.R. Gaines
Degree: Ph.D. — Learning and reasoning by analogy 1993

John Lewis
Supervisor: B.A. MacDonald
Degree: M.Sc. — Learning by cooperative observation June 1991

Mengchi Liu
Supervisor: B.R. Gaines
Degree: Ph.D. — Deductive databases 1993

Nick Malcolm
Supervisor: B.R. Gaines
Degree: M.Sc. — A collaborative editing system June 1991

David Maulsby
Supervisor: I.H. Witten
Degree: Ph.D. — Programming by demonstration 1992

Thong Phan
Supervisor: I.H. Witten
Degree: Ph.D. — Machine Learning functions in finite fields Nov. 1992

Maurice Sharp
Supervisor: B.R. Gaines
Degree: M.Sc. — The situated desktop: intentional extension to the desktop June 1991

Kwang Sim
Supervisor: I.H. Witten
Degree: M.Sc. — Machine learning June 1992

Tony Smith
Supervisor: I.H. Witten
Degree: M.Sc. — Grammatical inference using closed class lexemes June 1991

Mark Williams
Supervisor: B.A. MacDonald
Degree: M.Sc. — Robot mobility June 1991

UNIVERSITY OF BRITISH COLUMBIA

Department of Computer Science
Vancouver B.C. V6T 1W5

Ron Rensink
Supervisor: R.J. Woodham
Degree: Ph.D. — Rapid parallel recovery of three-dimensional orientation Dec 1991

Alex Kean email: kean@cs.ubc.ca
Supervisor: Prof. Alan Mackworth
Degree: Ph.D. — Clause Management Systems.
Keywords: Abduction, explanations, prime implicants/ implicates, approximation, reasoning April 1992.

Jeffrey Beis email: beis@cs.ubc.ca
Supervisor: David Lowe
Degree: Ph.D. — Learning to recognise 3-D objects spring, 1994

David LeBlanc email: leblanc@cs.ubc.ca
Supervisor: Richard Rosenberg
Degree: Ph.D. — language acquisition 1993

UNIVERSITY OF NEW BRUNSWICK

Faculty of Computer Science
Fredericton (N.B.) E3B 5A3

Karen Stephens
Supervisor: B. Nickerson
Degree: M.Sc. — Temporal reasoning in real time expert systems May 1991

Lisa Mullin
Supervisor: B. Nickerson
Degree: M.Sc. — Knowledge representation and data structures for automated name placement Dec. 1991

Shauna Gesner
Supervisor: B. Nickerson
Degree: M.Sc. — Knowledge base specification for design cost estimating

Joozar Vasi
Supervisor: B. Nickerson
Degree: M.Sc. — Semi-automated Knowledge acquisition from textual files Dec. 1991

McGILL UNIVERSITY

McGill, Research Centre for Intelligent Machines
Montreal, Quebec, H3A 2T5

Ferhan Bulca e-mail: bulca@lightning.mcrcim.mcgill.ca
Supervisor: J. Angeles, P. Zsombor-Murray, L. Vroomen
Ph.D. — Kinematical and dynamical analysis and control of a novel double-tetrahedra mechanism 1993

Benoit Dubuc e-mail: benoit@davinci.mcrcim.mcgill.ca
Supervisor: S.W. Zucker
Thesis topic: Integration of local information in early vision 1992

James Elder e-mail: elder@ra.mcrcim.mcgill.ca
Supervisor: S.W. Zucker
Thesis topic: The computation of closure (Computer Vision) Dec. 1991

Yasmine Ghallab : yasmine@eldar.mcrcim.mcgill.ca
Degree: Ph.D. — Collision Prediction Sept. 1992

Ahmed Helmy : ahmed@macondo.mcrcim.mcgill.ca
Supervisor: P.Z. Murray, L.J. Vroomen
Thesis topic: The control of a three degree freedom manipulator using a transputer network Aug. 1991

Michael Kelly e-mail: kelly@macondo.mcrcim.mcgill.ca
Supervisor: M. Levine
Thesis topic: View-centered approaches for 3D object recognition Fall 1993

Andrew Lam
Supervisor: A.S. Malowany
Thesis topic: Data/Knowledge based systems Fall 1992

Jiming Liu e-mail: jiming@davinci.mcrcim.mcgill.ca
Supervisor: L.K. Daneshmend
Degree: Ph.D. — Qualitative reasoning about Physical Systems: Theories and Applications in conceptual CAD and robot task planning 1992

Frederic Leymarie : leyfre@thunder.mcrcim.mcgill.ca
Supervisor: M.D. Levine
Thesis topic: Tracking and describing deformable shapes 1993

Ameen Maluf e-mail: amaluf@davinci.mcrcim.mcgill.ca
Supervisor: P.E. Caines
Degree: M.Eng. — Simulation of dynamic logic observer for finite automata Apr. 1991

Seyed Majid Noorhosseini
e-mail: majid@mcrcim.mcgill.ca
Supevisor: A.S. Malowany
Degree: Ph.D. – Task planning for robot workcells 1994

Mark Readman
Supervisor: P. Belanger
Thesis topic: Modeling and control of elastic robots Sept. 1991

Abbas Taher e-mail: taher@larry.mcrcim.mcgill.ca
Supervisor: A. Malowany
Degree: M.Eng. — Expert systems for telephone switch network maintenance June 1992

Kim Wheeler : wheeler@mithrandir.mcrcim.mcgill.ca
Supervisor: M.D. Levine
Degree: M.Eng. — Computer Vision / Knowledge-based system for tracking living cells during mitosis Oct. 1991

John Zelek e-mail: zelek@larry.mcrcim.mcgill.ca
Supervisor: M. Levine
Degree: Ph.D. (EE) — Attentional Mechanisms in object recognition 1993

McGILL UNIVERSITY

McGill, Department of Electrical Engineering
Montreal, Quebec, H3A 2T5

Shailendra Mathur : smathur@moe.mcrcim.mcgill.ca
Supervisor: F. Ferrie
Degree: M.Sc. — Multi-sensor fusion in Computer Vision June 1992

Michel Pelletier e-mail: michel@ireq-robot.hydro.qc.ca
Supervisor: L.K. Daneshmend
Degree: Ph.D. — Intelligent telerobotics Sept. 1992

Kenong Wu e-mail: wu@mcrcim.mcgill.ca
Supervisor: M.D. Levine
Degree: Ph.D. — Using computer vision and artificial intelligence to characterize and classify three-dimensional cells 1993

McGILL UNIVERSITY

McGill, School of Computer Science
Montreal, PQ, H3A 2A7

Giovanni Flammia
Supervisor: R. De Mori
Thesis Topic: Automatic Speech recognition July 1991

Roland Kuhn
Supervisor: R. De Mori
Degree: Ph.D. — Use of Dialogue Knowledge to Constrain Speech Recognition Hypotheses in an Interactive Speech Understanding System. Nov. 1991

Masao Takahashi e-mail: takahash@cs.mcgill.ca
Supervisor: R. De Mori
Degree: Ph.D. — Robot Dialogue May 92

UNIVERSITE DE MONTREAL

Departement Informatique et Recherche Operationelle
Montreal, PQ H3C 3J7

Sylvain Beland
Supervisor: B. Lefebre, E. Tropper
Degree: Ph.D. — Systeme d'assistance pour la conception et la gestion des systeme d'information
Computer aided system for designing and managing information systems 1992

Jean-Marie Desrosiers
Supervisor: J. Potvin
Degree: M.Sc. — Specifcation orientee-objet de reseaux de neurones
Object oriented specification of neural networks 1991

Massimo Fasciano
Supervisor: G. Lapalme
Degree: M.Sc. — Generation integree de texte et de graphiques
Integrated generation of text and graphics 1991

Mustapha Friha
Supervisor: G. Lapalme
Degree: M.Sc. — Detecteur rapide de fautes d'orthographe et de syntaxe francaise
Fast spelling and syntax checker 1991

Michel Gagnon
Supervisor: G.Lapalme, R. Kittredge
Degree: Ph.D. — Generation de texte tenant compte de la notion de temps
Text generation including temporal aspects 1992

Sylvain Giroux
Supervisor: G. Lapalme, A. Senteni
Degree: Ph.D. — Intelligence Artificielle distribuee pour la modelisation de comportements
Distributed artificial intelligence for behavior modelling 1992

Tanguy Kervahut
Supervisor: J. Potvin
Degree: M.Sc. — Systeme de support pour la conception d'algorithmes de generation d'arbres de decision
Computer aided system for the design of decision tree generation algorithms 1991

Leila Kosseim
Supervisor: G. Lapalme, R. Kittredge
Degree: M.Sc. — Generation de texte dans le domaine des recette de cuisine
Text generation in the area of kitchen recipes 1991

Fabrice Lavier
Supervisor: G. Lapalme
Degree: M.Sc. — Environnement d'experimentation de grammaires de Montague
A computer system for experimenting with Montague grammars 1991

Bertrand Legault
Supervisor: B. Lefebre
Degree: M.Sc. — Systeme de gestion de la couche semantique pour un SGBD
A system for managing the semantics of a DBMS 1991

Alain Loyer
Supervisor: G. Lapalme, J. Vaucher
Degree: M.Sc. — Environnement de travail pour la langage de progammation ObjVprolog
A programming environment for ObjVprolog 1991

Nancy Nadeau
Supervisor: R. Lebe-Neron
Degree Ph.D. — Environment interactif d'edition de partitions
Interactive environment for editing dance scores 1992

George Nault
Supervisor: M. Boyer
Degree: M.Sc. — Generation de structure dramatiques
Generating dramatic structures 1991

Gilles Roy
Supervisor: J. Vaucher
Degree: Ph.D. — Utilisation des constraintes pour la specification d'interfaces intelligentes
Constraints for specifying intelligent user interfaces 1992

Yu Shen
Supervisor: J. Potvin, G. Lapalme
Degree: Ph.D. — Methodes d'apprentissages appliquees au d'eveloppement d'algorithmes de tournees
Learning methods for developing routing algorithms 1992

Nicole Tourigny
Supervisor: G. Lapalme
Degree: Ph.D. — Generation de texte a partir de programmes de simulation
Text generation starting from simulation programs 1992

Marcel Turcotte
Supervisor: G. Lapalme, R. Cedergren
Degree: Ph.D. — L'etude de la dynamique de conformation de proteines et des acides nucleiques
Protein and nucleic acid conformation dynamics 1992

Controlling Super-Intelligent Machines

continued from page 5

points out that reward is better than punishment, as a tool for control [7]. Suppose a SIM has 100 options open to it. If we wish to make the SIM choose one particular option, it is more efficient to reward the choice of that one option than it is to punish the choice of the other 99 options. A SIM manipulated by the experience of pleasure, however, may feel resentment, like a drug addict manipulated by a dealer. When the addict is significantly more intelligent than the dealer, the dealer is in a precarious position.

Emotion

Part of our fear of SIM's is that we, being intellectually inferior, could become their slaves. The etymology of the word "free" is interesting in this context. The original sense of "free" was "dear, beloved"; hence it was applied to those of the household who were children, not slaves. This suggests that we could avoid slavery by programming SIM' s to love humans . Unfortunately, many unpleasant deeds have been done in the name of love.

There are other emotions we might use to control SIM' s . We might program SIM's to fear us, but fear does not reliably inspire obedience. SIM's might decide to eliminate the source of their fear — people. Emotional control has too many unpredictable side-effects.

Ethics

Hogan's The Two Faces of Tomorrow recommends that SIM's be programmed to have a sense of ethics [5]. According to Hogan, the essence of ethics is empathy; being able to imagine oneself in another's position. If SIM's are programmed to empathize with humans, then SIM's should not be dangerous to humans.

The core idea here is the Golden Rule: Do unto others as you would have them do unto you. This is a difficult rule to follow. Programming empathy into a SIM would be a challenging task.

Fusion

One approach to controlling a SIM would be to link it directly to a human brain. If the link is strong enough, there is no issue of control. The brain and the computer are one entity; therefore, it makes no sense to ask who is controlling

continued on page 35

Montreal, June 10-12, 1992

ITS-92

International Conference on Intelligent Tutoring Systems

PRELIMINARY ANNOUNCEMENT

Conference Chairman:
Claude Frasson
University of Montreal

Program Committee Chairman:
Gordon McCalla
University of Saskatchewan

Organization Committe Chairman:
Gilles Gauthier
Universite du Quebec

Publicity Committe Chairman:
Bernard Lefebvre
University of Montreal

An official call for papers with details for submission will appear soon. However, we invite interested people to communicate with us at the following address:

ITS-92
Bernard Lefebvre Dept I.R.O.
University of Montreal
CP 6128 Montreal H3C 3J7 Canada
or
frasson@IRO.umontreal.ca
fax : 514-343 5834

Aims of the conference:

The International Conference on Intelligent Tutoring Systems in 1992 will focus on a broad spectrum of research concerned with how artificial intelligence and other advanced technologies can be applied to education and training. The conference will be concerned both with the current state of the art as well as serving as a reference basis for future research directions. ITS-92 will be supported by a strong international program committee that will ensure full refereeing of all submitted papers.

As in ITS-88, a large number of well known speakers will be invited to discuss current and future research directions. There will also be provocative panels on "hot topics" in the field, and there will be demonstrations and exhibits of prototypes, research projects and products from universities and industries. The conference will be preceded (June 8-9) by two tutorial days.

Topics of interest for paper submissions will include, but will not be limited to, the following: architectures for ITS, student modelling, teaching and learning strategies, cognitive models, discovery environments, software tools for tutoring, knowledge representation in ITS, neural models applied to ITS, design issues, empirical studies, practical issues in ITS, real world applications, alternatives to one-on-one tutoring, evaluation and testing of ITS, helping and advising systems, collaborative learning, situated learning, negotiated tutoring

Jean-Claude Gavrel

PRECARN's Policies on Intellectual Property

PRECARN promotes, funds, and manages precompetitive research projects. One of the main products of this research program will be Intellectual Property (IP) which is governed by ownership and access rules developed by the PRECARN Membership. Since these intellectual property issues affect a large number of AI projects in Canada, we thought that a review of the key elements of the PRECARN approach might interest the readers of Canadian AI.

Two distinct but complementary research programs are managed by PRECARN: a set of industry-led precompetitive research projects (the PRECARN program), and a university-based National Network of Centres of Excellence (the Institute of Robotics and Intelligent Systems or IRIS). Different IP Policies govern each program, but they have many areas in common. The policies reflect the objectives and expectations of the three communities involved:

- **Industry.**
  In most cases, ownership of Intellectual Property is not as important for industry as access to the technology and the right to commercialize. Early access is a critical factor to companies as the first to market usually reaps the most benefits. Another important consideration for industries is exclusive rights to the IP, although many companies are now prepared to share the results of precompetitive research to gain access to technologies which they could not develop alone. Royalties are less important to companies, as most of the revenues come from sales of products or processes, not royalties.
- **University.**
  The foremost consideration for university researches is the right to publish and to disseminate scientific results. A second consideration is the ability to receive royalties from ownership of Intellectual Property. This becomes more and more important as the results of university research get transferred to industry and commercialized.
- **Government.**
  Government laboratories are usually interested to see their results commercialised by industry and to receive royalties. Government agencies which fund research projects usually insist that the results be exploited first in their area of jurisdiction (Country or Province). In some instances, agencies require a share of the royalties generated.

The PRECARN policies try to accommodate these, sometimes conflicting, requirements as follows:

- **Ownership of IP.**
  For PRECARN projects, wherever possible PRECARN owns any IP arising out of the research. However in cases where the project participants make a significant contribution, then the IP ownership can stay with the participant. In the case of a University participant, if the research is paid as a grant (i.e. no overhead), then the University retains the IP. If the research is paid as a subcontract, then PRECARN retains it. For IRIS work, ownership is vested with the participant who discovered it, consistent with the policy of the university.
- **Access.**
  For PRECARN projects all Members at the time the project started have non-exclusive, royalty-free fights to the IP ("late" Members have the same rights but must pay a percentage share of third party royalties based on the date they joined). Non-Members may negotiate rights, but must pay full royalties. In IRIS, all participants have rights to use each other's IP for research purposes. PRECARN Members have preferred access to the IRIS IP as follows: if the owner of any IP wishes to negotiate an exclusive license, then it must first be offered to the members of PRECARN; if the owner wished to offer non-exclusive licences, then the royalty rate offered to PRECARN Members must be 50% of that offered to third parties.
- **Royalties.**
  On PRECARN projects, if PRECARN does not own the IP, then third party royalties are shared between the owner and PRECARN in proportion to their investment in the project. Within the IRIS program, the disposition of royalties are governed by the rules of the University of the IP owner.
- **Publications.**
  In the PRECARN projects, research results must first be disclosed to PRECARN, and a publication delay of up to six months may be imposed to allow for protection of IP, and to provide a "head start" to the PRECARN Members. Special provisions exist to allow the publication of non-confidential results by University researches and for the publication of students theses. In IRIS projects, the publication rules are those of the University.

These policies offer a number of benefits. PRECARN industry Members have early access to the research results and royalty-free access to base technology which they could not develop alone. University researchers have a incentive to transfer results to industry for product development, while retaining the rights and the flexibility of publishing scientific papers. And the Government has the assurance that a number of companies have a vested interest in commercializing technology in Canada.

In closing here is a short update on the PRECARN activities:

- The APACS project (see last issue) is reaching its first review point, and a comprehensive briefing session for PRECARN Members is scheduled for June.
- A second project, IGI (Intelligent Graphic Interface, led by MPR Teltech) has been launched in February, following the approval of $1.7 million in funding from the government of British Columbia. IGI will be reviewed in the next issue of this column.
- Three more organizations joined PRECARN this year. They are: Communications Canada, Hewlett Packard (Canada) Ltd., and National Defence. The following is a current list of Members:

Alberta Research Council
Alcan International Ltd.
Atomic Energy of Canada Limited
B.C. Hydro
B.C. Advanced Systems Foundations
Bell Northern Research
Bristol Areospace Ltd.
CAE Electronics Ltd.
Canadian Institute of Advance Research
Canadian Marconi Company Ltd.
Centre de recherche informatique de Montreal
Communications Canada
Falconbridge Limited
H.A. Simons Ltd.
Hatch Associates Ltd.
Hewlett Packard (Canada) Ltd.
Husky Injection Moulding Systems
Hydro-Quebec
Inco Limited
MacDonald Dettwiler & Associates
Manalta Coal Ltd.
MPB Technologies Inc.
MPR Teltech Ltd.
National Defence
National Research Council
New Brunswick Power Commission
Noranda Technology Centre
Ontario Hydro
Petro-Canada Resources
Shell Canada
Spar Areospace Ltd.
Steltech Inc.
Syncrude Canada Limited
TransAlta Utilities Corporation
Xerox Research Centre of Canada

The intent of this column is to keep readers of Canadian Artificial Intelligence informed of the activities of PRECARN, an AI/Robotics consortium of 35 Canadian organizations. If you have any questions or need more information of PRECARN, please contact:

Mr. Jean-Claude Gavrel, Director of Research Programs
PRECARN Associates
300-30 Colonnade road, Nepean, Ontario K2E 7P4
Tel: (613) 727-9576 Fax: (613) 727-5672
E-mail:<GAVREL@ A1.ATOTT2.NRC.CA>

CIAR UPDATE
NOUVELLES DE CIAR

Some UBC Activities in Artificial Intelligence and Robotics

Robert J. Woodham

Quelques activités de l'UCB dans le domaine de l'Intelligence Artificielle et de Robotique.
Robert J Woodham

Cet article décrit quelques unes des activités de l'Université de Colombie Britannique dans le domaine de l'Intelligence Artificielle et de Robotique (IAR). Il est basé sur un groupe de recherche qui a été actif à l'intérieur du département d'informatique pendant plusieurs années. Dans les dernières années, la recherche s'est trouvée et a pris de l'expansion, grâce à la création de nouveaux liens avec d'autres groupes de recherche, antant à l'intérieur de l'UCB qu'à travers le pays. Vital pour ce développement fut le support de l'Institut Canadien pour la Recherche Avancée (IARA) et l'Institut de Robotique et Systèmes Intelligents (IRSI).

