Canadian Artificial Intelligence
Intelligence Artificielle au Canada

February 1991 No. 26 février 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

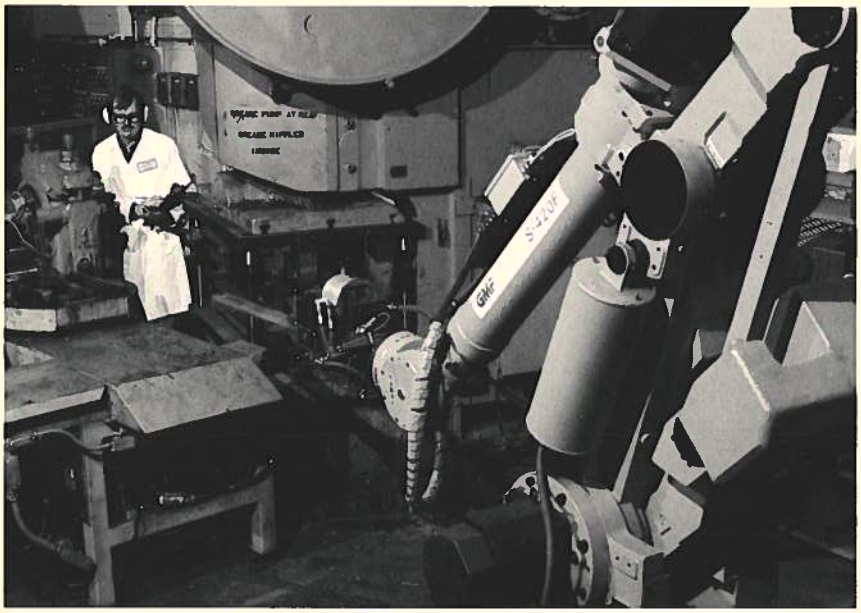

Robotic workcell at Standen's Ltd., Calgary performing hot forging tasks for manufacturing leafsprings

Photo courtesy Alberta Research Council

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

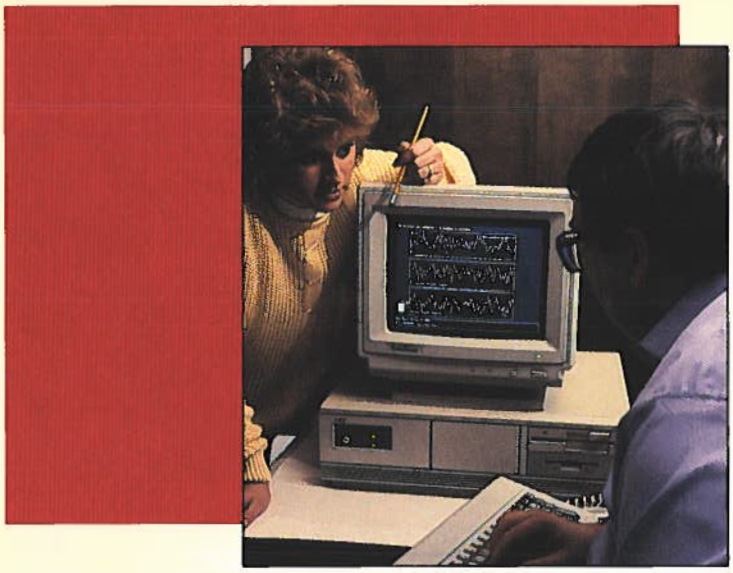

NEURALWARE, INC.

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

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

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

The sensible approach to neural network development.

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

User-Defined Neuro-Dynamics

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

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

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

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

APPLIED AI SYSTEMS, INC.

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

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

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

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

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

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

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

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

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

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

NeuralWare
INCORPORATED

Phone or write for more information

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

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

February 1991 | No. 26 | février 1991

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

COMMUNICATIONS
COMMUNICATIONS

Roy Masrani

Canadian Artificial Intelligence is brought to you each quarter by a dedicated team of volunteers who, in addition to their full-time jobs, put in many hours of work to get the magazine out to press. Our expansionist plans include a huge empire with newswires constantly churning out up-to-the-minute stories of AI action in Canada; reporters, on location (or in hotel rooms) fighting to get their copy past the section editors; phone calls from all over the world asking for our opinion on recent newsbreaking advances on the Canadian AI scene; thousands of letters to the editor lambasting our recent undercover operations; constant streams of reprint requests from second rate newspapers like the Globe & Mail and the New York Times . . . .

But in the mean time . . . we cannot afford the newswire services just yet, and are still waiting for our international bureau chiefs to submit interesting articles (most notably from Europe) . . . .

Nonetheless, the excitement and joy of seeing yet another issue of *Canadian Artificial Intelligence* in print is hard to contain, even when we get the occasional letter blasting us for the typos. All the scrambling to get the articles to print is more than offset by the great feedback we keep getting from our colleagues (most of whom have promised to send us a letter expressing this joy in writing).

You too can be involved. How, you ask? Simple. Just let me know. We are desperately in need of French translators and reporters. Although we can translate the magazine into numerous languages, we are having difficulties finding people to volunteer their time to translate to French for us, or even submit an article on the goings-on in Quebec.

And reporters? Haven't you always wanted to be in a meeting where people all of a sudden clam up because "there are members of the press here"? And you have to say "Is this for the record?". Don't you have a burning desire to say "Hold the presses" – or, more appropriately, "Interrupt the laser writer" – to get in a last minute news item? Well, this could be your chance. We may even throw in a reporter's hat and a press badge. All you have to do is keep your eyes and ears out for interesting news items, write them up and submit them with your photograph.

So . . . don't delay. Send me a note.

Your (volunteer) Editor
Roy Masrani

Announcements

Notice: Membership List

In the Summer of 1991, the Canadian Society for the Computational Studies of Intelligence (CSCSI), the American Association for Artificial Intelligence (AAAI), and the Association for Computing Machinery's Special Interest Group on Artificial Intelligence (SIGART) will release a joint membership directory. We hope to include the members of the Mexican AI society in the directory so that we have a directory for all the North American AI societies.

The purpose of this directory is not only to provide a listing of the names of the societies' members but also to serve as a general information resource about North American AI research centers, AI companies, and universities and colleges with AI studies. If you wish to include your company, university or college, or research center in this list, please send address, telephone, and contact information to me, preferably by electronic mail, to arrive before 31 March 1991.

We have also had requests from companies that would like to use our mailing list for advertising. For example, Morgan Kaufmann would like to send catalogs of new books to CSCSI members; others have contacted us for similar reasons. We will certainly not supply the mailing list indiscriminately but only in cases where we feel that what is offered is a genuine service to members. However, if you do not want your name to be used for these purposes, please send me email to that effect as well.

Peter F. Patel-Schneider
Secretary, CSCSI
AT&T Bell Labs
600 Mountain Avenue, Room 3C-410A
Murray Hill, New Jersey 07974 U.S.A.
pfps@research.att.com

Toronto CIPS Special Interest Group on Knowledge Based Systems

The Toronto chapter of CIPS has formed a special interest group in Knowledge Based Systems - SIGKBS. The objective of SIGKBS is to promote the understanding of Knowledge Based Systems to the Toronto community served by CIPS through presentations and discussions with leading experts in the field. An emphasis will be placed on practical issues of successfully deploying KBS in a variety of areas such as finance, manufacturing, etc.

SIGKBS is run by six co-chairpeople: Lew Baxter, who initiated the SIG, and Walter Berndl, Gordon Graham, Paul Mackay, Peter Priest and Andrzej Taramina. Meetings are held from 6:30 to 8:00 on the third Thursday of each month from September to May (except December). So far presentations have been given by Bill Tatham of Andersen Consulting on KBS Architecture and Development, Susan Gardner of The Herongate Group on Identifying KB Opportunities in Business, Denis Dorval of Eclin Systems Corporation on IFAX: A customer information KBS, and Rainer von Konigslow of Cognex on KBS in Manufacturing.

If you are interested in attending, or if you wish to make a presentation please contact Lew Baxter at (416) 598-0196. SIGKBS meetings are free to CIPS members; alternatively there is a fee of $45 to join SIGKBS or a fee of $10 per meeting.

Yours sincerely, Lew Baxter

Dr. Noel Lazure

A Banker of a Different Breed

Penelope Rokeby

Did you know you can go to a bank in Canada and withdraw, not money, but words?

Dr. Noel Lazure is a terminologist with the Government of Canada's terminology bank. He collects words. More specifically, he collects words related to artificial intelligence. His book, Vocabulaire de l'intelligence artificielle/Artificial Intelligence Vocabulary was published in 1988 and has sold over 800 copies. While it is used primarily as a tool for translation of artificial intelligence terminology between French and English, it is also useful for anyone from a scientist to a secretary who may require a reference book in artificial intelligence.

Dr. Lazure began work on the book in 1985, after research he was doing on an extensive Starwars lexicon. The Starwars research made him realize that artificial intelligence was large enough to merit a complete and separate lexicon of its own, and so he began work on the current volumes. Working with English as the source language, Dr. Lazure studied AI concepts in English, beginning by reading AI- related journals. This gave him a good idea of the scope of the field, and he was able to build on this with further research. He then turned to the French literature to find the same concept in French. He discovered a few notions that were specific to the English language, and some that were only found in the French, and so his own translation skills were put to use. For example, the term "chunking" did not have a parallel term in French, and so Dr. Lazure coined the term "memorisation par bloc." The book deliberately avoids robotics terminology because of an agreement between the Government of Canada and the Quebec Government to allow the Quebec government to continue work on a robotics vocabulary already underway.

Canada's terminology bank was founded in 1972 at the University of Montreal, and in 1975 the Government of Canada bought it. Dr. Lazure was hired as a terminologist in 1976 from Marquette University in Milwaukee, Wisconsin, where he had been a professor of religious studies for nine years. His background was in the Bible, and he studied in Ottawa, Rome and Jerusalem. With a total of 18 years experience as an academic in Canada and the U.S., his research skills were deemed a valuable asset to the terminology bank. As well, it is clear to anyone speaking with him that he is a man who possesses considerable energy and great curiosity about a variety of subjects, so one can understand how he could move with relative ease between such diverse topics as eschatology and heuristics.

Canada now has the largest group of terminologists of any country in the world, as well as the world's largest databank on terminology. The Secretary of State Terminology Bank is available on CD-ROM for a cost of $1,100 which includes three updates per year. More information can be obtained by phoning (819) 997-2067. In addition, the government operates a phone-in service which can be accessed by dialling (819) 953-6793. The terminology bank contains over two million entries in every field: sociology, economics, medicine, as well as computer science.

Dr. Lazure is currently doing research for future books on radar and laser terminology as well as a lexicon for the computer language known as ADA, and a future volume of the AI vocabulary which will be expanded to 25,000 entries from the current 12,000. He loves his work, explaining that he finds researching different topics interesting. The projects he works on are all related: he never starts from scratch anymore, which is why he has one little problem. For over thirteen years, he has been collecting information on little cards – he has over 150,000 of them now, and at the age of 67, he figures it's just too much trouble to enter all that data into a computer.

Canadian AI System hits US Market

RTES, the Real Time Expert System application development tool, previously available only in Canada, is now distributed in the US by Knowledge Systems (Schiller Park, IL). RTES uses a data-driven (forward chaining) inference engine and features a very fast development environment. Its intended use is for manufacturing automation and process control. The Canadian developer, Real Time Systems Inc. (Toronto, Ontario), also sells the Knowledge Processor, a backward and forward chaining inference tool, but this is still available only in Canada. The Knowledge Processor is designed to assist with decision-making for the manufacturing, medical and financial industries.

RTES combines rule-based knowledge representation with data acquisition and control techniques. The RTES shell uses rules that define the intended actions of the operations they control, and can be entered as rules of thumb, logical relationships, or mathematical calculations. On-line editing of rules can also be performed without recompiling and without stopping the monitoring or controlling of the equipment. "You could compare it to Lotus or dBase, in terms of the development cycle," said Frank Simtob, the shell's developer. RTES can also apply its decisions to the process under its direction directly, without a human operator. Canadian companies have installed between 100-200 RTES-based applications.

One of Simtob's clients put together a process control application for gem manufacturing in less than a month and installed it in a week during a scheduled holiday plant shutdown. Ontario Hydro has also developed RTES-based applications, using it to build complex substation control systems and small generating plant control systems. They can produce an application in a month to a month and a half.

"The reason development goes so fast is because rules are generated on-line; there is no editing process," Simtob explained. "When you run the shell, you can create rules, edit them, and test them at the same time." Users can add rules to build a knowledge base, and immediately find out how well the rule works in the context of the others already generated.

Simtob recently automated a grenade manufacturing plant in Mississippi, and arrived with only a sketch of what the knowledge base would contain. The rule generation was performed on-line and on-site. The entire system was developed and commissioned in three weeks. It runs on a network, with modules scattered throughout the plant, an approach Simtob felt made more sense than trying to put all of it in one place.

Development

Simtob began developing what would become RTES in 1975. He was only able to work on it part-time, as he was an employee of another company at the time. The first release of RTES ran on a Data General minicomputer in 1979, and the first application built was a production monitoring system for the US Army. With the arrival of the IBM PC in 1981, Simtob switched to PC development in 1983 and had a version for sale in late 1984.

The Knowledge Processor was developed from 1985-1987. Simtob has sold copies to Constellation Assurance and Prudential Assurance, two large Canadian insurance companies, but otherwise sales have been very disappointing. Constellation was one of the early supports of KP, with an initial bulk purchase, a yearly maintenance fee, and an unlimited site license.

Simtob is disappointed with the sales of Knowledge Processor, since it is the more robust solution of the two tools. KP is more of an expert system development tool for applications that will interact with humans, while RTES is directed at process automation. Given the success of RTES in Canada, he is hoping for similar results in the US, but he concedes that marketing AI is difficult. Recognition can take time and money to establish. "RTES was not successful when it first came out," said Simtob. "The typical reaction after giving a demonstration would be 'Well, what does it do?' after spending a couple hours showing it." KP uses the same multitasking features of RTES, but is able to backward and forward chain simultaneously. This gives KP quite a bit of speed, since the forward chaining will allow many questions to be answered simultaneously, while the backward chaining asks only one question." Its domain of application is insurance companies, particularly underwriting.

The interfaces of both products are not graphically very strong, Simtob admitted, and he is investing development effort into strengthening that. The problem is the graphic interface and the actual processing both compete for CPU cycles, meaning a colorful interface may distract the system from the job at hand.

Other applications

Intelimotion was developed by UTDC (Kingston, Ontario), and uses RTES as a central control shell managing a group of linear induction motor-driven passive vehicles. These vehicles are dispatched from production points to warehousing, and all control is run by RTES over a network. Intelimotion sends commands from its central processor to the local controllers and directs traffic, detects possible

malfunctions, and displays real-time system status, the location and current load of all vehicles.

WHE Process Systems (Weston, Ontario) is using its Bio-Trol application as an adaptive control system for a new type of vortex fermenter. All the user need do is give the machine the correct information about the substance to be processed, and Bio-Trol configures the controls accordingly.

Canmet (Elliot Lake, Ontario) is using RTES to check for the presence of radon and thoron gases in a radiation-monitoring application. The RTES application manages a network of monitors and a variable-speed aerosol pump to maintain the desired concentration of radon/thoron, through algorithm-directed trend prediction. The pump is regulated to dispense the gases as needed, based on information from the monitors and its own algorithm-derived results.

The city of Cornwall, Ontario, is using an RTES-based system to monitor its sewer system. A series of transducers sends readings to RTES, which converts them to hydraulic pressure, and displays the status of the line, taking into account the geometry of the system. It produces an alarm when conditions are abnormal, and enters a recording mode when the situation becomes critical.

RTES sells for CAN$3,200, and the Knowledge Processor sells for CAN$4,500. RTES runs on DOS-based microcomputers, and can perform minicomputer-like multitasking with its own multitasking software. Simtob recommends at least an AT, but a 386 or 486-based platform is not required. RTES can also be interfaced with data acquisition and control subsystems from Allen Bradley, Siemens, Modicon, Mitsubishi, Square D, and others.

Reprinted with permission from Intelligent Systems Report, November 1990

News Release

ISTC Minister announces federal funding for artificial intelligence development projects

OTTAWA, January 16, 1991 - Industry, Science and Technology Minister Benoît Bouchard today announced funding totalling $3.9 million, from the Artificial Intelligence Research and Development Fund. The projects will be undertaken by a number of federal departments and agencies on a cost-shared basis, with the major portion of each contracted-out to firms in the private sector with capabilities in artificial intelligence.

"I am particularly pleased that many of the projects will create new opportunities in the areas of environmental protection and training. By taking the lead in using advanced technology in these and other priority applications, the government will be a test bed for risky developments which have the potential to both improve the effectiveness of its operations and enhance industrial competitiveness," Minister Bouchard said.

Taken together with the first group of projects announced by the Minister in April, 1990, $6.3 million has now been allocated from the $10 million to a total of 20 active projects. This will result in about $13 million being devoted to leading-edge applications of AI technology. It is expected that over the five year life of the program, total project investments in these new technologies will be well over $20 million, because of the leverage exerted by the Fund.

The Artificial Intelligence R&D Fund is a procurement-based program which costshares projects within the federal government which offer high potential payoffs in both the public and private sectors. Sponsoring departments and agencies are responsible for the management of their own projects, with the contracting-out conducted according to standard government procedures. The resulting technologies are diffused to the private sector to the maximum extent possible.

The Fund is a component of ISTC's Strategic Technologies Program, which assists Canadian industry to respond to the competitive challenges being brought about by rapid technological change in biotechnology, advanced industrial materials, and information technology. "It is an example of the Department's new approaches in working with industry to build a competitive economy", the Minister added.
A summary of each of the projects is below.

ARTIFICIAL INTELLIGENCE

Artificial Intelligence (AI) is a set of computer-based technologies which seeks to approximate performance considered to require human intelligence. In most cases, AI is being used because it offers new ways to tackle complex and time-consuming problems which cannot be solved by conventional computer-based methods. AI also provides a means of capturing the knowledge of experts, especially those with specialized or hard to find experience in solving problems within a particular domain of knowledge. Other applications of AI include computerized systems for vision or speech understanding and robotic systems which can function in dangerous environments. The following is a description of the projects supported under the second round of the AI R&D Fund in cooperation with federal government departments.

PROJECTS CO-SPONSORED BY THE AI R&D FUND

"Expert Patent Search Assistant"

Sponsor: Automated Systems Branch
Intellectual Property Directorate (Patents)
Consumer and Corporate Affairs Canada
Contact: Mr. Ray Taylor
(819) 997-2186

The objective of this project is to create a PC-based, natural language interface that will make it far easier for non-experts to access Canada's patent documentation. The interface will enable people who are not familiar with either the patent classification scheme or database terminology to extract useful classification information from the Index, and, in turn, lead to greater exploitation by Canadian industry of the wealth of technological information inherent in patents.

