Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Winter 1994 | No. 33 | hiver 1994

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

Photo courtesy Daishowa-Marubeni International Ltd.

Daishowa-Marubeni International Ltd. Operators in control room monitoring pulp mill processes.

Call for Nominations

See page 4

Retrospective

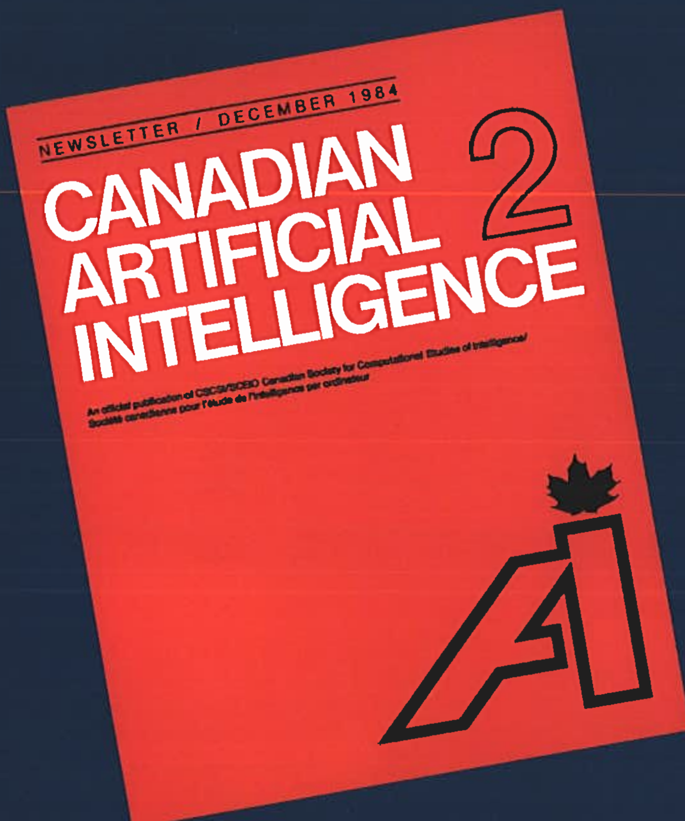

Contents

Cover graphic, contents from Issue Number 2, December 1984
Canadian Artificial Intelligence Newsletter

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Winter 1994 | No. 33 | hiver 1994

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

ISSN 0823-9339

Contents

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

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

Canadian Society for Computational Studies of Intelligence

Treasurer's Annual Report for Fiscal Years 1991 and 1992

Annual Report 1/1/91 Through 12/31/91

Balance as of 1/1/91	**33,326.93**
(not including monies on hand at CIPS)	
Income	
Advertising[1]	10,113. 66
AI-92 Sponsorship Income	8,000. 00
Credit Memo from Bank	13. 87
GST Collected	371. 70
GST Refund	510. 93
Bank Certificate Interest	3,211. 97
Signature Account Interest	619. 88
Proceeding Sales (M. K.)	2,299. 27
Total Income[2]	25,141. 28
Expenses	
Magazine Production[4]	26,538. 18
AI-92 Conference Expenses	2,678. 99
GST paid out[3]	1,250. 64
Bank Service Charges	64. 27
Provincial Taxes[5]	82.32
Total Expenses	30,614. 40
Year Total[6]	-5,473. 12
Monies on Hand with CIPS	17,000. 00
Balance as of 31/12/91	**44,853. 81**

Notes

1) The income from sponsorship for the AI-92 conference had accumulated to $8,000.00 by year end and only $2,678.99 of that had been expended. Therefore the year end deficit is actually greater by $5,321.01 as this represents the amount held by the society on behalf of the conference.

2) Income from membership fees did not appear within this year's annual report as it had not been transferred from CIPS by year end. The society's income from membership dues in the fiscal year was $14,360 but after paying all costs and charges to CIPS we had an effective income from this source of $9,803 for the year. At the year end, the society had on hand with CIPS some $17,000 which had been accumulated by them over more than one year. CIPS have changed their accounting method and from 1/1/92 we are in receipt of monthly fees from CIPS to the society's treasurer.

3) As will be noted, the society is now registered for GST and this registration took place during the latter part of 1991. The need to register was brought on by the requirement from CIPS, who operates our membership collection activity (for a charge!), that we come in line with all its other SIGS otherwise it would charge us an additional handling fee for a non-GST account. Although this year's accounts do not show a net benefit from registering, it is likely that we will be able to make substantial recovery on the GST that we do pay out.

4) The cost of producing the magazine during the year ($26,538.18) was considerably greater than the income that we were able to achieve from advertising ($10,113.66). This has resulted in a cost to the Society of $16,424.52 for the production of the magazine. In an attempt to keep these costs under control, there were only three editions of the magazine published in 1991. Over the last two years the average cost of providing the magazine has been approximately $7,050 per edition .

5) Some of our activities outside of Alberta have been subject to provincial taxes which we have attempted to track.

6) This is rather a crude measure of our activity in 1991 as it includes the inflow of sponsorship for the AI-92 conference and fails to account for the membership dues still on hand with CIPS. If we were to take out the sponsorship monies and expenses for the conference and add in the membership dues, the year balance would look slightly different (Total income [minus AI-92 sponsorship but including dues = $26,944.28, Total expenses [minus AI-92 expenses = $27,935.41, Year balance = -$991.13).

During the 1991 fiscal year, the society has tried to keep its expenses in line with its income. To this end we reduced the magazine to three editions for the year and the magazine editor will be presenting some recommendations on various alternatives for this publication.