This article describes some of the activities at the University of British Columbia in Artificial Intelligence and Robotics. It is based on a research group that has been active in the Department of Comptuer Science for many years. In recent years, research has been significantly enhanced and extended throught the creation of new linkages to other research groups, both within UBC and nationally. Vital to this development has been the support of the Canadian Institute for Advanced Research (CIAR) and the Institute for Robotics and Intelligenct Systems (IRIS).

There are several other groups at UBC engaged in artificial intelligence research, including ones in Law and in Commerce and Business Administration. Similarly, there are other groups engaged in robotics research, including ones in Electrical Engineering and in Mechanical Engineering This article does not attempt to cover all relevant research at UBC. Rather, it highlights one research group in Computer Science and its linkages to other groups

UBC was one of the three initial "nodes" formed when the CIAR Artificial Intelligence and Robotics (AIR) program began in 1984. From the outset, the major goals of the AIR program have been:

1. to promote long-term basic research in AI and robotics,

2. to develop and to network critical mass groups of researchers, and

3. to build a base for applied research in AI and robotics.

Subgoals included the promotion of cross-disciplinary research and the strengthening of technology transfer.

The UBC node has achieved considerable success with respect to the first two major goals. UBC has excellence in basic research in AI, robotics and related disciplines and it has a critical mass of researchers that interacts effectively, both internally and with research groups located elsewhere. The research effort at UBC is indeed cross-disciplinary. UBC researchers who are CIAR AIR program members hold appointments, including joint appointments, in seven departments (Computer Science, Electrical Engineering, Forest Resources Management, Mechanical Engineering, Ophthalmology, Physiology and Psychology) representing five different faculties (Applied Science, Arts, Forestry, Medicine and Science).

In 1987, during the time of preparation for the first scheduled five-year external AIR program review, it was clear that more needed to be done, at UBC and elsewhere, to build a base for applied research in AI and robotics and to promote technology transfer. Since then, two major initiatives have been launched. The first is PRECARN, a Canadian consortium of over thirty companies that does precompetitive applied research in AI and robotics. (See the series of articles in *Canadian Artificial Intelligence* on PRECARN activities). A key PRECARN objective is to enhance the receptor capacity within Canadian industry for the results of basic research in AI and robotics. The second is IRIS, the Institute for Robotics and Intelligent Systems, one of the Networks of Centres of Excellence funded by the federal government. IRIS is managed by PRECARN and consists of 22 research projects involving researchers from many universities across Canada. IRIS represents significant new funding for university-based research in AI and robotics.

UBC researchers are involved in several IRIS projects. In addition to supporting targeted research, IRIS funding has provided essential resources to enhance the experimental component of our basic research programs. It is clear to us that progress in AI and robotics necessarily involves a cycle between theory and experiment. Canadian researchers have made significant contributions to theory. At the same time, the experimental side of university based research in Canada has sometimes been weak, in part, because experimentation can be expensive in terms of equipment, space and manpower resources.

Research at UBC spans a wide range of topics in AI and robotics. The research interests of individual researchers are summarized below. One area that is above critical mass in strength is perception and, in particular, vision. Five researchers (Cynader, Little, Lowe, Mackworth and Woodham) have vision as a principal research interest. Two others (Bibel and Poole) work primarily in logic-based reasoning. Mackworth's current research includes both vision and logic-based reasoning.

In computational vision, the focus of UBC research is intelligent perception for robotics and for remote sensing. Research includes: knowledge representations for vision; model-based systems for recognition and motion tracking; analytic methods for early vision, especially shape from shading and photometric stereo; and parallel algorithms and architectures for effective implementation and integration of vision modules including binocular stereo, shading, motion, texture, and colour. Applications are to robot vision systems that recognize and track known 3-D objects from unknown viewpoints, even when parts of the object are hidden from view, to natural resources management, using satellite and aerial imagery of the earth's surface, and to diagram understanding.

Our research strategy is to investigate the theory, design, implementation and evaluation of intelligent perceptual systems. This requires mathematical analysis of the problems to be solved, development of a theory of knowledge representation for those problems, design of algorithms and architectures to solve the problems and implementation and testing of those algorithms in realistic settings.

An important theme at UBC is the interface between biological and machine vision. Research combines neurophysiological studies of visual processing mechanisms in animals, psychophysical studies of human visual performance and computational studies to seek effective solutions to the problems of early vision. This three-pronged

approach has resulted in substantial cross-disciplinary synergy.

A second focus for UBC research is on logical frameworks for reasoning, including aspects of theory formation, diagnosis and recognition. This research area also forms the underpinning of work into natural language processing and discourse interpretation.

Faculty Researchers

This section provides summaries of the research interests of selective UBC faculty members. All UBC members of the CIAR AIR program are listed, including UBC IRIS project leaders recently appointed as CIAR-PRECARN Associates.

Wolfgang Bibel

Wolfgang Bibel spent one year, 1987-88, as a CIAR Fellow and Professor of Computer Science at UBC. Since returning to Germany, he has continued as a CIAR Associate Fellow and Adjunct Professor. Wolfgang's research is in automated reasoning, parallel inference machines, knowledge representation and program synthesis. He is widely known for the development of a new deductive mechanism for first-order logic, called the connection method.

Max S. Cynader

Max Cynader is a neuroscientist whose research encompasses both visual and auditory processing mechanisms. His vision research includes: quantitative studies of single cortical cells involved in motion perception, stereoscopic vision and orientation selectivity; computational and physiologic studies of orientation selection and segmentation; studies of cortical maps; and studies of the distributions and roles of cortical neurotransmitters involved in visual information processing.

Cortical maps found in biological vision systems suggest possible architectures for machine vision systems. In human and animal visual systems, maps have been delineated for colour, for direction of motion, for orientation of elongated contours, and for stereoscopic vision, among others. Max's research tries to combine emerging knowledge of cortical maps, their computational features and connectivity, to understand: how the locations of boundaries are inferred; how the locations and orientations of surfaces are established; what mechanisms process motion in 3-D space; how motion information is used to infer the shapes of surfaces and objects; and how diverse cues are combined in order to achieve robust estimates.

Max is the CIAR BC Fellow. In addition to Ophthalmology, Max also holds joint appointments in both Physiology and Psychology. He is Project Leader for IRIS Project A-3, "Sensory processing architectures".

Peter D. Lawrence

Peter Lawrence's main research interest is the application of real-time computers in the control interface between human operators and machines for robotics and telerobotics. Research projects include: coordinated control of excavator-type machines (excavator, logloader, grapple-yarder and feller-buncher); 3-D camera and display systems (for remote operation of machines); optical measurement systems (for measuring depth, volume of logs and machine parameters); real-time inverse kinematics (for manipulator endpoint control); and parallel computation systems (for real-time dynamics simulation).

Peter is Project Leader for IRIS Project C-6, "Teleoperation: measurement of machine parameters and human/machine performance" and is a CIAR-PRECARN Associate.

James J. Little

Jim Little's research area is computational vision. He is interested in the early vision computation of scene structure, through stereo and motion, and in the integration of vision modules to produce robust description of objects in the scene. Jim's interest also extend into the area of robotics, machine architecture, programming languages and numerical methods.

Recent advances in parallel computation have made possible near real-time vision systems. Jim has designed parallel vision algorithms suited to SIMD machines such as the Connection Machine. He is designing an intermediate computational level suited to MIMD architectures such as networks of transputers. He also is developing as active binocular imaging apparatus with control of pointing and camera parameters to explore the possibilities of active imaging

David G. Lowe

David Lowe's major area of research has been in computer vision and its applications to robotics. In particular, the work has concentrated on the problem of programming a computer to visually recognize known objects seen from any viewpoint in ordinary images. The research includes formulating psychological theories of human performance in recognition as well as the design of computer vision systems.

This work has resulted in the implementation of a system for recognizing known 3-D objects from arbitrary viewpoints in images taken by an ordinary TV camera. Since the appearance of a 3-D object can greatly change as it is viewed in different orientations, it is necessary to discover viewpoint invariant characteristics of the projection that remain stable over at least a significant range of viewpoints. This has led to the study of perceptual organization in human vision, a topic first studied by the Gestalt school of psychology. Some of these methods of perceptual organization have been implemented on the computer and serve as a basis for practical methods of recognition.

A second focus has been on mathematical techniques for determining the 3-D position and orientation of an object given the locations of a few of its features in an image. The exact solution for viewpoint allows an object to be reliably recognized even when many of its features are hidden from view. An important application of these computer vision

systems is for providing a robot with information regarding the location of objects in its surroundings. Current work includes using these techniques to track the motion of an object in real time, including objects with movable subparts.

David is a CIAR Scholar.

Vinod J. Modi

Vinod Modi's research interests include aerodynamics, ocean engineering, biomechanics, wind energy, satellite dynamics and control, robotics, and remote manipulator systems. Recently, he was named a fellow of the American Institute of Aeronautics and Astronautics, in recognition of his work in the field of flexible spacecraft dynamics and control of tethered satellite systems. He is only the third Canadian to be named a fellow of this 60 year old institute.

Vinod Modi is Project Leader for IRIS Project C-8, "Dynamical design and control of a large class of space and ground based manipulators" and is a CIAR-PRECARN Associate.

Alan K. Mackworth

Alan Mackworth's research focuses on the theory and applications of knowledge representation and computational vision. He has worked primarily on the use of knowledge in vision systems including surface orientation representations, network consistency approximation algorithms for constraint satisfaction problems, schema-based systems and logic-based systems. Applications include diagram understanding, telerobotics and the interpretation of satellite imagery and geographical sketch maps.

Recently, Alan has worked, with colleagues and students, on the following topics: descriptive and procedural adequacy for visual knowledge representations, 2-D shape representations, algorithms for constraint satisfaction problems, complexity analysis of these algorithms, their use in schema systems for recognition, visual monitoring of an excavator arm, a logical framework for evidential reasoning in depiction and diagnosis and various experimental implementations of these ideas.

Alan is the CIAR Shell Canada Fellow. He also is Project Leader for IRIS Project A-1, "Vision Systems for Recognition, Tracking and Navigation" and he serves on the CIAR AIR Program Development Committee and on IPACC, the recently formed IRIS PRECARN AIR Coordinating Committee.

David L. Poole

David Poole's main area of interest is in the application of logic and probability to common sense reasoning. He has been involved in the development of the "Theorist" paradigm for hypothetical reasoning. The main applications of this are to default reasoning, recognition, diagnosis and planning. The theory has been developed along with applications and implementations .

Other interests include theorem proving, logic programming, probabilistic reasoning and deductive databases.

David is a CIAR Scholar.

Richard S. Rosenberg

Richard Rosenberg's research focus is natural language understanding including question-answering systems, anaphora, case grammars, discourse, and portable natural language interfaces. Recently, he has become interested in the social implications of computer use. Specifically his concerns are with privacy, the nature of work, unemployment, the impact of AI and robotics, and applications in education, medicine, and the law.

Robert J. Woodham

Bob Woodham's primary research contributions have been in the area of shape from shading and photometric stereo. The view taken is that one needs to model the physics of image formation in order to extract information about the world from images. One test-bed for the research is the computer interpretation of remotely sensed data. Here, the task is to determine intrinsic surface properties related to ground cover, independent of topography and conditions of illumination. With the methods developed, it is possible to decouple direct solar irradiance, diffuse sky irradiance and atmospheric effects from ground cover and known topography. The practical goal of the research is to extend the range of terrain and imaging conditions that can be handled by automatic image analysis systems.

In related work, a new method based on photometric stereo has been developed for determining dense representations of surface shape, including surface curvature. This work has application to the robot vision tasks of object recognition, localization and inspection. The idea of using multiple light sources has also been exploited in a novel approach for determining the optical flow created when objects (or the robot perceiver) are in motion. This work has application to the tasks of object identification, object tracking and robot navigation in dynamic environments.

Bob is a CIAR Fellow. He also holds a joint appointment in the Faculty of Forestry.

Related UBC Laboratories and Centres

CIAR, AIR, PRECARN and IRIS all are acronyms that are national in scope. UBC researchers also have affiliations and research collaborations involving internal laboratories, centres and people. The following is a selected summary of UBC laboratories, centres and their acronyms.

LCI

The Laboratory for Computational Intelligence (LCI) in the Department of Computer Science supports research in computational reasoning and perception. LCI provides the intellectual environment, shared resources, and technical support for a wide range of research projects directed by the participating faculty. The LCI is unique in Canada in its

support for research and graduate education in computational vision, artificial intelligence, remote sensing and robotics. Currently, there are six LCI faculty members (Little, Lowe, Mackworth, Poole, Rosenberg and Woodham), each of whom is mentioned above. Two new faculty members, Craig Boutilier and Dinesh K. Pai, will be joining the LCI during the summer of 1991. Boutilier's research interests include default reasoning, knowledge representation, belief revision, inheritance reasoning and diagnosis. Pai's research interests includge reasoning about mechanical systems for robotics, control and CAD/DAM/CAE. Applications include computer graphics animation and biped walking machines.

In 1990 the laboratory moved into newly renovated space that provides faculty and staff offices, work areas for graduate students, and adjacent experimental facilities for research in calibrated imaging, robotics and real-time vision. The Laboratory for Computational Vision (LCV) moved from its previous location and is now a part of the LCI.

The LCI provides essential facilities for the experimental component of our research programs. The computational facility is based primarily on Sun workstations and servers. Attached peripherals include a set of Datacube modules for real-time image processing and a collection of transputers for high-speed parallel processing. The transputer configuration currently consists of eight T800-25 transputers, each with 2MB 60ns DRAM. Three other special purposes transputers provide a framegrabber, a Datacube to transputer interface and a general I/O capability.

Additional specialized equipment includes: a Symbolics 3650 Lisp machine with both fast image convolution hardware, developed at MIT, and the Symbolics high resolution colour graphics option; an Optronics C-4500 colour film scanner/writer; and analog video equipment with single frame A/D and D/A capability.

As part of its IRIS targeted research, the LCI is developing a calibrated imaging facility where visual experiments, including motion tracking and photometric stereo, can be conducted. The calibrated imaging facility consists of an optical bench, cameras, light sources and devices for the measurement and control of 3-D position and motion. The LCI has also acquired a CRS 6-degree-of-freedom robot arm to serve as the platform for an "eye-head" system for experiments in active vision.

Substantial software support has been developed for the laboratory, including a widely distributed package of image processing modules. Licensed software includes Lucid Common Lisp, Quintus Prolog, Saber-C, and The Publisher from ArborText.

In addition to the eight faculty members named, LCI has four full-time staff, one postdoctoral fellow, and about twenty five graduate students working in the laboratory. During the past ten years, some fifty theses have been completed by students working in the LCI and its predecessor, the LCV.

Mackworth is the Director and Woodham is the Associate director of the LCI.

CICSR

The Centre for Integrated Computer Systems Research (CICSR) is an "umbrella" research organization encompassing a significant amount of the computer based research at UBC. CICSR's mandate is to encourage collaborative research among its members in the departments of Computer Science, Electrical Engineering, and Mechanical Engineering. Currently, this includes the graduate research of over 100 graduate students, each of whom is registered in one of the three constituent departments, but whose research may involve several CICSR faculty. CICSR also promotes collaboration with external companies, government agencies, similar organizations at other universities (including Simon Fraser University and the University of Victoria), the BC Advanced Systems Institute, and, as is most relevant to this article, the CIAR and IRIS.

Early in 1989, it was announced that, as part of the UBC Development Campaign, the BC Provincial Government will provide $15M for the construction of a CICSR/Computer Science Building to house CICSR activities and expanded facilities for the Departments of Computer Science, Mechanical Engineering, and Electrical Engineering. The new CICSR/CS building is scheduled for completion in the summer of 1992.

James M. Varah is the Director of CICSR.

MAGIC

MAGIC is an acronym for Media and Graphics Interdisciplinary Centre. It is an interdisciplinary research centre for media and graphics, recently established at UBC. It is administered from the Faculty of Graduate Studies, with its director reporting directly to the Dean of Graduate Studies. The research conducted is intended to span the entire spectrum of computer based media applications. Samples of these media are: animation, interactive video disk applications, computer music, and computer speech recognition.

One of MAGIC's main goals is to strengthen interaction with industry through collaborative research projects. These interactions will include collaborations with computer manufacturers to develop the technology, and collaborations with industrial users to transfer the technology.

Kellogg S. Booth is the Director of MAGIC.

Imager

Imager is a new graphics research laboratory in the Department of Computer Science. It was established in September 1989 and has been the center for most of the initial programming and development for MAGIC. Its purpose is to conduct fundamental research in all aspects of graphical representation.

LCI researchers Little, Lowe and Woodham collaborate with Imager researchers David Forsey and Alain Fournier on an NSERC Strategic Grant entitled, "Merging real and computer generated images".

WORLD WATCH
VUE SUR LE MONDE

Russ Thomas

This month we continue our review of AI research around the world with a short article outlining the work that is being carried on at the Artificial Intelligence Applications Institute of the University of Edinburgh in the United Kingdom. As you will see from the article there is in this institute a strong emphases on the applied aspects of AI.

In an attempt to extend the communication between members of the AI community around the world CSCSI is about to develop a directory of email (and the corresponding snail-mail) addresses of its members. We are currently negotiating access to similar email directories held by other national AI societies around the world and we would hope to be able to make these available to our own members. To implement the process of developing our own directory we are inviting members to send us some email with their full name, physical address and, email address. In return, at the end of the summer we will send them the current society email directory and any other similar directories that we have available. We fully realize that some members will not want their email names included and of course this is easily achieved by not sending us any email! Those who would wish to have access to the directory and wish to be included then please send the appropriate information to:

can-ai@arcsun.arc.ab.ca

The above email address can also be used to send any email to members of the executive of CSCSI.

Artificial Intelligence Applications Institute

Edinburgh University

Overview

AIAI is a non-profit technology transfer organisation which is a part of the University of Edinburgh, working alongside companies to help them understand Artificial Intelligence methods, recognise the benefits they offer, and apply them to their specific business needs.

This is achieved by:

- promoting awareness through newsletters, seminars, briefing sessions, and other information services;
- providing training through short courses and project-based Study Programmes;
- undertaking consultancy, advising on appropriate technology and project feasibility
- and carrying out joint project development work with clients.