"Development and Demonstration of a Comprehensive Computer-based Pollution Control Planning System"

Sponsor: Wastewater Technology Centre
Environment Canada
Contact: Dr. Tom Constable
(416) 336-4716

Canadian municipalities are currently spending billions of dollars to maintain and upgrade sewage collection and treatment systems because of aging, environmental impairment and municipal growth. These systems are very complex, and it is often difficult to determine the causes of their problems and to identify the most appropriate control measures. Extensive data collection and analysis studies are often undertaken to determine the causes, but these studies can be unnecessarily expensive unless a systematic approach is used prior to their initiation to eliminate inappropriate control measures and identify specific data needs.

The purpose of this project is to develop a computer software package which will allow systematic planning and evaluation of municipal sewage collection and treatment systems. One of the key components of the package will be an expert system to manage and integrate the various mathematical models and data requirements of the entire system, and to recommend data collection needs and appropriate control measures. The software will help municipalities to improve the performance and extend the life of existing collection and treatment facilities, and to realize major capital cost savings in the construction of new facilities. The software will be tested using data from the collection system and treatment facility of the City of Port Colborne.

"Spill Emergency Artificial Intelligence System"

Sponsor: Environmental Emergencies Technology Division
Environment Canada
Contact: Mr. Merv Fingas
(613) 998-9622

There is at present, no unified program to evaluate spills and predict their consequences and ecological effects. Poor decision making at the time of spills results in lost time, greater health and safety risks, environmental damage, and more extensive cleanup costs. The project will lead to the development of a system for predicting the effects and consequences of oil, chemical, and other spills resulting from most accidents. It will result in reduced cost of spill cleanup both in terms of financial and human resources and in improved health and safety for the public, as a result of improved response and corrective actions. The system will also be used to train Environment Canada and provincial government employees who respond to spill situations, enhancing their ability to respond to environmental emergencies.

"Automated Recognition and Prediction of Meteorological Features"

Sponsor: Atmospheric Environment Service
Environment Canada
Contact: Mr. Don Dueck, Project Manager
(416) 739-4944

The science of meteorology relies heavily on "conceptual models" of the atmosphere to organize, explain, and predict its behaviour. These models incorporate many complex physical processes within the atmosphere. This integration of many processes into a single model or feature is an essential ingredient in the forecast process. In order to do so, the forecaster must evaluate and integrate many large and diverse data sets of surface weather observations, weather balloon ascents, radar and satellite imagery, and complex computer model outputs, often dealing with contradictory and incomplete information. The project will lead to the development of a system that will allow the forecaster to do a more complete analysis of the available data, which at the same time will remove the variability arising from the assessments by different forecasters, which are dependent upon their personal abilities and knowledge. It is intended to provide more rapid and accurate forecasts of threatening weather conditions, especially explosively deepening marine storms. Improved forecasts also support the management and exploitation of renewable and non-renewable industries such as forestry, agriculture, fishing and mining, in which weather information and forecasts are an integral part of daily operations. The system will also help to train new meteorologists.

"An Environmentally Sensitive Investment System (ESIS)"

Sponsor: Federal Environmental Assessment Review Office
Contact: Mr. Patrice J. LeBlanc
(819) 997-2253

The increasing demand to include the economic and social impacts, as well as the biophysical impacts of proposed development projects during the Environmental Impact Assessment (EIA) of proposed development projects has led to a need to assess trade-offs between environment and

economics. This has resulted in increased difficulty in determining the financial viability of a project or program. It is important today to ensure that the full cost of compliance with requirements established during environmental assessment is evaluated. The broadened scope and coverage of EIA increases the amount of knowledge and information required to ensure proper financial and environmental planning. The ESIS project will lead to the development of a system that permits the integration of the full cost of requirements associated with environmental assessments, directly into financial investment analysis on an industry-specific basis. The system could be used by government to simulate the appropriateness of environmental requirements and to assess industry responses to regulations and guidelines as well as environmental strategies. The private sector will be able to use the system as an integral part of their investment analysis and planning. FEARO, with Dalhousie University, UMA Engineering Limited, the Pulp and Paper Research Institute of Canada (PPRIC), as well as other organizations from various sectors, will cooperate in this initiative. Future developments could lead to similar applications in other industry sectors.

"Radio Interference Advisor"

Sponsor: Engineering Programs Branch
Department of Communications
Contact: Mr. Glen Lockwood
(613) 990-4800

To manage the radio spectrum, Department of Communications (DOC) radio inspectors across the country perform about 15,000 radio investigations per year to resolve radio interference complaints from the public. The inspectors diagnose the cause of the radio interference using knowledge gained through formal education, specialized in-house training, and experience gained on the job. Given the large variety of problems encountered, each of some 150 inspectors has some unique expertise. The project will lead to the development of an expert system to consolidate the collective expertise of DOC inspectors in the domain of radio interference diagnosis. Initially, the knowledge base will be built on existing case histories. Over time, new case histories will be added to the knowledge base making the system more expert. The Radio Interference Advisor will be used as a tool for improving the training of new inspectors, and for sharing experience and knowledge gained by one inspector with all inspectors, thereby increasing their productivity.

"Intelligent Investigators Assistant"

Sponsor: Informatics Directorate
Royal Canadian Mounted Police
Contact: Mr. David Gordon
(613) 993-1521

The project will lead to the development of a system to build and enhance a knowledge base for investigation. The system will store and use specialized investigative and legal knowledge, and apply the knowledge base of legal precedent and investigative procedures, to assist in supporting criminal charges. On an ongoing basis, the system will capture the experience of experts and retain that knowledge for all investigators. The system will also be used to train police investigators. The successful development of the technology for this application will demonstrate that it could be applied in other situations where decision-making is based upon large amounts of data and complex rules.

"Expert System for Control of Hot Rolling of Steel"

Sponsor: Metals Technology Laboratories
Energy, Mines and Resources Canada
Contact: Dr. R. Thomson
(613) 995-9541

The need to meet stringent property specifications, conserve energy and to reduce scrap rates is critical for the competitiveness of Canadian steel companies in world markets. An expert system capable of storing and analyzing data will be a powerful tool for the design, optimization, and control of rolling operations, to produce products with specified properties. In addition to improving product quality, the system is expected to result in significant savings of both time and money. Ultimately, such a system will be capable of predicting rolling schedules for various steel grades to obtain specified properties, and would be used on-line to control and modify production operations. The project will be carried out in collaboration with Canadian steel companies.

"Photo Interpretation Keys Expert System"

Sponsor: Canada Centre for Remote Sensing
Energy, Mines and Resources Canada
Contact: Dr. David G. Goodenough
(613) 952-2760

Photo-interpretation, or human interpretation of imagery acquired from aircraft and space platforms, is used widely in Canada by governments and industry, for map making, resource inventorying, land-use studies and environmental monitoring. Photo-interpreters, because of the high level of skill required, are in short supply. As well, even experienced photo-interpreters often vary in their interpretations. This project will lead to the development of an expert system that will result in increased speed and accuracy of photo interpretation. The system will be built on the experience of expert photo-interpreters, and will shorten the time needed to train photo interpreters. This project is expected to enhance the competitive position of Canadian industry in the rapidly growing international remote sensing market - estimated at $20 billion.

"EXCLASS - Expert Job Evaluation Assistant"

Sponsor: Human Resources
Treasury Board Secretariat
Contact: Mrs. Lise Ouimet
(613) 952-3150)

The objective of the project is to develop an AI-based system to support the job evaluation function, including the processes of position description, position analysis and classification as they apply within the federal government. The project will involve the development of an expert system using knowledge-based technology to evaluate standardized job descriptions and provide the user with a recommended classification and an explanation of the reasoning used in the process. Natural language technology will be applied to support the interactive entry of job specification details and for the generation of standardized job descriptions which will be "understood" by the system.

The project offers very significant potential for high productivity gain within government operations. The overall cost to the federal government of the job evaluation function has been estimated at about $50 million annually. The system could ultimately lead to savings of $25 million per annum. A consortium of departments led by the Treasury Board Secretariat will oversee the development of the technology and its application.

"Trade Ref: Artificial Intelligence Project in International Trade Regulation"

Sponsor: Customs and Excise
Revenue Canada
Contact: Mr. John Shearer
(613) 954-6990

The objective of the project is to develop an expert system that will apply state-of-the-art AI technologies to the tariff classification process. The project will also develop a prototype document assembler, which will assist in the generation of correctly completed forms that accompany import/export transactions. The system is expected to simplify and reduce the costs of importing and exporting, thereby increasing the competitiveness of Canadian companies. It is also expected to be of particular assistance to importers in determining the correct tariff classification and, consequently, dutiable status of imports, and in particular, in determining whether goods qualify for the benefits of the Canada U.S. Free Trade Agreement. The project builds on earlier work carried out by the non-profit Legal Information Systems and Technologies (LIST) Foundation.

"Expert System Advisor for Air Traffic Conflict Resolution"

Sponsor: Transport Development Centre
Transport Canada
Contact: Mr. Yves Le Borgne
(514) 283-0024

Air traffic control activities provide for safe and efficient aviation by preventing airspace conflicts between users. Traditionally, air traffic controllers have used radar, communications systems and computers, and the identification, prediction and resolution of potential air space conflicts lies with the air traffic-controller. The project deals with the application of AI and Expert Systems to conflict prediction and resolution advice. The project is expected to result in improved efficiency and capacity of aviation movements in Canadian airspace, and to provide tactical assistance to controllers and produce consistent and safe procedures under all situations. The system will also be used as a training tool for automated self-instruction.

"Intelligent Instructor Tutor"

Sponsor: Canadian Workplace Automation Research Centre
Department of Communications
Contact: Dr. Jocelyne Picot
(514) 686-1990

The objective of this project is to develop an AI-based system, the Intelligent Instruction Tutor, which will advise and assist trainers in the design, delivery and evaluation of training programs. The three major modules of the system are the Training Design Advisor, for assistance in designing instructional interventions, the Training Delivery Advisor, from which to develop a pedagogical model of how to train, and the Training Evaluation Advisor, which is intended to help trainers in evaluating training systems and diagnosing student learning. The system will also include multi-media components such as video, a library of cases, databases of relevant information and simulations. The system is intended primarily to effectively train trainers, but may also serve as a job aid for more experienced users. The project reflects the importance of training to innovation and competitiveness. The improvement in the quality of training that is expected to result will become a model for training practices in other sectors of the Canadian economy.

"Intelligent Tutoring Systems Strategy"

Sponsor: Quality Assurance Program
Public Service Commission of Canada
Contact: Mr. Ian Jackson
(613) 957-2730

The objective of this project is to develop an overall intelligent tutoring systems strategy and technology development program for use as a training pilot within the federal Public Service. The project recognizes the major need in Canada for greater efforts in the development of training technology, and will use the federal government as a catalyst and test-bed to develop new knowledge-based training systems in conjunction with the private sector. Potential opportunities will be identified in which AI could be the basis for addressing a number of specific training requirements, and a plan of action will be established.

"Automated, Integrated Inspection System"

Sponsor: Food Production and Inspection Branch
Agriculture Canada
Contact: Mr. Jon Hutton
(613) 992-2114

The objective of this project is to develop a proof-of-concept prototype to demonstrate how meat and poultry inspection can be automated. An automated detection, decision, disposition, system using AI for inspection is expected to have an impact on, and provide the opportunity for: improving the cost performance of slaughter inspection; eliminating poor working environments for inspectors; enhancing inspection detection capabilities; increasing slaughter productivity; and reducing the unit cost of meat processing.

Software products industry sector campaign launched

MONTREAL, October 29, 1990 – "A strong, vibrant software products sector can provide the foundation for future high technology development in Canada," Mr. Benoît Bouchard, Minister of Industry, Science and Technology, said today announcing the launch of an ISTC-sponsored sector campaign for the software. The Minister made the announcement at the annual meeting of the Information Technology Association of Canada (ITAC).

This latest initiative, which builds upon the Software Agenda for Action announced by the department in February 1988, is aimed at building an internationally competitive software products industry in Canada.

"We already have many software products companies with world-class capabilities, developed through strong linkages with our post-secondary institutions. The initial thrust of this initiative is to better understand the environment under which Canadian companies can continue to build upon their successful record in what is recognized as a major growth sector in advanced economies," the Minister said.

A key element of the Campaign will be the formation of a Special Advisory Panel, comprised of representatives from industry and academia, to provide advice and guidance during the campaign process.

The studies which will be undertaken as part of the Campaign activities will attempt to isolate the conditions which determine the success of a software products company, focusing on such issues as management, financing, human resources and marketing. Particular attention will be paid to the business environment of the U.S. industry, as it is the acknowledged world leader, and Canada's major market and competitor.

The information gained through these studies will assist ISTC and industry in developing further special initiatives, targeting particular constraints and opportunities facing the industry.

Graeme Hirst

Planning the future of natural language research (even in Canada)

Graeme Hirst
Department of Computer Science,
University of Toronto, Toronto, Ontario MSS 1A4

At CSCSI's Eighth Canadian Conference on Artificial Intelligence, held in Ottawa in May, Graeme Hirst was invited to give a talk on natural language processing. Following is an abridged version of his text.

For this talk, I've tried to include something to offend everybody. I will cover:

- What research in natural language has done – and is likely to do.
- What it could do for Canada.
- Why it hasn't done it – yet.
- What we need for it to happen.

Like other parts of AI, Canadian natural language research has made significant contributions to the field. But many of our best people have left the field, or gone south. Natural language doesn't have the small but distinguished nucleus in Canada that some other AI subfields have.

Like other parts of AI, Canadian natural language research has been relatively lucky compared to many other sciences. AI is still a "glamour field", even if not as glamorous as putting a Canadian in space. But natural language hasn't been quite as lucky – it's had little support from CIAR, for example. (I won't speculate on the cause/effect relationship between this and the lack of a distinguished nucleus.)

That's all rather unfortunate. Natural language research could be important for Canada, as it already is in Europe (a point I'll come back to shortly) – perhaps even more important than putting a Canadian in space.

Graeme Hirst is an Associate Professor of Computer Science at the University of Toronto, where he is head of the computational linguistics research group.

1. Where NLP is now

I'd better start by defining the area that I'm going to cover, with an apology for the terminology. "Natural language processing", or "NLP", or "computational linguistics", is the subfield of AI that's concerned with the use of human languages – natural languages ('NL' for short). That means, in particular, building systems that can deal with the structure and content of language the way a human would – not just processing language the way a word processor, say, does. The field is sometimes called "NLU", "natural language understanding", but I'll avoid that term because I don't want to limit myself, or the field, just to the *comprehension* of language.

NLP, perhaps even more than other parts of AI, is an incremental science. That is, it's not given to sudden breakthroughs or elegant theorems; rather, it is an accumulation of smaller ideas, techniques, and formalizations that together build up a system. This reflects the nature of language itself. Language has evolved to take advantage of many different parts of human cognition in its operation. Language understanding draws on many different kinds of knowledge. It seems to use both hardware in the brain that's specialized for language and also the brain's more general-purpose hardware. Language pervades cognition. So it's hardly a surprise that language understanding programs often look like an agglomeration of bits and pieces.

NLP has come a long way since the 50s and 60s – thanks in part to the co-evolution of modern linguistics. Not only has the development of modern syntactic and semantic theories strongly influenced NL research, but the concerns of NLP for process-oriented theories of language have had a strong influence on theoretical linguistics. The two fields

are now closer than ever before. The same is true of psycholinguistics. Computational work on NLP has influenced the development of psycholinguistic models of how people process language, and experimental work on how people understand has influenced the development of techniques in NLP.

Let me give some examples of where the field is today. We now have fairly large, robust systems for parsing – not just toys – that are, for example, used in corpus studies in which large bodies of text are parsed or processed. Semantic analysis is also doing well. We have new approaches to semantic interpretation, to knowledge-based ambiguity resolution, and to fitting compositional semantics into unification-based systems. We are rapidly learning to use on-line dictionaries and reference works. There's an awful lot of useful knowledge in such books. They were originally intended for use by people, but methods that enable NLP programs to use them are a matter of current research. Language generation and user modelling are starting to come into their own. These used to be just small areas within NLP. In general, there is a greater emphasis in the field on "real" language use. That is, we are worried about actual texts written by actual people, and not just silly artificial examples like "John saw the spy with the telescope".

Applications are filtering into the world:

- Grammar checkers, although presently modest, are starting to appear for personal systems.
- NL database interfaces, both for large systems and for personal systems, are now widely available.
- News- and message-routing systems, which use some understanding of the content of the item to decide where to send it, are in regular production use.
- There are many machine-aided translation systems, some quite sophisticated, now available.
- Intelligent computer-assisted language teaching, although lagging behind the other areas, is starting to produce some interesting systems.

In general, we can expect to see much greater use of NL in information systems – both for communication with the systems, and in the knowledge base of the systems themselves. NL research applications will have use in just about any domain or industry that uses language or information systems, and that's just about everywhere and anywhere that communication takes place.

2. Who is doing NLP research?

To think about the potential benefits of NLP research for Canada, let's look at what's happening in other countries. Although the basic research is largely centred in the U.S., much of the action in applied research and development is in Europe and Japan.

Europeans have always been more conscious of language than North Americans – even Canadians. Any educated European is assumed to speak two or three languages. (Interestingly, the exceptions to that seem to be the English and the French.) Many European countries are bilingual or multilingual – Belgium, Switzerland. Many have languages that are spoken by few outsiders – the Netherlands, Denmark, most central European countries – so it is essential to learn the more widely spoken languages. And the EC, which has nine official languages, and the single European market coming in 1992, are making an emphasis on language even more important.

Thus Europe has what are called "language industries", while we in North America hardly even have the term. There's an emphasis on translation, and computerization of all aspects of language – writing, translating, managing multi-lingual documentation. The EC is supporting Eurotra, a huge machine translation project.

Likewise, in Japan. The Japanese Fifth Generation project includes a large amount of NL research. Enormous effort is going into machine translation. And the Japanese Electronic Dictionary project is, I understand, more far-reaching than any comparable Western project. And language also features prominently in the applications proposed for the successor project in massively parallel computing.

3. What could have happened in Canada?

This is all rather relevant to Canada. We also have a large need for translation in Canada. And it's expensive and there's always a shortage of qualified translators. We have a great need for teaching languages in Canada. And not just teaching a second official language, but also, for many immigrants, a first official language. Anything that improves translation or language teaching – anything that can make it easier, anything that can make it cheaper – is surely of benefit to Canada, both economically and socially.

And, on a larger scale, language is central to the future economy of Canada. We keep hearing about how Canada has to lower its dependence on a resource-based economy. About how so much of the world economy will be – is already! – based on information and services. About how the typical worker of the future will be a "knowledge worker". Well, language is how people represent knowledge.