We are providing the organizers of the AI-92 conference with funds to support the setting up and running costs of the conference. It is expected that the society will achieve a profit from the conference although the society is providing surety for losses.

At the present time, the society is having to pay the cost of GST on our membership dues as they are currently being treated as if they include GST. This results in a reduced income to the society so I would like to recommend that we adopt one of two options:

a) keep the membership dues at its current level ($35) but put GST on top, or,

b) raise the membership dues to $40 but with GST included.

Option (a) would not increase our revenue but only reduce some loss due to GST (we do recover much of it by being able to claim

GST on our costs). Option (b) would increase our revenue and provide us with some much needed breathing space while we find additional sources of income. I would therefore recommend that we adopt option (b) with a base membership fee of $40 including GST.

We are about to elect a new executive committee for the society and we all hope that they will be able to come up with new ways to increase the society's income in order to maintain its present services to its members and at the same time to have the financial flexibility to undertake new initiatives on behalf of the membership. We had hoped this year to obtain some long term sponsorship of the society's activities from a variety of private and governmental sources. To date, this has not been achieved and we may have to re-think the society's priorities for its activities in future years.

Russ Thomas,
Calgary, May 1992

Annual Report 1/1/92 Through 12/31/92

Balance as of 1/1/92	**44853.81**[1]
Income	
Advertising	7,386.54
AI-92 Sponsorship Income	7,488.79
AI-92 Conference Income[3]	4,296.11
Workshop Income	975.00
Membership Income[2]	10,135.76
GST Collected	1,140.48
GST Refund	1,259.23
Interest	987.42
Total Income	33,669.33
Expenses	
Magazine production	11,560.12
AI-92 Conference	21,714.62
GST paid out	1,097.07
GST to Government	738.52
Postage costs	9.10
Supplies	647.15
Service Charges	16.93
Total Expenses	35,783.51
Year Total[4]	-2,114.18
Balance as of 31/12/92	**42,739.63**

Notes

1) The new balance forward includes approximately $2,000 still on hand with CIPS.

2) In 1992, we received cheques for $15,000 more than reported here from CIPS. However, this $15,000 was not 1992 membership income but prior membership income on hand with CIPS returned to us.

3) Accounting is not yet in for AI/GI/VI-92 shared expenses. There may be further net income as little as $0 and as much as $4,000 more in conference income. I understand it is closer to the latter figure, but details are not finalized.

4) The net deficit of $2,114.98 varies with the date transactions are attributed. Figures here are based on date posted rather than date of activity. Some alternate views of the financial picture:

a) Last year the treasurer reported a $5,473.12 deficit prior to accounting for dues collected in 1991 having a value of $9,803 but still on hand with CIPS. (We received in 1992 $15,000 to cover these and other funds for 1991 and before. Note $5,197 in dues is for years previous to 1991.) Because the society held $5,321.01 net sponsorship income for AI-92, the treasurer suggested rolling forward these transactions for a net deficit of -$991.13 in 1991. Using this reasoning, we could add $5321.01 to 1992 AI-92 sponsorship income for an overall positive balance of $3207.03.

b) As noted, accounting for shared expenses of AI/GI/VI is not in. We might receive as much as $4,000 further net income for a positive balance of about $7,300. This is a "best case" view of the financial picture and is about one half of the net income during the last conference year.

c) Magazine production ($11,560.12) exceeded advertising income of $7,386.54. However several large bills incurred last fall were not paid until early 1993. Magazine expenses from Jan 1/93 to May 24/93 are already $10,000. Using an estimated cost of $7,000/issue for each of three issues, we might add about $10,000 to magazine production costs. This gives a deficit of about $2,700 in a year we previously enjoyed a substantial earning. (As a side effect, this would mean we would be enjoying a small earning as of mid 1993.)

In 1990, the last conference year, the treasurer reported a surplus of about $15,000, which included a net profit on the conference of $13,000 reported for 1990. Over 1992, we are looking at a net loss on AI-92 of about $6,000-$10,000 (although the conference broke even when 1991 sponsorship income is included). There appear to be other differences with the 1990 methods of accounting. Cost of magazine production remains constant and advertising income has dropped $5,000. Membership income, after dues increases, remains about the same as 1990.

CSCSI has two major activities as far as financial transactions go — the conference, which is biannual, and the magazine, which appears three times a year. The conference creates spikes in expenditure and income that complicate the overall picture. I propose to keep accounting for the conference completely separate from the rest of CSCSI, and furthermore to track magazine expenses and income on a per issue basis. This would mean two treasurer's reports, but at the expense of slightly more work, this would simplify the treasurer's job by making it easier to get an understanding of the organization

Eric Neufeld
Saskatoon, August 1993

President's Notes

CSCSI/SCEIO ELECTION 1994

Recommendations are solicited for officers (President, Vice-President, Secretary, and Treasurer) of CSCSI/SCEIO for the period from 1 June 1994 to 31 May 1996. Recommendees must be members of CSCSI/SCEIO in good standing, and must be willing to serve as officers for the two-year period. Nominations should be sent to me (preferably e-mailed) to be received by 1 February 1994.

The election for officers will be held by mail. Ballots will be sent out in late February 1994, and must be returned to me at the address below by the end of March.