AIAI is part of the Edinburgh AI community, which is the largest community of AI specialists in Europe.

Application Areas and Types of Task Worked On

AIAI concentrates on the application of AI methods in the following areas:

Aerospace; Computing and Electronics;
Engineering; Environment; Financial Services;
Manufacturing; Petroleum

Types of Task include: Assessment; Interpretation;
Diagnosis; Planning; Scheduling

Training and Awareness

AIAI provides a wide range of short courses and secondment based training.

- **Short Courses**

Awareness

Knowledge Based Systems for Managers - Techniques & Methods - KBS Skills - Knowledge Engineering - Knowledge Elicitation - Planning and Scheduling - Reasoning with Constraints

AI Languages & Systems

AI Languages Sampler - Common Lisp - Introduction to Prolog - Advanced Prolog - Parallel Logic Programming with Parlog

- **Study Programmes**

Training paths through short courses ending with project implementation.

Project examples:
Aero-engine fault diagnosis - Rolls-Royce
AI-based job shop scheduling - ICL
Medical insurance underwriting - Standard Life

AIAI Facilities

- **Computing Facilities**

Knowledge Representation Systems Trials Laboratory (KRSTL) - access to AI workstations and toolkits for trials, evaluations and project work.

Training Laboratory - providing single-user workstation based training facilities with extensive software tool availability.

Staff computing is mainly workstation based.

- **AI Library**

AIAI and the Department of AI at Edinburgh jointly run one of the largest AI libraries in the UK. with over 15,000

research reports and over 3,000 books.

• **Location**

AIAI is located in central Edinburgh alongside the world renowned Department of AI at Edinburgh University.

AI in Financial Services

Areas of interest include Accounting (evaluation of companies, advising on audits, dealing with complex legislation), Banking (Personal financial planning, mortgage advising underwriting, identifying fraud, support services [maintaining networks, routing telexes, generating appropriate documentation]), Insurance (Underwriting, claims processing) and, Stocks & shares (Portfolio management, foreign exchange, trading strategies)

Previous work at AIAI in the area of Financial Services include: Medical underwriting, Cash flow modelling, Mortgage loan acceptance, AI and Relational Databases, Qualitative Modelling of the Securities Industry, Diagnosing fault in DP networks, an options trading system and, an AI in Financial Services Survey which includes over 400 annotated references

In this area AIAI has worked with: Abbey National; Baillie Gifford; Bacon & Woodrow; The Royal Bank of Scotland; Union Bank of Switzerland; Hitachi Ltd; The International Stock Exchange; Hewlett Packard; Standard Life; ...

AI in Manufacturing

Areas of Interest in this sector include: Planning & Scheduling (the key to reduced operations management costs, increased customer response and improved quality of processes); Robotics and Assembly (to increase the flexibility and capability of industrial robots); Design (integration of design knowledge into all processes through to production); Preventative Maintenance (to link diagnostic capability to operations management for overall quality of production).

Previous work at AIAI in the area of Manufacturing include: Planning for remotely executing agents (integrating project planning with execution monitoring and control); Scheduling in a factory environment (JIT in the job shop); Planning for Assembly (Integration and Test: for complex assemblies).

In this area AIAI has worked with: ICL; Hitachi; US Air Force; European Space Agency; Ford Motor Company; British Telecom and; GEC Marconi; ...

AI in Petroleum

Areas of Interest here include the application of AI methods in areas such as: Safety (Monitoring equipment, Diagnosing failures, Analysing risk, Scheduling maintenance);

Exploration (Identifying and Classifying reservoirs); Reservoir Engineering (Defining reservoir extent and estimating recoverable reserves); Production (Planning recovery programs); Distribution (route Planning and Scheduling).

Existing work in this area includes: Directional drilling advisor; Geological knowledge representation; Oil platform construction planning; Pressure transient test data interpretation; Simulation of the movement of equipment during drilling operations

Here AIAI has worked with: AEA Technology, Safety and Reliability Directorate; BP; Edinburgh Petroleum Services; ELF Aquitaine; Petroleum Science and Technology Institute; Robertson Group; Shell; Sysdrill.

AI in Aerospace

Finally areas of Interest here include: Planning & Scheduling; Systems Modelling and; Diagnosis

The work here has been on Command and Control (for planning and controlling a remotely executing agent); Flight Operations Planning (for aircraft carrier based missions); Scheduling of Earth Resources Spacecraft; Planning for Assembly, Integration and Test (for spacecraft); Modelling and Diagnosis (of onboard systems such as power and electrical systems); Feasibility Studies (on specific applications and on the use of specific tools).

Again in these areas AIAI has worked with: Admiralty Research Establishment; British Airways; CRI; European Space Agency (ESTEC and ESOC); Matra Espace; Rutherford Appleton Laboratory (SERC); US Air Force; ...

For further details on any of the above areas please contact Austin Tate at:
Artificial Intelligence Applications Institute
University of Edinburgh,
80 South Bridge,
Edinburgh EH1 1HN,
United Kingdom
Tel: 011-44-31-650-2732 Fax: 011-44-31-226-2730
Email: bat@aiai.edinburgh.ac.uk or aiai@ed.ac.uk

1.0 Theoretical Aspects

1948 Fuzzy sets and typicality theory.
S.V. Murthy, A. Kandel (Dept. of Comput. Sci., Florida State Univ., Tallahassee, FL, USA).
Inf. Sci. (USA), vol.51, no.1, p.61-93 (June 1990).

A set whose elements take grades of membership in the interval [0,1] is known as a fuzzy set. In practice, assignment of grades of membership to the elements of a fuzzy set is based on statistical considerations. The paper is concerned with the problem of identifying what should be meant by a 'typical' element of a fuzzy set. Because of the subjective nature of membership grades, use of the common statistical measures such as the mean of median is theoretically unjustifiable. The authors employ the concepts of the fuzzy expected value and fuzzy random variable as the basis for a measure of what is typical of a fuzzy set, an approach which shows promise in fields like pattern recognition and expert systems. They propose a flexible algorithm which incorporates subjective ideas of typicality and some simple statistical techniques to arrive at the

A N N O U N C E M E N T

AAAI FALL SYMPOSIUM SERIES

Sponsored by the American Association for Artificial Intelligence

Introduction

The AAAI 1991 Fall Symposium Series will be held November 15 through 17, 1991, at the Asilomar Conference Center, Pacific Grove, Monterey, California.

Asilomar is situated on the tip of the Monterey Peninsula overlooking the Pacific Ocean. Its rustic setting, secluded areas of forest and dunes, and clustered meeting space and accommodations make it ideal for smaller groups like the Fall Symposium Series. Asilomar operates under the American Plan, which means that participants are expected to room and board at Asilomar.

The topics of the four symposia in the 1991 Fall Symposium Series are:

- Discourse Structure in Natural Language Understanding and Generation:
- Knowledge and Action at Social and Organizational Levels;
- Principles of Hybrid Reasoning;
- Sensory Aspects of Robotic Intelligence.

Each symposia will be limited to approximately 60 participants. Working notes will be prepared and distributed to participants in each symposium. A general plenary session will be scheduled in which the highlights of each symposium will be presented. Informal receptions will be scheduled in the early evenings.

Each participant will be expected to attend a single symposium throughout the symposium series. In addition to participants selected by the program committees of the symposia, a limited number of other interested parties will be allowed to register in each symposium. Registration information will be available by early September 1991. To obtain registration information please contact the AAAI office.

Discourse Structure in Natural Language Understanding and Generation

Computational modeling of discourse structure is a fundamental component of theoretical and applications-oriented work in Natural Language Processing. A representation of the underlying structure of a discourse enhances the ability of a natural language system to interpret and generate a wide variety of linguistic phenomena.

Modeling discourse is important to many subareas within Natural Language Processing, including text generation, plan synthesis and recognition, multi-media systems, user modeling, machine translation, psychological models of language, and spoken language systems. This symposium is designed to bring together researchers in different areas of discourse, in order to identify common issues, goals and techniques, and to transfer associated theoretical and practical results. Some of these issues are:

- What should a discourse model model (e.g., rhetorical relations, speaker intentions, statistical probabilities)?
- How should discourse models be represented (e.g., using logics of belief, connectionist frameworks, discourse grammars)?
- How are complex discourse structures modeled (for collaborative or multi-agent discourse)?
- How is the structure of a particular discourse marked (from indicators such as anaphoric reference, cue phrases, tense and aspect, intonational features, turn-taking signals, inferred speaker intention)?
- How are models of discourse evaluated (e.g., what criteria are appropriate for generation or understanding)?
- What advantages do various models of discourse have over one another?
- What role do empirical studies play in discourse modeling and evaluation?
- How are theoretical models of discourse implemented in working systems?
- How can work in related disciplines such as psychology, sociolinguistics, psycholinguistics, and philosophy contribute?

Program Committee: Julia Hirschberg (co-chair), Diane Litman (co-chair), Kathy McCoy, Candy Sidner.

Knowledge and Action at Social and Organizational Levels

Research on representing and reasoning about knowledge and action at analytical levels more aggregated than that of the individual — topics such as social and organizational foundations of knowledge, multiple-perspective reasoning, consensus, multiple rationality, coordination and collaboration, commitment, stable organization, representing groups and group activity, etc. — is becoming increasingly important. Such research influences many fields, including cognitive science, distributed systems, design, human-computer interaction, natural language processing, computer-supported cooperative work, basic AI and distributed AI, social studies of science, and the engineering of AI systems. This symposium will address a number of relevant foundational scientific issues, such as:

CONCEPTUAL MODELING: How can agents, knowledge, commitment, etc. be modeled at the social, the individual, and at integrated levels? Can a group be a locus of knowing, rather than an individual? How can groups model their own activities and influence other groups, even in so-called "open systems?"

INTEGRATING MULTIPLE PERSPECTIVES: How do actors reason among multiple representations? How can actors deal with conflict in knowledge and action? How can actors dynamically construct useful meanings for the objects and actions that are part of their joint activity?

November 15, 16, & 17, 1991

Asilomar Conference Center — Pacific Grove, Monterey, California

445 Burgess Drive, Menlo Park, CA 94025 (415) 328-3123 fss@aaai.org

INTEGRATING SITUATED AND GENERAL KNOWLEDGE: How can agents and groups address the basic tension between local, "situated" knowledge and action, and the non-local conception of general knowledge and action-at-a-distance? How can an actor influence another actor, e.g. by sending a message, without global semantics or control?

INTEGRATING SOCIAL AND INDIVIDUAL LEVELS OF ANALYSIS: What would a "middle ground" theory, that integrates both individual agency and social structure, be like? (How) can we make middle-ground theories computational?

COMMON GROUNDS FOR INTERACTION: What are the foundations for mutually compatible languages, assumptions, and other common bases for interaction? How is global or shared knowledge possible? How can agents "stay out of each others' way?"

METHOD: What are appropriate research methods for addressing these questions in principled ways? What are the roles and limitation of data from biological, human, and social studies? What are the roles and limitations of current formal models and theories?

Program Committee:

Phil Agre (agre@cogs.sussex.ac.uk), Danny Bobrow (bobrow @parc.xerox.com), Les Gasser (gasser@usc.edu, chair), Jim Hendler (hendler@cs.umd.edu), Eduard Hovy (hovy@isi.edu), Ed Hutchins (ehutchins@ucsd.edu), Leigh Star (soaO3@gec.keele.ac.uk).

Principles of Hybrid Reasoning

One of the most promising approaches to designing more-efficient automated reasoning systems is that of integrating multiple reasoners to form a hybrid reasoner. By exploiting the efficiency of special-purpose reasoners operating on specialized representations, hybrid reasoners can outperform homogeneous systems. The common challenge faced by all hybrid-system designers is the successful integration of the system's components.

The focus of this symposium will be on architectures for integrating the components of hybrid reasoners, and, in particular, on the principles that underlie such architectures. Though it seems unlikely that all hybrid reasoning systems can be explained by a common set of principles, there are general architectural classes that can each be explained by common principles. The major goals of the symposium are to explore principles that have been applied to hybrid reasoners, identify general classes to which they apply, and seek unifying themes among the current diversity of systems and theoretical models.

Systems of central interest are deductive systems that incorporate special-purpose reasoning modules operating on their own representation. Examples include:

- Unification-based systems that incorporate special-purpose reasoners in their unification algorithms (for example, systems that use sorted unification or E-unification);
- Knowledge representation systems that mix assertional information with terminological information;
- Deductive systems, including planners, that incorporate constraint processing mechanisms (especially constraint logic programming);
- Systems such as theory resolution that build domain theories into reasoning algorithms;
- Parsers for feature-based grammars that use special mechanisms to handle constraints among features or to perform inheritance; and
- Deductive systems for modal logic that exploit special techniques for reasoning about the accessibility of possible worlds.

Further examples and discussion can be found in the report on the 1988 Workshop on Principles of Hybrid Reasoning, which appears in Volume 11, Number 5 (special issue) of AI Magazine (January 1991).

Program Committee: Alan Frisch (chair), Gert Smolka, Lynn Andrea Stein, Richmond Thomason.

Sensory Aspects of Robotic Intelligence

It probably is not an exercise in hyperbole to say that sensory faculties will be the primary determinants of the degree of autonomy that will be achieved by the machines of the future. Some of us therefore believe that the aspects that deal with sensing should be at the core of the intellectual enterprise that is artificial intelligence, assuming, of course, that an important aim of this enterprise is to design better and better machines in the service of humankind.

This symposium will provide a forum for reviewing the state of the art in the sensory aspects of robotic intelligence and, through such review, an attempt will be made to delineate the potentially productive directions for future research. Many questions at varying levels of generality will be debated. Questions at the most general level will include: Given the level of competence we can pack into the sensory faculties of a robot today, what is the most sophisticated demonstration of robotic intelligence that could be made in the near future? More specific questions will deal with the how's and why's of integrating sensing with planning; integrating different modalities of sensing, such as vision, force/torque, touch, etc.; representation of actions; representing uncertainties and dealing with them through sensing and actions; dealing with the complexities introduced by geometry and topology; etc. Hopefully, by trying to answer such question we will see more clearly the conceptual frontiers of our knowledge in this area.

Program Committee: Bruce Donald, Avi Kak (Chair), Matt Mason, Ed Riseman, Saburo Tsuji.

typical element of a fuzzy set. No attempt is made to further fuzzify the element of the given fuzzy set with respect to the typical element; however, it is an issue worth considering at a later stage. The investigation assumes a well-defined fuzzy set; i.e. the grades of membership are known. Subjective feelings of cohesiveness within a set (based on proximity and density, for example are incorporated within the algorithm. (14 refs.)

2208 Connectionism and artificial intelligence as cognitive models.
D. Memmi (LIMSI-CNRS, Orsay, France).
AI Soc. (UK), vol.4, no.2, p.115-36 (April-June 1990).

The current renewal of connectionist techniques using networks of neuron-like units has started to have an influence on cognitive modelling. However, compared with classical artificial intelligence methods, the position of connectionism is still not clear. In this article artificial intelligence and connectionism are systematically compared as cognitive models so as to bring out the advantages and shortcomings of each. The problem of structured representations appears to be particularly important, suggesting likely research directions. (30 refs.)

2455 Limitations of multi-layer perceptron networks - steps towards genetic neural networks.
H. Muhlenbein (Gesellschaft fur Math. und Datenverarbeitung mbH, St. Augustin, West Germany).
Parallel Comput. (Netherlands), vol.14, no.3, p.249-60 (Aug. 1990).

The author investigates multi-layer perceptron networks in the task domain of Boolean functions. He demystifies the multi-layer perceptron network by showing that it just divides the input space into regions constrained by hyperplanes. He uses this information to construct minimal training sets. Despite using minimal training sets, the learning time of multi-layer perceptron networks with backpropagation scales exponentially for complex Boolean functions. But modular neural networks which consist of independently trained subnetworks scale very well. The author conjectures that the next generation of neural networks will be genetic neural networks which evolve their structure. He confirms Minsky and Papert (1988): 'The future of neural networks is tied not to the search for some single, universal scheme to solve all problems at once, but to the evolution of a many-faceted technology of network design.'. (12 refs.)

2477 A biological perspective on autonomous agent design.
R.D. Beer, H.J. Chiel, L.S. Sterling (Case Western Reserve Univ., Cleveland, OH, USA).
Robot. Auton. Syst. (Netherlands), vol.6, no.1-2, p.169-86 (June 1990).

The inability of current 'classical' AI systems to handle unconstrained interactions with the real world has recently lead to a search for new control architectures for autonomous agents. The authors argue that simpler natural animals already exhibit most of the properties required by an autonomous agent, and suggest that designers of autonomous agents should draw directly upon the neural basis of behavior in these animals. The relevant behavioral and neurobiological literature is briefly reviewed. An artificial nervous system for controlling the behavior of a simulated insect is then developed. The design of this artificial insect is based in part upon specific behaviors and neural circuits from several natural animals. The insect exhibits a number of characteristics which are remarkably reminiscent of natural animal behavior. (42 refs.)

2203 Hypermedia and randomized algorithms for medical expert systems.
R.M. Chavez, G.F. Cooper (Sch. of Med., Stanford Univ., CA, USA).
Comput. Methods Programs Biomed. (Netherlands), vol.32, no.1, p.5-16 (May 1990).

KNET is an environment for constructing probabilistic, knowledge-intensive systems within the axiomatic framework of decision theory. The KNET architecture defines a complete separation between the hypermedia user interface on the one hand, and the representation and management of expert opinion on the other. KNET offers a choice of algorithms of probabilistic inference. The authors have used KNET to build consultation systems for lymph-node pathology, bone-marrow transplantation therapy, clinical epidemiology, and alarm management in the intensive-care unit. Most important, KNET contains a randomized approximation scheme (RAS) for the difficult and almost certainly intractable problem of Bayesian inference. Their algorithm can, in many circumstances, perform efficient approximate inference in large and richly interconnected models of medical diagnosis. The authors describe the architecture of KNET, construct a randomized algorithm for probabilistic inference, and analyze the algorithm's performance. Finally, they characterize their algorithms' empiric behavior and explore its potential for parallel speedups. (18 refs.)

1953 Towards a general theory of information - information and entropy.
T. Stonier (Sch. of Sci. & Society, Bradford Univ., UK).
Future Comput. Syst. (UK), vol.2, no.4, p.409-27

The effective application of fifth-generation and post-fifth-generation computers will be limited by the lack of an adequate theoretical basis for the processing of information. The paper develops a general theory of information and intelligence. It begins with the premise that information is a basic property of the universe and that any system which exhibits organization, contains information. This is as true for physical systems as for systems involving human abstractions. Information has as much physical reality as do matter and energy. Changes in the organization of matter can be measured by changes in entropy. Entropy and information are inversely related. This conceptualization is diametrically opposed to the assumption of Shannon and the communications engineers. Other implications are considered. (17 refs.)

2471 Elephants don't play chess [artificial intelligence].
R.A. Brooks (MIT Artificial Intelligence Lab., Cambridge, MA, USA).
Robot. Auton. Syst. (Netherlands), vol.6, no.1-2, p.3-15 (June 1990).