So if computers and automation are to be involved in this knowledge work – and I think we'd all agree that for efficiency they must be – then it would be a good idea if the computers involved could process language. And process it not just as text, like a word processor does, without regard to its meaning, but process it as a repository of knowledge.

Even in artificial intelligence, people seem to have lost sight of this. People rightly see AI as having great potential, and are willing to spend time and money on work in expert systems and formalisms for knowledge representation and reasoning. But what they've forgotten is that knowledge comes from people, and is for the benefit of people. So we need to also be concerned with people's knowledge representation formalisms – natural language – as well as those for computers, and we need to worry about mapping between the two kinds of representation.

After all, there's a vast amount of knowledge out there in the world. About 10 to the minus 87 percent of it is presently in a form suitable for use in any AI system. Of the other 99.9999 percent, a good slab, maybe half, is in natural language. The rest is in people's heads in a non-linguistic form, and when it comes out it does so in the form of actions or, again, language.

So for all these reasons, NLP research could be of great benefit to Canada. What's more, we could have been a leader!

The TAUM METEO project (at the University of Montreal, in the 70s) was a world leader in machine translation. I've found that the project and the people who were on it still command enormous respect in the MT community – outside Canada, that is.

In the latter part of the 1970s, Canada was almost unique in having good AI people, good MT people, and, as it happens, a significant number of good people working in computer applications in the linguistic humanities. With a base of such people working together, Canada could have been a leader in machine translation, in intelligent computer-assisted language teaching, in multi-lingual processing, in NLP in general. But Canada isn't any of those things. Why?

Funding was withdrawn from the TAUM group because MT wasn't found to be immediately cost-effective in the short term! How incredibly short-sighted – and typically Canadian. We had something good and we blew it. What should have been thought of as basic research was evaluated as if it were product development.

And, in general, Canada couldn't – or at least didn't – match the resources and opportunities available to researchers in the south.

4. Current work

So what are we doing in NL research now in Canada? I can't be a spokesperson for other groups, but I'll briefly mention some of the projects we have worked on at Toronto in the past five years.*

One of our projects is what we've called "theoretical machine translation". We don't have the resources to work on real MT projects, but some of our theoretical work has been explicitly directed towards developing ideas that could be used by someone else in machine translation systems.

In prototype systems for intelligent instruction in a second language, again, our group doesn't have the resources to be building large, complete CAI systems, let alone developing courseware and evaluating the systems with real language learners. So again we have concentrated on theoretical work directed toward application in such systems.

In knowledge acquisition as a problem in language understanding, we have been considering what's involved in knowledge acquisition for knowledge-based systems just as a manual task in which an analyst – or "knowledge engineer" as we say these days – has to convert an expert's utterances into some formal knowledge representation.

And besides all this, we've also been doing basic research in various other aspects of NL, including the special problems of knowledge representation for language understanding.

5. Prospects for the future

Well, so far I've told you that NL research is steaming ahead in the rest of the world, while Canada missed the boat. But some of us are still pushing ahead, rowing as hard as we can. The question is, what we'd need in Canada in order to catch up with the boat.

What we need are several NL projects, at universities or elsewhere, that would be big enough to have an effect: to get researchers together, to do pre-competitive research (as it's called these days), to show what can be done – with enough infrastructure to get the job done properly, to train graduate students, and to give the students somewhere to work when they graduate.

I'm thinking of projects comparable to, say, the Center for Machine Translation at Carnegie Mellon University in Pittsburgh. The Center has long-term goals in MT applications, and pursues both those applications and basic research directed toward them. It has external funding, provides a place for grad students to learn and work, and has an active visitors program to promote the exchange of ideas. Canada, unfortunately, has little tradition of large project centres like this at universities.

Ideally, we'd want several such projects. We don't want to get the whole country committed to just one group's paradigm or one group's approach.

What are the obstacles to realizing this?

- A lack of money.
  Large projects need not just researchers, but space, equipment, programmers, assistants. Building a "real" MT system is 90% software hackwork.
- A lack of momentum, a long up-to-speed time, and already being quite a way behind.
- A lack of industrial and government interest.

We all know the sad story of the state of R&D in Canada. It's so well known that even Maclean's had an article on it last week [21 May 1990] – "Canada's spending on research is falling far behind that of its competitors" – with the obligatory picture of the Canadarm, as if that's the only thing Canadian science has ever produced. We spend a much smaller fraction of our GNP on R&D than other industrialized nations. And the present government has repeatedly broken its promises to change that. Indeed, the NRC itself has a doubtful future. And in Ontario and elsewhere, university funds have been cut.

The federal and Ontario governments and the CIAR have initiated their respective "centres of excellence" programs – which are better than nothing, but by their nature are an admission of failure of the research funding system. And

---

**An overview of NLP research at the University of Toronto is available from the author upon request.*

neither "centres of excellence" programs nor NSERC strategic grants, as they presently work, are intended for the kind of bootstrapping that I'm talking about here of a research field that's at the pre-pre-competitive stage. Centre-of-excellence programs imply existing momentum and "excellence"; strategic grants imply existing industrial interest.

And even where there is some interest in AI, there seems to be little appreciation of NLP research.

Could it still happen? Could we develop NLP research in Canada? What are our resources?

- A few NLP people in universities and elsewhere. And many more who could be repatriated, including those who have moved to other sub-fields of AI. Good people in other areas of AI and in computing in the linguistic humanities. And lots of keen, good students.
- A few companies that could apply some of the research.
- Larger R&D companies that might be induced to support long-term research, like companies such as Bell Labs and Bellcore do in the U.S. And perhaps CIAR might be too. But, what else do we need?
- We, as a natural language community, need to effectively communicate our visions and our commitment.
- Active interest from government and industry to support this vision and commitment: A realization that there could be long-term economic benefit from support of NLP research.
- Money.

But we have to be wary of the problem of inadequate half-measures. That is, being offered just a little money, and then written off when we fail to perform miracles on a shoestring.

6. Conclusion

You might wonder why I'm telling you all this. It sounds a bit like a grant proposal. Perhaps I just want a sympathetic ear. I think we're doing some good NL research in Canada, but we're not getting the resources we need.

Questions

Why were you critical of special programmes giving funding for "centres of excellence"?

Centres-of-excellence programmes are like food banks. Almost by definition, they are an admission that our research-funding system has failed! – Not supporting adequately the best people in the field, so a special band-aid programme is needed to bring their resources up to (near) adequacy – to what a proper system would have given them in the first place. The U.S. doesn't need such programmes! To the extent centres of excellence exist in the U.S., e.g., MIT, Stanford, they tend to be automatic by-products of the system.

Now, such programmes are certainly welcome as better than nothing. But for each Centre of Excellence we need half-a-dozen Centres of Okayness. A research field doesn't move just by supporting the few superstars. They can't do all the work by themselves; other competent researchers need adequate funding too.

It's not all sweetness and light in the U.S., you know.

That's true. In fairness, let me say what we have in Canada that's better than in the U.S.

The main thing is the way NSERC operating grants work. They permit research that is open-ended and curiosity-driven, not project-oriented. The money (what little NSERC has to distribute) is "more real" than that from U.S. grants because no overheads, salaries of principal investigators, etc, are taken out; but of course, that means that those costs have to be paid from other funds. NSERC grantees are not dependent on making their work suit the military or government policy fashions the way DARPA grantees are in the U.S. And with NSERC, one is able to spend more time on the research and less just writing grant proposals.

The only other advantages we have in Canada seem to be cleaner streets and some kind of general moral superiority.

Instead of setting up any kind of rival project, perhaps Canada or Canadian researchers should ask to be included in Eurotra or other projects?

That might be nice. But we'd have to have something special that they would want enough to be willing to make us a co-developer rather than a customer. It's not clear that we do. And I would imagine that participation in Eurotra would require some pretty special high-level agreements between the Canadian government and the EC.

I think you've been unnecessarily hard on the government. While it's true that our R&D spending is low, the government's share is nearly as big as that of other countries. The shortfall is due to industry not doing its part.

But why isn't industry doing its part? Because it has little incentive to do so. Because research is done at companies' head offices in the U.S. Because there's no tradition of any need for much research in what was once a more resource-based economy – just dig it up or cut it down and ship it out. Because we don't have the tradition of venture capitalists willing to gamble on small leading-edge research-based companies.

It *is* the government's job to try to change this. After the debacle of the Scientific Research Tax Credit program, one can understand a reluctance to try something new. But they did promise that they would. And they haven't.

You say you are doing "theoretical MT" work. Perhaps that's what our contribution to the world should be. We can do this sort of thing well exactly because we aren't distracted by large projects requiring real deliverables!

Even if that's true, it would still be nice to have a bit more recognition and support for the work. If Canadian government policy is to be simply that Canadian research, to the extent it exists at all, is just our altruistic gift to the world of science for others to exploit at will, then what can I say? It seems to me that perhaps we can be just as altruistic *and* benefit more ourselves.

ACADEMIA
ACADÉMIQUES

Jan Mulder

Academia editor's notes

As researchers and teachers of AI in Canada, we are probably all aware that the AI job market in Canada, especially in academics, is currently not good. This is of major concern to those of us who are supervising graduate students who are to be released into the job market during the next two years. It was this concern that caused me to announce a new initiative in the November issue of *Canadian Artificial Intelligence*, namely the regular publication of a list of AI graduate students in Canada providing people's names, affiliation, thesis topic or title, supervisor, and (optionally?!) expected completion date.

If you have read the November issue of *Canadian Artificial Intelligence* then you will have noticed a major new initiative with the objective to broaden *Canadian Artificial Intelligence's* readers and contributors base. Not only has an attempt been made to get regular contributions from all sources in Canada where AI research and development takes place, we also have undertaken a major membership drive to increase our membership especially in the industrial sector. It is my hope that the regular publication of an AI graduate student list, apart from making an informative contribution, will facilitate the communication between students and potential employers.

If you are currently supervising graduate students then you (and your students) can contribute to this initiative by sending (preferably e-mail) me a list of student names, affiliation, supervisor, thesis topic, and expected completion date. I will assume that all information I receive is given with both student and supervisor consent and no information other than what I receive from you will be published.

I would like to publish a first list in the May issue. A speedy response would therefore be appreciated. If for one reason or another you are reluctant to participate in this initiative, then please let me know as well.

A new initiative started in this issue is the publication of titles and abstracts of AI tech reports published at Canadian universities. The listing in this issue is somewhat meager. I am still in the process of setting up communication lines with computer science departments across Canada.

This issue contains two contributions from academia: a report on the Robotics and Intelligent Systems Centre at the University of Saskatchewan and an abridged version of the invited address given by Graeme Hirst at the Canadian Artificial Intelligence conference in Ottawa held in May of last year. In that address Graeme shared his concerns about the state of natural language research in Canada with the conference participants. I would like to share his concerns with you as well.

On Minimax Game Tree Search Pathology and Node-Value Dependence

Liwu Li and T.A. Marsland
Computing Science Department, University of Alberta, Edmonton, Canada T6G 2H1 – August 1, 1990

Abstract

Here we are concerned with the mysterious phenomenon of minimax game-tree search pathology, where it appears and how it happens. It is commonly believed that the strategy of searching-deeper for computer game-playing programs can enhance the accuracy of position evaluation and increase the possibility of detecting the correct move. The strategy has been successfully implemented in practice; but it has not been justified theoretically, and various previous investigations of the phenomenon were based on uniform trees whose terminal node values are independent from each other. Those trees fail to account for the apparent relationship between the node values for common games and, therefore, cannot be used to explain this pathology. Here, we present a new method to introduce node-value dependence into board-splitting games and relate the pathological phenomenon to the dependence. In particular, we study the effect of minimax searching-deeper for the games by using an evaluation of invariant accuracy with respect to the search depths; we also examine the relationship between the quality of minimax search and the node value dependence of the game trees by assuming a real evaluation function described by Nau (8). The results of both approaches reveal that the pathological phenomenon is related to weak node value dependence and confirm that searching -deeper is effective for game trees with strong node-value dependence.

Robotics and Intelligent Systems

Of the 14 centres of excellence networks, this one is likely the most diverse. This diversity is reflected by the U of S participants, all of whom are in the College of Engineering: Professors Jean-Paul Tremblay, Gordon McCalla, Jim Greer, and Eric Neufeld, of the Department of Computational Science; Madan Gupta, director of the Intelligent Systems Research Laboratory; and Jim Wilson, chairman of the Division of Control Engineering.

The Computational Science participants, led by Professor Paul Sorenson of the University of Alberta (and until 1989, of the U of S), are supporting the development of large software systems. These are required for such complex undertakings as the control of airline reservations, industrial inventories, and space stations.

Teams of software engineering specialists, with as many as several dozen members, develop the system. But to do so they need specialized sets of computer tools, which vary markedly from one system to another.

This is where the Saskatchewan and Alberta computational scientists come in. They are working to devise a computer system (metasystem) that will assist in generating sets of tools tailor made for particular application areas. They are also developing computer-based training systems to help the software specialists learn how to generate and use the tools.

Artificial intelligence (AI) techniques are being applied to overcome some of the difficulties in generating these diverse sets of tools. With AI, the computer can make the software engineers' work easier by ensuring that necessary ground rules are not violated. This requires a "reasoning" capacity similar to that of humans.

AI techniques are also being used to develop intelligent training systems. These integrate knowledge from a variety of contexts with strategies for imparting it in a flexible and responsive manner. The objective is to have the computer "advise" the operator on how to make the most complete and consistent use of the metasystem.

Professor Gupta and his associates are working on the development of machine vision. The objective of their computer-based systems is to emulate the visual processes of the human eye and brain. Potential applications of their research are far reaching, including assistance to the sightless, improvements to industrial robots and to remote sensing technology, and support for space and ocean exploration.

The key is to translate the biological functions of the human eye into mathematical language computers can understand. But the work is challenging and difficult because of the complexity and efficiency of human vision, which is much more than simply seeing. Vision involves neural functions of the retina and cerebral cortex and encompasses such things as perception, memory, feedback, and decision making. This is what enables human

vision to distinguish color, texture, shape, location, in relation to the surrounding environment.

His colleagues in the research group include Dean Peter Nikiforuk, Professors Hugh Wood (Electrical Engineering) and Kailash Prasad (Physiology), two research scholars, and 12 graduate students.

Professor Wilson is the principal investigator for two projects of importance to the agriculture industry: 1) a computer-vision system to assist in guiding field implements and 2) a sensor and control system to monitor and automatically adjust combine harvesters for optimum performance.

The computer-vision system will be developed to detect the demarcation between successive field operations. It will show the farmer, perhaps by means of a simple gauge and pointer, the extent of any overlap, or of any missed strips, during such operations as tillage, seeding, or harvesting. The gauge will indicate steering adjustments that should be made to eliminate overlap or missed strips and enable the farmer to reduce energy, labor, seed, and maintenance costs.

By monitoring the field at the edge of the implement, the system detects differences in surface reflectivity that indicate the demarcation between successive operations.

It's ultimately intended to incorporate the vision system into the steering system. This will provide a degree of automatic control in making appropriate steering adjustments.

To accommodate different crops, as well as crop conditions that vary within and between fields, combines have been provided with a variety of manual adjustments. Professor Wilson's second project will develop sensors for measuring changes in performance that result from these adjustments.

Among the key variables for which adjustments are required are feed rate, cylinder loss, and straw moisture content. The rate at which crop is fed into the combine, for example, affects the amount of threshed grain lost in the straw and chaff effluent. An optimum balance is maintained by adjustments that aim at maximizing feed rate while minimizing grain loss.

Sensors will enable farmers to operate the combine more efficiently. Eventually, the sensors will be incorporated into a control system which makes adjustments continuously and automatically.

Reprinted with permission from The Green & White, University of Saskatchewan Alumni Association Magazine.

AI Capabilities at Thomson Systems

Dr. Serdar Kalaycioglu

1.0 General AI Capabilities

Thomson Systems recognized the great potential of AI in the development of complex systems and in 1987 made a serious commitment by establishing a program to develop AI capabilities.

Presently, there are five people in this field with extensive backgrounds in the company and this number will increase substantially in the near future as a result of continuous in-house training efforts.

One of the major efforts in this field at Thomson Systems is technology transfer from universities and research centres. To start with, a survey of the Canadian capabilities in the field of AI was performed. Then, some visits were made to establish links with appropriate institutions and several information-exchange seminars were held. In order to successfully implement technology transfer, various collaborations took place with, for example McRCIM (McGill University Research Centre for Intelligent Machines) and PAMI (Pattern Analysis and Machine Intelligence of University of Waterloo), etc.

Currently, Thomson Systems has the following tools:

- PC-Based AI Tools
  - Rulemaster
  - Prolog
  - Smalltalk V/286
  - Smalltalk-80
- HP-9000 Based AI Tools
  - Ojectworks-80
  - Humble

2.0 Research and Development Work

Thomson Systems has initiated some research and development work in AI in order to strengthen its capabilities in this field. In the following, these investigations are briefly described:

- **Real-Time AI**

Thomson-DCS (the parent company) is involved in the development of a parallel transputer-based inference engine – Parallel Associative Development Machine as a vehicle for AI (PADMAVATI). Thomson Systems is transferring this technology to Canada and also is identifying new application areas in the field of AI.

PADMAVATI is a tool kit for designing a parallel transputer-based system suited for time critical applications, with numeric and symbolic computing facilities in a real-time context. PADMAVATI is composed of a high performance processor - 20 MIPS which uses transputer links and a dynamic network with add on features up to 5,000 MIPS. It also supports high level languages (such as LISP, PROLOG, C++, ADA).

Among the various identified applications are pattern recognition, speech recognition, system identification, diagnosis, trouble shooting, data fusion, threat evaluation and weapon assignment, simulation, task planning, intelligent control, and decision making.

- **AI and ADA**

Thomson Systems is developing a strong real-time Ada capability on 286 machines operating under OS/2. To support this, the Ada Inference Engine was developed and forward and backward chaining were incorporated. Recent work investigates the uncertainty factors associated with the tasks to be carried out and efforts will be made to incorporate these uncertainties.

- **Neural Networks**

Thomson Systems has developed a Hopfield Network for Pattern Recognition. This system is useful especially for classification problems (radar target recognition), optimal routing/scheduling, prediction problems (e.g. collision avoidance) and intelligent control.

- **Functional Analysis Case Tool (FACT)**

Thomson Systems has developed a case tool called "FACT" which is a Turbo-Prolog based program designed specifically for the task of functional analysis (Functional Definition/Requirement Definition Tool). FACT is an expert system based tool and tracks all the inputs and outputs in an analysis and ensures consistency throughout.

- **Computer Aided Maintenance System (CAMS)**

Thomson Systems has initiated a research and development project to develop a Computer Aided Maintenance System (CAMS) to guide operator/technicians through the fault-finding/diagnosis and the maintenance cycle.