Des mises en candidature sont sollicitées pour les postes d'officiers (Président, Vice-Président, Secrétaire et Trésorier de CSCSI/SCEIO pour la période du 1er Juin 1994 au 31 Mai 1996. Les candidats doivent être des membres en règle du CSCSI/SCEIO et s'engagent à siéger comme officiers pour une période de deux ans. Les nominations devraient m'être envoyées personnellement (de préférence par courrier électronique), au plus tard le 1er Février 1994.

Les élections pour les officiers auront lieu par la poste. Les bulletins de vote seront envoyés vers la fin de Février 1994, et doivent m'être retourné à l'adresse ci-dessous vers la fin de Mars 1994.

CALL FOR NOMINATIONS CSCSI Distinguished Service Award

This award is presented biannually to an individual who has made outstanding contributions to the Canadian AI community in one or more of the following areas:

- Community service • Research • Training of students
- Research/Industry interaction

The award, which will be presented at the CSCSI conference, provides:

- honorary lifetime membership in CSCSI
- conference fees when the award is presented
- a token gift.

The first award was presented in 1992 to Professor John Mylopoulos of the University of Toronto.

Recommendations for the award should be addressed to me to arrive by 31 March 1994. They should include a brief (1 page) summary of the nominee's qualifications for receiving the award. The final decision will be made by the CSCSI executive.

It is very much hoped that the winner will be present to receive the award in person, and we ask nominators to make discreet enquiries in advance to ensure that there is a reasonable likelihood of their nominee attending the conference.

APPEL DE CANDIDATURES Prix de distinction de la SCEIO pour service rendu

Cette récompense est présentée deux fois par année à un individu qui aura fait une contribution hors de l'ordinaire à la communauté d'Intelligence Artificielle au Canada dans au moins un des secteurs suivants:

- Service à la communauté • Recherche
- Formation d'étudiants • Intéraction Recherche/Industrie

La première récompense fut présentéé en 1992 au professeur John Mylopoulos de l'université de Toronto.

La récompense qui sera présentée à la conférence de la SCEIO, comprend:

- Carte de membre à vie à la SCEIO à titre de membre honoraire
- Accès gratuit à la conférence à laquelle la récompense sera attribuée
- Cadeau d'appréciation

Les nominations devraient m'être adressé personnellement au plus tard le 31 Mars 1994. Elles devraient inclure un bref (1 page) sommaire des qualifications du candidat. La décision finale sera rendue par l'exécutif de la SCEIO.

Il serait souhaitable que le gagnant puisse venir chercher sa récompense en personne, et nous demandons aux parrains de s'informer auparavant de la disponibilité de leur candidat, quant à la possibilité que le candidat assiste à la conférence en vue de recevoir son prix.

Nominations for officers of CSCSI/SCEIO and for CSCSI Distinguished Service Award are to be addressed to:

Janice Glasgow , President, CSCSI
Department of Computing and Information Services
Queen's University, Kingston, Ontario K7L 3N6
email: janice@qucis.queensu.ca

Integrated Distributed Intelligent System for On-Line Monitoring and Control of Pulp Processes


Ming Rao


Qijun Xia

Ming Rao and Qijun Xia
Intelligence Engineering Laboratory
Department of Chemical Engineering
University of Alberta
Edmonton, Alberta T6G 2G6

Abstract

This paper introduces an integrated intelligent system, namely intelligent on-line monitoring and control system (IOMCS) for pulp processes. IOMCS is built upon the existing distributed computer system (DCS) and information management system (IMS). It has functions such as process upset monitoring, fault diagnosis, decision support, quality prediction and operation optimization. The objective of developing IOMCS is to increase the "intelligence" of the information management systems by applying artificial intelligence technology to enhance pulp process operation from information management to decision automation.

Résumé

Cet article introduit un système intelligent intégré, nommément le système intelligent pour la surveillance et contrôle en temps réel des procédés de transformation de la pulpe (IOMCS). IOMCS est construit sur le système distribué existant (DCS) et sur le système de gestion de l'information (IMS). Ces principales fonctions sont: la surveillance des procédés, diagnostique des défauts, aide à la décision, prévision de la qualité et optimisation des opérations. L'objectif principal d'IOMCS est d'augmenter l'intelligence des systèmes de gestion de l'information en appliquant la technologie d'intelligence artificielle pour améliorer les opérations de transformation de la pulpe de la gestion de l'information à l'automation des décisions.

Introduction

Pulp and Paper industry is resource and energy intensive. Successful control of the related processes will result in significant rewards (Matson, 1989).

Slave Lake Pulp Corporation and Daishowa-Marubeni International Ltd. Peace River Pulp Division in Alberta, Canada, are utilizing the best available technology to produce a high quality pulp product. A distributed computer system (DCS) and a mill wide information system (MOPS) have been successfully installed in the mill. MOPS has such functions as display handling, trend handling, material tracking, statistical process control, cost reports, etc. These functions help operators as well as managers to check the status of the mill quickly, make decisions efficiently, and access the operating conditions for new production of repeated grades. The implementation of MOPS has provided substantial benefits (Frith and Henriksson, 1992).

However, the mill operation still relies on individual operators' experience. Even with consistent and thorough operator training, individual start up experience has led operators to react differently to similar situations. To deal with these problems, the Intelligence Engineering Laboratory (IEL) at the University of Alberta, SLPC, DPRPD and MoDo Chemetics (MC) are working together to develop an intelligent on-line monitoring and control system (IOMCS).