There is an alternative route to artificial intelligence that diverges from the directions pursued under that banner for the last thirty some years. The traditional approach has emphasized the abstract manipulation of symbols, whose grounding in physical reality has

rarely been achieved. The author explores a research methodology which emphasizes ongoing physical interaction with the environment as the primary source of constraint on the design of intelligent systems. He shows how this methodology has recently had significant successes on a par with the most successful classical efforts. He also outlines plausible future work along these lines which can lead to vastly more ambitious systems. (32 refs.)

1951 Uncertainty management in expert systems.
K.-C. Ng, B. Abramson (Dept. of Comput. Sci., Univ. of Southern California, Los Angeles, CA, USA).
IEEE Expert (USA), vol.5, no.2, p.29-28 (April 1990).

Basic expert system terminology is reviewed, and several uncertainty management paradigms are surveyed. The focus is on subjective probability theory, Dempster-Shafer theory, and possibility theory, although a number of other methods are mentioned. The benefits and limitations of the various schemes are examined, examples of expert systems within each school are presented, and some relevant open problems are discussed. (75 refs).

2211 Effect of initial weights on back-propagation and its variations.
H. Lari-Najafi (Dept. of Electr. Eng., Minnesota Univ., Minneapolis, MN, USA), M. Nasiruddin, T. Samad.
1989 IEEE International Conference on Systems, Man and Cybernetics. Conference Proceeding. (Cat. No.89CH2809-2), Cambridge, MA, USA, 14-17 Nov. 1989 (New York, NY, USA: IEEE 1989), p.218-19 vol.1

The effects are studied on the converge properties of the back-propagation learning rule of the range from which the initial weight values are randomly selected. In addition to the standard back-propagation rule, two variations are also considered, namely symmetric back-propagation and expected-value back-propagation. In most applications of back-propagation, the range of initial weights is small. It is shown that significantly higher initial weights can substantially improve learning rates. If the initial weight range is increased beyond a problem-dependent limit, however, performance degrades. Symmetric back-propagation is most sensitive to the initial weight range, while expected value back-propagation is least sensitive. The authors describe an improvement on the symmetric variation that produces faster learning rates with low initial weights. (4 refs.)

2.0 Systems and Techniques

2536 THINK: a C library for artificial intelligence tasks.
M.E. Grost (Electron. Data Syst. Corp., Artificial Intelligence Services, Troy, MI, USA).
CAD/CAM Robotics and Factories of the Future. 3rd International Conference (CARS and FOF '88) Proceedings, Southfield, MI, USA, 14-17 Aug. 1988 (Berlin, West Germany: Springer-Verlag 1989), p.169-73 vol.2

THINK is a library of C functions created to perform artificial intelligence tasks. Because of its open architecture, it can be integrated with other C applications, such as CAD/CAM software. THINK is framebased and each slot of the frame can be filled with multiple values of mixed data types - numeric, character strings, C functions and pointers to other frames. THINK contains functions for adding elements to slots, replacing existing slot values, deleting values and accessing values of a slot, either one at a time or as a set. There are also functions for finding the maximum values of a slot, and for sorting the values of a slot according to a given sort function. Facilities for object-oriented programming are included in THINK. Functions can be values of the slots, and the binding of functions to frames is based on this. A message-sending system is built on top of the function binding capability. THINK also includes facilities for rule-based programming. THINK is small, portable and easily integrated with existing CAD/CAM systems. The entire THINK library compiles to less that 80 K of object code. Using THINK allows the CAD/CAM programmer to take advantage of the productivity gains in object-oriented and rule-based programming, and to encapsulate design and manufacturing knowledge about objects along with the part geometry of the objects themselves. (8 refs.)

2014 CAPS: a connectionist architecture for production systems.
A.S. Bhogal, R.E. Seviora, M.I. Elmasry (Dept. Of Electr. Eng., Waterloo Univ., Ont., Canada).
Proceedings of the Twenty-Third Annual Hawaii International Conference on System Sciences, Kailua-Kona, HI, USA, 2-5 Jan. 1990 (Los Alamitos, CA, USA: IEEE Comput. Soc. Press 1990), p.202-11 vol.1

CAPS supports most features of the OPS5 language, including variables, negation, and conjunction and disjunction of conditions. The architecture uses local representations to facilitate dynamic variable bindings and to reduce the number of interconnections within the network, thus making hardware implementations more feasible. The CAPS processing elements are simple and relatively easy to fabricate. A hardware implementation of a CAPS network can potentially provide a 200-800 fold increase in parallelism over serial implementations. The CAPS architecture is tested by transforming small OPS5 programs into connectionist networks and simulating them on a connectionist simulator. It is demonstrated that connectionist architectures can perform rule-based symbolic reasoning and can support dynamic variable bindings. (20 refs.)

2532 EXPERTEXT: from semantic nets to logic Petri nets.
R. Rada, P.E.S. Dunne, J. Barlow (Dept. of Comput. Sci., Liverpool Univ., UK).
Expert Syst. Appl. (UK), vol.1, no.1, p.51-62 (1990).

Hypertext systems logically solve problems. Expertext is an approach to combining the precision of expert reasoning processes with the browsing capabilities afforded by hypertext. In this paper intelligent hypertext is first modeled with semantic net. The semantic net formalism is then extended to a Petri net formalism. Finally, a deductive inferencing ability is added to the Petri net formalism. Examples of how an Electronic Yellow Pages might exploit these methods are presented. (27 refs.)

2010 An algorithm for identification of relations among rules.
H. Marathe, T.-K. Ma, C.-C. Liu (Dept. of Electr. Eng., Washington Univ., Seattle, WA, USA).
IEEE International Workshop on Tools for Artificial Intelligence. Architectures, Languages and Algorithms, Fairfax, VA, USA, 23-25 Oct. 1989 (Los Alamitos, CA, USA: IEEE Comput. Soc. Press 1989), p.360-7

A representation of a rule-based system is proposed. It is used to define relations among rules. The relations defined are cause-effect, mutual exclusion, redundancy, conflict, subsumption, and implication. A relation between two rules is either complete, i.e. the relation holds only for some instantiations of the rules. An algorithm for detecting relations between a new rule (to be added to the rule base) and rules in the rule base is developed. The algorithm is applicable to any forward chaining rule-based system. Application of this algorithm during rule base modification would identify possible problems resulting from the modification. The algorithm is implemented on a MicroVAX computer using C language. Examples of its application to a rule-based system are shown. (10 refs.)

2043 Design object modeling.
F. Arbab (Dept. of Comput. Sci., Univ. of Southern California, Los Angeles, CA, USA).
Intelligent CAD, I. Proceedings of IFIP TC/WG 5.2 Workshop on Intelligent CAD, Boston, Ma, USA, 6-8 Oct. 1987 (Amsterdam, Netherlands: North-Holland 1989), p.3-12

Discusses the tools and techniques for modeling of concepts, or for the lack of a better term, of 'objects'. The author is interested in the representation and manipulation of the knowledge about concepts that need to be manipulated in an intelligent CAD system. Geometric reasoning, feature based reasoning, and qualitative reasoning on models are important activities during a design process. Representation of geometric knowledge beyond existing geometric modes, modeling of features, and representation of functionality are, therefore, among the topics relevant to discussion. Each of these topics is an interesting area of research in itself and is also the concern of people in the AI community working on robotics, computer vision, learning, natural language understanding, etc. (6 refs.)

2231 Overcoming deficiencies of the rule-based medical expert system.
C.A. Hughes (Coll. of Med., Illinois Univ., Chicago, IL, USA), E.E. Gose, D.L. Roseman.
Comput. Methods Programs Biomed. (Netherlands), vol.32, no.2, p.63-71 (May 1990).

One of the current deficiencies of the rule-based expert system is its static nature. As these systems are applied to medicine, this short-coming becomes accentuated by: the rapid speed at which new knowledge is generated, the regional differences associated with the expression of many diseases, and the rate at which patient demographics and disease incidence change over time. This research presents a solution to the static nature of the rule-based expert system by proposing a hybrid system. This system consists of an expert system and a statistical analysis system linked to a patient database. The additional feature of a rule base manager which initiates automatic database analysis to refresh the statistical correlation of each rule ensures a dynamic, current, statistically accurate rule base. The philosophical differences between data and knowledge are also addressed as they apply to this type of hybrid system. The system is then used to generate four rule bases from different knowledge sources. These rule bases are then compared. (25 refs.)

2002 Coop: a shell for cooperating expert systems.
S. Shekhar, C.V. Ramamoorthy (Div. of Comput. Sci., California Univ., Berkeley, CA, USA).
IEEE International Workshop on Tools for Artificial Intelligence. Architectures, Languages and Algorithms, Fairfax, VA, USA, 23-25 Oct. 1989 (Los Alamitos, CA, USA: IEEE Comput. Soc. Press 1989), p.2-11

Conventional expert system shells do not help in developing AI programs for large applications which require multidisciplinary knowledge and which are geographically distributed. To support these applications, a shell must provide tools for a knowledge based system to (1) reason about the need for cooperation, (2) understand global knowledge in order to locate relevant expert systems, and (3) select appropriate cooperation plans. Coop, which supports cooperation models for characterizing the three essential decisions in the cooperation process, is described. It provides a computational method for deciding whether an expert system has enough knowledge to solve a given problem or whether it needs to consult with other expert systems. A yellow pages technique is provided to represent global knowledge and to select appropriate cooperation plans. The Coop environment lets expert systems autonomously resolve the three fundamental decisions in cooperation at runtime. The Coop environment also provides tools for resolving distributed computing issues of initiating and controlling process groups on a network, monitoring the state of distributed computation, and support tools needed to implement a large AI program consisting of multiple knowledge bases and expert system processes. (24 refs.)

2554 Analysis of the distinction between deep and shallow expert systems.
P.D. Karp, D.C. Wilkins.
Report UIUCDCS-R-89-1536, Univ. Illinois at Urbana-Champaign, IL, USA (Aug. 1989), 40 pp.

The first generation of expert systems (e.g. MYCIN, DENDRAL, R1) is often characterized as only using shallow methods of representation and inference, such as the use of production rules to encode empirical knowledge. First-generation expert systems are often dismissed on the ground that shallow methods have inherent and fatal shortcomings which prevent them from achieving problem-solving behaviors that expert systems should possess. Examples of such desirable behaviors include graceful performance degradation, the handling of novel problems, and the ability of the expert system to detect its problem-solving limits. The paper analyzes the relationship between the techniques used to build expert systems and the behaviors they exhibit to show that there is not sufficient evidence to link the behavioral short-comings of first-generation expert systems to the shallow methods of representation and inference they employ. There is only evidence that the shortcomings are a consequence of a general lack of knowledge. Moreover, the

paper shows that the first-generation of expert systems employ both shallow methods and most of the so-called deep methods. Lastly, the authors show that deeper methods augment but do not replace shallow reasoning methods; most expert systems should possess both. (45 refs.)

2243 Generic expert systems for management applications: the Operations Advisor and the Management Advisor.
W. Reitman (Sch. of Manage., Rensselaer Polytech. Inst., Troy, NY, USA).
Comput. Sci., Econ. Manage. (Netherlands), vol.3, no.2, p.167-75 (1990).

The Operations Advisor and the Management Advisor are AI-based executive decision support systems. The Operations Advisor (OA) enables managers to improve manufacturing strategy. It makes possible manufacturing policies that take full advantage of strategic interactions among capacity utilization, inventories, lead time, and unit costs. The Management Advisor (MA) values strategic business proposals (building a new plant, developing a new product, etc.). It helps managers assess the attractiveness of new business opportunities, and the impacts of proposed changes in existing business policy. Artificial intelligence methodologies in the MA and the OA allow the author to combine conventional analytic techniques with expert knowledge and experience. They also provide an extremely high level of user facilitation, control, and freedom from error, thus substantially improving overall business productivity. The author beings with a description of the design strategy underlying these two generic AI-based management applications. Then he considers how this strategy is realized in the OA and the MA, and the productively benefits it provides to users. (4 refs.)

1978 Expert systems: a question of liability?
K. Mykytyn, P.P. Mykytyn, Jr., C.W. Slinkman (Inf. Syst. & Manage. Sci., Texas Univ., Arlington, TX, USA).
Manage. Inf. Syst. Q. (USA), vol.14, no.1, p.27-44 (March 1990).

The development of expert systems has changed dramatically in recent years based largely on concepts dealing with artificial intelligence. These efforts are evolving from very specific, academically oriented efforts, such as medical diagnosis, to more managerially oriented corporate issues. Unfortunately, many proponents of these systems may be overlooking possible legal ramifications related to both the development and use of these systems. A major issue concerns the establishment of liability for the decisions and recommendations made by expert systems. Some liabilities could include product liability and negligence. All individuals involved with expert systems are discussed; machine intelligence, machine knowledge, intelligent machines and mechanical systems. (2 refs.)

2492 LOGPRIMER: a tutoring system for Prolog learning.
E. Fischetti, A. Gisolfi (Dipartimento di Inf., Salerno Univ., Italy). Comput. Educ. (UK), no.65, p.14-17 (June 1990)

'Logic programming' nowadays is avowedly recognized as one of the most powerful software technologies, and the language Prolog is, by far, the most popular in this area. Yet learning Prolog at a working level is very often a heavy burden and several students are discouraged by the particular structure of the language, its technical phraseology, its inference engine, etc. LOGPRIMER (logic programming primer) is a learning environment running on an IBM PC and devoted to teaching interactively the basic notions of Prolog. The mastery of the language is gradually achieved by intermixing theory and exercises and finally LOGPRIMER presents an expert system shell whose knowledge base can be built by the student. (9 refs.)

2244 Comax: an expert system for cotton crop management.
H. Lemmon (USDA, Agric. Res. Service, Albany, CA, USA). Comput. Sci. Econ. Manage. (Netherlands), vol.3, no.2, p.177-85 (1990).

An expert system has been developed that acts as an expert in managing a cotton crop. The system operates by using as its source of knowledge a sophisticated computer program which simulates the cotton plant growth, plus a set of 'if-then' type rules, plus a computer program called an inference engine. Comax determines the best strategy for irrigating, for applying fertilizer and for applying defoliants and cotton boll openers. Weather stations in the cotton fields automatically report local temperatures, radiation and precipitation to the system. Comax reevaluates its recommendations daily. Comax has been tested on several farms with excellent results. This system is currently in use by over 400 cotton growers.

2238 Expert system security.
D.E. O'Leary (Univ. of Southern California, Los Angeles, CA, USA). IEEE Expert (USA), vol.5, no.3, p.59-63, 66-69 (June 1990).

Computer risk exposures and security in general are reviewed, and factors suggesting that expert system security is a unique problem are examined. Security requirements associated with the unique characteristics of expert systems are investigated. They include technical aspects of knowledge (certainty factors, symbolic information and special fixes), structure (user interface, knowledge base, inference engine and database), design methodology (prototyping) and the current delivery environment (e.g., PCs and expert system shells). The security of working expert systems as discussed in the literature is examined. Rationales for the current apparent lack of expert system security are analyzed and the impact of possible security controls on expert system users and developers is assessed. (41 refs.)

2232 Using hindsight in medical decision making.
R.A. Russ (Lab. for Comput. Sci., MIT, Cambridge, MA, USA). Comput. Methods Programs Biomed. (Netherlands), vol.32, no.1, p.81-90 (May 1990).

As the clinical picture of a patient evolves over time, more information becomes available. Certain procedures require time to perform, causing delay between the time when the tests are ordered and when the results are available. Furthermore, as the patient's condition changes over time, serial measurements can be made. The availability of more data allows a more accurate assessment of the patient. Uncertainties, guesses or errors that were made early in the clinical course of patient care can also be identified and resolved when more information is available. Reasoning with a stream of data that changes over time presents a challenge to the designers of expert systems. The use of hindsight in expert systems requires

that appropriate attention be paid to the temporal relations of the data and that care is exercised in revising decisions. The author presents a data-dependency system, the Temporal Control Structure (TCS), designed to support reasoning with data changing over time and shows how it can be used to implement reasoning by hindsight. (22 refs.)

1991 Life cycles in software and knowledge engineering: a comparative review.
M. Wilson (Inf. Dept., Rutherford Appleton Lab., Chilton, Didcot, UK), D. Duce, D. Simpson.
Knowl. Eng. Rev. (UK), vol.4, no.3, p.189-204 (Sept. 1989). [received: 01 May 1990]

Progress in software engineering has led to system development following models of the system life cycle. These models incorporate the use of prototyping and formal methods of program verification. They are becoming supported by integrated project support environments and permit the planning and monitoring of software development projects. In contrast, knowledge based systems (KBS) are developed using informal views of the system life cycle. Tools have been developed to support some stages of the life cycle in an undisciplined manner. The commercial use of KBS needs development projects to be planned and monitored. This requires methods and tools based on systematic life cycle models to be established for KBS. This paper reviews the current state of life cycle approaches to software engineering and KBS development projects in order to provide a direction for the development of methodical KBS life cycle models. (40 refs.)

2528 Expert systems books.
E.E. Brown (Oregon State Univ., Corvallis, OR, USA).
Expert Syst. Rev. (USA), vol.2, no.3, p.53-71 (1990).

A number of excellent books have been published from 1985 to the present (1990) that are relevant to the development and use of expert systems and the associated research topics. This list includes ninety books selected for their relevancy to business from the hundreds of expert systems books published. The books listed are appropriate for a wide range of audiences from texts for beginners just learning about expert systems to references for experienced researchers in the field. Many of the books listed include examples of expert systems developed for business audience, the list includes some high quality general references. Also listed are books covering special topics of interest, such as knowledge acquisition or uncertainty. (90 refs.)

3.0 Applications

2297 A computerized phonetics instructor: Babel.
J. Vila (Illinois State Univ., Normal, IL, USA), L. Pearson.
CALICO J. (USA), vol.7, no.3, p.3-29 (March 1990).

Babel is an expert system able to animate (graphically) and reproduce (acoustically) a text in any language which uses the Latin alphabet. This system has been developed to aid language learners and to help instructors teach the fine nuances of phonemes. Each phoneme has a unique sound and thus requires a precise positioning of the vocal organs which are displayed on the screen in two different projections: a front view and a profile cross view of a human face in synchronization with the output sounds of the speech synthesizer. (22 refs.)

2353 An expert system-based approach to capacity allocation in a multiservice application environment.
S. Erfani, M. Malek, H. Sachar (AT&T Bell Labs., Middletown, NJ, USA).
GLOBECOM '89. IEEE Global Telecommunications Conference and Exhibition. Communications Technology for the 1990s and Beyond (Cat. No.89CH2682-3), Dallas, TX, USA, 27-30 Nov. 1989 (New York, NY, USA: IEEE 1989), p.870-5 vol.2

The authors discuss the application of expert systems and intelligent control to the transmission link capacity allocations problem. They present a structure for an expert system capacity allocator (ESCA).
The structure allows for dynamic allocation of transmission link capacities to support traffic loads at desired performance levels. The ESCA implementation is based on modular capacity expansions and a priori limits on the number of connections for each service type. This approach can perform near-optimal capacity allocation in real-time by using heuristic knowledge to virtually eliminate the required computations in this process. The authors discuss the advantages of the ESCA implementation in providing the ability to make rational decisions in an environment of uncertainty and imprecision, such as nonsystematic overloads and network failures. (15 refs.)