The system will contain an expert system capable of diagnosing a specific piece of equipment, and then planning repair actions. In both processes, the expert system will require the co-operation of the human maintainer to carry out physical actions on its behalf.

CAMS could be hosted on a micro computer (e.g. a 386

based machine) with multi-media capabilities. It would be capable of accessing multi-gigabytes of laser disk storage, display high resolution images, and synthesize speech or other sounds. In turn, it would be operable by maintainers through voice, keyboard or simple controller inputs.

CAMS's knowledge representation will address the complexities inherent in any diagnosis problem. Its reasoning algorithm will utilize a fault-tree analysis. The user will be able to interact wih the system through an effective user interface which supports the display of contextual advice, schematics, hypertext-based glossary information and the database of the specific parts such as a combat system or a jet engine, etc.

3.0 AI Contracts

- **Application of AI to Air Defence**

This contract was awarded to Thomson Systems in 1988 by DND/DREV to investigate the applicability of AI technologies to future AAW systems (e.g. NATO Anti-air Warfare Systems). In this study, Thomson Systems identified the components of a modern AAW system which are best suited to AI in order to support the following requirements:

- reliable detection using a variety of sensors;
- processing of a large amount of information;
- automation to meet short reaction times.

Comparisons of the merits of AI relative to more conventional technologies were made and the recommendations for the implementation of AI in NAAWS were also provided as results of this investigation.

- **TEWA Expert System**

This contract was awarded to Thomson Systems in 1989 by DND/DREV to develop a "Threat Evaluation and Weapon Assignment System" Expert System Simulator. In this study, Thomson Systems developed an expert system which prototypes the environment and testbed that supports experimentation with TEWA concepts. Assessment of the viability of AI technology in building a TEWA system was also made.

The developed testbed environment for the evaluation of TEWA concepts and designs incorporates all elements and attributes of an AAW combat environment.

The TEWA Testbed System consists of two main components: the TEWA Stimulator and the TEWA Simulator. The purpose of the TEWA Stimulator is to produce disk files which emulate the output of the ship's sensor data fusion process. These files contain a series of coherent data frames which will ultimately be used to stimulate the TEWA function with information about the scenario that is unfolding in the airspace around the ship.

The TEWA Simulator, which comprises the second component, will read and process the files produced by the Stimulator and will render judgement, in accordance with its built-in strategies, as to which weapons to use against the current threats to the ship. The evaluation and decisions made are logged and an expert system explanation facility can be used to discover the reasoning that led to a particular TEWA evaluation/decision.

As a part of the TEWA Expert System contract, a complete AAW simulation has been specified. Such a simulator would support the development and evaluation of TEWA algorithms modeling the complete air defence process. This simulator will incorporate AI technology including Expert System Techniques.

- **Intelligent Task Planning**

Thomson Systems was awarded a contract to define and develop a concept for a limited or supervised autonomous robotic system (MSS Autonomous Robotics System) by the Canadian Space Agency in 1988. As a part of this study, Thomson Systems identified several strategic technologies to be developed in the second phase of this program. Intelligent Task Planning was identified as one such strategic technology. The concept definition and the survey of this technology were performed during the first phase of this program.

In the Second Phase of this program, Thomson Systems and PAMI (of University of Waterloo) will develop an AI-based Intelligent Task Planning System which will be responsible for the decomposition of high level mission goals into low level tasks in the temporal as well as spatial domain. The human operator and the robot manipulator will be considered as agents and assigned to certain task elements in order carry out the decomposed low level tasks in a coordinated fashion.

4.0 Concluding Remarks

Thomson Systems is looking forward to applying its expertise in AI and Expert Systems to various applications such as Anti-air Warfare Systems, Trouble-shooting, Fault Analysis and Diagnosis, Robotics, Automation, Simulation and others.

Thomson Systems appreciates both the importance and the applicability of this technology and is committed to develop more capabilities in this field by involvement in other contracts, possible collaboration and technology transfer.

Space Robotics at Thomson Systems

Introduction

Thomson Systems believes that the Space Station Freedom (SSF) is the most significant international space project of this century and the largest international technology development project ever undertaken.

Canada is an important partner in this great venture and will develop the Mobile Servicing System (MSS) - a

moveable "workshop" containing special tools, the Remote Manipulator System, the Special Purpose Dexterous Manipulator, as well as vision equipment. Development of the MSS will stimulate the advancement of strategic technologies such as automation, robotics, machine vision and artificial intelligence.

To begin with, the STEAR (Strategic Technologies in Automation and Robotics) program was created in support of Canada's involvement in the Space Station Freedom Program.

The STEAR Program has two main objectives: advancing strategic robotics and automation technologies for the MSS; identifying terrestrial spin-off ideas, developing Canadian industry in these areas and capturing significant returns for investment.

Thomson Systems' Role

Thomson Systems has made a serious committment to be part of STEAR program and to apply its systems expertise to robotics in Space. Space Robotics Programs and their terrestrial spin-offs represent the initial committment of Thomson Systems to a future in space. This future will be intellectually rigorous, technically demanding, and commercially profitable.

MSS Autonomous Robotics Program - Phase I

Thomson Systems' first involvement in STEAR Program was through participation in the MSS Autonomous Robotics Program - Phase I for the Canadian Space Agency in 1989.

Thomson Systems has developed a Supervised Autonomous Robotic Concept and prepared a plan for full development of the Supervised Autonomous Robotic System Concept at the end of Phase I. During this program, the functional requirements were defined, a control hierarchy and a technology breakdown structure were developed, and assessments of technologies were made in terms of Canadian industrial capability as well as the strategic importance to Canada. Also, a Strawman Task - Orbital Replacement Unit (ORU) Change will be used for the maintenance of the Space Station. This task will be the basis for proof-of-principle demonstration in the following phases.

The McGill University Research Centre for Intelligent Machines (McRCIM) collaborated with Thomson Systems during the first phase and the project was completed on time with excellent results.

MSS Autonomous Robotics Program - Phase II

Thomson Systems was recently awarded the Second Phase of the MSS Autonomous Robotics Program. This phase will take two years in order to develop the defined Supervised Autonomous Robot Concept and the identified strategic technologies (telerobotics-telesensing considering time delays and bandwidth restrictions). At the end of Phase II, there will be some basic proof of principle demonstrations for the developed concept and the technologies. Thomson Systems will also develop a plan and proposal for full implementation of the concept in Phase III. A "Go or No-go" decision for Phase III will be made at this stage by the Canadian Space Agency Project Office.

Thomson Systems has an excellent team in this Program all of whom are recognized as leaders and world class experts in the field of robotics. Thomson Systems Robotics Team is comprised of the University of Waterloo (PAMI), McGill University, the University of Ottawa and the Alberta Research Council (ARC). The team is strongly committed to bring this program to successful completion.

Main Achievements From This Program

There is very close collaboration and understanding among these team members which provides an excellent technology transfer between companies and the universities. Thomson Systems and the STEAR Program will benefit by the leverage that is obtained by this technology transfer.

Thomson Systems, as a result of this program has identified some spin-off ideas and developed a commercialization plan to develop these ideas. This opportunity opened new doors to Thomson Systems to apply its expertise and capabilities. One of these spin-off ideas is "Computer Aided Maintenance System (CAMS)". This system will utilize intelligent task planning and human-system interaction technologies which will be mainly transferred from the MSS Autonomous Robotics Program. Our preliminary market survey results show that there is a great deal of interest in such an application. Thomson Systems is pursuing this opportunity and looks forward to some excellent results.

Thomson Systems developed a generic hierarchical functional system architecture for both fully autonomous and supervised autonomous robots which can be applied to any telerobotics system for dexterous tasks. Also, a supervised autonomous robotic concept which has a great potential for providing effective automation and supervisory control capabilities for hazardous, difficult and dexterous tasks has been developed. The great advantage of this concept is that the concept can be applied to an unstructured environment where there is no "a priori" information available. Some parts of this work were presented at the Canadian Conference on Electrical and Computer Engineering and also published in their proceedings. The developed "Supervised Autonomous Robot Concept" will be presented at the Canadian Astronautics and Space Institute Conference.

General Capabilities

Thomson Systems Robotics Team has conducted research and has excellent capabilities in the areas of: Dynamic Modelling, Dual-arm motion coordination, Force-moment accommodation, Motion control, Path/trajectory planning and Object avoidance, Automated task planning, Hierarchical control and Telerobotics.

Concluding Remarks

Thomson Systems views the Robotics, Automation and Artificial Intelligence technologies, and programs such as STEAR which stimulate these technologies, as being of strategic importance to its existence. Thomson Systems has made a serious investment of personnel and resources to pursue R&D in these areas and is looking forward to applying its robotics, automation and AI capabilities to provide Systems Engineering solutions to challenging problems.

Robotic Vision Technology

Thomson Systems is involved in a technology development project as a subcontractor to VIRTEK to provide technical and marketing support in the field of robotic vision. Currently, the proposal for the Second Phase of this program is under evaluation with the Canadian Space Agency.

During this program, VIRTEK and Thomson Systems will investigate the following problems:

- Automatic real time determination of object position and attitude (for closed-loop control of manipulator systems);
  - design of photogrammetry targets and lighting schemes for reliable tracking in a space environment
  - development of schemes to automatically identify known target arrays in cluttered scenes
  - integration of additional sensor information, e.g. LIDAR range data, into a photogrammetric solution
- Three dimensional surveying of the robotic work site in order to build up a computer world model database of the surroundings (for collision avoidance and path planning purposes);
  - further development of laser scanning, depth profiling imaging systems
  - further development of stereo processing of images from CCTV camera pairs
  - further development of structured light CCTV systems
  - development of schemes to identify a known object in a depth profile image
  - development of schemes to build up world model data bases from depth profile images
- Object identification using natural and special features of the object (for automation of robotic operations).
  - identification of high contrast markings or features in a general visual scene
  - identification of known objects in a visual scene
  - interpretation of special coding or marking (e.g., bar codes) for object identification

At the end of First Phase, VIRTEK and Thomson Systems:

- proposed a development plan for some of the above mentioned vision technologies;
- identified some terrestrial robotics and automation applications where the proposed vision system concept can be used; and
- developed a commercialization plan.

The MSS Robotic System will make use of many CCD cameras, laser scanners, wrist force/moment sensors, and tactile sensor, etc. One of the challenging activities of this program will be the incorporation of aspects related to the sensor data fusion problems into the MSS robotic system. Therefore, there is a great deal of analogy, support, and technology transfer between the Robotic Vision Technology Program and the Multi-Sensor Data Fusion Project.

MSS Autonomous Robotics - Phase II

Thomson Systems has recently been awarded a contract to develop a ground-based telerobotics system to control robot manipulators (Special Purpose Dexterous Manipulator -SPDM on the Mobile Servicing System of the Space Station Freedom). This work is a continuation of the first phase of the same program. Thomson Systems will design and develop a supervisory autonomous robotics system (telerobotics) considering tele-link time delays and bandwidth restrictions. This system will consist of a ground control station, a remote station with robot manipulators and the communication interfaces. The system will be used to implement a strawman task "Orbital Unit Changeout" for the Space Station Freedom.

Some of the technology development activities are listed below:

- Supervisory Control System;
- Dual-arm Coordination:
  - kinematic and dynamic coordination of multi-arm robotic systems;
  - trajectory control and collision avoidance;
  - optimal load distribution for two manipulators; and
  - hybrid position / force and impedance control for multiple manipulators.
- Development of Robot Skills; and
- Development of a Simulator.

The Thomson System's Team consisted of the University of Waterloo - PAMI, the National Research Council - IIT, McGill University, the University of Ottawa, and Alberta Research Council.

In Phase III of this Program, Thomson Systems will do a system integration to implement the developed strawman task for the proof-of-principle demonstration. The sensor data fusion technologies will play very important roles to provide a robust telemanipulation capabilities in this project.

Trajectory Planning and Object Avoidance - Phase I

Thomson Systems has been awarded a contract by the Canadian Space Agency to develop and demonstrate the application of techniques in trajectory planning, object avoidance and object contact in a robot system (especially a

multi-arm robot system) which provide a real-time capability to automatically generate manipulator paths around objects in the workspace, including moving objects, and the capability to be able to automatically modulate the forces of contact between the end-effector and payload in a stable manner.

The identified techniques in this phase, will be developed to provide robust automatic trajectory planning and collision avoidance capabilities for the Special Purpose Dexterous Manipulator on the MSS (Space Station Freedom). Real-time collision detection and avoidance significantly depends on the on-line sensing/perception capabilities. These capabilities depend on the techniques of interpretation and process of the sensory data. Thus, sensor data fusion is an important part of this project which will determine the robustness of the system.

Dr. Kalaycioglu is a Senior Systems Engineer for Thomson – CSF Systems Canada Inc.

Research Report: Machine Learning Applied to Industrial Process Planning

P. Turney and A. Famili
Information Technology Section, Laboratory for Intelligent Systems, Division of Electrical Engineering
National Research Council Canada, Ottawa, Ontario, Canada KIA OR6
PETER@AI.DEE.NRC.CA FAZEL@AI.DEE.NRC .CA

May 8, 1990

The function of a process planning system is to determine the methods by which a product is to be manufactured economically and competitively. In a modern manufacturing environment, a process planning system consists of highly trained people and complex software. The plans prepared by a process planning system are not always executed as planned. The aim of the project described here is to design and implement a computer program which will help human planners in a process planning system to understand why plans fail.

Introduction

Industrial process planning is the translation of design and production information into instructions for production of either parts or complete products. The plans prepared by a process planning system are not always executed as planned. The system must monitor the execution of its plans, learn from unsuccessful plans, and replan the process. It must use what it learned to improve future process plans (Wilkins, 1988). The process planning system must analyze a number of plans, both successful and unsuccessful, to find patterns in the failures of plans.

Abstractly stated, the problem is to take a set of events which have been divided into two classes (successful and unsuccessful plans), and find a way to characterize the division. Machine learning can be applied to this problem. Michalski (1986) has provided a taxonomy of machine learning techniques. The relevant technique for this type of problem is *learning from examples*. In learning from examples, a learning system is presented with a set of examples *(positive* examples) and counterexamples *(negative* examples) of a concept (the *target concept)*. The learning system then attempts to induce a general concept description. The aim of the project described here is to design and implement a computer program which will assist human planners in a process planning system, by helping them to understand why plans fail. To this end, a program called IMAFO (Intelligent Manufacturing Foreman) has been written, in Common Lisp. There are several approaches to learning from examples. IMAFO uses *decision tree induction*, because this approach has been extensively studied and is relatively well understood. This paper discusses some of the design decisions involved in building IMAFO.

2 The Role of IMAFO

Hierarchical planning is central to information processing and decision-making in industrial environments (Albus *et al.*, 1981)(McLean, 1986). IMAFO is designed to advise a planner working near the middle level of the hierarchy.

Figure 1 illustrates the flow of information between IMAFO and the manufacturing planning and decision making hierarchy.

IMAFO takes *job commands* (plans from the level above), *task commands* (plans sent to the level below), and *execution results* (reports from the level below) as input. A simulator (written in Common Lisp) provides test data for IMAFO. IMAFO produces a report of the patterns it finds in the plan failures.

3 The Design of IMAFO

IMAFO contains two main modules. One module performs *initialization* and the other *analysis*. Initialization consists of

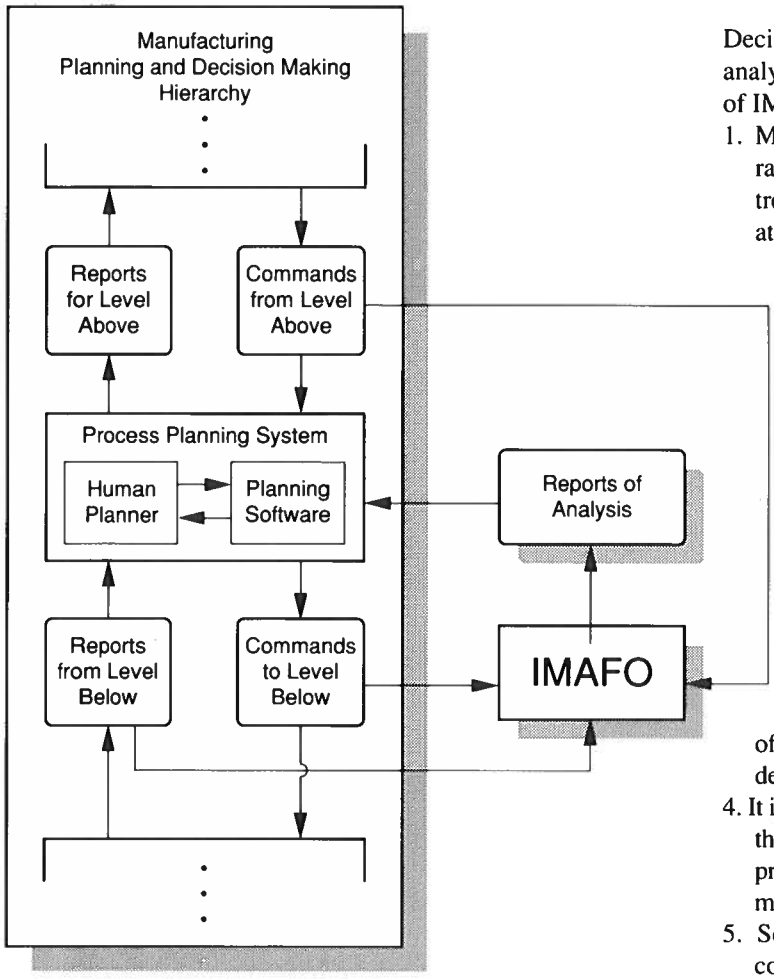

Figure 1: The flow of information between IMAFO and the manufacturing planning and decision making hierarchy.

specifying the format of the input files, defining a measurement space, defining subspaces of the measurement space, and controlling the use of dimensions in descriptions. IMAFO reads data in three files, which contain job commands, task commands, and execution results. IMAFO can handle data in many different formats and types, which are specified by the user during initialization. This specification can be saved in a fourth type of file, an *initialization* file.

The three sources of input must be consolidated before analysis can begin. The inputs are mapped into a *measurement space*, using functions that the user constructs from IMAFO's built-in functions. Each point in the measurement space is called an *observation*. The observations in a measurement space consist of both successful and unsuccessful plans. Observations are classified as successful or unsuccessful by defining a subspace of the measurement space.