IOMCS is a real-time intelligent system which links with the Fisher Provox DCS through MOPS. It takes advantage of the DCS's value-added data in MOPS. The main functions are: *monitoring process variables and advising evasive or corrective actions to the operator to recover the production process from undesirable situations; predicting changes in the final quality of pulp; and, decision support for normal and upset conditions.* IOMCS is an integrated distributed artificial intelligence application for knowledge-intensive automation.

The project is financially supported by Canada-Alberta Partnership in Forestry, SLPC, DPRPD, MC and Perde Enterprises (PE). MC is the company which developed and commercialized MOPS. Therefore, this is a collaboration project among government, university, pulp companies, an information system research company and a private company.

The research team consists of five research associates and graduate students from IEL and three engineers from SLPC, DPRPD and MC, respectively. IEL develops the intelligent system. SLPC and DPRPD provide process information and knowledge. MC provides MOPS training course and ensures that all interfacing to the MOPS is trouble-free.

Functional Specifications of IOMCS

IOMCS is an integrated distributed intelligent system which fulfills on-line monitoring and control tasks (Rao, 1991). The system aids process operation in two ways: (1) it generates optimal operating conditions in normal situations and corrective actions in abnormal situations and interfaces with MOPS and DCS to influence the pulp processes; (2) it makes feasible operation recommendations to operators who make final decisions and take actions (Rao and Corbin, 1992).

One of the main issues considered in the design of IOMCS is protecting the companies' previous investment. The objective of developing IOMCS is to make full use of the available computer facilities in mills and enhance pulp production from information management to decision automation (Xia, et al., 1993). The intelligent system does not intend to replace the existing DCS and information management system, but to enhance the functionality of the existing systems. IOMCS is a computer software system implemented based on the existing MOPS system. It acquires data from MOPS CVD (current value data base) and HDB (historical data base) and uses the "value-added" data for automatic decision making. The companies need to invest very little in hardware to install IOMCS. Figure 1 presents the interconnection of IOMCS with MOPS and DCS.

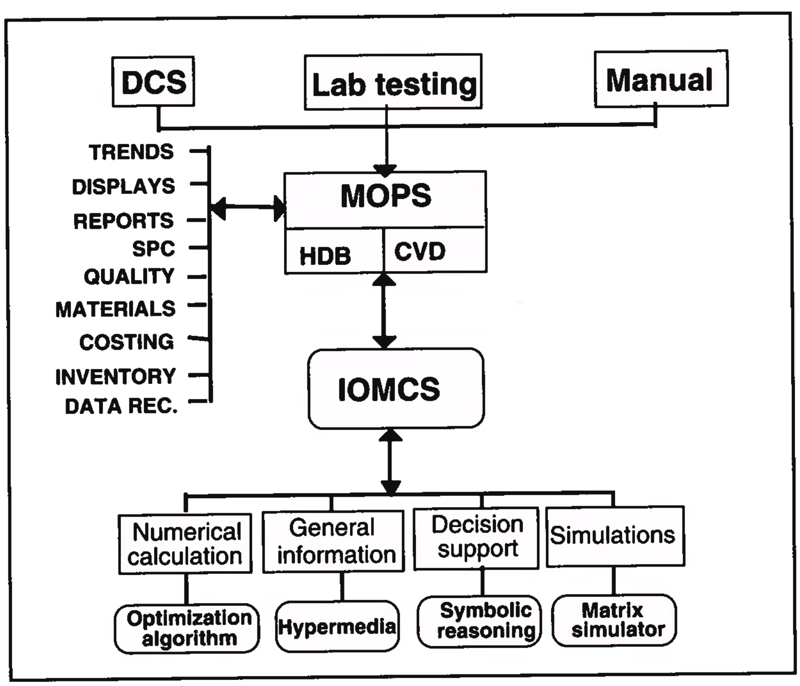

Figure 1. Interconnection of IOMCS with MOPS

Another issue considered is keeping operators in decision making and control loops. IOMCS does not intend to replace human operators in decision making, but to help them by providing them with highly concentrated and understandable information from vast amounts of data. The multimedia operator interface can provide information and explanation with natural language, numerical values, or graphical display. Thus, operators get "closer" to the production processes and perform better.

SLPC and DPRPD required the IOMCS to have the following features:

- standardization of operator changes,
- more timely and knowledgeable decisions,
- identification of complex process situations that have the potential of causing process upsets,
- automatic learning of optimal operation conditions for different product grades, and
- having an open structure to implement the local systems independently and then integrate them into a mill-wide system.

Through a multimedia operator interface, operators are able to:

- view and acknowledge situations and advisories;
- view "how & why" type information and have access to further process information if required,
- scroll through old messages and sort by process area,
- provide an interface for process engineers to program the knowledge base.

Design of IOMCS

IOMCS possesses the following desired properties in order to satisfy pulp production requirements:

Process generality: IOMCS is general and flexible in nature, and not process-specific. The methodology applied enables the systems to easily accommodate the changes in the plant configurations.

Application-specific shells: The use of IOMCS does not require much knowledge in computer science or artificial intelligence. Users can implement IOMCS on a specific process just by making revisions in process-dependent knowledge bases.

Open architecture: Developing a mill-wide IOMCS in one phase is very difficult. IOMCS has an open structure such that different applications can be implemented step by step, independently. In our project, we develop the system for the bleach plant and final mill quality. Next will be a similar system for effluent treatment system. Finally, all these local systems will be integrated into a mill-wide system.

Multiple knowledge representations: To represent different knowledge involved in monitoring and control of pulp processes, IOMCS facilitates multiple knowledge representation techniques, such as facts, rules, numerical models, as well as neural networks.