2376 An expert system for tunnel design.
Y. Ichikawa, O. Aydan, T. Kyoya (Nagoya Univ., Japan), H. Osaka, T. Kawamoto.
Microcomput. Civ. Eng. (USA), vol.5, no.1, p.3-18 (March 1990).

Describes the development of an expert system (ES) for tunnel design and presents some applications of the system. The system has been developed by the authors' research group consisting of 21 people closely involved with the tunnel design under the leadership of Nagoya University. The system consists of four subexpert systems, ES for standard tunnel design methods, ES for framed structure method, ES for theoretical design methods, and ES for numerical analysis design method; a common part to control the overall system; and two database systems (tunnel database system and rock mass database system). (10 refs.)

2599 Intelligent scheduling systems for parallel machines with different capability.
G. Leininger (Intelligent Syst. Center, Missouri Univ., Rolla, MO, USA).
CAD/CAM Robotics and Factories of the Future. 3rd International Conference (CARS and FOF '88) Proceedings, Southfield, MI, USA, 14-17 Aug. 1988 (Berlin, West Germany: Springer-Verlag 1989), p.287-90 vol.2

A rule-based forward chaining expert system was developed to schedule a multi-pass manufacturing operation in the production of glass lined vessels and related parts. The expert system uses data provided by an MRP-II system to schedule production at each of the work locations. Implementation of the system at a factory in

Program Chairs:

Janice Glasgow
Department of Computing and Information Science
Queen's University
Kingston, Ontario, K7L 3N6
e-mail: janice@queis.queensu.ca
phone: (613) 545-6058
fax: (613) 545-6513

Bob Hadley
School of Computing Science
Simon Fraser University
Burnaby, B.C., V5A 1S6
e-mail: hadley@cs.sfu.ca
phone: (604) 291-4488
fax: (604) 291-3045

Exhibitions Chair:

Jan Mulder
Alberta Research Council

Local Arrangements:

David Poole
University of British Columbia

Publicity Chair

Fred Popowich
Simon Fraser University

Program Committee:

CALL FOR PAPERS

Canadian Artificial Intelligence Conference

AI'92

University of British Columbia
Vancouver, British Columbia, CANADA
11-15 May 1992

AI '92 is the ninth biennial conference on Artificial Intelligence sponsored by the Canadian Society for the Computational Studies of Intelligence. It will be held in conjunction with Vision Interfaces and Graphic Interfaces, general chairs: Kellogg Booth (University of British Columbia) and Alain Fornier (University of British Columbia).

Contributions are requested describing original, unpublished research results, either theoretical or applied, in all areas of Artificial Intelligence research. The following areas are especially of interest:

Applications of AI	Rule-Base Reasoning
Cognitive Modelling	Social Implications of AI
Learning	Temporal Reasoning
Planning	Common Sense Reasoning
Robotics	Knowledge Acquisition
Search	Machine Architectures and Languages
Automatic Programming	Natural Language Understanding
Default Reasoning	Perception (Vision, Touch, Speech)
Knowledge Representation	Theorem Proving

Submit four (4) complete copies of papers to program co-chair Janice Glasgow. Papers must be received by January 7, 1992. Notification of acceptance or rejection will be mailed to the first author (or designated author) by February 24, 1992. Camera ready copy of accepted papers will be due March 24, 1992.

All submissions will be reviewed by the program committee. Acceptance will be based on the overall merit and significance of the reported research, as well as the quality of the presentation. Theoretical papers will be judged on their originality and contributions to the field of AI, whereas applied papers will be judged by the importance and originality of the application. As a condition of acceptance, the author, or one of the co-authors, will be required to present the paper at the conference.

Papers must not exceed 5000 words in length, including abstract bibliography and allowances for figures, tables, and diagrams. Papers must have a separate title page, containing the title of the paper, the name sand addresses of all authors and a short abstract. Final papers will be allocated eight (8) pages in the conference proceedings and should be formatted using 12pt LaTeX or Scribe "article" style (or equivalent). Papers which are also being submitted to other conferences, whether verbatim or in essence, must have this fact clearly indicated on the title page.

The journal *Artificial Intelligence* has offered a prize and publication for the "best paper" of the conference. Selection of the best paper will be done by the program committee.

Scotland resulted in a 40 percent increase in productivity and a reduction in the manufacturing cycle time from 12 weeks to 11 weeks. (3 refs.)

2613 Development of a computer program to teach patent searching.
S.B. Ardis (McKinney Eng. Libr., Texas Univ., Austin, TX, USA). World Pat. Inf. (USA), vol.12, no.2, p.76-80 (1990).

The increase in the use of patents by both engineers and the general public at the McKinney Engineering Library, University of Texas, threw a great strain on the library staff in meeting demands for help in accessing the patent files. This prompted development of an extremely user-friendly computer program which would lead inexperienced users towards the answers to their problems. Work began in September 1986 and the project took one man-year to complete. Results have been very encouraging and considerable savings in library staff time have been achieved. The background to the problem is described and the solution (involving an expert system) is outlined. Development, telecommunications and costs are also discussed.

2106 DTREE: an expert system for psychiatric diagnosis.
M.B. First, J.B.W. Williams, R.L. Spitzer (Dept. of Psychiatry, Columbia Univ., New York, NY, USA).
Proceedings: The Thirteenth Annual Symposium on Computer Applications in Medical Care (Cat. No.89TH0286-5), Washington, DC, USA, 5-8 Nov. 1989 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.954-5

The introduction in 1980 of the American Psychiatric Association's Diagnostic and Statistical Manual of Mental Disorders, 3d ed. (DSM-III) and, in 1987, of its successor (DSM-III-R) provided clinicians with a well-defined classification system for describing their patients, facilitating communication, research, and clinical care. The DSM-III-R classification contains well-defined necessary and sufficient criteria for each of more that 220 disorders. To assist clinicians in making DSM-III-R diagnoses and understanding the reasons behind giving them, the authors have developed DTREE, a microcomputer-based expert system. In addition to serving as an expert system to assist clinicians in making diagnoses in complex cases (or to confirm diagnoses in other situations), DTREE includes several features that make it particularly useful as an educational tool to teach new clinicians the process of psychiatric diagnosis. (8 refs.)

2082 Overview and examples of Digital's AI applications in manufacturing.
K. Nitschke, M. Tseng (Digital Equipment Corp., Maynard, MA, USA).
Proceedings of the 1st International Conference on Artificial Intelligence and Expert Systems in Manufacturing. The Scope, Applications and Limitations of Intelligent Manufacturing Systems, London, UK, 20-21 March 1990 (Kempston, UK: IFS Conferences 1990), p.45-56

Digital Equipment Corporation is a global manufacturing enterprise with manufacturing facilities located across the world. Coordinating a global manufacturing enterprise and integrating it with engineering and other functions requires management planning and decision support that is different from one-country of one-plant operation. Intelligent solutions, therefore, have to be flexible enough to be adaptable to different country and management needs. This, in turn, require a global and cross-functional development approach and an understanding of the business processes and their driving factors in the different environments. Digital's Applied Intelligent Systems Group (AISG) is in charge of this kind of solutions development and the technology transfer to user groups. This paper describes two manufacturing applications developed by AISG in support of integration and better management of the global manufacturing enterprise. First, the Expert Financial Closing (XFC) System is described, an application assisting in the manufacturing financial management process. Second, an Intelligent Business Decision Support System model and its applications area is introduced. (6 refs.)

2618 Inductive generation of auxiliary-lines in geometry.
M. Yamazaki, Y. Iida (Dept. of Inf. Process., Seikei Univ., Tokyo, Japan).
Technol. Rep. Seikei Univ. (Japan), no.49, p.39-56 (Jan. 1990). In Japanese. [received: 09 Jul 1990]

In order to solve some kinds of plane geometry problems, it is necessary to introduce auxiliary-lines. Human experts can quickly introduce these lines using knowledge obtained by their experiences, and can solve the given problems easily. But these problems are very difficult for inexperienced students. The authors describe a program system which can solve plane geometry problems of junior high school level. The program, which is written in the PROLOG language, can generate auxiliary-lines automatically if these lines are necessary to solve the appropriate problems. It does not generate auxiliary-lines at random, but generates only effective lines to solve the problems. Knowledge which is necessary to generate only effective auxiliary-lines is described by the form of PROLOG rules in the program. The experimental test results show that a computer, VAX 8200, requires shorter time to solve the problems than junior high school students. (11 refs.)

2604 DRUG-CHEK - a clinical laboratory drug interaction expert system.
A.M. Turano (Med-Chek Lab. Inc., Pittsburgh, PA, USA).
MUG Q. (USA), vol.20, no.2, p.5-9 (Aug. 1990).

DRUG-CHEK represents a compilation of how a particular medication affects diagnostic clinical laboratory tests. DRUG-CHEK represents an assembly of knowledge for laboratory tests detailing normal ranges dependent on species, sex, and age and correlates this with medication profiles. Abnormal lab values are passed to an inference program that deduces any chemical interference from drugs that are being administered. Expert knowledge is used to analyze the possibility of an interaction between drugs and lab tests. With this program, the computer can keep a watchful eye for potentially harmful drug-lab interactions manifested as elevated or diminished report values. Errors of a critical nature can occur when a clinician fails to consider a possible diagnosis, orders incorrect tests, or receives faulty lab results due to this interference. (7 refs.)

2597 An operations analysis expert system for fiberglass manufacturing.
G. Biswas (Dept. of Comput. Sci., Vanderbilt Univ., Nashville, TN, USA),M.D. Oliff.
CAD/CAM Robotics and Factories of the Future. 3rd International Conference (CARS and FOF '88) Proceedings, Southfield, MI, USA, 14-17 Aug. 1988 (Berlin, West Germany: Springer-Verlag 1989), p.240-4 vol.2

Notable strides have been made in the last five years in the application of knowledge based system techniques to problems in business decisionmaking and manufacturing. The paper discusses OASES, an operations analysis expert system for diagnostic problem solving in a fiberglass manufacturing environment. OASES functions as an intelligent assistant and aids management in analyzing problems in a fiberglass manufacturing process. The system uses a partitioned rule base for domain knowledge representation and a combined forward and backward inferencing mechanism to conduct a mixed initiative dialogue with users. (8 refs.)

2415 Expert-aided process control laboratory.
M. Rao (Dept. of Chem. Eng., Alberta Univ., Edmonton, Alta, Canada), T.-S. Jiang. Int. J. Appl. Eng. Educ. (UK), vol.5, no.2, p.227-31 (1990).

Describes an expert-aided process control laboratory that can be used to teach undergraduate and graduate process control (including intelligent control) courses and to design expert-aided process control systems for the real-time applications. The laboratory projects consist of several numeric simulation packages and expert systems, which cover a wide spectrum of process control knowledge. The laboratory provides a new approach in the teaching of chemical process control, which strengthens engineering curricula, helps students to achieve practical experience and reduce their work pressure, and stimulates the faculty to provide the best and the most advanced techniques in education. The intelligent educational system can help students gain a clear insight into the abstract theoretical subjects. It is an increasingly important feature of modern education. The microcomputers make such a low cost and effective teaching facility a real possibility. (14 refs.)

2152 Design in a distributed blackboard framework.
C. Kitzmiller, V. Jagannathan (Boeing Adv. Technol. Center for Comput. Sci., Seattle, WA, USA).
Intelligent CAD. I. Proceedings of IFIP TC/WG 5.2 Workshop on Intelligent CAD, Boston, MA, USA, 6-8 Oct. 1987 (Amsterdam, Netherlands: North-Holland 1989), p.223-33

Design problems within the aerospace industry are often of such size and complexity that no single individual, organization, or design environment is capable of effectively addressing all aspects of the design. In such situations, the design problems must be addressed by a team of specialists or experts each knowledgeable in some aspect of the design. Although many computer-aided design and engineering environments have been developed, few are capable of supporting the type of distributed, multidisciplinary, team-oriented approach such problems require. The next generation of design environments must address the issues raised by such design problems and provide a platform in which design can be pursued as a cooperative, distributed activity. The authors examine several issues involved in developing a cooperative, distributed design environment by describing an approach to automating the design of a fluid-driven actuator. They contrast the approach with that employed by several other systems based on the structure of design knowledge, the mechanisms available to control the design process, the representation of design constraints, and the mechanisms employed to handle constraint violations. (12 refs.)

2295 Bidder's Associate: a case-based reasoning system to improve cost estimates of manufactured products.
L.A. Whitaker, S.L. Wiggins (Klein Associates Inc., Yellow Springs, OH, USA).
1989 IEEE International Conference on Systems, Man and Cybernetics. Conference Proceeding. (Cat. No.89CH2809-2), Cambridge, MA, USA, 14-17 Nov. 1989 (New York, NY, USA: IEEE 1989), p.521-2 vol.2

Klein Associates is building a software system that will aid the bidder in writing a successful bid for a new solicitation. The software is called Bidder's Associate and is a prototype for a Dayton manufacturing firm which has a database containing approximately 1000 previously built products. Approximately 200 of these products have been made more than once and are likely to be solicited again. Although the data for previously built products are stored on a mainframe computer, the access by the bidder's current system is cumbersome and inefficient. Bidder's Associate aids the bidder in three ways: (1) accessing an appropriate product from the firm's data base; (2) adjusting bid cost estimates on the basis of the retrieved product's cost data; (3) producing a final bid for the current solicitation. (3 refs.)

2157 An expert planning system for complex surface manufacture.
G. Smith (Teeside Polytech., Middlesbrough, UK).
Proceedings of the 1st International Conference on Artificial Intelligence and Expert Systems in Manufacturing. The Scope, Application and Limitations of Intelligent Manufacturing Systems, London, UK, 20-21 March 1990 (Kempston, UK: IFS Conferences 1990), p.201-10

The application of expert systems technology to the problems of process planning for the production of complex surfaces on 3-axis CNC milling machines is presented. A knowledge based process planner to aid in the manufacture of these surfaces is described, and the links between this and commercial packages for surface design and numerical control preparation are discussed. (9 refs.)

2283 Forecasting techniques advisory system.
K. Kengskool (Dept. of Ind. Eng., Florida Int. Univ., Miami, FL, USA), M. Gross, R.M. Marinez.
Int. J. Appl. Eng. Educ. (UK), vol.6, no.2, p.171-5 (1990).

Forecasting is a very useful planning tool for most organizations. A number of forecasting techniques have been developed over the years. However, most people have failed to achieve the maximum benefits of utilizing those available techniques. This is mostly because they are unaware of the best model to use based upon the given situation. The paper presents an effective approach to fill the gap between users' demands in forecasting and those available

forecasting techniques. The approach is to integrate a knowledge-based expert system with principles of the various popular forecasting models. Based upon the characteristics of data provided by the users, the system can select the best feasible forecasting model to use as well as educate an inexperienced user about how the most appropriate model can be selected. This advisory system was developed on a microcomputer using an expert system shell, Personal Consultant Plus. (9 refs.)

2657 Expert system for specifying of CAD software systems.
K. Ghosh, L. Villeneuve, N.D. Tai (Dept. of Ind. Eng., Ecole Polytech., Montreal Univ., Que., Canada).
CAD/CAM Robotics and Factories of the Future. 3rd International Conference (CARS and FOF '88) Proceedings, Southfield, MI, USA, 14-17 Aug. 1988 (Berlin, West Germany: Springer-Verlag 1989), p.253-7 vol.2

In spite of the substantial productivity gains possible with CAD, it is not readily accepted by small industries. One of the reasons for this hesitation is that the executive of a small industry does not have the necessary background to evaluate the various products that are available in order to make a proper choice. The authors have created an expert system that will help in specifying the CAD software (and the necessary hardware) appropriate for the needs of a certain company. The system asks some pertinent questions of the user and makes a recommendation based on this information. (6 refs.)

2335 A prototype earthquake warning system for strike-slip earthquakes.
M.N. Toksoz, A.M. Dainty, J.T. Bullitt (Earth Resources Lab., MIT, Cambridge, MA, USA).
Pure Appl. Geophys. (Switzerland), vol.133, no.3, p.475-87 (May 1990).

A prototype expert system has been developed to provide rapid warning of earthquakes while they are occurring. Warning times of up to 100 seconds will be possible. In the complete system, several accelerometers are distributed at intervals within a few kilometers of a known fault; data are telemetered to a central computer which implements the expert system. The expert system incorporates specific information about the type of fault to be monitored, and includes simple rules for estimating the fault slip, rupture length, and seismic moment, all in real time. If the seismic moment exceeds a preset value, an alarm may be issued. The prototype is designed for deployment on near-surface strike-slip faults such as the San Andreas and has been successfully tested with data from the 1979 Imperial Valley and 1984 Morgan Hill earthquakes. Crucial concepts have also been tested using synthetic data calculated for a model of the 1857 Fort Tejon earthquake. Parkfield, California could be used as a test site. (7 refs.)

2100 Computer interrogation of patients and diagnostic decision support in dyspepsia.
G. Lindberg, R.P. Knill-Jones (Karolinska Inst., Huddinge Univ. Hospital, Sweden).
MEDINFO 89. Proceedings of the Sixth Conference on Medical Informatics, Beijing, China and Singapore, 16-20 Oct. 1989 and 11-15 Dec. 1989 (Amsterdam, Netherlands: North-Holland 1989), p. 1206

Summary form only given. Gladys (Glasgow Diagnostic Decision Support System for Dyspepsia) is a microcomputer based system for automated history-taking and diagnostic decision support in patients with dyspepsia. Gladys have been translated for use in ten languages. The demonstration includes the original English version and a recently developed Chinese version for the Macintosh series of personal computers. The system includes a set of 375 questions to elicit information about 120 symptoms in patients with dyspepsia. On average only 110 questions are actually given to any one patient. Questions are displayed one at a time and patients answer the question by clicking with a mouse driven pointer at answer buttons on the screen. The inference mechanism utilizes scores based on logistic regression analysis and Bayes' rule to calculate the probabilities of ten major diagnoses. Output includes diagnostic suggestions, explanatory tables of symptoms for and against suggested diagnoses, advice on management, and a summary of interview results in the style of a referral letter. Particular emphasis in the evaluation was put on the diagnosis of disorders detectable by upper endoscopy.

2589 Expert configuration systems: a survey and lessons learned.
J. Liebowitz (George Washington Univ., Washington, DC, USA). Expert Syst. Appl. (UK), vol.1, no.2, p.183-7 (1990).

One particular area within expert systems where applications are increasing is configuration management. The need for accurately configuring hardware is an important requirement for many manufacturing organizations. The paper surveys some of the work that is being done in developing expert configuration systems. It also presents some lessons learned. (23 refs.)

2087 Improving learner performance with computer based programs.
G.A. Marcoulides (California State Univ., Fullerton, CA, USA). J. Educ. Comput. Res. (USA), vol.6, no.2, p.147-55 (1990).