The search for descriptions of unsuccessful plans uses a variation on Quinlan's *ID3* algorithm for induction of decision trees (Quinlan, *1983*, *1986a*, 1986b, 1987a, 1987b). Decision tree induction is not immediately applicable to the analysis of unsuccessful industrial process plans. The design of IMAFO required addressing the following issues:

1. Much of the data in industrial process planning is numeric, rather than symbolic. The standard algorithm for decision tree induction (Quinlan, 1983) only handles symbolic attributes. IMAFO readily handles several different data types. IMAFO builds binary decision trees, as evidence suggests that they are more accurate than other types of decision trees (Bratko and Kononenko, 1986). Thus IMAFO is designed to automatically convert all data to Boolean.
2. Decision tree induction works well only when the attributes of the examples are appropriately selected. IMAFO uses a measurement space, which allows data from distinct sources to be consolidated; it lets the user represent the data more abstractly; and it facilitates experimentation with different sets of complex attributes.
3. Decision trees are difficult to understand. IMAFO translates decision trees into a form which is designed to be more readily used and understood by a human planner. The method of improving the understandability of the tree is based on Quinlan's method for converting decision trees to production rules (Quinlan, 1987b).
4. It is important to minimize the number of examples of plans that are required in order to learn the target concept, to provide rapid feedback to the planner. IMAFO attempts to minimize the number of examples required to learn a concept.
5. Some plan failures may be random. These failures constitute *noise*, which IMAFO must be able to ignore. IMAFO prunes the decision trees to reduce their sensitivity to noise. IMAFO uses Quinlan's reduced-error pruning algorithm (Quinlan, *1987a*), since it improves the accuracy of decision trees at least as much as do the other four principal methods of pruning, and it is relatively easy to compute (Mingers, *1989b*).
6. The standard algorithm for decision tree induction requires a *teacher* to manually classify objects or events as positive or negative examples of the target concept. This would mean extra work for the planner. IMAFO automatically finds positive and negative examples of the target concept, using criteria supplied by the planner.

4 The Output of IMAFO

Table 1 gives an example of output from IMAFO. Rules 1 and 2 attempt to predict which plans will have a quantity-shortage. Rule 1 means "The current task for a station is stamp-bumper-support and the prior task for the same station was not stamp-bumper-support." Rule 2 means "The current task for a station is stamp-bumper-bar and less than three hours were allocated for this task." A coverage of 90% means that a rule applies to 90% of the positive examples of quantity-shortage.

The error rate is the ratio of the number of negative examples of the target concept that satisfy the rule to the

Subspace quantity-shortage:
Rule 1:
task-type is sta m p-bu m per-support prior-task-same-stn is not stamp-bumper-support Coverage: 90 % (train) 85 % (test) Error rate: 3 % (train) 5 % (test)
Rule 2:
task-type is stamp-bumper-bar planned-task-duration is less than 3 hours Coverage: 10 % (train) 10 % (test) Error rate: 6 % (train) 5 % (test)
Combined rules:
Coverage: 100 % (train) 95 % (test)
Error rate: 4 % (train) 5 % (test)

Table 1: An example of output from IMAFO.

number of positive examples of the target concept that satisfy the rule. Rules 1 and 2 cover different parts of the subspace. The coverage and error rate are reported for both the training set and the test set. The combination of the rules is the disjunction of rule 1 and rule 2.

It is useful to report the coverage of each rule, so that the planner can concentrate on rules which cover more of the subspace, and to report the error rate, so that the planner knows the reliability of the rules. It is interesting to note that rules that cover only a small part of the subspace tend to have higher error rates (Holte et al., 1989).

There are many possible responses that the planner might make, given the output presented in Table 1. For example, the planner might allocate more time for the situations described by the rules, try to avoid the situations, examine the situations to determine whether it is possible to eliminate the difficulty, or ignore the problem.

5 Conclusions and Future Work

The nature of industrial process planning has been outlined and the problem of learning from unsuccessful plans has been described. It is relatively straightforward to find positive and negative examples of unsuccessful plans, but the task of finding patterns in these examples is difficult for people, due to the quantity of data involved and time constraints. Given this problem, decision tree induction is a natural approach.

IMAFO has been fully implemented and tested using data generated by the simulator. The next step is to test IMAFO in a real process planning environment. This will require an intermediate module to provide the data which are now provided by the process simulator.

We are seeking industrial collaborators to participate in testing and refining IMAFO, both for process planning in general and for the problems presented by their particular industry. The goal is to investigate the practicality of integrating IMAFO with a real process planning environment. For more information about this system, please contact A. Famili at the Laboratory for Intelligent Systems, Division of Electrical Engineering, National Research Council, (613) 993-8554.

References

Albus, J., Barbera, B., & Nagel, R. (1981). Theory and practice of hierarchical control. *Proceedings of the 23rd IEEE Computer Society International Conference*, pp. 18-39. IEEE Computer Society.

Bratko, I. & Kononenko, I. (1986). Learning diagnostic rules from incomplete and noisy data. In AI *Methods in Statistics*. London: UNICOM Seminar.

Holte, R.C., Acker, L.E., & Porter, B.W. (1989). Concept learning and the problem of small disjuncts. *Proceedings of the Eleventh International*
Joint Conference on Artificial Intelligence, pp. *813-818*. Los Altos, California: Morgan Kaufmann.

McLean, C.R. *(1986)*. Information architecture of the automated manufacturing research facility. *Proceedings of the Information Technologies for Manufacturing Conference*. Troy, New York: Rensselaer Polytechnic Institute.

Michalski, R.S. *(1986)*. Understanding the nature of learning: issues and research directions. In J.G. Carbonell, R.S. Michalski & T.M. Mitchell, Eds. *Machine Learning: An Artificial Intelligence Approach, Volume II, pp. 3-25*. Los Altos, California: Morgan Kaufmann.

Mingers, J. *(1989b)*. An empirical comparison of pruning methods for decision tree induction. *Machine Learning, 4, 227-243*.

Quinlan, J.R. *(1983)*. Learning efficient classification procedures and their application to chess end games. In J.G. Carbonell, R.S. Michalski & T.M. Mitchell, Eds. *Machine Learning: An Artificial Intelligence Approach, Volume I*, pp. *463-482*. Los Altos, California: Morgan Kaufmann.

Quinlan, J.R. *(1986a)*. The effect of noise on concept learning. In J.G. Carbonell, R.S. Michalski & T.M. Mitchell, Eds. *Machine Learning: An Artificial Intelligence Approach, Volume II*, pp. *149-166*. Los Altos, California: Morgan Kaufmann.

Quinlan, J.R. *(1986b)*. Induction of decision trees. *Machine Learning, 1, 81-106*.

Quinlan, J.R. *(1987a)*. Simplifying decision trees. *International Journal of Man-Machine Studies, 27, 221-234*.

Quinlan, J.R. *(1987b)*. Generating production rules from decision trees. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pp. *304-307*. Los Altos, California: Morgan Kaufmann.

Wilkins, D.E. *(1988). Practical Planning: Extending the Classical AI Planning Paradigm*. Los Altos, California: Morgan Kaufmann.

Call for Papers

Final Version

INTERNATIONAL SYMPOSIUM ON ARTIFICIAL INTELLIGENCE

APPLICATIONS IN INFORMATICS

Software Engineering, Data Base Systems,
Computer Networks,
Programming Environments,
Management Information Systems,
Decision Support Systems

November 13-15, 1991
Sheraton Cancun Resort & Towers
Cancún, México

The Fourth International Symposium on Artificial Intelligence will be held in Cancún México on November 13-15, 1991. The Symposium is sponsored by the **ITESM** (Instituto Tecnológico y de Estudios Superiores de Monterrey) and supported by the International Joint Conferences on Artificial Intelligence Inc., in cooperation with the American Association for Artificial Intelligence, the Canadian Society for Computational Studies of Intelligence, the International Association of Knowledge Engineers, the Sociedad Mexicana de Inteligencia Artificial and IBM of México.

Papers from all countries are sought that **(1)** present **applications** of artificial intelligence technology to the solution of problems in Software Engineering, Data Base Systems, Computer Networks, Programming Environments, Management Information Systems, Decision Support Systems and other Informatics technologies; **(2)** describe research on techniques to accomplish such applications, (3) Address the problem of transferring the AI Technology in different socio-economic contexts and environments.

Areas of application include but are not limited to:
Software development, software design, software testing and validation, computer-aided software engineering, programming environments, structured techniques, intelligent databases, operating systems, intelligent compilers, local networks, computer network design, satellite and telecommunications, MIS and data processing applications, intelligent decision support systems.

AI techniques include but are not limited to:
Expert systems, knowledge acquisition and representation, natural language processing, computer vision, neural networks and genetic algorithms, machine learning, automated reasoning, search and problem solving, knowledge engineering tools and methodologies.

Persons wishing to submit a paper should send **five** copies written in English to:

Hugo Terashima, Program Chair
Centro de Inteligencia Artificial, ITESM.
Ave. Eugenio Garza Sada #2501, Col. Tecnológico, C.P. 64849 Monterrey, N.L. México.
Tel. (52-83) 58•2000 Ext. 5134
Telefax (52-83) 58-1400 Dial Ext. 5143 or 58•2000 Ask Ext. 5143
Net address isai @tecmtyvm.bitnet or isai @tecmtyvm.mty.itesm.mx

The paper should identify the area and technique to which it belongs. Extended abstract is not required. Use a serif type font, size 10, single-spaced with a maximum of 10 pages. No papers will be accepted by electronic means.

Important dates:
Papers must be received by **April 30, 1991**. Authors will be notified of acceptance or rejection by **June 15,1991.** A final copy of each accepted paper, camera ready for inclusion in the Symposium proceedings, will be due by **July 15, 1991.**

PRECARN UPDATE
NOUVELLES DE PRECARN

Jean-Claude Gavrel

Already the second issue of the "PRECARN UPDATE" column! Unfortunately some of you may still wonder what PRECARN stands for, as little gremlins removed some of the introductory remarks of my first column! PRECARN is the acronym for PRECompetitive Applied Research Network, and was chosen as the name for a Canadian consortium of 34 companies formed in 1987 to do precompetitive AI and robotics research. We will use this column, kindly offered to us by the editor of *Canadian Artificial Intelligence*, to keep you informed of our activities.

We believe that PRECARN is a good mechanism to help bring new AI and robotics technologies to Canadian industries, faster. By putting together strong teams of industry, university and government researchers, and by sharing the costs between the Member companies and Canadian governments, the PRECARN research projects can address problems of longer term and larger scope than would normally be carried out by Canadian industry in these new technology areas.

An example of such a project is APACS (Applications of Expert Systems to Advanced Process Analysis and Control), the feature of this month's column. At a ceremony held in December to mark the launching of APACS, federal Science Minister William Wineguard said that this project illustrates companies' increasing recognition that they must combine resources to undertake costly research. APACS involves four companies and one university and is jointly funded by the participants, PRECARN, the National Research Council's IRAP program and Industry Science and Technology Canada. It is a five-year, $8.7 million undertaking. The following describes the need for, the objectives and the approach of this research project; it has been prepared by Connie Bryson, a science writer well known to the readers of *Canadian Artificial Intelligence.*

Modern process plants incorporate computer control in almost every phase of their operations. Automation is the name of the game. Automation, that is, until something goes wrong when in most cases a human operator must take charge of the situation. But having an operator step in at the first sign of trouble is a far cry from a fully automated system, one that is capable of diagnosing and correcting problems on its own. PRECARN'S Advanced Process Analysis and Control System project is a first step towards this goal. The APACS project involves the use of artificial intelligence and expert system technology to automate the diagnostic analysis and control of the feedwater system at Ontario Hydro's Bruce Nuclear Generating Station. It is intended as a demonstration of the applicability of AI technology. The main participants in the project are Ontario Hydro, CAE Electronics, STELCO, Shell Canada, PRECARN and a research group headed by John Mylopoulos at the University of Toronto.

Although several task-specific expert systems have proven successful in industrial application, APACS project lead applicant Jordan Chou from Ontario Hydro says there are still many challenges in the plant-wide application of technology and the integration of expect systems with conventional plant control systems.

"Many commercially available expert systems are small rule-based systems with limited application," he says. "If we were to apply them to a large systems like a nuclear plant, we'd be faced with a rule explosion. The large number of rules would render the system impractical, and difficult to verify and validate."

The project researchers are using model-based reasoning and simulation technology to reduce the number of rules and improve confidence levels. "Faster than real-time" simulation is capable of supporting on-line, real-time diagnosis of system faults and the prediction of plant behaviour. The project builds on Canadian expertise in simulation technology. CAE Electronics and Ontario Hydro have already sold over $100 million-worth of operator training simulators to the United States. The four objectives of the APACS project are:

- design and develop knowledge-based systems to deal with operation and control problems within a process plant;
- design and develop an advanced simulation facility, the Plant Analyzer, to support the real-time diagnosis of plant malfunctions and the prediction of plant behaviour;
- design and develop an operator-machine interface in compliance with the operating requirements of a control room environment; and
- integrate these systems within the Advanced Process and Control System applicable to power and process plants.

A working APACS promises to increase the efficiency, reliability and availability of generating station. Benefits include more timely diagnosis of process faults, more precise identification of process trends indicative of malfunctions or upsets, and more efficient and informative methods for filtering and presenting alarms to operators. Application of the system goes beyond nuclear plants to any process industry, as evidenced by the participation of Shell Canada and STELCO in the APACS project.

Of course Canada isn't the only player in this game. Japan has already set a national target to have a fully automated

nuclear plant by 2030 and is pouring million of dollars into the effort. They plan to push the application of AI technology beyond self-diagnosis and correction, to self-maintenance.

"Canada is a small country. We can't compete in many areas but real-time simulation and computer control are areas where we have significant expertise. It's critical that we keep in the forefront of technology and maintain our lead," says Chou. "The APACS project is only a first step. But if we succeed, this work will revolutionize the future design of our process plants."

For more information on APACS or PRECARN, 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

Giving disabled a hand

Brainy computer sought to translate hand signs into speech

Carolyn Leitch

Researchers at the University of Toronto say a computer modelled on the human brain may eventually provide voices for hearing- and speech-disabled people.

Computer science professor Geoffrey Hinton and PhD student Sidney Fels have trained a computer to recognize hand movements and speak simple words in response.

Their ultimate goal is a machine that can translate hand signs into speech.

Prof. Hinton, a fellow at the Canadian Institute for Advanced Research, said people could learn specialized hand movements much in the way they learn to play a musical instrument.

"You learn to play this thing, but instead of producing music you produce speech," he said.

The wide variations in human speech, handwriting and movement have tripped up computers for years because the machines rely on precise digital codes to process information.

But a system of pattern recognition that learns much the same way as do humans could be the breakthrough that Prof. Hinton and Mr. Fels need.

Neural networks recognize patterns after being shown many examples. Researchers don't understand how, but the machines actually get better with practice. Their uses range from reading handwriting to detecting abnormal human cells among thousands of normal ones.

To communicate with his neural net, Mr. Fels dons a black lycra glove covered with fibre-optic cables that trail away to a black box. As he runs through a series of choreographed signs, a disembodied voice replies "short, shorts, shorter."

A different set of manoeuvres produces "sister, mother and brother."

Mr. Fels ran through his hand signals 9,000 times before the computer learned to recognize them.

"It got everything wrong – it didn't have a clue what it was doing," he said of early attempts.

Now Mr. Fels has taught the computer 66 hand shapes that represent words. By moving his hand in different directions, he can increase its vocabulary to 203 words.

The machine teaches itself what information is reliable and what is not. If, for example, Mr. Fels has trouble keeping his baby finger straight, the computer makes allowances for variation in that finger's position.

The neural net also adjusts quickly to the quirks of a new operator.

The project is part of a trend in research to build computers that learn to communicate like humans rather than requiring people to learn the computer's language.

Prof. Hinton said a practical device is still several years away, but he hopes the system will eventually recognize either conventional sign language or signs that denote phonetic sounds.

Such a system, Mr. Fels said, would be more flexible in making conversation.

"With around 50 phonetic sounds, you can make any word in the English language," he said.

Prof. Hinton acknowledges that the system is cumbersome today, requiring a desktop computer and a suitcase-sized processor, but he predicts a portable model could be available in a few years.

"Look what's happening with computers and look down the road five years – it could be made cheap and small," he said.

Reprinted with permission
– The Globe and Mail, January 2, 1991

Russ Thomas

AI Research In Finland

There are a number of centers carrying out AI related work in Finland. In the 1980's much of this work was carried out under the FINPRIT (Finnish Program for R & D in Information Technology) a program similar in concept to the European Communities ESPRIT programme. The largest of the five subprograms of FINPRIT was in the area of AI and was aimed at bringing Finnish competence in applied AI to an international level. The work reported here is largely focused on projects carried out under the FINPRIT programme. These projects was generally based either at VTT (Technical Research Center of Finland) or at HUT (Helsinki University of Technology).

Intelligent Interfaces

The goal of this project was to develop methods and software to create intelligent interfaces to complex systems. One of the successful prototypes to come out of this project was an intelligent intermediary system for information retrieval. Others provided support for ship designers to carry out the design process without having to worry about the underlying mathematical software, and provided advice on optimal travel selection between Finnish Cities given various time constraints. These projects were undertaken in the Laboratory for Information Processing at VTT.

Expert Systems

HUT carried out research into both real-time embedded expert systems (these included fault diagnosis, maintenance of telephone exchanges, optimal control of elevator groups) and systems for production management. Some of the work on production management which looked at distributed scheduling was carried out in co-operation with Carnegie Mellon University.

Three prototype expert systems in the area of medicine were developed by the Medical Engineering Laboratory of VTT. One system supported the doctor in prescribing antibiotics for serious but common bacterial infections. Another system diagnosed operational disorders in thyroid disfunction based on hormone analysis and other background information. The third system provided support for the treatment of fluid and electrolytic balance disorders in an intensive care environment.

Finally, the Computer Technology Laboratory at VTT carried out work in the development of tools, based on a knowledge engineering approach, to support structured specification and the design process. The work focused upon developing embedded products where both hardware and software must be designed simultaneously. A prototype was developed which was connected to the design of control software for elevators.

Other AI related projects have been carried out in Finland outside of the FINPRIT program although again these have also been focused upon the VTT and HUT centers. VTT has also developed a number of expert systems in the construction, medical and process sectors. These include expert systems for checking compliance with building regulations, control of energy distribution, cancer treatment and safety analysis of industrial processes. At the same time researchers at HUT have been carrying out work into natural language and text processing including machine translation, speech recognition, pattern recognition, self-adapting systems and neural networks.

This short report is in part based upon a longer article by Seppo Linnainmaa of VTT. Those interested in learning more about the AI work in Finland or seeking further details about specific projects could try contacting Seppo on internet as linnainmaa@tik.vtt.fi.

1.0 Theoretical Aspects

1260 An approximate analogical reasoning schema based on similarity measures and interval-valued fuzzy sets.
I.B. Turksen, Zhao Zhong
(Dept. of Ind. Eng., Toronto Univ., Ont., Canada).
Fuzzy Sets Syst. (Netherlands), vol.34, no.3, p.323-46
(26 Feb. 1990).