Based on considerations above, IOMCS is designed using general function modules. Users can select functions from the available function modules and implement them according to the equipment and process layout. Modular technology makes IOMCS more flexible and easier to implement (Soucek, 1991). A number of function modules are identified to be important for pulp processes:

- Matrix simulator,
- Quality prediction,
- Operation optimization,
- Emergency support,
- General information,
- Operation interface,
- Engineering interface, and
- System interface.

IOMCS has a three layer architecture, as depicted in Figure 2.

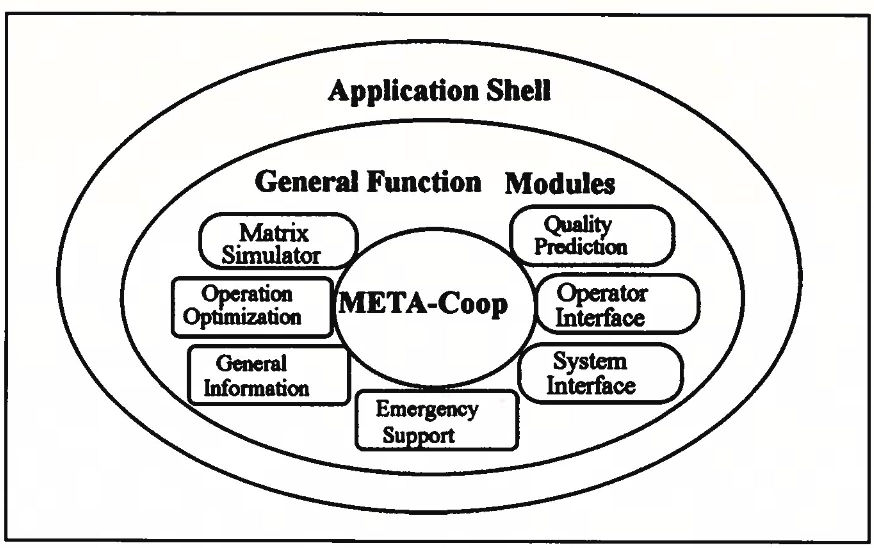

Figure 2. Three layer architecture for IOMCS

The first layer, the core of IOMCS, is an integrated distributed intelligent system shell, Meta-COOP. It provides an inference engine, and facilitates knowledge representations, acquisition, editing, and compiling.

The second layer provides general function modules for on-line monitoring and control. It is a medium between the system core (Meta-COOP) and specific applications. In this layer, a number of modules are developed by implementing the general monitoring and control technologies into computer software supported by Meta-COOP. Each module fulfills a typical monitoring and control function which is necessary in pulp mill operations. Every module can be run independently or incorporated with several other modules.

The third layer is an application-specific shell for users in SLPC and DPRPD to implement IOMCS for some specific processes. It applies the function modules provided in the second layer. The users can simply select some of the function modules, and enter knowledge by using the engineering interface. Therefore, not much experience in computer science and artificial intelligence is required for implementing specific applications. The process operators and process control engineers in pulp mills are able to do it.

The General Function Modules

The function modules fulfill the on-line monitoring and control tasks.

The ***General Information*** module provides general mill information including text-like and graphical information by applying hypertext method. Hypertext is a promising technology for organizing heavily referenced information bases in a more efficient way. It links a keyword or phase to text or graphics further explain the original keyword of phases (Gessner, 1990). In the General Information module, all the mill information is organized into discrete blocks or nodes of information. Links are then used to join these blocks to one another to form the document as a whole. The information includes:

- main process equipment in the mill,
- brief introduction of the process,
- architecture and functions of DCS, MOPS and IOMCS,
- hazard alerts,
- grade recipes,
- mode of operation for all grades,
- procedures for startup and shutdown operation, and
- procedures for emergency operation.

Operators can easily find subjects of interest by going through a menu system or searching for keywords in text or graphics. This module can serve as a training tool for new operators to become familiar with operations. It can also be used for on-line help.

The ***Matrix Simulator*** is a process simulation package and graphical tool built into IOMCS. In the Matrix Simulator, process variables are divided into two parts: action variables and result variables. Action variables are those that affect the product quality and can be changed to correct the process conditions. Result variables, which are affected by the action variables, are either product qualities or those used to evaluate the performance of the production processes. The Matrix Simulator provides the relationship between result variables and action variables. Unlike conventional simulation packages, the Matrix Simulator is built based on the expert system building tool, Meta-COOP.

The Matrix Simulator receives process data from MOPS and identifies the normal operating conditions (product grades and production ranges). According to the normal operating conditions, one relationship between the action variables and result variables is selected. The operator can introduce changes in action variables (perform an operation action) and observe the corresponding changes in result variables from the graphical displays. The Matrix Simulator also displays the trend curve, target, current, and final change of the result variables.

The Matrix Simulator display is produced using an object oriented drawing tool, PICAD, which is a function of MOPS and is PC based. It is displayed in PCs using MOPS' operator station interface program, EDE/2.

Other function modules can call the Matrix Simulator for different purposes.

The ***Quality Prediction*** module applies the Matrix Simulator to predict quality changes. The accurate measurement of some important quality variables such as freeness and brightness is not instantly available. This module tracks the model's state (using forward reasoning in the matrix simulator) with observations from the physical system. The laboratory test results are used to correct the prediction. Due to the complexity of pulp processes, it is impossible to obtain complete process models using conventional methods. On the other hand, there is a lot of historical data available in the pulp mills. Neural network technology is thus applied as a complement of conventional modeling techniques for quality prediction.