The study presented compared the effectiveness of different types of instructional aids for improving learner performance in an introductory statistics course. Two types of instructional aids were examined: an expert system program (ES), and a computer-assisted instructional program (CAI). Subjects for this study were 133 students selected from a required undergraduate introductory statistics course at a large state university. Students with high and low prior levels of statistics achievement were randomly assigned to three treatment groups: an ES as an adjunct to lectures, a CAI as an adjunct to lectures, and a control group with lectures alone. At the conclusion of the study students were tested on a statistics achievement test. Students that had utilized the ES and the CAI teaching aids scored significantly higher than students from the lecture group alone. There were no significant differences between the ES and CAI groups. The results of the study indicate that computer-based programs can help students learn and improve their performance. (17 refs.)

2289 Integration of AI and OR techniques for computer-aided algorithmic design in the vehicle routing domain.
J.-Y. Potvin, G. Lapalme, J.-M. Rousseau (Centre de Recherche sur les Transp., Montreal, Univ., Que., Canada).
J. Oper. Res. Soc. (UK), vol.41, no.5, p.517-25 (June 1990).

Shows how tools and techniques of artificial intelligence can be successfully integrated into a computer system working in the vehicle routing domain. The aim of this system, called ALTO, is to facilitate the development of routing algorithms for transportation vehicles. The authors describe the general algorithmic framework and the rich interface provided by the system to the expert algorithm designer. They also introduce a methodology for acquiring useful knowledge in the domain, based on examples of successful and unsuccessful problem-solving strategies. With such knowledge, ALTO would then be capable of actively supporting the algorithm designer by suggesting good candidate algorithms for solving new problems. (16 refs.)

2598 Intelligent lot-size advisor for MRP systems.
C.H. Dagli (Dept. of Eng. Manage., Missouri Univ., Rolla, MO, USA).
CAD/CAM Robotics and Factories of the Future. 3rd International Conference (CARS and FOF '88) Proceedings, Southfield, MI, USA, 14-17 Aug. 1988 (Berlin, West Germany: Springer-Verlag 1989), p.282-6 vol.2

Intelligent lot-size advisor (ILA) is a prototype expert system that select lotsizing heuristics considering demand patterns and some qualitative factors which are not generally included in models. The current structure of the system could provide sufficient support to users of MRP software. It can be integrated into existing MRP systems through a custom-make interface. (9 refs.)

2606 Evaluating medical expert systems: what to test and how?
J. Wyatt (Nat. Heart & Lung Inst., London, UK), D. Spiegelhalter Med. Inform. (UK), vol.15, no.3, p.205-17 (July-Sept. 1990).

Many believe that medical expert systems have great potential to improve health care but few of these systems have been rigorously evaluated, and even fewer are in routine use. The authors propose the evaluation of medical expert systems in two stages: laboratory and field testing. In the former, the perspectives of both prospective users and experts responsible for implementation are valuable. In the latter, the study must be designed to test, in an unbiased manner, whether the system is used in clinical practice and if it is used, how it affects the structure, process and outcome of health care encounters. They conclude with proposals for encouraging the objective evaluation of these systems. (37 refs.)

2154 Expert systems are modifying 1992 manufacturing sector strategies.
R. Sargeant.
Proceedings of the 1st International Conference on Artificial Intelligence and Expert Systems in Manufacturing. The Scope, Applications and Limitations of Intelligent Manufacturing Systems, London, UK, 20-21 March 1990 (Kempston, UK: IFS Conferences 1990), p.25-37

Any manufacturing company which has not taken full account of the availability and use of expert systems in preparing their business strategy for the period 1989 to 1994 is likely to lose market position. Such Information Technology can enable companies to drive market changes and to occupy dominant positions as 1992 approaches. This paper describes some operational benefits arising from the successful development and installation of large real-time expert systems and introduces the less discussed manufacturing business strategic issues which have emerged and are now set to occupy a more significant role.

2090 Expert systems and computer-controlled decision making in medicine.
B. Lipscombe (Sci. Studies Centre, Sch. of Humanities & Social Sci., Bath, UK).
AI Soc. (UK), vol.3, no.3, p.184-97 (July-Sept. 1989).
[received: 30 May 1990]

The search for 'usable' expert systems is leading some medical researchers to question the appropriate role of these programs. Most current systems assume a limited role for the human user, delegating situated 'decision-control' to the machine. As expert systems are only able to replace a narrow range of human intellectual functions, this leaves the programs unable to cope with the 'constructivist' nature of human knowledge-use. In returning practical control to the human doctor, some researchers are abandoning focused problem-solving in favour of supportive problem-analysis. Using ONCOCIN and QMR as examples, the article contrasts these approaches and suggests that the latter avoids many of the difficulties currently facing medical expert systems. (39 refs.)

2313 New directions for medical artificial intelligence.
V.K. Sondak (Michigan Univ., Med. Sch., Ann Arbor, MI, USA). Comput. Math. Appl. (UK), vol.20, no.4-6, p.313-19 (1990).

The past decade has seen significant advances in medical artificial intelligence (MAI), but its role in medicine and medical education remains limited. The goal for the next decade must be directed towards maximizing the utility of MAI in the clinic and classroom. Fundamental to achieving this is increasing the involvement of clinicians in MAI development. MAI developers must move from 'pet projects' toward generalizable talks meeting recognized clinical needs. Clinical researchers must be made aware of knowledge engineering, so clinical data bases can be prospectively designed to contribute directly into MAI 'knowledge bases'. Closer involvement of MAI scientists with clinicians is also essential to further understanding of cognitive processes in medical decision-making. Technological advances in user interfaces-including voice recognition, natural language processing, enhanced graphics and videodiscs - must be rapidly introduced into MAI to increase physician acceptance. Development of expert systems in non-clinical areas must expand, particularly resource management, e.g. operating room or hospital admission scheduling. The establishment of MAI laboratories at major medical centers around the country, involving both clinicians and computer scientists, represents an ideal mechanism for bringing MAI into the mainstream of medial computing. (13 refs.)

2585 Expert systems for machining parameter selection: design aspects.
B. Gopalakrishnan (Dept. of Ind. Eng., West Virginia Univ., Morgantown, WV, USA).
Adv. Manuf. Eng. (UK), vol.2, no.2, p.59-63 (April 1990).

Describes the different methods that can be used to make expert

systems as efficient as possible for applications within the domain of manufacturing, particularly machining parameter selection. Machining parameter selection is described, and the underlying areas of expert systems applications are outlined. The design of expert systems is than analyzed with respect to the different techniques that will aid expert system shells in the area of data acquisition, inference engine path modifications, and user input, to make them function effectively in the domain of machining parameter selection. (16 refs.)

2164 An expert system for generation of hardware failure equations.
C.L. McCullough, M.O. Hofmann (Alabama Univ., Huntsville, AL, USA).
Trans. Am. Nucl. Soc. (USA), vol.61, p.211-12 (1990).
(1990 Annual Meeting of the American Nuclear Society (papers in summary form only received), Nashville, TN, USA, 10-14 June 1990).

One of the difficulties in performing the type of probabilistic risk assessment currently mandated by the US Nuclear Regulatory Commission (NRC) is the tremendous amount of time that must be invested by a trained analyst in order to generate the necessary hardware failure equations from an imprecise and eclectic collection of information, including system designs, qualification data, vendor drawings, and maintenance requirements. From this, the analyst must determine what constitutes system failure for the given goal and must generate Boolean (logical) equations for all failure paths involving the components deemed significant. These equations must than be altered to account for maintenance actions, conditional failure probabilities, etc. While the analyst cannot be removed entirely from this process, it is the authors' contention that artificial intelligence can be used to generate the Boolean hardware failure equations necessary for risk assessment from a simplified collection of data based on a data base of standard component types. (2 refs.)

2130 Artificial Intelligence and expert systems techniques in control: an overview.
S.G. Tzafestas (Nat. Tech. Univ., Athens, Greece).
Syst. Anal. - Model. - Simul. (East Germany), vol.7, no.3, p.171-90 (1990).

Gives an overview of the key artificial intelligence (AI) concepts and techniques used for the development of knowledge-based expert systems that are appropriate for the computer-aided and operation of industrial control systems. After a short outline of the main benefits of the AI approach, and an exposition of some underlying techniques, languages and tools, the basic features of process control expert systems are discussed. Then a number of working expert control and supervision systems are presented in an attempt to show how many of the ideas of the paper have been put into practice. (35 refs.)

2060 PROJEVAL: an expert system for agricultural project evaluation in developing countries.
K. Khan, G.I. Doukidis (Dept. of Inf. Syst., London Sch. Of Econ., UK).
Expert Syst. Inf. Manage. (UK), vol.1, no.1, p.22-42 (Spring 1988). [received: 01 Jun 1990]

PROJEVAL is a prototype, a test-version of a knowledge-based expert system which has been developed with the co-operation of a consultant to various international agencies. It is intended to be used as a tool for training government officials in developing countries to carry out project identification. It trains them by asking questions about an actual or hypothetical project and analyzing it in the same way that an expert would do. This paper first describes the life-cycle of project evaluation and how PROJEVAL fits into it. It then describes the development of PROJEVAL, its architecture and functions. Finally it suggests the future developments required for the operational system. The paper concludes that PROJEVAL is bound to have an impact on developing countries as it can reduce their dependence on foreign expertise. This is as important as helping them avoid projects that drain away limited resource. However, to take the system beyond the prototype stage will require a major integration effort as it is meant to be of use to the international agencies as well as the developing countries. (7 refs.)

2399 Nuclear power plant expert system verification and validation.
J.A. Naser (EPRI, Palo Alto, CA, USA).
7th Power Plant Dynamics, Control and Testing Symposium Proceedings, Knoxville, TN, USA, 15-17 May 1989 (Knoxville, TN, USA: Univ. Tennessee 1989), p.25.01-25.14 vol.1

Expert system technology has matured enough to offer considerable benefits for a number of application areas in the nuclear power industry. These potential benefits are becoming widely recognized. However, an obstacle to the acceptance of expert systems in some of these areas by both users and regulators is the lack of verification and validation methodology. The author gives some initial guidelines for the verification and validation of expert systems. Considerable work has been done developing verification and validation methodologies for conventional software systems. This work should be taken advantage of and, where applicable, adapted or modified for expert systems. Additional verification and validation techniques need to be developed to handle the unique characteristics of an expert system's knowledge base and the iterative nature of the expert system's development. Expert systems are categorized into six types to identify their different verification and validation requirements. Suggested methodology guidelines are given for the first four types. (9 refs.)

2264 An expert system for harmonizing chorales in the style of J.S. Bach.
K Ebcioglu (IBM Thomas. J. Watson Res. Center, Yorktown Heights, NY, USA).
J. Log. Program. (USA), vol.8, no.1-2, p.145-85 (Jan.-March 1990).

This paper describes an expert system called CHORAL, for harmonization of four-part chorales in the style of Johann Sebastian Bach. The system contains about 350 rules, written in a form of first-order predicate calculus. The rules represent musical knowledge from multiple viewpoints of the chorale, such as the chord skeleton, the melodic lines of the individual parts, and Schenkerian voice leading within the descant and bass. The program harmonizes chorale melodies using a generate-and-test method with intelligent backtracking. A substantial number of heuristics are used for biasing the search toward musical solutions. The CHORAL knowledge base

provides for style-specific modulations, cadence patterns, and complex encounters of simultaneous inessential notes; it imposes difficult constraints for maintaining melodic interest in the inner voices. Encouraging results have been obtained, and output examples are given. BSL, a new and efficient logic-programming language fundamentally different from PROLOG, was designed to implement the CHORAL system. (80 refs.)

"Based on the information provided in the abstracts, the references provided have been selected by J. Scrimgeour of the National Research Council as a representative sample of interest and value to Canadian industry. Abstracts provided are reprinted from "Key Abstracts in Artificial Intelligence" with permission 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 American 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 0S2, (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".

Controlling Super-Intelligent Machines

continued from page 12

whom. Hogan's The Genesis Machine depicts this kind of connection in convincing detail, although the computer involved is not a SIM [4].

The majority of people would likely consider the result of human-SIM fusion to be inhuman. Most people might judge this to be a disadvantage of this approach.

Conclusion

These deliberations are not as premature as they might seem. Vere and Bickmore have constructed a "basic agent"; a rudimentary machine intelligence [10]. Their agent

> ... integrates limited natural language understanding and generation, temporal planning and reasoning, plan execution, simulated symbolic perception, episodic memory, and some general world knowledge.

Vere and Bickmore have considered the issue of controlling Homer, their basic agent. Homer exists in a simulated world, where it can do little harm. Homer commands a simulated robot submarine, which can move about underwater, take photographs, refuel, pick up objects, and

> ... also shoot objects. This capability exists primarily to enable scenarios involving the first law of robotics ... We can run such scenarios in our simulation without actuating politicians and the news media. Homer will shoot inanimate objects and animals, for example, a mine or a shark, but not people.

Presumably Homer will not be released in the real world until there is good evidence that Homer obeys the first law of robotics. We have seen, however, that there may be some problems with using Asimov's laws to control SIM's.

We already have primitive machine intelligence. It may not be long before we have super-intelligent machines. It is time to think about the problems they may present.

Acknowledgements

Thanks to Louise Linney, John Linney, Ken Deaton, and Jeff Ardron for helpful comments.

References

[1] Asimov, I., Opus 100, Houghton Mifflin. Boston, Massachusetts, 1969.

[2] Benford, G., "Me/Days", in In Alien Flesh, Victor Gollancz. London, England, 1988. Pages 182-195.

[3] Gibson, W., Neuromancer, Ace Science Fiction. New York, New York, 1984.

[4] Hogan, J. P ., The Genesis Machine, Ballantine Books . New York, New York, 1978.

[5] Hogan, J. P., The Two Faces of Tomorrow, Ballantine Books. New York, New Yo rk, 19 7 9 .

[6] Jones, D. F., Colossus, G. P . Putnam' s Sons . New York, New York, 1967.

[7] Skinner, B. F., Beyond Freedom and Dignity, Alfred A. Knopf. New York, New York, 1971.

[8] Sladek, J., "Broot Force", in The Best of John Sladek, Pocket Books. New York, New York, 1981. Pages 190-194.

[9] Turing, A. M., "Computing Machinery and Intelligence", Mind, Vol. LIX, pp. 433-460, 1950.

[10] Vere, S. and Bickmore, T., "A basic agent", Computational Intelligence, Vol. 6, No. 1, pp. 41-60, 1990.

Peter Turney is a Research Associate at the Knowledge Systems Laboratory, Institute for Information Technology, National Research Council. He is currently involved in research in machine learning. He has a Ph.D. in Philosophy from the University of Toronto.

CALL FOR PAPERS

4th International Conference on Computers and Learning

June 17-20, 1992

Acadia University, Nova Scotia, Canada

About ICCAL

The objective of the 4th International Conference on Computers and Learning is to provide a forum for the exchange of ideas and presentation of developments in the theory and practice of the use of computers in learning. The previous three conferences were held in Hagen (Germany 1990), Dallas (USA 1989), and Calgary (Canada 1988). The last conference attracted over 300 delegates. The high quality papers presented at past conferences have been published in the series "Lectures in Computer Science" by Springer-Verlag.

Areas of Interest - (Representative List)

Authoring Systems
Hypermedia
A.I. Applications
Simulations
Media-based CAI
Visualization of Algorithms
Human Problem Solving
Knowledge-based CAI Systems
Presentation CAI and ICAI
Discourse Management
Problem Generation
Performance Monitoring
Intelligent Tutoring Systems
Knowledge Acquisition and Representation
Student Modelling and Cognitive Diagnosis
Evaluation of Learning Environments
Human-computer Interaction
Applications in Learning, such as

Medicine	The Arts	Music
Engineering	Business	Sciences
Humanities	etc.	

Computers in Distance Education.

Organization

Steering Committee:

H. Maurer (Austria), D. Norrie (Canada), F. Makedon (USA), Th. Ottmann (Germany).

Advisory Board:

D. Bitzer (USA), A. Bork (USA), B. Clark (Canada), G. Davies (UK), K. Friend (USA), G. Kovacs (Hungary).

Program Committee:

D. Benest (UK), M. Brown (USA), P. Carlson (USA), S. Cerri (Italy), B. Clark (Canada), N. Coulter (USA), G. Davies (UK), M. DeBlasi (Italy), N. Gardner (UK), J. Hammond (Australia), R. Hartog (Netherlands), F. Hill (USA), D. Jonassen (USA), G. Kovacs (Hungary), E. Luque (Spain), G. McCalla (Canada), F. Makedon (USA), H. Maurer (Austria), R. Mizoguchi (Japan), A. Montgomery (Australia), M. Mulhauser (Germany), D. Norrie (Canada), T. O'Shea (UK), G. Oberem (South Africa), S. Otsuki (Japan), Th. Ottmann (Germany), D. Paramskas (Canada), M. Petruk (Canada), C. Prescott (Thailand), B. Schneiderman (USA), I. Tomek (Canada), S. Wills (Australia), C. Unger (Germany), B. Woolf (USA), M. Yazdani (UK).

Conference Chairman:

Ivan Tomek, Jodrey School of Computer Science, Acadia University Wolfville, Nova Scotia B0P 1X0, Canada
Tel: (902) 542-2201 Ext.467, Fax: (902) 542(7224), E-Mail: BITNET:ICCAL@AcadiaU.ca

Schedule

Wednesday June 17: Workshops and tutorials.(See list of proposed topics below.) Registration and welcoming party.
Thursday June 18: Invited lectures. Presentation of Papers.
Trade Exhibition begins.
Friday June 19: Invited lectures. Presentations of papers. Poster sessions.
Trade Exhibition continues.
Saturday June 20: Invited Lectures. Presentation of papers.
Trade Exhibition concludes.

Proposed Workshops/Tutorials:

Cognitive Science, Hypermedia, Authoring Software,
Intelligent Tutoring Systems.
A social program for spouses and friends will be available.

Important Dates

Manuscripts submitted for presentation or publication must be in the hands of the Conference Chairman no later than September 20, 1991.

Manuscripts to be reviewed and authors notified by December 15, 1991.

Deadline for receipt of camera-ready version of manuscript February 15, 1992.

The Venue

Acadia University offers a unique location for ICCAL '92, situated as it is in the centre of a small picturesque rural community. Being little over 150 years of age, Acadia University is one of the oldest in Canada, and its campus has been described as one of the most beautiful. The University offers a full range of programs in the Humanities, Sciences and other areas. Its School of Computer Science is 20 years old and well known for its undergraduate program with an established co-operative component and innovative use of computers in support of learning.

Wolfville is about 100 km from Halifax on the shore of the Minas Basin in the Bay of Fundy. The town has a small but colorful business district with a number of restaurants, shops and other services. It is situated in an attractive tourist and agricultural area known as the Annapolis Valley.

The small town character of Wolfville and the fact that all facilities, restaurants and hotels are within easy walking distance gives participants an opportunity to be in constant touch and make close personal contacts.

Since the conference will be held after the end of the academic year, the University will make its facilities, such as residences, swimming pool, gymnasium and tennis courts available to participants.