An approximate analogical reasoning schema (AARS) is proposed which exhibits the advantages of fuzzy sets theory and analogical reasoning in expert system development. The AARS avoids going through the conceptually complicated compositional rule of inference. It uses a similarity measure of fuzzy sets and a threshold T_0 to determine whether a rule should be fired and a modification function inferred from a similarity measure to deduce a consequence. Linguistic variables and terms, rather than quantitative variables, are used to represent decision rules and/or decision variables. The schema is extended to interval-valued fuzzy sets where additional uncertainty can be handled in a formal model. Finally, the proposed AARS is implemented using Texas Instrument's PC * Expert System Shell. The expert system AARS is applied to an aggregate production planning prototype study, where intuitive decision-making seems to be common practice.

The system's performance is compared with other approaches and the result is encouraging. (28 refs).

1277 An approach to combining explanation-based and neural learning algorithms.
J.W. Shavlik, G.G. Towell
(Dept. of Comput. Sci., Wisconsin Univ., Madison, WI, USA). *Connect. Sci. (UK)*, vol.1, no.3, p.231-53 (1989).

A hybrid system that combines the symbolically-oriented explanation-based learning paradigm with the neural backpropagation algorithm is described. In the presented EBL-ANN algorithm, the initial neural network configuration is determined by the generalized explanation of the solution to a specific classification task. This approach overcomes problems that arise when using imperfect theories to build explanations and addresses the problem of choosing a good initial neural network configuration. Empirical results show that the hybrid system more accurately learns a concept than the explanation-based system by itself and learns faster and generalizes better than the neural learning system by itself. (55 refs.)

1484 Applications of nonmonotonic logic to diagnosis.
P. Jackson
(McDonnel Douglas Res. Labs., St. Louis, Mo, USA).
Knowl. Eng. Rev. (UK), vol.4, no.2, p.97-117 (June 1989).
[received: 01 May 1990]

This paper attempts to assess the practical utility of nonmonotonic logic in diagnostic problem solving. It begins with a brief review of the main assumptions which motivate work in this area, and discusses two logic-based approaches which involve nonmonotonic arguments. It then considers two recent proposals for the application of default logic to diagnosis, as well as a proposal based on counterfactual logic. In conclusion, these methods are briefly compared with other diagnostic reasoning paradigms found in the artificial intelligence literature. (31 refs.)

1489 Numerical and symbolic approaches to uncertainty management in AI.
D.A. Clark
(Imperial Cancer Res. Fund, London, UK).
Artif. Intell. Rev. (UK), vol.4, no.2, p.109-46 (1990).

Dealing with uncertainty is part of most intelligent behavior and therefore techniques for managing uncertainty are a critical step in producing intelligent behavior in machines. The paper discusses the concept of uncertainty and approaches that have been devised for its management in AI and expert systems. These are classified as quantitative (numeric) or symbolic techniques. Each is discussed, illustrated, and assessed in relation to various criteria which illustrate the relative advantages and disadvantages of each technique. The discussion summarizes some of the criteria relevant to selecting the most appropriate uncertainty management technique for a particular application, emphasizes the differing functionality of the approaches, and outlines directions for future research.

1501 Contribution analysis: a technique for assigning responsibilities to hidden units in connectionist networks.
D. Sanger
(AT&T Bell Labs., Holmdel, NJ, USA).
Connect. Sci. (UK), vol.1, no.2, p.115-38 (1989).

Contributions, the products of hidden unit activations and weights, are presented as a valuable tool for investigating the inner workings of neural nets. Using a scaled-down version of NETtalk, a fully automated method for summarizing in a compact form both local and distributed hidden-unit responsibilities is demonstrated. Contributions are shown to be more useful for ascertaining hidden-unit responsibilities than either weights or hidden-unit activations. Among the results yielded by contribution analysis: for the example net, redundant output units are handled by identical patterns of hidden units, and the amount of responsibility a hidden unit takes on is inversely proportional to the number of hidden units. (7 refs.)

1508 Generation, local receptive fields and global convergence improve perceptual learning in connectionist networks.
V. Honavar, L. Uhr
(Dept. of Comput. Sci., Wisconsin Univ., Madison, WI, USA). IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.180-5 vol. 1

Presents and compares results for three types of connectionist networks on perceptual learning tasks: (1) multi-layered converging networks of neuron-like units, with each unit connected to a small randomly chosen subset of units in the adjacent layers, that learn by re-weighting of their links; (2) networks of neuron-like units structured into successively larger modules under brain-like topological constraints that learn by re-weighting of their links; and (3) networks with brain-like structures that learn by generation-discovery, which involves the growth of links and recruiting of units in addition to reweighting of links. Preliminary empirical results from simulation of these networks for perceptual recognition tasks show significant improvements in learning from using brain-like structures over networks that lack such structure; further improvements in learning result from the use of generation in addition to reweighting of links. (22 refs.)

1527 Minimal change and maximal coherence: a basis for belief revision and reasoning about actions.
A.S. Rao
(Australian AI Inst., Carlton, Vic., Australia), N.Y. Foo. IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.966-71 vol.2

The study of belief revision and reasoning about actions have been two of the most active areas of research in AI. Both these areas involve reasoning about change. However, very little work has been done in analyzing the principles common to both these areas. This paper presents a formal characterization of belief revision, based on the principles of minimal change and maximal coherence. This formal theory provides an elegant solution to the conceptual frame and ramification problems. It also facilitates reasoning in dynamic situations where the world changes during the execution of an action. The principles of minimal change and maximal coherence seem to unify belief revision and reasoning about actions and may form a fundamental core for reasoning about other dynamic processes that involve change. (20 refs.)

1538 A modest, but semantically well founded, inheritance reasoner.
F. Bacchus
(Dept. of Comput. Sci., Waterloo Univ., Ont., Canada). IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.1104-9 vol.2

A modest-exception-allowing inheritance reasoner is presented. The reasoner allows restricted, but semantically well-founded, defeasible property inheritance. Furthermore, it gives a well-defined and easily understood semantic interpretation to all of the assertions encoded in it. The semantics allows a knowledge engineer to decide what knowledge can be encoded in the system, and gives him understandable formal guarantees about the quality of the conclusions that will be generated. For this reason the system is a more practical, usable inheritance reasoner than others that have appeared in the literature. The system has been fully implemented in a short (<75 lines) PROLOG program which executes all the examples presented, among others. Furthermore, although the system performs a restricted form of inheritance reasoning it can still represent and solve most of the inheritance 'puzzles' that have appeared in the literature, including the recent heterogeneous inheritance problems. (16 refs.)

1728 Dynamic backward reasoning systems.
A. Ligeza (Inst. of Autom., Acad. of Min. & Metall., Cracow, Poland). *Artif. Intell. (Netherlands)*, vol.43, no.2, p.127-52 (May 1990).

Theoretical backgrounds for dynamic backward reasoning systems are presented. By dynamic reasoning is meant reasoning which allows for dynamic changes of the knowledge base, i.e. that new facts are added to it while certain facts that are no longer true are deleted from it. It is argued that many problems can be more easily solved by reasoning backwards that by reasoning forwards. A classification of problems into constructive and destructive ones, with regard to the potential efficiency of systematic backward and forward search, is introduced. A formal method for describing physical states and sets of states as well as transformation rules (operators) is given briefly; metarules for performing forward- and backward-chained reasoning are also provided. Moreover, the potential equivalence of reasoning forwards and backwards is proved and the advantages of reasoning backwards are pointed out. This paper, primarily aimed at developing a deeper theoretical understanding of reasoning backwards, can provide basic concepts to be incorporated into a specialized planning or expert system for support of knowledge processing in realistic domains. (21 refs.)

1747 Information fractals in evidential reasoning.
A.M. Erkmen, H.E. Stephanou
(Sch. of Inf. Technol. & Eng., George Mason Univ., Fairfax, VA, USA). Proceedings IEEE International Symposium on Intelligent Control 1988, Arlington, VA, USA, 24-26 Aug. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.34-9

Evidential reasoning based on a fractal model of belief is outlined. The specific focus is on the fractal modeling of belief functions. After a qualitative justification and interpretation of this model, several concepts and tools needed for its incorporation into evidential reasoning are formally defined. A particularly important concept is that of conductivity, as it provides the basis of partial evidential matching in the present approach to reasoning by analogy. A conductivity analysis algorithm is derived, and it is illustrated by an application to a simple object classification problem. The fractal model provides potentially powerful mechanisms for a quantitative measure of relevance of a piece of evidence to a knowledge base, and a systematic approach to the coarsening and refining of frames of discernment. The proposed model is motivated by applications to the design of intelligent systems, such as sensor-based dexterous manipulators that must operate in unstructured environments in the presence of high levels of uncertainty. (16 refs.)

1752 An extended framework for default reasoning.
M.A. Nait Abdallah
(Dept. of Comput. Sci., Univ., of Western Ontario, London, Ont., Canada). Fundamentals on Computation Theory, International Conference FCT '89 Proceedings, Szeged, Hungary, 21-25 Aug. 1989 (Berlin, West Germany: Springer-Verlag 1989), P.339-48

Investigates the proof theory of default reasoning. The author generalizes Reiter's (1980) framework to a monotonic reasoning system, and in particular allows formulae with nested defaults. He gives proof rules for this extended default logic, called default ionic logic, and gives deduction theorems. He also gives examples of applications of his framework to some well-known problems: weak implication, disjunctive information, default transformations, and normal versus non-normal defaults. (7 refs.)

2.0 Systems and Techniques

1303 Parallelism in knowledge-based machines.
A.N. Refenes
(Dept. of Comput. Sci., Univ. Coll., London, UK).
Knowl. Eng. Rev. (UK), vol.4, no.1, p.53-71 (March 1989).
[received: 19 Mar 1990]

The application area of knowledge-based expert systems is currently providing the main stimulus for developing powerful, parallel computer architectures. Languages for programming knowledge-based applications divide into four broad classes: functional languages, logic languages, rule-based languages, and, what is referred to as self-organizing networks. Despite their many differences, a common problem for all language classes and their supporting machine architectures is parallelism: how to de-compose a single computation into a number of parallel tasks that can be distributed across an ensemble of processors. The aim of the paper is to review the four types of language for programming knowledge-based expert systems, and their supporting parallel machine architectures. In doing so the author analyzes the concepts and relationships that exist between the programming languages and their parallel machine architectures in terms of their strengths and limitations for exploiting parallelization. (40 refs.)

1319 SimKit: knowledge-based simulation tools.
D. Silverman, M. Stelzner
(IntelliCorp., Mountain View, CA, USA).
Modelling and Methodology. Knowledge Systems' Paradigms, Tucson, AZ, USA, 19-23 Jan. 1987 (Amsterdam, Netherlands: North-Holland 1989), p.189-96

Describes the SimKit system, an integrated set of general-purpose simulation and modelling tools built on the Knowledge Engineering Environment (KEE) software development system of J.C. Kunz et al. (1984). KEE provides tools used in the construction of expert systems, many of which are explicitly designed to facilitate modelling of domain knowledge, a goal that expert systems share with simulation. Because SimKit is built on KEE, it takes full advantage of the expressive representation, reasoning, and interface tools that KEE provides. SimKit is explicitly designed for two classes of users: library developers and model builders. Library developers are programmers who build libraries - sets of domain-specific object classes, relationships, and behaviors. Individual instances of these classes are placed via a graphic interface in a model. A library may have many models. Model builders are frequently experts in the simulation domain, but no experienced programmers. However, once a library has been constructed, they can use SimKit's constructor-kit interface to configure and run models of specific objects and the relationships among those objects. (2 refs.)

1325 A technique for customizing object-oriented knowledge representation systems, with an application to network problem management.
L. Goffaux, R. Mathonet
(CIG-Intersys Syst., Brussels, Belgium).
IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.97-103 vol.1

Over the last few years, object-oriented techniques have gained an increasing recognition both in software engineering and in AI. Object-oriented systems present undisputable advantages and provide features that make them really suitable to represent knowledge. However, from practical experience in using these techniques for knowledge representation the authors suggest that object-oriented systems also present serious drawbacks, essentially due to their lack of expressive power. These drawbacks are severe when modelling domains characterized by a wide variety of knowledge. This paper introduces the notion of representation cluster, by which any object-oriented system can be provided with customizable knowledge representation formalisms. These formalisms permit the expression and handling of diverse kinds of knowledge using only their natural high-level properties. Such an approach speeds up development of knowledge bases, makes them clearer and more natural, and avoids a great deal of code writing. Its advantages are illustrated on two realistic examples extracted from the knowledge representation system of an expert system shell dedicated to real-time network troubleshooting. (9 refs.)

1327 On the road to automatic knowledge engineering.
J. Patel
(Human Cognition Res. Lab., Open Univ., Milton Keynes, UK). IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.628-32 vol.1

Represents a scheme for categorizing knowledge engineering tools. The classification of knowledge acquisition systems has revealed some interesting facts about these systems. It seems that systems which are able to work on multiple tasks produce very shallow knowledge bases. On the other hand, systems which produce expert-level knowledge bases function on a single task. These insights have led to the design of ASKE, a knowledge acquisition system which can be used to build expert-level knowledge bases in several domains and for different task-types. The knowledge acquisition process is based on the notion of templates, the knowledge-bearing units of ASKE. Templates provide a convenient way of representing domain knowledge (8 refs.)

1580 Experiences implementing a parallel ATMS on a shared-memory multiprocessor.
E. Rothberg, A. Gupta
(Dept. of Comput. Sci., Stanford Univ., CA, USA).
IJCAI-89 Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI, USA, 20-25 Aug. 1989 (Palo Alto, CA, USA: Morgan Kaufmann 1989), p.199-205 vol.1

The assumption-based truth maintenance system (ATMS) is an important tool in AI. So far its wider use has been limited due to the enormous computational resources which it requires. The authors investigate the possibility of speeding it up by using a modest number of processors in parallel. They begin with a highly efficient sequential version written in C and then extend this version to allow parallel execution on the Encore Multimax, a 16 node share-memory multiprocessor. They describe their experiences in implementing this shared-memory parallel version of the ATMS, present detailed results of its execution, and discuss the factors which limit the available speedup. (7 refs.)

1766 AI planning: systems and techniques.
J. Hendler, A. Tate, M. Drummond
(Maryland Univ., College Park, MD, USA).
AI Mag. (USA), vol.11, no.2, p.61-77 (Summer 1990).

Reviews research in the development of plan generation systems. The goal is to familiarize the reader with some of the important problems that have arisen in the design of planning systems and to discuss some of the many solutions that have been developed in the over 30 years of research in this area. The authors broadly cover the major ideas in the field of AI planning and show the direction in which some current research is going. They define some of the terms commonly used in the planning literature, describe some of the basic issues coming from the design of planning systems, and survey results in the area. Because such tasks are virtually never ending, and thus any finite document must be incomplete, the authors provide references to connect each idea to the appropriate literature and allow readers access to the work most relevant to their own research or applications (137 refs.)

1782 Issues in the design of expert systems for business.
E.A. Edmonds, L. Candy
(Dept. of Comput. Studies, Loughborough Univ. of Technol., UK), P. Slatter, S. Lunn.
Expert Syst. Inf. Manage. (UK), vol.2, no.1, p.1-22 (1989).
[received: 01 Jun 1990]

Critical issues which have emerged from experience in the design and implementation of expert systems are discussed. The key elements of the design process which are vital to the success of the system are identified and discussed with reference to a collaborative project. Some current approaches to system design, which arise from differing perspectives, are surveyed briefly. It is suggested that, in the context of expert systems applications, the conventional phased, incremental approach has significant limitations. A design and development model is described which provides a framework that places user involvement and the establishment of system acceptability in a prime position. The authors also address the organization of collaborative, team work in a complex and dynamic situation. In particular, a view of the role of prototyping and its relationship to system evaluation by users within the whole process is discussed. (35 refs.)

1792 A new paradigm for explanation in medical causal reasoning systems.
P.W. Jamieson
(Decision Syst. Lab., Pittsburgh Univ., PA, 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.17-21

A significant obstacle confronting medical expert systems that employ causal reasoning is providing explanations without prohibitive computational expense. Another problem is the construction of a suitable user interface for the explanatory system. The interface should allow a physician to explore in detail the explanation of some concepts while avoiding lengthy explanations with conclusions generated by a causal reasoning system. This approach allows the reasoning module to structure explanations without excessive computational overhead. It also allows the user to obtain additional detail of a causal explanation. (23 refs.)

1797 A graphical expert system for microfossil identification for use in the petroleum industry.
P.A. Swaby
(BP Res. Int., Sunbury-on-Thames, UK). IEE Colloquium on 'AI in the User Interface' (Digest No. 118), London, UK, 27 April 1990 (London, UK: IEE 1990), p.6/1-7

Describes the design and development of a graphical expert system for the identification of microfossils. The main components of the system are a graphical expert system shell and a set of knowledge bases. The process of knowledge elicitation is highly structured, and partly automated for maximum efficiency. The elicited knowledge for microfossils is stored in specially designed metafiles which are easily modified using the automated elicitation process. A knowledge base building tool generates knowledge bases automatically from the metafiles which makes the expert system easy to construct and extend. The graphical expert system shell displays generic pictures of microfossils and their attributes. The user compares these generic pictures to the actual observation and selects the most appropriate values in each case. On the basis of this information the expert system suggests possible fossils. Combining extensive domain knowledge with a graphical interface makes this system powerful, flexible and effective (18 refs.)

3.0 Applications

1356 Machine layout: an optimization and knowledge-based approach.
S.S. Heragu
(Dept. of Manage. & State Univ. of New York, Plattsburgh, NY, USA), A. Kusiak. *Int. J. Prod. Res. (UK)*, vol.28, no.4, p.615-35 (April 1990).

The machine layout problem in an automated manufacturing system is addressed. Two models for solving the machine layout problem are developed. A knowledge-based system for machine layout, called KBML, is presented. KBML combines the optimization and expert system approaches and considers quantitative as well as qualitative factors while solving the machine layout problem. The system is coded in Common LISP and implemented on a Symbolics 3650 machine. It is illustrated with a numerical example. (24 refs.)

1357 An integrated approach to facilities layout using expert systems.
G. Abdou, S.P. Dutta
(Dept. of Ind. Eng., Windsor Univ., Ont., Canada). *Int. J. Prod. Res. (UK)*, vol.28, no.4, p.685-708 (April 1990).

The paper develops an expert systems approach to define appropriate layouts of machining facilities under specific combinations of manufacturing and materials handling systems. The knowledge base incorporates six factors relating product variety and quantity, degrees of flexibility, level of automation, materials handling systems, work-in-process and environmental considerations. The EXSYS system program is used to manipulate the knowledge base. The program operates in tandem mode and interfaces with algorithms to optimize the selection of materials handling equipment and to generate appropriate layouts. Application of the program is illustrated for an existing facility. (27 refs.)