The ***Operation Optimization*** module applies backward reasoning in the Matrix Simulator and the real-time value-added information in MOPS to learn the optimal operation conditions which result in the best product quality and lowest costs. An optimization package called SWIFT is developed using FORTRAN language. SWIFT can do numerical optimization with equality and inequality constraints. The optimization criterion is to produce pulp economically, which implies: making the best use of raw materials, controlling waste, limiting fiber losses, and better quality.

The ***Emergency Support*** module monitors undesired conditions and diagnoses the original source of the process upsets (Kowalski, 1991; Kramer and Lenoard, 1990). The sequence of corrective actions to recover the process from the undesired situations will be automatically generated after an upset is detected. The on-line monitoring function is carried out in the following five steps:

- monitoring the process for undesirable conditions (fault detection),
- identifying the initial cause of the situations (fault isolation),
- analyzing the criticality of the situation (fault analysis),
- deciding the corrective actions to be taken (fault handling), and
- report the alarm, cause and criticality of the situation and the suggested actions to operators.

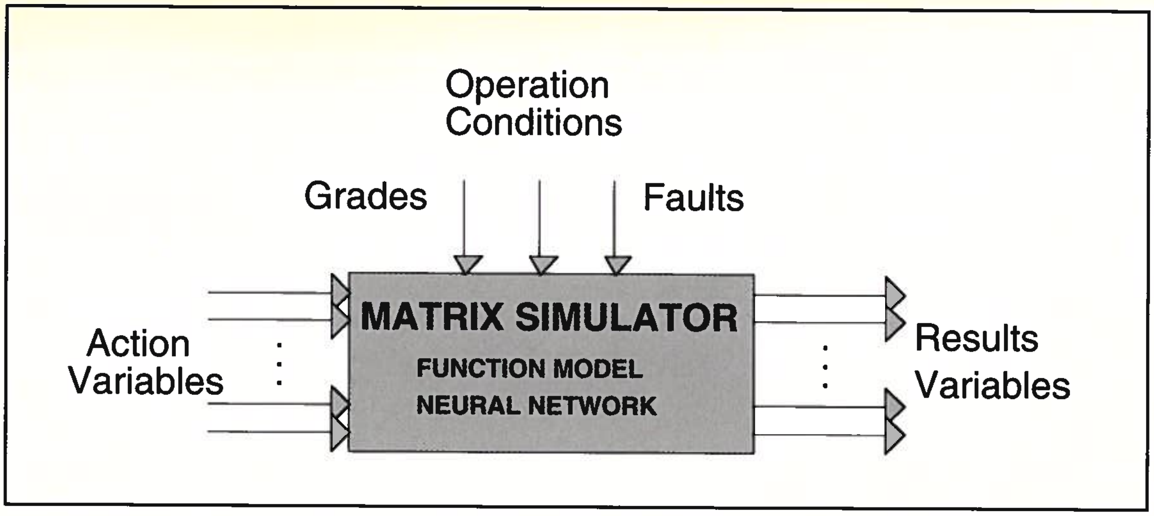

Figure 3. Illustration of matrix simulator

Structure and function modeling, fault propagation modeling, and process hierarchical decomposition techniques are applied in the Emergency Support module.

A failure is defined as a violation of expected plant behavior or certain function constraints. The pulp process is modeled based on both function and structure and built in the Matrix Simulators. A fault propagation model is a fault propagation graph used to express the dynamics of fault evolution. Using a forward direction search of the graph from a component, the system can simulate a failure of that component in terms of a time sequence of symptoms. Using a backward search from the symptoms, we can locate fault sources. The hierarchical decomposition of pulp processes is used to improve the efficiency of fault diagnosis.

The fault detection, analysis and handling basically use forward and backward searching of fault propagation models as qualitative knowledge and the Matrix Simulator as quantitative knowledge. The principle of the fault diagnostic technique can be represented by Figure 4.

The ***System Interface*** provides IOMCS with the ability to communicate with MOPS. By interfacing with the MOPS database, IOMCS can access more than 2000 points per minutes of data collected.

IOMCS communicates with MOPS in two steps: communication with VAX, and with PCs. The communication between IOMCS and VAX is achieved by adding two function routines, "getdata" and "putdata," to IOMCS. The communication between IOMCS and PCs is achieved by using the function routine" matdsp," " usinit," " usgepn," " usput," and " usgtcm." In this way, IOMCS is embedded into MOPS.

The ***Operator Interface*** is used for communication between operators and IOMCS. The ***Engineering interface*** is used by engineers to edit and reorganize the knowledge bases.

System Implementation

The IOMCS for the bleach plant is implemented by using Meta-COOP. The distinct characteristics of Meta-COOP make the application of the above techniques possible.

Meta-COOP distributes its knowledge into a number of knowledge bases. Each knowledge base is a basic object within the Meta-COOP environment called a unit. Each unit has an arbitrary number of slots, in which the attributes of the unit are described. A slot represents one attribute of the unit and has several facets, where the attributes is specified in more detail. There are two types of units: class units and member units. Member units describe individual objects; class units group several objects with common attributes into a single class. Class units can be defined as a subclass or superclass of another unit. The knowledge units are organized in a hierarchy with inheritance properties. With this characteristics, we can apply process hierarchical decomposition techniques to reduce a complex engineering problem to a number of less complex problems.

As a hybrid system, Meta-COOP allows the integration of various knowledge representations and inference methods, such as frame-based, rule based and method-based, external procedures written in any other language, and internal subroutines written in C.