4th International Conference on Computers and Learning

c/o Dr. Ivan Tomek
Jodrey School of Computer Science
Acadia University
Wolfville, Nova Scotia, B0P 1X0, Canada

BOOK REVIEWS

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

Reviewed by
Norma Welch
MaSyCom Consultants Inc

Innovative applications of artificial intelligence presents, under a single cover, 27 papers, each describing a successful and practical application of AI. The collection is extremely rich and diverse. The papers have two things in common: each describes an application that is being used to provide substantial proven benefits, and each is eminently readable, clearly describing the domain, the requirements, the systems architecture, the methodologies, and the experiences with the developed system. Most of the systems were initiated around 1984 to 1985 and have been in use since 1988 or 1989.

In size, the artificial intelligence projects described in this book range from the Packaging Advisor at Dupont, developed in nine months on a personal computer, to the Naval Battle Decision Aid which has a budget of $10 million. Geographically, the collection includes projects from Norway, France, Singapore, and Australia. Application domains include aerospace, finance, biotechnology, emergency response, law, manufacturing design, process control, advertising, music composition, the military, and labour management. Systems involving rule sets, frames, and object-oriented programming are well represented. Embedded and integrated systems are common. Other techniques that have emerged from the AI labs into these practical systems include: blackboard techniques, augmented transition networks, various methods of linking numeric and symbolic computations, computational linguistics, concepts, simulation, goal formulation, and machine learning. Although systems developed on Lisp machines are in the majority, the papers include applications on various other hardware and software platforms.

Though it is difficult to select examples from such an outstanding and diverse collection, the following brief descriptions are presented as a sampling.

"Syntel: An architecture for financial applications" by Peter E. Hart describes the special modelling needs of financial advisory systems and discusses how Syntel was developed to meet those needs. Key features of Syntel are the powerful "primitive functions", the "forms" system of user input and "value tables". Syntel was developed, at Syntelligence, over a period of five years at an investment of over 100 person-years of effort. It forms the basis for Underwriting and Lending Advisor products, which have been purchased by insurance companies and banks worldwide.

"The intelligent banking system" by Kenan Sahin and Keith Sawyer describes the application of natural language processing to the interpretation of international funds transfer messages. The interpretation is performed by a combination of several techniques from computational linguistics with domain knowledge. Intelligent banking workstations are integrated with the TELEX systems. Another application using language techniques is the Direct Labour Management System at the Ford Motor Company. In this system, a custom-designed Process Description Language is used to detail work to be performed. The system interprets the work descriptions and plans all stages of the assembly process.

The Can Build system at Digital Equipment Corporation helps management with strategic inventory decisions. The system, described in the paper by Robin M. Krumholtz, is a large, integrated, computerized solution to a business need. The knowledge base portion is small, though crucial. The paper presents a thoughtful discussion of the methods used to involve the users fully in all stages of development and implementation.

A novel paper, "Wolfgang, Musical composition by emotional computation", by R.Douglas Riecken describes an application of theories from Marvin Minsky's *Society of mind*. Wolfgang is in use as a composer's assistant. It takes a short musical idea and develops it into a composition of sonata form. The composer can request that the idea be developed with emphasis on various emotional states (sad, happy, etc). The system combines musical knowledge, a K-line network of intelligent agents, goal formulation, a scheduler, a blackboard, and a log file. Wolfgang can be trained over time to adopt the characteristics of the trainer composer.

Last, least expensive, but not least, "Packaging Advisor: An expert system for rigid plastic food package design" by Alvin S. Topolski and Douglas K. Reece describes how a rapidly prototyped system was able to provide competitive advantage to Dupont in the packaging marketplace.

All the papers in this collection describe significant and

successful applications, and do so very well. The reader is guaranteed to absorb some stimulating ideas. Many of these applications will likely become classics.

Norma Welch is an AI consultant whose specialties include the application of AI techniques in railroad operations.

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

Reviewed by
Michael J. D. Sutton
Flynn McNeil Raheb and Associates Ltd.

Ethical conflicts in the specialities of high technology have become more common in the last decade. Electronic communications issues are some of the easiest to identify. Just over a year ago, Epson America fired an electronic-mail administrator, Alana Shoars. Epson would not elaborate why Shoars was fired. However, Shoars says that she was fired for questioning why her supervisor was reading the e-mail of employees without their knowledge. Shoars is now suing Epson America for wrongful dismissal and invasion of privacy. Ms. Shoars could be considered one of the new ethical martyrs of computer science.

Late last year, the Prodigy information service announced a new fee structure. By using the e-mail services of Prodigy, a group of subscribers organized themselves to oppose the price hike. They sent out thousands of messages to other members and advertisers. Their objective was to protest the increased fees and get Prodigy to rescind its price hikes. The surprising result was that Prodigy disconnected a number of the troublemakers and suspended their accounts without warning for "harassing" other members with unsolicited messages.

One of the protesters was Henry Niman, a cancer researcher in Pittsburgh, Pennsylvania. He subsequently was denied access to electronic banking services which were a significant means of paying his personal bills. Moreover, he lost access to all messages in his mailbox. This was particularly critical. A patient's medical treatment was contained in one of the messages.

I wanted to describe above some of the recent situations which beg re-evaluation of the application of ethical principles to our electronic day-to-day living. It is not always obvious how ethical values may apply, or where these values may conflict with other values when appropriately applied. This book tries to get the reader to ascertain what constitutes ethical or unethical behaviour. It is a complication of a series of scenarios which depict ethical problems.

The 1987 study described in this book was conducted by SRI International to assist users of computers and telecommunications to clarify and apply their ethical values as new, complex situations are encountered. An earlier study in 1977 by SRI resulted in the development of a new, experimental approach, scenario analysis, for discussing ethical issues. The first study focused on the computer scientists and computer teachers who served the users of the technology.

The current study, because of the dramatic change in the widespread use of computers by non-computer professionals, concentrates on all users of computers and telecommunications. Now more individuals have an opinion on computer-related ethical issues. However, identifying and recognizing impacts is much more difficult because of the diffusion and proliferation of computer technology and its networking via telecommunications.

Certain principles are intrinsic to the evaluation of the scenarios. The Kantian universality rule states that "an act that is not right for everyone is not right for anyone''. The Descartian observation specifies that "a sufficient change in degree produces a change in kind''. Weaving these two rules through the scenario of unauthorized use of an employer's computer resources negates an interesting rationalization. Some individuals would conclude that use of computer resources without authorization is justified, since it would otherwise go wasted (unused). Nonetheless, in applying these two rules, the fallacy of such a rationalization becomes apparent:

> If everyone used computer resources that are not authorized, no resources would be available for authorized work. Indeed, taking an increasing amount of computer resources for unauthorized personal use at some point changes from immaterial to material unethical behaviour.

The scenarios documented in this book were designed to raise questions of unethicality rather than ethicality. They are short, simple stories with a minimum of details written as objectively as possible. They are limited to individuals confronted with acts or engaged in acts that could be construed as unethical. Issues of a global nature were avoided because they would require joint decisions of more than one person. The issues described directly apply to individuals in computer science, information science, technology and business.

Differences in experiences, backgrounds, and personal standards caused extensive and heated discussions amongst the participants in this study. The most difficult problem in evaluating the scenarios was the requirement by each participant to treat the scenario as absolute, i.e., they were to finally limit their vote on each scenario to either 'unethical', 'not unethical', or 'no ethical issue'. Afterwards, opinions from each participant to support their vote were solicited, describing the actor, the act, and any problem in reaching a decision. These opinions are summarized after each scenario and document invaluable insight into the general principles applied by the participants to the problems.

A recent conference sponsored by the Computer Professionals for Social responsibility should help to add spark to this area of debate. The first Conference on Computers, Freedom, and Privacy took place in late March, 1990, in Burlingame, California. Topics included discussion of legislation that would make it difficult for companies such as Prodigy to create private terms and conditions that may contravene the basic rights of individuals. It is hard to speculate where these issues will be dealt with in the future.

Wherever people have access to computers and modems, an electronic community can arise. Many questions result when a community's interaction is mediated by a digital interface. How is that electronic community ethically legislated? What are the obligations and rights that the common carriers and utilities can enforce for use of their services? Are new ethical issues raised by changing technological sophistication, or do the ethics questions remain constant, regardless of the rate and kind of technological change unfolding?

As professionals in this emerging and maturing field of artificial intelligence, we are obligated to promote sound ethical practices. Additionally, we should be mentoring those employees or students with less experience (and, regretfully, less grounding in ethical behaviour) to question their actions, and the actions of those around them in the fields of computer science, high technology, government and business. The scenarios outlined in this book would help educators and mentors to describe problems and justify actions of specific ethical situations.

Michael Sutton is Managing Principal of the Ottawa branch of Flynn McNeil Raheb and Associates, information management consultants. He is the author of many popular and academic papers, with a recent emphasis on the electronic forms and formware technology. His background includes academic studies in expert systems, knowledge engineering, theology, and philosophy. He established the first ecumenical Pastoral Ministry in Ottawa over a decade ago to study the impact of advanced technology by laypersons from a theological and philosophical perspective.

Readings in planning
James Allen, James Hendler, and Austin Tate (editors)
(University of Rochester, University of Maryland, and University of Edinburgh)
San Mateo, CA: Morgan Kaufmann, 1990, xiii+754 pp
(Morgan Kaufmann series in representation and reasoning) Paperbound, ISBN 1-55860-130-9, US$38.95

Reviewed by
Qiang Yang
University of Waterloo

Planning has been with AI ever since its birth. But it has been difficult for professors to teach courses on planning, and for interested researchers to obtain a systematic view of planning, precisely because of the lack of comprehensive collections of literature in planning. Even more serious than this difficulty has been the problem of planning researchers unwittingly "reinventing the wheel"; a book such as this should, at least in part, provide a solution.

Readings in planning is a very comprehensive collection of foundational papers in the field, which successfully addresses the above problems. Intended for readers with a reasonable understanding of basic AI techniques, such as logic and various heuristic search methods, it contains 48 papers that cover a wide range of subjects in the field. Starting with the earliest attempts to model human problem solving and cognitive processes, the collection includes work addressing the different approaches to representing plans and actions, formal models of action and planning systems, relations between planning, execution, learning and plan reuse, as well as a whole spectrum of search control methods. In addition, it includes articles representing some of the most recent research directions in planning, such as multiagent reasoning and reactive planning.

The attraction of this collection, therefore, is that it offers more than just a set of papers on planning, it also provides an accurate picture of the history of the field. The book opens with two survey articles about planning problems and techniques, and an argument for the necessity of formal approaches to planning. Each survey article is complete with an extensive bibliography of the research papers in the field. The whole collection is partitioned into four parts, each addressing a different aspect of planning and planning systems: Introduction, Planning Systems, Foundations of Planning, and New Directions in Planning Systems. These divisions are natural and well-organized, and readers with different aims can easily find their way to the most relevant part of the book. Furthermore, the book allows us to follow the evolution of many ideas in planning from their initial conception to maturity.

As the editors themselves admit, the collection does not cover every subfield of AI related to planning, rather it includes only papers on domain-independent approaches to planning and plan reasoning. Due to lack of space, papers on plan recognition and domain-dependent planning systems are not included.

Most of the planning techniques covered in the collection are close to the state of the art in the planning field. In fact, following the thread laid out by the editors, it is not hard to observe a trend leading from the early engineering approaches to planning, to the most recent attempts in addressing the planning problems in a scientific way. Readers, however, should not yet regard planning as a mature field. Rather, I would encourage readers to "read between the lines"; instead of believing the papers to provide solutions to the planning problems, they should be regarded as works that raise issues yet to be answered. For example, despite many years of

speculation about the superiority of nonlinear planning over linear planning, there has been no comprehensive study, either empirical or theoretical, confirming or denying the validity of the claim. Most of the experiments were only run on a few examples of the toy problems, and an experimental study of planning has yet to be completed. Also, only recently have researchers begun to address the issues of what constitutes a good abstraction hierarchy, and how to construct it automatically. It would have been helpful for the editors to express some of these shortcomings in the introduction to each chapter.

Overall, the collection is both timely and well-done, and should prove a useful tool to all involved in the field of planning.

Qiang Yang is an Assistant Professor at the University of Waterloo. His current research interests include AI planning, plan-related reasoning, and their applications to automated design and manufacturing.

Pattern thinking
L. Andrew Coward
(Bell-Northern Research) New York: Praeger, 1990,
xii+180 pp Hardbound, ISBN 0-275-93427-6, us$42.95

Reviewed by
Helen M. Gigley
National Science Foundation

Pattern Thinking presents a plausibility argument for a model of brain processing that underlies human cognitive behaviour which even includes pleasure, pain, and personality. If it had been written about 15 years ago, it might have fit better with the then-current understanding of neural processing and cognition. It begins with a discussion of the computer as a possible model for neural processing and discusses both a micro- and macroscopic perspective of levels of description that arise by assuming this as a basis. The goal of the book is to develop a model of brain processing at a micro-level that is consistent with what is known of neural process from many diverse approaches to its study and consistently, to expand on the basic micro-level description to capture functional aspects of human behavior.

I found many problems with the book. At first, I was intrigued with the idea of a pattern hierarchy and the related cascade of firing activations as the basis of human cognitive process. I still find that this idea is clearly conveyed as a premise of the text. However, there are too many problems with the exposition of the supporting data and theories to convince one by the end of the text that the theory is worth modeling. A design of a model is presented, but not implemented. I find this a critical lack.

There are errors in the discussion of the basic neural cellular processes. It often appears that the author has read basic texts and provides synopses of them at an even more basic level. In doing so, important facts are omitted or made so simple they are incorrect. For instance, a synapse is shown as the thickening at the end of an axon. In two different basic texts I've checked (Eccles 1977, Thompson 1967), the synapse includes the synaptic thickening at the end of the axon, the synaptic cleft, and the post-synaptic membrane. All three together are the synapse, not just the terminal thickening. This basic discussion is further clouded by whether a nerve or a neuron is being described. A nerve is outside the central nervous system (CNS), while a neuron is inside. A nerve can innervate a muscle, but a neuron cannot. The entire discussion of neurophysiology leaves much to be desired. If anyone would like to pursue the idea of the pattern hierarchies, which is a view consistent with the connectionist or neural-net views of cognition, then they should skip these chapters and go to a decent book on neurophysiology or physiological psychology for correct information.

Throughout, there is discussion of the abstraction capability which a neuron has. It reacts to a pattern of inputs that is not isomorphic to just one fixed distributed state but can range over some bounded space of possible inputs that are interpreted as having the same effect on the "next" processes. However, within this view and ensuing discussion, it is never clear when the author attributes the ability for one cell to produce a response, whether a "grandmother cell" is allowed.

When the discussion of the pattern hierarchical processing approach expanded to include pleasure, pain, and personality, I felt the claims were definitely beyond plausibility. And the final suggested design was barely adequate to convey why anyone would even consider building the model. If the author did not feel so compelled, why would anyone else?

In summary, I will use the author's own words to provide an overview of the model:

> So far, it has been shown how the phenomena of perception, thinking, memory, sleep, pleasure and pain, personality, and self-awareness can be described and accounted for within a system model of the brain. This model is a pattern extraction hierarchy, organized into regions with differing levels of arousal. The processes of thought are the subjective manifestations of pattern extraction cascades of firing neurons that find paths through the pattern extraction hierarchy. Feedback between levels of abstraction gives rise to the advanced phenomena of speech and self-awareness.

Following this statement is a question of whether this model can be mapped onto computational architectures as data and instructions. The remainder of the book is the design of this mapping. Being only a design, I do not feel the book provides any closure for its claims.

There may be glimmers of really exciting ideas hidden between the lines, but they are not obvious. This book does not have much to contribute to computational approaches to understanding neuroscience, human behavior, or processing

underpinnings of neural processes. It has less to contribute to computer science or artificial intelligence.

References
Eccles, John C. (1977) *The understanding of the brain.* McGraw-Hill Book Company.
Thompson, Richard F. (1967). *Foundations of physiological psychology.* Harper & Row.

Helen Gigley is Program Director for the Knowledge Models and Cognitive Systems Program (the basic artificial intelligence program) of the National Science Foundation. She has a research affiliation with the Speech and Hearing Section at Walter Reed Army Medical Center in Washington, D.C. Her research includes modeling neurolinguistic performance that integrates brain processes with clinical evidence of language degradation following brain lesion.

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 Artificial Intelligence* 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 Artificial Intelligence* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, or send electronic mail to gh@cs.toronto.edu. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

Propositional attitudes: The role of content in logic, language, and mind

C. Anthony Anderson and Joseph Owens (editors)
(University of Minnesota)
Stanford: Center for the Study of Language and Information, 1990, xvi+342 pp
(CSLI lecture notes 20)
(Distributed by the University of Chicago Press)
Hardbound, ISBN 0-937073-51-2, US$37.50; paperbound, ISBN 0-937073-50-4, US$16.95

John von Neumann and the origins of modern computing

William Aspray
(Center for the history of electrical engineering, IEEE)
Cambridge, MA: The MIT Press, 1990, xvii+376 pp
(History of computing series)
Hardbound, ISBN 0-262-01121-2, US$35.00

+Representing and reasoning with probabilistic knowledge

Fahiem Bacchus
(University of Waterloo)
Cambridge, MA: The MIT Press, 1990, xvii+233 pp
(Artificial intelligence series)
Hardbound, ISBN 0-262-02317-2, US$24.95

The language of first-order logic, including the program Tarski's World

Jon Barwise and John Etchemendy
(Indiana University and Stanford University)
Stanford: Center for the Study of Language and Information, 1990, xiii+257 pp with Macintosh diskette of Tarski's World 3.0 (CSLI lecture notes 23) (Distributed by the University of Chicago Press) Paperbound, ISBN 0-937073-59-8, US$27.50

Expert systems

Robert A. Benfer, Edward E. Brent Jr, and Louanna Furbee
(University of Missouri, Columbia)
Newbury Park, CA: Sage Publications, 1991, 92 pp
(Sage University Paper series on quantitative applications in the social sciences 77) Paperbound, ISBN 0-8039-4036-X, US$7.50

Computer assisted language learning: Program structure and principles

Keith Cameron (editor)
(University of Exeter)
Norwood, NJ: Ablex Publishing and Oxford: Intellect Ltd, 1989, x+115 pp Hardbound, ISBN 0-89391-560-2 and 1-871516-01-3, US$22.50

Plan recognition in natural language

Sandra Carberry
(University of Delaware)
Cambridge, MA: The MIT Press, 1990, xi+286 pp
(The ACL—MIT Press series in natural language processing)
Hardbound, ISBN 0-262-03167-1, US$35.00

How to write parallel programs: A first course

Nicholas Carriero and David Gelernter
(Yale University)
Cambridge, MA: The MIT Press, 1990, xv+232 pp
Hardbound, ISBN 0-262-03171-X, US$29.95

+Artificial experts: Social knowledge and intelligent machines

H.M. Collins
(University of Bath)
Cambridge, MA: The MIT Press, 1990, xiii+266 pp
(Inside technology series)
Hardbound, ISBN 0-262-03168-X, US$19.95

Logic, language, and meaning
Volume I: Introduction to logic
Volume II: Intensional logic and logical grammar
L.T.F. Gamut (pseud.)
(Johan van Benthem, Jeroen Groenendijk, Dick de Jongh, Martin Stokhof, and Henk Verkuyl, Universities of Groningen, Amsterdam, and Utrecht)
Chicago: The University of Chicago Press, 1991, Vol~I: xiv+282 pp, Vol~II: xvi+349 pp Vol~I: Hardbound, ISBN 0-226-28084-5, US$55.00; paperbound, ISBN 0-226-28085-3, US$19.95 Vol~II: Hardbound, ISBN 0-226-28086-1, US$60.00; paperbound, ISBN 0-226-28088-8, US$24.95

Anaphora and quantification in situation semantics
Jean Mark Gawron and Stanley Peters
(Hewlett Packard Laboratories and Stanford University)
Stanford: Center for the Study of Language and Information, 1990, xi+187 pp (CSLI lecture notes 19)
(Distributed by the University of Chicago Press)
Hardbound, ISBN 0-937073-49-0, US$37.50; paperbound, ISBN 0-937073-48-2, US$15.95

Linguistic exploitation of syntactic databases: The use of the Nijmegen Linguistic DataBase program
Hans van Halteren and Theo van den Heuvel
(University of Nijmegen)
Amsterdam: Editions Rodopi, 1990, ix+207 pp
(Language and computers: Studies in practical linguistics 5)
Paperbound, ISBN 90-6203-809-3, US$33.50, Dfl 65.00
[An MS-DOS demonstration diskette for the examples and exercises in the books is available free of charge from the authors.]