1377 Expert/knowledge-based systems for ceramic glaze technology.
S. Sengupta
(Dept. of Ceramic En., Clemson Univ., SC, USA).
Am. Ceram. Soc. Bull. (USA), vol.69, no.1, p.83-4, 90 (Jan. 1990).

A ceramist working with ceramic glasses might be called upon to make a decision in the formulation and diagnosis of glass flaws. Glaze flaws afford the opportunity to explore the diagnosis aspect of a KBS. The calculation of a glaze recipe using electronic spreadsheets warrants the construction of an intelligent front end to facilitate communication between and ease of use of such software packages. An expert system has been developed as a demonstration of the concept for the diagnosis of glaze defects and the formulation of glaze compositions. The development of this system involves

the following steps: a feasibility study, which concerns itself with questions like suitability of the project, cost, long-term benefits, etc.; the extraction and construction of knowledge to build the knowledge base; implementation of the system; and finally, validation of the expert system. The two application areas are discussed separately, following the guidelines outlined above.

1398 Automatic color segmentation of images with application to detection to variegated coloring in skin tumors.
S.E. Umbaugh, R.H. Moss
(Dept. of Electr. Eng., Missouri Univ., Rolla, MO, USA), W.V.Stoecker. *IEEE Eng. Med. Biol. Mag. (USA)*, vol.8, no.4, p.43-50 (Dec. 1989).

A description is given of a computer vision system, developed to serve as the front-end of a medical expert system, that automates visual feature identification for skin tumor evaluation. The general approach is to create different software modules that detect the presence or absence of critical features. Image analysis with artificial intelligence (AI) techniques, such as the use of heuristics incorporated into image processing algorithms, is the primary approach. On a broad scale, this research addressed the problem of segmentation of a digital image based on color information. The algorithm that was developed to segment the image strictly on the basis of color information was shown to be a useful aid in the identification of tumor border, ulcer, and other features of interest. As a specific application example, the method was applied to 200 digitized skin tumor images to identify the feature called variegated coloring. Extensive background information is provided, and the development of the algorithm is described. (25 refs.)

1440 Computer generation of geometrical error equations applicable for improvement of robots' positioning accuracy.
N. Vira, E. Tunstel
(Dept. of Mech. Eng., Howard Univ., Washington, DC, USA). Proceedings of an International Conference. Intelligent Autonomous Systems 2, Amsterdam, Netherlands, 11-14 Dec. 1989 (Amsterdam, Netherlands: IOS 1990), p.650-60 vol.2

A symbolic manipulation software package has been developed to automatically generate geometrical error model equations applicable for robot error compensation and calibration. The software package named AREEM (Automatic Robot Error Equation Modeler) utilizes MACSYMA, a LISP-based artificial intelligence language program, to output scalar algebraic equation representing the positioning error correction in world coordinates for N degree-of-freedom robots. At present, AREEM incorporates three kinematic error models based on the Denavit-Hartenberg representation (DH) and anon-DH representation. The AREEM program is user-friendly, interactively menu-driven, and has been tested on numerous robots. (11 refs.)

1449 A qualitative model for reasoning about shape and fit.
S.P. Carney
(Symbolics Inc., Cambridge, MA, USA), D.C. Brown. Artificial Intelligence in Design: Proceedings of the Fourth International Conference on the Applications of Artificial Intelligence in Engineering, Cambridge, UK, July 1989 (Southampton, UK: Comput. Mech. Publications 1989), p.251-72

Previous research on qualitative reasoning about shape and fit laid the foundations to determine whether two objects fit together. Continued investigation has refined the theory and produced a functioning implementation. The paper describes the theory and some aspects of the implementation. The reasoning process has been divided into five layers: grouping, topology, orientation, matching, and confirmation. The grouping layer clusters feature into groups for each surface of an object. The topology layer recognizes patterns formed by the groups on each surface, and describes the pattern in terms of topological structures. The orientation layer selects promising surfaces from the two objects and attempts to align them. If the topological structures on the two surfaces are aligned, the matching layer tries to pair the features between those topological structures. The confirmation layer inspects paired features to determine whether the surfaces are compatible. (17 refs.)

1461 Coupling symbolic and numerical computation for intelligent simulation.
G.O. Beale
(Dept. of Electr. & Comput. Eng., George Mason Univ., Fairfax, VA, USA), K. Kawamura.
In book: *Knowledge-based system diagnosis, supervision and control*, S.G. Tzafestas [Ed.], p.137-52. New York, NY, USA: Plenum (1989), xii+305 pp. [0 306 43036 3]

Expert systems technology has been gaining in popularity and acceptance in the engineering community. Due to the need for results generation and results interpretation in complex problems, there is a need to couple symbolic and numerical techniques. Such coupled systems promise to integrate the explanation and problem solving capabilities of knowledge-based systems with the precision of traditional numerical computing. One application area where this coupling of symbolic and numerical computing will have great benefit is the digital simulation of complex dynamic systems. The authors describe the design principle and architecture of one such coupled system, named NESS (NASA Expert Simulation System). NESS assists the user in running digital simulations of dynamic systems; interprets the output data to determine system characteristics; and, if the output does not meet the performance specifications, recommends a suitable series compensator to be added to the simulation model. (14 refs.)

1465 A survey of expert systems for equipment maintenance and diagnostics.
W.T. Scherer, C.C. White, III
(Dept. of Syst. Eng., Sch. of Eng. & Appl. Sci., Virginia Univ., Charlottesville, VA, USA).
In book: *Knowledge-based system diagnosis, supervision and control*, S.G. Tzafestas [Ed.], p.285-300. New York, NY, USA: Plenum (1989), xii+305 pp. [0 306 43036 3]

Considerable research is being conducted in the area of expert systems for diagnosis. Early work was concentrated in medical diagnostic systems. Current efforts are expanding to the area of equipment maintenance and diagnostics with numerous systems having been built during the past several years. The authors discuss a survey on expert systems for diagnosis. They remark that it is common for diagnostic systems to integrate concepts from artificial intelligence (expert systems), decision theory, and operations research. (110 refs.)

1603 Integrating a knowledge-based meat-grading system with a voice-input device.
Y.R. Chen, S.A. Robinson
(Roman L. Hruska US Meat Animal Res. Center, US Dept. of Agric., Clay Center, NE, USA). *Comput. Electron. Agric. (Netherlands)*, vol.4, no.4, p.303-13 (March 1990).
(Integration of Problem-Solving Techniques in Agriculture Workshop, San Antonio, TX, USA, 10-12 Aug. 1988).

The US livestock and meat industries depend heavily on the USDA to provide the rating of their.meat for marketing. Since the carcass-beef grading system was put in place in 1927, carcasses have been graded by human graders. To achieve better consistency and equity, a knowledge-based expert system for assisting graders to grade carcass beef has been developed. The paper presents the development of the knowledge-based expert system and discusses how the expert system was integrated with a voice-recognition subsystem that provides the characteristics of the carcasses to the expert system. (6 refs.)

1618 An expert system for aluminum-alloy welding procedure specification.
C.S. Wu, J. Xu, L. Wu
(Dept. of Welding Eng., Harbin Inst. of Technol., China).
Eng. Appl. Artif. Intell. (UK), vol.2, no.4, p.321-3
(Dec. 1989).

It is of great technical and economic significance to develop expert systems in the welding field. The paper describes an expert system for Aluminum-alloy welding procedure specification. The system involves 17 welding processes and 16 aluminum-alloy materials. Its output consists of expert advice on process selection, a colour graphic display of the suggested joint design, and recommendations for the main welding process parameters. It is written in Turbo Prolog, and is designed for running on an IBM PC, or compatible computers.

1645 Expert systems for the prediction of ovulation: comparison of an expert system shell (Expertech Xi Plus) with a program written in a traditional language (BASIC).
A. Schreiner, T. Chard
(Dept of Reproductive Physiol., London Hospital Med. Coll., UK). *Methods Inf. Med. (West Germany)*, vol.29, no.2, p.140-5 (March 1990).

The use of an expert system shell (Expertech Xi Plus) in the construction of an expert system for the diagnosis of infertility has been evaluated. A module was devised for predicting ovulation from the medical history alone. Two versions of this system were constructed, one using the expert system shell, and the other using Quick BASIC. The two systems have been compared with respect: (1) ease of construction; (2) ease of knowledge base update; (3) help and explanation facilities; (4) diagnostic accuracy; (5) acceptability to patients and clinicians; (6) user-friendliness and ease of use; (7) use of memory space; and (8) run time. The responses of patients and clinicians were evaluated by questionnaires. The predictions made by the computer systems were compared to the conclusions reached by clinicians and to the 'gold standard' of day 21 progesterone. (24 refs.)

1647 Database issues for a veterinary medical expert system.
M. McLeish, M. Cecile, A. Lopez-Suarez
(Dept. of Comput. & Inf. Sci., Guelph Univ., Ont., Canada). Statistical and Scientific Database Management. Fourth International Working Conference SSDBM Proceedings, Rome, Italy, 27-30 June 1988 (Berlin, West Germany: Springer-Verlag 1989), p.177-92

A large project to build an expert system for veterinary medicine was recently begun at the University of Guelph in collaboration with the Ontario Veterinary College. The hospital database system has been collecting online medical data for almost ten years. The prototype being developed is comparing several knowledge acquisition techniques from data which limit the number of rules used for diagnosis. Along with this approach, a fuzzy relational system involving all test results is being implemented. This paper shows how a combination of INGRES and the C programming language can be used as a knowledge base to support these different methodologies. (30 refs.)

1653 Expert system and interactive videodisc for lymphnode pathology.
B.N. Nathwani, D.E. Heckerman, E.J. Horvitz.
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. 1184

Summary form only given. The expert system makes diagnosis on 65 diseases of lymph nodes based on histopathologic features, clinical and immunologic findings. The system uses the method of sequential diagnosis and recommends to the user additional histologic features that should be evaluated to narrow the differential diagnosis. Moreover, the system explains its recommendations graphically. The information retrieval portion of the program provides hypertext-style access to over 1 megabyte of text information including criteria for diseases, definition of histologic features, the value of immunology and gene rearrangement studies, and relevant journal references. The computer controls an analog interactive videodisc containing 6700 slides on lymph node pathology. The slides are organized by diseases, by histologic features, and in comprehensive lecture format.

1699 A simulation study of expert control system for a phosphate flotation process.
S.-L. Jamsa-Jounela
(Kemira Oy, Helsinki, Finland), A. Niemi. Automation in Mining, Mineral and Metal Processing 1989. Selected Papers from the 6th IFAC Symposium, Buenos Aires, Argentina, 4-8 Sept. 1989 (Oxford, UK: Pergamon 1990), p.45-52

The parameters of a detailed phenomenological flotation model are evaluated by steady state testing. The dynamic model has been used in simulation of the operation of a rougher flotation cell, as a step in development of expert control of a larger flotation system. The simulator has been interfaced with the ONSPEC control software package run under the IBM DOS environment. The control strategy aims to maximization of the recovery, while the concentrate grade is maintained above a given minimum level. The hierarchy

of manipulated variables of the cell is described and the simulated results of expert control are presented.

1666 Regional analysis of watershed acidification using the expert systems approach.
D.C.L. Lam, A.S. Fraser
(Nat. Water Res. Inst., Burlington, Ont., Canada), D.A. Swayne, J. Storey, I. Wong. *Environ. Softw. (UK)*, vol.3, no.3, p.127-34 (Sept. 1988). [received: 22 Mar 1990]

This paper describes a workstation environment for the regional analysis by intelligent system on a microcomputer (RAISON) for studying watershed acidification and other environmental problems. Data is processed with a map-based query language and a spreadsheet paradigm. Spatial and temporal data retrieval is made possible for any chosen region and time frame, so that information on water, air and soil is combined in a number of watershed acidification models. A simple prototype expert system uses each of these models and determines which is the most appropriate according to a set of decision rules defined by the user. Based on the data from 53 watersheds in Southern Quebec, Canada, regionalization of the expert model to 91 watersheds is achieved by using a Monte Carlo procedure. A mean relative error of 19.1% is obtained by the Monte Carlo results when compared to observed data. (10 refs.)

1685 A real time data acquisition and monitor system for integrated circuit visual inspection.
P. Pinchera, Y.-K. Yang
(Harris Semicond., Melbourne, FL, USA).
Seventh IEEE/CHMT International Electronic Manufacturing Technology Symposium. Proceedings 1989 (Cat. No.89CH2720-1), San Francisco, CA, USA, 25-27 Sept. 1989 (New York, NY, USA: IEEE 1989), p.88-95

The design and implementation of a system used for real-time acquisition and monitoring of data collected at an IC manufacturing process are described. How it is used for process improvement and control is discussed. The approach taken to solve this problem can be divided into three functional areas: acquisition, monitor, and integration. The data acquisition talk utilizes a touch screen terminal which is used by each inspector to enter his inspection results. Icons represent each defect category and, when a unit is rejected, the inspector touches the icon and a counter is incremented. The data monitor is an expert system which accesses these data and displays them dynamically in a status board fashion, including statistics on each lot. It checks if any inspector has rejected or accepted too many die or if any defect rate for a reject category is too high. The data integration includes an INGRES-based relational database and many report generator programs. At the end of the lot inspection, all data and warning messages related to the lot are stored in the database for later analysis.

1718 Expert database systems: knowledge/data management environments for intelligent information systems.
L. Kerschberg
(Dept. of Inf. Syst. & Syst. Eng., George Mason Univ., Fairfax, VA, USA).
Inf. Syst. (UK), vol. 15, no.1, p.151-60 (1990).

Expert database systems (EDS) are database management systems (DBMS) endowed with knowledge and expertise to support knowledge-based applications which access large shared databases. The architectures, tools and techniques needed to build such systems are varied, and draw upon such fields as artificial intelligence, database management and logic programming. The goal of EDS research and development is to provide tools and techniques to make databases active agents that can reason, and to allow database systems to support artificial intelligence applications that manage and access large knowledge bases and databases. Expert database systems allow the specification, prototyping and implementation of knowledge-based information systems that represent a vertical extension beyond well-defined, transaction-oriented systems to those with knowledge-directed reasoning and interpretation. (42 refs.)

1816 IBM Burlington's logistics management system.
G. Sullivan
(IBM Corp., Essex Junction, VT, USA), K. Fordyce.
Interfaces (USA), vol.20, no.1, p.43-64 (Jan.-Feb. 1990).

Describes IBM's logistics management system (LMS), a real-time imbedded transaction-based integrated decision and knowledge-based expert support system, which serves as a dispatcher, monitoring and controlling the manufacturing flow of IBM's semiconductor facility near Burlington, Vermont. Burlington produces various logic and memory microelectronic chips and modules that are used throughout the IBM product line. LMS helps improve manufacturing performance and is critical to running major areas of the manufacturing facility. It was developed by the advanced industrial engineering department. (33 refs.)

1826 The educational potential of linking a videodisc and an expert system.
S.S. Stensaas, O. Bouhaddou, L. Hardy, D.K. Sorenson, N.E. Dougherty, S.J. Altman
(Utah Univ., Salt Lake City, UT, 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.864-8

A medical expert system (Iliad) and a videodisc (Slice of Life), have been successfully interfaced. The expert system was designed to support medical students in their clinical diagnosis experience. The videodisc was also designed for undergraduate medical students. The systems were developed independently and are being used in medical education. Modifications to the expert system were minimal and designed to support the linkage of dictionary terms in a knowledge base to images and textural descriptions of videodisc images. These pictorial links enhance the informational capacity and utility of the expert system. It is hoped that images will increase retention, problem-solving ability, and enjoyment of the expert system. (7 refs.)

1830 Lessons learnt from the field trial of ACORN, an expert system to advise on chest pain.
J. Wyatt
(Dept. of Med. Inf., Nat. Heart & Lung Inst., London, UK).
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. 111-15

Chest pain presents diagnostic problems, particularly in the accident and emergency department. A decision-aid, ACORN, was built to reduce doctors' false negative rates for admission of patients to the cardiac care unit, and to speed up decision-making. After a laboratory evaluation, the system was submitted to a randomised controlled field trial. ACORN was well accepted by doctors and nurses, but failed to alter the management of patients. The reasons for this failure are analysed and the changes made are described. The conclusion from the study is that field testing is a vital component of the evaluation process, necessary not only to ensure adequate performance but also to confirm that the system is usable in the intended clinical environment. (14 refs.)

1831 A connectionist expert system for diagnosing hepatobiliary disorders.
K. Yoshida
(Dept. of Preventive Med. & Public Health, Keio Univ., Tokyo, Japan), Y. Hayashi, A. Imura.
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.116-20

Expert systems that have connectionist (neural) networks for their knowledge bases are sometimes called connectionist expert systems. A connectionist expert system used to diagnose hepatobiliary disorders has been developed. The purpose of this study is to compare the diagnostic capabilities provided by the connectionist approach using the Pocket Algorithm, and that provided by the statistical approach. The authors used a medical database containing the results of nine biochemical items of four hepatobiliary disorders. After learning by using training data (334 patients), the connectionist model correctly diagnosed 71.8% of external (test) data from 163 previously unseen patients and correctly diagnosed 91% of the training data. Conversely, the diagnostic accuracy of the linear discriminant analysis was 63.2% of the external (test) data and 67% of the training data. The connectionist expert system showed higher diagnostic accuracy than that of linear discriminant analysis. The authors suggest that the connectionist approach will be useful technique for generating a knowledge base from medical databases. (12 refs.)

1836 Expert system for common diseases in pediatrics.
Wang Quiyan
(Dept. of Pediatrics, Xian Med. Univ., China), Wang Qiuhe, Shi Hongbao. 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.180-2

This system applies the theory of both artificial intelligence and fuzzy mathematics and sets up a diagnostic inference model with simulating properties of human thinking. The whole system includes 7 subsystems which are respiratory system, digestive system, circulatory system, urinary system, nervous system, hematic system and newborn diseases. Through checking and clinic practice of 2485 cases, the average accuracy of computer diagnosis was 95.86% which raised 8% more that clinic attending did. (6 refs.)

1839 The Iliad expert system for medicine.
Chinli Fan, M.J. Lincoln, H. Warner, Jr., Hong Yu, H.R. Warner
(Dept. of Med. Inf. & Internal Med., Utah Univ., Salt Lake City, UT, USA). 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.196-200

Iliad is an expert system for medicine jointly developed by the University of Utah Applied Informatics Inc. The paper describes the design and function of Iliad. Iliad is a stand-alone system that utilizes Apple Macintosh computers. The system is unique because it combines statistical and rule-based diagnostic algorithms so as to simultaneously increase diagnostic and teaching efficiency. Iliad's inference engine, coupled with its knowledge base, provides expert diagnostic consultations in the field of internal medicine for over 200 diseases. This internal medicine knowledge base also serves as a medical reference for users. The knowledge base will be expanded to include 600 diseases by 1990. (17 refs.)