The knowledge units of IOMCS for the bleach plant are organized into four layers. The operator interface of the system provides illustrative and graphical explanations, such as trend curve of important operating conditions and quality variables, for operators to easily understand the current operation situations and the action to be taken.

Conclusions

The intelligent on-line monitoring and control of pulp processes is an important and challenging research field. The presented system provides powerful operation support

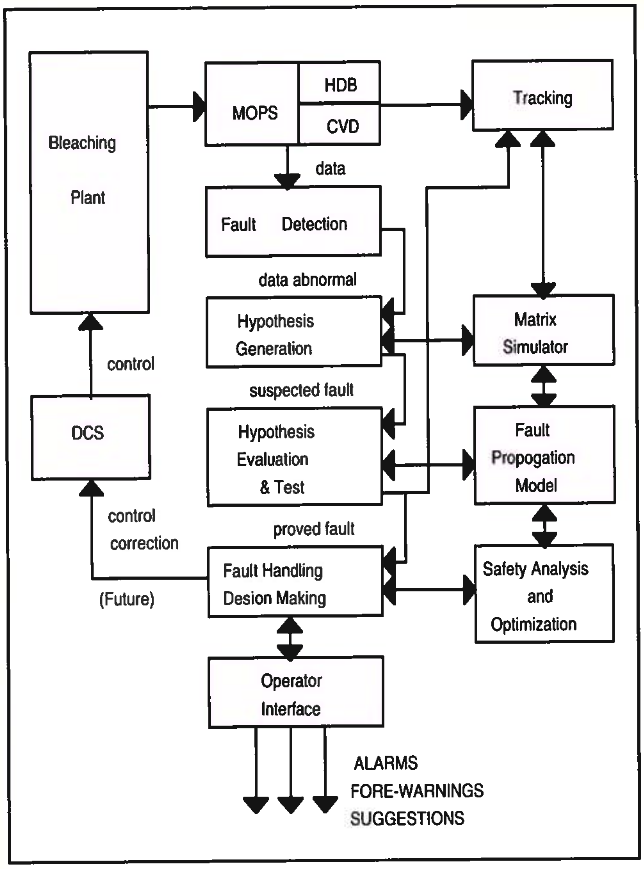

Figure 4. Principle diagram of emergency support system

and on-line advice functions. The techniques developed can be also applied to develop high-performance intelligent systems for many complicated engineering applications.

References

Frith, M.D. and Henriksson, C. (1992). Integration of Distributed Control and Mill Wide Information Systems at the Slave Lake Pulp Corporation Plant. *CPPA, 1992 Spring Conference*, Jasper, Alberta, Canada.

Gessner, R. (1990). Building a Hypertext System. *Dr. Dobb's Journal*, June, 22-33.

Kowalski, A. (1991). Diagnostic Expert System For Solving Pitch Problems. *Pulp and Paper Expert Systems Workshop*, November, Pointe Claire, Quebec.

Kramer, M.A. and Leonard, J.A. (1990). Diagnosis Using Backpropagation Neural Networks—Analysis and Criticism. *Computers and Chemical Engineering*, 12, 1323-1338.

Matson, W. (1989). The Outlook for the Canadian Pulp and Paper Industry. *Pulp and Paper Canada*, 90, 297-302.

Rao, M. (1991). *Integrated System for Intelligent Control*, Springer-Verlag, Berlin.

Rao, M. and Corbin, J. (1992). Intelligent System for Batch Digester Operation. *78th Annual Meeting, Technical Section, CPPA*, B81-85, Montreal, Canada.

Soucek, B. (1991). From Modules to Application-Oriented Integrated Systems. in *Neural and Intelligent Systems Integration: Fifth and Sixth Generation Integrated Reasoning Information Systems*, pp. 1-36, John Wiley & Sons, Inc.

Xia, Q., Farzadeh, H., Rao, M., Henriksson, C., Danielson, K. and Olofsson J. (1993). Integrated Intelligent Control System for Peroxide Bleaching Processes. *Proceedings 1993 IEEE Conference on Control Application*, Vancouver, B.C., Sept. 13-16.

***Ming Rao** is an associate professor of chemical enginneering/computing science, and the director of the Intelligence Engineering Laboratory, at the University of Alberta. He received his B.S. in Chemical Engineering from Kunming Institute of Technology, M.S. in Computer Science from University of Illinois, and Ph.D. in Interdisciplinary Engineering from Rutgers University. He has 5 years of industry experience, and is a consultant for several industrial companies. In 1989, he worked in FAA (Federal Aviation Administration, U.S.A.) Technical Center, on intelligent control. Dr. Rao has published 5 books and over 150 technical papers. He is a member of AAAI, AIChE, CPPA, CSChE, IEEE, Sigma Xi and Tau Beta Pi.*

***Qijun Xia** is a research associate in the Department of Chemical Engineering at the University of Alberta. He received his B.Sc., M.Sc., and Ph.D. in Chemical Engineering from Zhejiang University, P.R. China. His research interests include process modeling and control, fault diagnosis and fault-tolerant control, and intelligent control systems for industrial processes.*

AAAI Calendar

AAAI Spring Symposium Series 1994, March 21-23, 1994, Stanford University, California. Contact: Spring Symposium Series 1994, AAAI, 445 Burgess Drive, Menlo Park, CA94025; 415-328-3123; FAX: 415-321-4457; e-mail: sss@aaai.org