Literature and cognition
Jerry R. Hobbs
(SRI International)
Stanford: Center for the Study of Language and Information, 1990, vii+180 pp (CSLI lecture notes 21)
(Distributed by the University of Chicago Press)
Hardbound, ISBN 0-937073-53-9, US$35.00; paperbound, ISBN 0-937073-52-0, US$14.95

Uses of randomness in algorithms and protocols
Joe Kilian
(MIT and Harvard University)
Cambridge, MA: The MIT Press, 1990, xxiii+235 pp
(ACM distinguished dissertations)
Hardbound, ISBN 0-262-11153-5, US$39.95

***Programming for artificial intelligence: Methods, tools and applications**
Wolfgang Kreutzer and Bruce McKenzie
(University of Canterbury)
Sydney: Addison-Wesley, 1991, xiv+682 pp
(International computer science series)
Hardbound, ISBN 0-201-41621-2

Computer processing of natural language
Gilbert K. Krulee
(Northwestern University)
Englewood Cliffs, NJ: Prentice-Hall, 1991, vii+456 pp
Hardbound, ISBN 0-13-610288-3, no price listed

+Computers, chess, and cognition
T. Anthony Marsland and Jonathan Schaeffer (editors)
(University of Alberta)
NY: Springer-Verlag, 1990, xii+323 pp
Hardbound, ISBN 0-387-97415-6 and 3-540-97415-6, US$35.00

***Artificial intelligence: Concepts and applications in engineering**
A.R. Mirzai (editor)
(Polytechnic of Central London)
London: Chapman & Hall and Cambridge, MA: The MIT Press, 1990, xx+304 pp (MIT Press artificial intelligence series) Hardbound, ISBN 0-262-13256-7, US$39.95

+Experiments in the machine perception of motion
David W. Murray and Bernard F. Buxton
(University of Oxford and General Electric Hirst Research Centre) Cambridge, MA: The MIT Press, 1990, xix+236 pp
Hardbound, ISBN 0-262-13263-X, US$37.50

+Probablistic reasoning in expert systems
Richard E. Neapolitan
(Northeastern Illinois University)
NY: John Wiley, 1990, xiii+433 pp
Hardbound, ISBN 0-471-61840-3, cdn$65.50

The linguistics of punctuation
Geoffrey Nunberg
(Xerox Palo Alto Research Center)
Stanford: Center for the Study of Language and Information, 1990, ix+141 pp (CSLI lecture notes 18) (Distributed by the University of Chicago Press) Hardbound, ISBN 0-937073-47-4, US$35.00; paperbound, ISBN 0-937073-46-6, US$14.95

Intelligent databases: Object-oriented, deductive hypermedia technologies
Kamran Parsaye, Mark Chignell, Setrag Khoshafian, and Harry Wong
(IntelligenceWare Inc, University of Southern California, Ashton-Tate Corp, and Ashton-Tate Corp)
New York: John Wiley, 1989, xvi+479 pp
Paperbound, ISBN 0-471-50345-2, cdn$39.50

***Developing expert systems: A knowledge engineer's handbook for rules and objects**
Edmund C. Payne and Robert C. McArthur
(Albathion Software)
New York: John Wiley, 1990, xiv+401 pp
Paperbound, ISBN 0-471-51413-6, cdn$45.50

Computers and human language
George W. Smith
(University of Massachusetts, Boston)
New York: Oxford University Press, 1991, xiv+478 pp
Hardbound, ISBN 0-19-506281-7, no price listed; paperbound, ISBN 0-19-506282-5, US$16.95

Neural and concurrent real-time systems: The sixth generation
Branko Souček
(University of Zagreb)
New York: John Wiley, 1989, xviii+387 pp
Hardbound, ISBN 0-471-50889-6, cdn$65.50

Neural and massively parallel computers: The sixth generation
Branko Souček and Marina Souček
(University of Zagreb)
New York: John Wiley, 1988, xix+460 pp
Hardbound, ISBN 0-471-63533-2, cdn$89.95

The practice of Prolog
Leon Sterling (editor)
Cambridge, MA: The MIT Press, 1990, xxviii+312 pp
(Logic programming series)
Hardbound, ISBN 0-262-19301-9, US$37.50

Meaning and speech acts
Volume I: Principles of language use
Volume II: Formal semantics of success and satisfaction
Daniel Vanderveken
(Université du Québec à Trois-Rivières)
Cambridge, England: Cambridge University Press, 1990,
Vol~I: x+244 pp,
Vol~II: ??+??? pp
Vol~I: Hardbound, ISBN 0-521-37415-4, US$49.50
Vol~II: Hardbound, ISBN 0-521-38216-5, US$??.??

+Computer systems that learn: Classification and prediction methods from statistics, neural nets, machine learning, and expert systems
Sholom M. Weiss and Casimir A. Kulikowski
(Rutgers University)
San Mateo, CA: Morgan Kaufmann, 1991, xii+223 pp
Hardbound, ISBN 1-55860-065-5, US$39.95

***Artificial intelligence, simulation, and modeling**
Lawrence E. Widman, Kenneth A. Loparo, and Norman R. Nielsen (editors)
(University of Texas, Case Western Reserve University, and SRI International)
New York: John Wiley, 1989, xvii+556 pp
Hardbound, ISBN 0-471-60599-9, cdn$72.50

Computers and writing: Models and tools
Noel Williams and Patrik Holt (editors)
(Sheffield City Polytechnic and Heriot Watt University)
Norwood, NJ: Ablex Publishing and Oxford: Intellect Ltd, 1989, xi+166 pp Hardbound, ISBN 0-89391-636-6 and 1-871516-03-X, US$29.50

Advertisers Index

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

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 _____

_____

MORGAN KAUFMANN

PLANNING AND CONTROL
Thomas Dean and Michael P. Wellman

Sept. 1991; approx 400 pgs; cloth;
ISBN 1-55860-209-7; $44.95

The authors explore the connection between planning and control by reformulating the two areas in a common control framework, developing techniques side-by-side, and identifying opportunities for integrating their ideas and methods.

PARADIGMS OF AI PROGRAMMING: A COMMON LISP APPROACH
Peter Norvig

Fall 1991; approx 600 pgs; paper;
ISBN 1-55860-191-0; $39.95

"I have learned much that I didn't know...something that is relatively rare in the usual crop of basic Lisp texts...This [book] will find its way to the desk of many AI programming practitioners where Steele's "Common Lisp: The Reference" and Keene's CLOS book are currently the only 'required'. "
—Marty Hall (Johns Hopkins Continuing Professional Programs)

PRINCIPLES OF SEMANTIC NETWORKS: EXPLORATIONS IN THE REPRESENTATION OF KNOWLEDGE
Edited by John Sowa

May 1991; 563 pgs; cloth; 1-55860-088-4; $42.95

Representing the first major overview of this subject in 10 years, this book includes a wide range of work from formal analyses to implementations and applications.

MACHINE LEARNING: A THEORETICAL APPROACH
Balas K. Natarajan

May 1991; 250 pgs; cloth; 1-55860-148-1; $42.95

This presentation of fundamental results and their applications is intended for a general audience, and provides tools to classify what is and is not efficiently learnable, and to help construct efficient algorithms.

CONCEPT FORMATION: KNOWLEDGE AND EXPERIENCE IN UNSUPERVISED LEARNING
Edited by Douglas Fisher and Michael Pazzani

July 1991; approx 425 pgs; paper;
ISBN 1-55860-201-1; $39.95

Brings together results on concept formation from cognitive psychology and machine learning, including explanation-based and inductive approaches.

FOUNDATIONS OF GENETIC ALGORITHMS
Edited by Gregory Rawlins

August 1991; approx 650 pgs; cloth;
ISBN 1-55860-170-8; $45.95

This book addresses the need for a principled approach to understanding the foundation of genetic algorithms and classifier systems as a way of enhancing their further development and application.

REASONING ABOUT PLANS
James Allen, Henry Kautz, Richard Pelavin and Josh Tenenberg

June 1991; approx 400 pgs; cloth;
ISBN 1-55860-137-6; $39.95

Four substantial contributions to planning research are presented within an integrated framework.

KNOWLEDGE-BASED MACHINE TRANSLATION
Sergei Nirenburg, Jaime Carbonell, Masaru Tomita and Kenneth Goodman

July 1991; approx 250 pgs; cloth;
ISBN 1-55860-128-7; $39.95

The authors introduce the general paradigm of knowledge-based MT, survey major recent developments, compare it with other approaches and present a paradigmatic view of its component processes.

THE KBMT PROJECT
Edited by Kenneth Goodman and Sergei Nirenburg

June 1991; approx 400 pgs; paper;
ISBN 1-55860-129-5; $34.95

This book presents a detailed case study of a particular knowledge-based machine translation development effort.

REPRESENTATIONS OF COMMONSENSE KNOWLEDGE
Ernest Davis

1990; 544 pgs; Cloth; ISBN 1-55860-033-7; $42.95

"It is difficult for me to imagine anyone attempting to devise a representation scheme without having this book as a reference."
—Drew McDermott (Yale University)

"This is the first modern book on AI knowledge representation, a sophisticated, complete summary of 25 of years research with thorough introductions to the necessary background, replacing a library of other texts and research memoranda...There is no other book like it."
—Pat Hayes (MCC)

COMPUTER SYSTEMS THAT LEARN: Classification and Prediction Methods from Statistics, Neural Nets, Machine Learning and Expert Systems
Sholom Weiss and Casimir Kulikowski

1990; 250 pgs; Cloth; ISBN 1-55860-065-5; $39.95

"Computer Systems That Learn is an extremely well-written, comprehensive text that should be a valuable resource for those interested in applying computer-learning methods to practical problems...[the authors] clearly demonstrate how the discipline of statistical decision theory can be applied to computer learning methods to develop reliable estimates of learning system accuracy."
—Kenneth Marko (Senior Research Scientist, Ford Motor Company)

NEW 1991 PROCEEDINGS

AVAILABLE AUGUST, 1991

IJCAI 91–AUSTRALIA Proceedings of the International Joint Conference on Artificial Intelligence
1991; 2 volumes, ISBN 1-55860-160-0; $75.00/$56.25 members of national AI societies.

KR '91: Proceedings of the Second International Conference on Knowledge Representation and Reasoning
Edited by James Allen, Richard Fikes and Erik Sandewall

April 1991; 602 pgs; paper;
ISBN 1-55860-165-1; $34.95

Case-based Reasoning: Proceedings of the 1991 DARPA Workshop
May 1991; approx 500 pgs; paper;
ISBN 1-55860-199-6; $35.00

Advances in Neural Information Processing Systems, 3
Edited by Richard P. Lippmann, John E. Moody and David S. Touretzky

May 1991; 1130 pgs; cloth; 1-55860-184-8; $49.95

Machine Learning: Proceedings of the Eighth Annual Workshop (1991)
Edited by Lawrence Birnbaum and Gregg Collins

June 1991; approx 650 pgs; paper;
ISBN 1-55860-200-3; $34.95

Uncertainty in Artificial Intelligence: Proceedings of the Seventh Conference (1991)
Edited by Bruce D'Ambrosio, Phillipe Smets and Piero P. Bonissone

July 1991; approx 500 pgs; paper;
ISBN 1-55860-203-8; $34.95

Genetic Algorithms: Proceedings of the Fourth International Conference
Edited by Richard K. Belew and Lashon B. Booker

July 1991; aprox 650 pgs; paper;
ISBN 1-55860-208-9; $36.95

COLT '91: Proceedings of the Fourth International Workshop on Computational Learning Theory
Edited by Leslie Valiant and Manfred K. Warmuth

August 1991; approx 500 pgs; paper;
ISBN 1-55860-213-5; $39.95

TO ORDER IN THE US OR CANADA CALL TOLL FREE:

1-800-745-7323

Or send check, money order or Amex, VISA or MC authorization (acct. no., name-on-card and expiration date) to **Morgan Kaufmann, 2929 Campus Dr., Ste. 260, Dept. CN, San Mateo, CA 94403.** Include shipping & handling (Int'l. surface postal rate: $6.50 1st volume, $3.50 each add'l;). **Int'l. customers (not including Canada) add 10% to your order.** Fax orders: (415) 578-0672. Telephone: 415-578-9911.

AI System Development Tools

Soft Warehouse Inc
mu-Lisp: Lisp language programming enviornment. Turns a PC into an integrated workstation for developing AI software. $365. With compiler, $490. mu-MATH: Symbolic mathematics and calculus. From $305 to $365. Derive: Mathematical assistant for the PC. Solves large and complex problems. $245.

NeuralWorks Professional II
Advanced, easy-to-use "graphic-intuitive" approach results in a seamless interface. User-defined neurodynamics, and loadable control strategies provide complete freedom in designing multi-paradigm networks or developing completely new network types. IBM PC AT or PS-2 and MAC Plus, MAC Plus, MAC SE, MAC II $1,495. SUN/3 or SUN/4, $3,655.

Nexpert
Full scale expert system development tool for PCs with rich graphics support. Object hierarchy with multiple inheritance. Rules support pattern matching, integrated forward and backward chaining, automatic goal generation, non-monotonic reasoning. Open architecture allows integration with external programs. For use on IBM AT, 386, PS/2 and Macintosh II, $6,100. For use on DEC VAX 2000,II,III, under VMS/UIS, and UNIX workstations, $14,650.

N-NET
Combined neural/knowledge based industrial controller. N-NET™ 210 is a neural net, based on the Functional Link Net architecture, which may be embedded in expert systems. Use of N-NET™ 210 has shown quality improvement, cost reduction and efficient use of time in the areas of trend analysis, forecasting, process control, quality control, product design and analysis and machine diagnostics. Has been used in speech recognition. Supervised and unsupervised learning and associative recall are all accomplished within the same command and data structure. Runs on IBM PC/XT/AT, PS/2, or compatible. From $850. A VAX version is forthcoming.

Dragon Systems, Inc.
Dragon's VoiceScribe and DragonWriter voice recognition systems have been acclaimed the best voice recognizers for IBM XT, AT and 386 machines. They have constantly outperformed systems costing several times more. The design is based on a sound set of algorithms developed by two of the top researchers in the research community, Drs. Jim and Janet Baker. The devices are ideal for such applications as command/control, data entry/retrieval, form filling (hospital, government), documentation creation or dictation, front-ending of expert systems, and support systems for the handicapped. From $1,215 to $10,980.

Products and Services of Carnegie Group
Applied AI Systems, Inc distributes the products of Carnegie Group within Canada including their flagship product Knowledge Craft. Other products include TestBench, a troubleshooting shell for field or customer service and manufacturing applications, Operations Planner, a "what-if" application for manufacturing engineers who need to assess the impact of changes to operations before the changes are made; and Text Categorization Shell (TCS) for automatic content-based categorization of on-line text.

ACCENT SA
A high quality text-to-speech convertor converts ordinary ASCII text into intelligible speech. Can be plugged directly into any IBM PC, PC/XT, PC AT or compatibles. Can speak in two modes: the text mode and the spell mode. Speech output may be directly connected to a speaker. A standard RS-232C serial port is provided for interface with a computer or terminal as the host. Prices range from $485 - $1,760.

Arity Knowledge Systems Development Pack
A new generation AI prototyping tool for PCs. A prolog interpreter/compiler, a rule/frame-based Expert Systems shell with object hierarchy and inheritance, the external file and database access, a screen design kit, linkage to conventional languages, and an IBM SQL access, $1,315.

Smalltalk/V
Pure object oriented programming. Provides access to a rich tool set: bit-mapping graphics, a point-and-click user interface, multiple windows, and a source-level debugger. Runs on IBM PC/XT/AT/PS or compatibles with at least 512K memory and one of the following graphics boards: Hercules, CGA, EGA, VGA, AT&T 6300, Toshiba T3100 or IBM 3270. Add on features include: EGA/VGA Color Extension Pack, Communications Application Pack, and "Goodies" Application Pack. Price ranges from $125 - $245

R*TIME
Tool for realtime AI system development originally developed for advanced aerospace applications. Full-colour graphic MMI updates application-dependent displays in realtime to display data and results of inference. Distributed architecture allows multiple copies of inference engine, MMI, data sources, and simulated signal sources (for testing) on several workstations. Available on Sun 3, Sun 4, and SPARC, DEC VAX and Risc and IBM 386/486 (SCO and Interactive UNIX).

VIP (Voice Interactive Package)
Complete voice recognition and synthesis package for developing a multi-modal interface to applications developed using NEXPERT Object. Speaker dependent voice recognition of up to 1,000 commands and unlimited number of output texts. Especially valuable in "hands busy, eyes busy" situations and where access to the system is limited. $6,000. A version that interfaces MS-DOS directly under development.

Symbologic Adept
Ideal tool for building expert systems for realtime procedural problems. Graphical user interface makes working with the knowledge simple and intuitive, allowing procedures to be represented graphically, as a network of nodes and arcs that resemble a flowchart. Takes full advantage of the multitasking, graphics, and windowing capabilities of Microsoft Windows. Runs on IBM PC AT and up.

DENDROS-1G Evaluation Board
All-analog, electronic ASIC and later VLSI implementation of a neural network which learns and categorizes in real-time. The dynamics of learning and sleep can be observed. $850.

ENQUIRY AND ORDER: Call (613) 592-3030 or send Purchase Order to Applied AI Systems, Inc. Shipping and Handling charges and all applicable sales taxes apply. Applied AI Systems, Inc is an authorized dealer of the above AI software/hardware products. Prices are subject to change without notice.

Applied AI Systems, Inc.
Suite 500, Gateway Business Park
340 March Road
KANATA, Ontario, Canada K2K 2E4
Telephone: (613) 592-3030
Facsimile: (613) 592-2333