1865 RXPERT: a decision support expert system for drug product interchangeability assessment.
M.L. Greer
(Coll. of Pharmacy, Saskatchewan Univ., Saskatoon, Sask., Canada), T.J. Quinn, J.E. Greer.
Proceedings: The Thirteenth Annual Symposium on Computer Application in Medical Care (Cat. No.89TH0286-5), Washington, DC, USA, 5-8 Nov. 1989 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.157-61

A knowledge-based expert system has been developed to support human decision makers who assess drug products for inclusion in the Saskatchewan Formulary. Formulary inclusion and deletion decisions are made by two committees. The Drug Quality Assessment Committee (DQAC) performs the initial clinical evaluation and makes recommendations to the Formulary Committee (FC). The FC then considers the practical and administrative implications of the DQAC recommendations to the Minister of Health. The RXPERT formulary advisor is designed to model and support the decision processes of the two committees. The prototype expert system attempts to model committee decision making, based on multiple experts with overlapping, but distinctly focused, areas of expertise. The major anticipated contribution is screening of routine decisions. (9 refs.)

1872 Research on interactive knowledge-based indexing: the MedIndEx prototype.
S.M. Humphrey (Nat. Libr. of Med., Bethesda, MD, 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.527-33

A description is given of the MedIndEx (medical indexing expert) system for interactive knowledge-based indexing of the medical literature. The general purpose of the project is to design, develop, and test interactive knowledge-based systems for computer-assisted indexing of literature in the MEDLINE database using terms from the MeSH (Medical Subject Headings) thesaurus.

In conventional MEDLINE indexing, although indexers enter MeSH descriptors at computer terminals, they consult the thesaurus, indexing manual, and other tools in published form. In the MedIndEx research prototype, the thesaurus and indexing rules are incorporated into a computerized knowledge base which provides specific assistance not possible in the conventional indexing system. The system, which combines principles and methods of artificial intelligence and information retrieval, is expected to facilitate expert indexing that takes place at the National Library of Medicine. (9 refs.)

1874 Patient monitoring in the operating room: validation of instrument reading by artificial intelligence methods.
D. Garfinkel, P.V. Matsiras, T. Mavrides, J. McAdams, S.J. Aukburg
(Pennsylvania Univ., Philadelphia, PA, 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.575-9

During physiological monitoring in the operating room, the large number of alarms when there is no hazard to the patient (false alarms) is a considerable problem. A method is described for comparing and validating instrument reading in this situation. The method involves a knowledge base whose core is a set of 36 rules. This was applied to 7803 warnings (6287 cautions and 1516 alarms) from 68 day surgery patients undergoing 115 h of surgery. Most of the cautions were validated by the analysis, but 734 of the 1516 alarms were invalidated., while 419 were validated and 363 left indeterminate. This translates to a potential reduction from one alarm every 4 min to one every 16 min. (9 refs.)

1917 Design of self-learning controllers using expert system techniques.
Z. Geng, M. Jamshidi
(Dept. of Electr. Eng. & Comput. Sc., George Washington Univ., WA, USA), J. Liebowitz. Proceedings IEEE International Symposium on Intelligent Control 1988, Arlington, VA, USA, 24-26 Aug. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.551-8

The use of real-time expert system techniques to control systems, including robot manipulator control systems, is discussed. A novel type of intelligent controller structure, the expert learning controller prototype called ELEC (expert learning controller), is developed for the control series optimization of trajectory tracking problems in repeat operations. The ELEC, acting as an intelligent real-time controller in a closed-loop system, can modify the control series in a human-expert-like way using the experience of previous operations in order to force the system output to converge to the previously given desired trajectory. With the self-learning functions, the ELEC does not require the knowledge of system models; thus, it can be used in a wide range of control problems, especially in robot control. Numerical examples and simulation results of nonlinear, time-varying and multiple-variable robot systems are given to show the satisfactory performance of ELEC. (7 refs.)

1938 Static analysis of program source code using EDSA.
L.I. Vanek, M.N. Culp
(Array Syst. Comput., Downsview, Ont., Canada).
Proceedings: Conference on Software Maintenance (Cat. No.89CH2744-1), Miami, FL, USA, 16-19 Oct. 1989 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.192-9

EDSA (expert data flow and static analysis tool) which uses static analysis of source code to help gain an understanding of existing code, is described. This may be for the purpose of tracking down a bug or to determine in advance whether an intended change will have any undesirable side effects. In either case, the phase of the development life cycle that is most likely to benefit from a tool like EDSA is the maintenance phase. EDSA provides three kinds of facilities. It helps to browse through code, following either the control flow of data flow, rather than the order in which the code happens to be written. It displays code with unimportant source lines eluded, so that the user can get a more global view of the program. Finally, it provides search management to make it easier to examine all possibilities when browsing. The mechanism used to create and display various views of programs is described and the approach to the user interface is discussed. The state of the current implementation is described. (14 refs.)

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

Persons wishing to obtain copies of references cited should contact their nearest technical library or the Canada Institute for Scientific and Technical Information, National Research Council, Building M-55, Montreal Road, Ottawa, Ontario, CDA, K2A 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".

International Conference

Cognitive science: Tools for the Development of Organizations

April 29 to May 2, 1991

Montréal

Conference Objective

The objective of this conference is to provide researchers and representatives from industry and government with an opportunity to:

- Present and discuss their research results
- Review applications of artificial intelligence and expert systems in organizations
- Discuss of future directions for research and development Ico'91 will include tutorials, and a three day scientific conference.

Topics of Interest

The conference will primarily focus on the following areas:

- Impact of AI and CS in organizations and in industry
- Natural language processing
- Knowledge-based tools for management and decision support
- Knowledge representation and acquisition
- Intelligent tutoring systems for training
- Computer-aided manufacturing and industrial robotics
- Evaluation of working systems.

acm Association for Computing Machinery

IEEE COMPUTER SOCIETY

A simultaneous translation service will be provided for the participants.

Ico '91

Organized by:
GIRICO
Université du Québec à Montréal

Sponsored by:
ACI Association canadienne d'informatique
AFCET Association française pour la cybernetique économique et technique
BCS The British Computer Society
CEFRIQ Centre francophone de recherche en informatisation des organizations
CRIM Centre de recherche informatique de Montréal
FIQ Fédération de l'informatique du Québec
INRIA Institut national de recherche en informatique et en automatique

In cooperation with:
ACM Association for Computing Machinery
SIGART, SIGOIS
CSCSI Canadian Society for Computational Studies of Intelligence
IEEE-CS IEEE Computer Society

Conference Chairperson:
Marc Bouisset
(Université du Québec à Montréal, Canada)

Program Co-Chairpersons:
Lorne H. Bouchard
(Université du Québec à Montréal, Canada)
Gilles Gauthier
(Université du Québec à Montréal, Canada)

Organizing Committee Chairperson:
Wilfried Probst
(Université du Québec à Montréal, Canada)

Conference Treasurer:
Robert Dupuis
(Université du Québec à Montréal, Canada)

Invited Speakers

Alain Colmerauer
GIA Luminy, France
Jean-Pierre Desclés
Université Paris-Sorbonne, France
Mark Fox
Carnegie-Millon University, USA
Brian Gaines
University of Calgary, Canada
Jean-Guy Meunier
Université du Québec à Montréal, Canada
Beverly S. Woolf
University of Massachusetts, USA

Tutorials

April 29th, 9h-12h a.m.

1. Text processing in UNIX™ (in english)
   Kenneth W. Church
   AT & T BELL Labs & USC/ISI, USA
2. Verification of Expert Systems (in english)
   Shinghai Rajjan and Alun D. Preece
   Concordia University, Canada

April 29th, 2h-5h p.m.

3. Designing Intelligent Tutoring System (in english)
   Jim Greer
   University of Saskatchewan, Canada
4. Traitement de l'information en réseaux neuronaux: une approche connexioniste (in french)
   Alex Friedmann and Robert Proulx
   Université du Québec à Montréal, Canada

Accommodation

Rooms are reserved at the Hôtel des Gouverneurs, near the conference location. These rooms will only be held until March 28th, after which time availability cannot be guaranteed. Participants are asked to reserve as soon as possible and mention ICO'91.

Hôtel des Gouverneurs
1418, rue St-Hubert
Montréal, Québec
H2L 3Y9
For other holds, contact:
Montréal Reservation Center 1-(800) 567-8687

Rate: 90 Can.$/day
Single or Double
Phone: 1 (514) 842-4881
Fax: 1 (514) 842-1584

Registration

Conference	Before April 1st	After
Members of sponsoring or cooperating organizations	290 (Can.$)	345(Can.$)
Non-members	325	390
Students	80	80
Banquet (May 1st)	25	25
Proceedings (Suppl. Copy)	35	35
Tutorials (each 1/2 day)	110	130

Conference fees cover admission for conference sessions, simultaneous translation, lunches (except for students), coffee-breaks, and a copy of the proceedings. Tutorial fees include admission to tutorial, documentation available and coffee-breaks. Official taxes are included.

For Further Information

Please contact:
Martine Fuentes
Fax: 1-(514) 987-8477
E-mail ICO91 @INFO.UQAM.CA
Phone: 1-(514) 987-3239

Registration Form

Last name _____
First name _____
Company _____
Address _____
City _____ Prov _____ Country _____
Postal _____ Phone _____
Affiliation _____ # Member _____

❑ The Conference _____ $
❑ The Banquet _____ $
❑ (...) Suppl. Proceeding(s) _____ $
❑ Tutorials _____ $
Precise which: 1 ❑ or 2 ❑ ; 3 ❑ or 4 ❑
Total (Can$): _____ $

Make cheque or money order payable to Ico'91
(Can. $ or equivalent US $)

Ico '91

Ico'91 Conference
Dept. de maths/info,
Université du Québec
à Montréal
Case postale 8888, Succ. A
Montréal, Quebec, Canada
H3C 3P8

BOOK REVIEWS
CRITIQUES DE LIVRES

BOOK REVIEWS

Artificial neural systems: Foundations, paradigms, applications, and implementations
Patrick K. Simpson
[General Dynamics Electronics Division]
New York: Pergamon Press, 1990,
145 pp+36 pp bibliography
Hardbound, ISBN 0-08-0378951, us$39.50
Softbound, ISBN 0-08-0378943, us$19.50

Reviewed by
Lynn Sutherland
Alberta Research Council

Patrick Simpson has produced a very valuable reference source for anyone interested in neural networks. He presents the nomenclature for the domain, a taxonomy of neural architectures, and detailed descriptions of 27 artificial neural system paradigms. The book presents artificial neural systems from a computational viewpoint. It does not address psychological, biological, or philosophical issues. The book would be a good companion for anyone learning about, experimenting with, or just interested in, neural systems as a computational technique. It is primarily a reference book. It needs to be supplemented by more detailed presentations in areas such as theory and applications.

The subtitle of the book, "Foundations, paradigms, applications, and implementations", is a bit too ambitious. The book provides a comprehensive presentation of neural systems paradigms, includes a brief history of the field, and an extensive bibliography. It does not include much in the area of mathematical foundations or details of applications or implementations. References to these areas are provided.

The chapter on the foundations of artificial neural systems characterizes processing elements, activation functions, network topologies, memory, recall, learning, and stability and convergence. The section on learning describes some examples of supervised and unsupervised learning, including error-correction, reinforcement, stochastic, Hebbian, competitive, and cooperative, learning. The section on stability and convergence presents three stability theorems for autoassociative and heteroassociative networks. This section is particularly weak in breadth. An introduction to the foundations of neural networks should cover at least some results in complexity and scope of representation and learning.

The most significant contribution of this book is the collection of descriptions of a wide range of artificial neural systems paradigms. Each paradigm is described by its encoding, its learning and recall mechanisms, stability properties, and a brief discussion of strengths, weaknesses, applications, and implementations. Each description provides the reader with a general overview of the paradigm and with enough information to decide whether the paradigm is potentially useful for a specific application. At this point, the reader is provided with references for further information.

Artificial neural systems is a useful contribution to the synthesis of current work in neural networks, connectionist systems and parallel distributed processing. It meets its goals of providing a reader with a rapid introduction to the concepts and terminology of artificial neural systems and providing a valuable reference guide to most of the primary artificial neural systems paradigms.

Lynn Sutherland is an assistant research officer at the Alberta Research Council. Her research interests include parallel processing and neural networks. She is currently A.I. project leader for a STEAR project for intelligent fault detection, diagnosis, monitoring, and control.

The age of intelligent machines
Raymond Kurzweil
Cambridge, MA: The MIT Press, 1990, xiii+565 pp
Hardbound, ISBN 0-262-11121, US$39.95

Hermes and the golden thinking machine
Alexander Tzonis
(Technical University, Delft)
Cambridge, MA: The MIT Press, 1990, xiv+284 pp
Hardbound, ISBN 0-262-20076-7, US$19.95

Reviewed by
Graeme Hirst and Diana Latent

The MIT Press, never too bound to academic convention, has recently released two new AI books rather different from the normal research monograph or textbook. One is a coffee-table introduction to the field; the other is a murder mystery.

The age of intelligent machines is a large, popularly-styled introduction to artificial intelligence and its milieu, as seen by Raymond Kurzweil. Kurzweil is best known as the maverick engineer responsible for the Kurzweil Reading Machine, the Kurzweil K250 music synthesizer, and other Things They Said Could Never Be Done. His book, while intended as a general *Scientific American*-level introduction for the lay reader, very much reflects his own perspective on the field.

Although the title seems to emphasize the modern social context of AI, the emphasis is on AI itself and how it came into being. That is, the book is a history of ideas, not a social history. There is a strong emphasis on the mathematical and philosophical background, and modern AI techniques don't make an appearance until page 223. The primary discussion of social issues is in the final chapter, speculating on the future.

Kurzweil's book is attractively designed. Its pages are large (20 cm by 28 cm; that's 8 by 11 inches), and full of pictures, many in colour. In addition to Kurzweil's text, many of the chapters feature extra essays by well-known AI researchers. (But setting these essays in a vibrant boldface is a failure in the design of the book.) There are about 460 pages of text, plus another hundred of reference material: glossary, bibliography, notes, etc. Printed on heavy matte paper, it is a bargain at US$40 (with a $10 discount if you were smart enough to order before 31 December 1990); prices may be slightly higher in Canada, as they say.

Less successful, but undoubtedly fun for true AI aficionados, is *Hermes and the golden thinking machine*, which is an AI-based murder mystery that smugly prides itself on having a technical bibliography at the end. Hermes is a professor of archaeology and artificial intelligence who finds himself a suspect in the murder of his uncle, from whom an ancient Greek computing device has been stolen. With his amazing laptop computer and Nina, his "attractive young disciple", Hermes investigates the crime.

The characters are cardboard, the dialogue is more stilted than Charlotte Brontë's, and the text has the feel of a translation from another language. Stretches such as that in which Hermes gives Nina a quick tutorial in search-space pruning have all the excitement of a quick tutorial in search-space pruning. On page 38, Hermes tells Nina that Lisp is "a computer language for manipulating symbols in a humanlike manner". Still, the interweaving of the principles of AI with a murder mystery is an interesting, perhaps unique, achievement.

Graeme Hirst, self-styled expert on coffee tables, is a former editor of this magazine. Diana Latent is the pseudonym of his attractive young disciple.

BOOKS RECEIVED

Books listed below that are marked with a + will be reviewed in a future issue. Reviewers are still sought for those marked with a *. Authors and publishers who wish their books to be considered for review in *Canadian AI* should send a copy to the book review editor at the address below. All books received will be listed, but not all can be reviewed.

Readers who wish to review books for *Canadian AI* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

+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

Artificial intelligence and learning environments

William J. Clancey and Elliot Soloway
(Institute for Research on Learning and University of Michigan) Cambridge, MA: The MIT Press, 1990, 162 pp
(A reprint of *Artificial intelligence*, 42(1), 1990, special issue)
Paperbound, ISBN 0-262-53090-2, US$19.95

Current research in natural language generation

Robert Dale, Chris Mellish, and Michael Zock (editors)
(University of Edinburgh and CNRS)
London: Academic Press, 1990, ix+356 pp
(Cognitive science series)
Hardbound, ISBN 0-12-200735-2, $45.00, pounds 23.50

+Representations of commonsense knowledge

Ernest Davis
(New York University) San Mateo, CA: Morgan Kaufmann, 1990, xxiii+515 pp (Morgan Kaufmann series in representation and reasoning) Hardbound, ISBN 1-55860-033-7, US$42.95

Antilinguistics: A critical assessment of modern linguistic theory and practice

Amorey Gethin
Oxford: Intellect Ltd, 1990, iii+275 pp Hardbound, ISBN 1-871516-00-5, pounds 14.95

Ada programmer's handbook

Dean W. Gonzalez
Redwood City, CA: Benjamin/Cummings, 1991, 230 pp
Paperbound, ISBN 0-8053-2529-8, $13.95

Machine learning, volume III

Yves Kodratoff and Ryszard Michalski (editors)
(CNRS and George Mason University) San Mateo, CA: Morgan Kaufmann, 1990, x+825 pp Hardbound, ISBN 1-55860-119-8, US$49.95

Computers and conversation

Paul Luff, Nigel Gilbert, and David Frohlich (editors)
(University of Surrey) London: Academic Press, 1990, 284 pp (Computers and people series) Hardbound, ISBN 0-12-459560-X, $45.00, pounds 23.95

The logical structure of English: Computing semantic content
Allan Ramsay
(University College Dublin)
London: Pitman, 1990, iv+209 pp
Paperbound, ISBN 0-273-03287-9, pounds 19.99

Artificial intelligence: An engineering approach
Robert J. Schalkoff
(Clemson University) New York: McGraw-Hill, 1990, xxiv+646 pp (Distributed in Canada by McGraw-Hill Ryerson Ltd) Hardbound, ISBN 0-07-055084-0

+ Readings in uncertain reasoning
Glenn Shafer and Judea Pearl (editors)
(University of Kansas and University of California, Los Angeles) San Mateo, CA: Morgan Kaufmann, 1990, x+768 pp (Morgan Kaufmann series in representation and reasoning) Paperbound, ISBN 1-55860-125-2, US$38.95

Computational models of scientific discovery and theory formation
Jeff Shrager and Pat Langley (editors)
(Xerox PARC and NASA Ames Research Center)
San Mateo, CA: Morgan Kaufmann, 1990, xi+498 pp
(Morgan Kaufmann series in machine learning)
Hardbound, ISBN 1-55860-131-7, US$39.95

Meaning-text theory: Linguistics, lexicography, and implications
James Steele (editor)
(Carleton University) Ottawa: University of Ottawa Press, 1990, xi+439 pp Paperbound, ISBN 0-7766-0289-6, $35.00

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 _____

_____

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