AAAI-94, Twelfth National Conference on Artificial Intelligence July 31 - August 4, 1994, Seattle, Washington. Contact: AAAI-94, 445 Burgess Drive, Menlo Park, CA94025; 415-328-3123; FAX 415-321-4457; e-mail: ncai@aaai.org

IAAI-94, Sixth Annual Conference on Innovative Applications of Artificial Intelligence, July 31 - August 4, 1994, collocated with AAAI-94, Seattle, Washington. Contact: IAAA-94, 445 Burgess Drive, Menlo Park, CA94025; 415-328-3123; FAX 415-321-4457; electronic mail: iaai@aaai.org

DataLogic/R: A Tool for Mining Knowledge in Databases

Adam J. Szladow
REDUCT Systems Inc.,
Regina, Saskatchewan, Canada

Introduction

DataLogic/R est un logiciel d'exploitation de base de donnée pour travailler avec des bases de données incomplètes et imprécises. Ce logiciel est basé sur des théories d'interprétation de connaissance et de logique inductive. Ce qui rend DataLogic/R unique est qu'il analyse des modèles logiques de données à différents niveaux d'interprétation de connaissance. Ce qui signifie qu'il peut découvrir des faits et des relations qui ne sont pas disponibles avec d'autres méthodes.

Datalogic/R possède plusieurs caractéristiques uniques en ce qui a trait à l'acquisition de connaissance. Il peut aider les analystes et les ingénieurs de connaissance dans des tâches comme la classification, conception de système expert.

DataLogic/R peut déceler automatiquement des modèles logiques dans d'immenses volumes de données et générer le plus adéquatement possible des règles de prédiction basés sur ces modèles. Ces règles peuvent être utiliséés pour l'aide à la décision, prévision. ...Les possibilités analytiques de ce logiciel amène une puissante conceptualisation des relations entre les données et les modèles.

Introduction

DataLogic/R is a database "mining" software for working with incomplete and imprecise databases. The software is based on the theories of knowledge representation, inductive logic and most importantly, rough sets. What makes DataLogic/R unique is that it analyzes logical patterns in data at different levels of knowledge representation. This means that it can discover facts and relationships that are not accessible with any other method.

DataLogic/R provides several unique and powerful features for knowledge acquisition and machine learning. It can assist analysts and knowledge engineers in tasks like classification, predictive modeling, and expert system building.

DataLogic/R can automatically discern logical patterns in large volumes of data and generate the best possible prediction rules based on those patterns. The rules can be used for decision support, forecasting, and more. The analytical capabilities of the software provide strong conceptualization of data relationships and patterns.

History

REDUCT's "sister" company, Lobbe Technologies Ltd, has worked with advanced process control and engineering simulators since the mid 80's. In our work, we were looking for alternative technologies that could uncover patterns in process data and present them as simple rules. We sought methodology that could:

- work with imprecise and incomplete data
- work with numeric and symbolic variables
- represent patterns in an easy to understand manner

The methodology also had to be easily integrated with other AI methods for development of computer assisted decision systems.

That's when we came across rough sets. The rough sets method appealed to us because it satisfied all the above requirements for pattern recognition and representation, plus it provided strong data analysis and modeling capabilities.

We started with a selected application of rough sets in process control, specifically, adaptive control and hybrid expert systems development. After a while, we realized that the technology would be valuable to many users who have large amounts of data and need tools for reasoning from data.

We recognized that rough sets technology would have many applications in business, industry, and science. Our objective was to develop a tool for knowledge discovery from databases. The knowledge or organized information was to be presented in the form of simple rules. We also added more requirements concerning what the technology should do. The rules derived had to be adequate for use in predictive modeling or forecasting with capabilities comparable to neural networks or other available machine learning methods. We also combined the approaches of knowledge acquisition and machine learning in DataLogic/R.

In early 1991, we started a new company, REDUCT Systems Inc., with the objective of providing packaged and custom software for knowledge acquisition and decision support. We began by developing DataLogic/R, a tool for decision makers who want to "mine" as much knowledge as possible from data in their computers.

The Team

After a period of exploratory work on rough sets, a team was formed at REDUCT Systems to work on software design, development, and marketing. The team consisted of four core professionals responsible for rough sets research, programming, technical support (quality control and customers), marketing, and project management. REDUCT staff covered all areas except the theory and research on rough sets, which has been provided by Dr. W. Ziarko of the University of Regina. It was clear to us that continued exploratory and applied research will have to be carried out to build stronger products and to identify potential new applications for rough sets.

Throughout the project, we have insisted that there be some overlap of the responsibilities and expertise among the key team members. For example, programmers became familiar with the theory of rough sets, marketing and customer support staff exchanged daily information on customers, etc. The existence of the team, its close working relationship and the application of formal project management techniques where appropriate have been key to timely and successful development of DataLogic/R.

A broad understanding of all issues by the team members has been emphasized throughout the project to achieve maximum synergism from each individual's knowledge and expertise. The team was also supported by other professionals when required.

At REDUCT, DataLogic/R means to have the capacity to acquire knowledge from large databases and to express this knowledge in form of explicit rules. The way in which DataLogic/R develops these rules is explained below.

Rough Sets

The theory of rough sets is a new approach to reasoning from data. It provides a collection of tools for data analysis, pattern discovery, and characterization of regularities in databases. Since the early 1980's, over 400 technical papers have been published on rough sets all over the world.

The key idea behind rough sets is that the data which we collect is too refined for a clear representation of knowledge. To change data into knowledge, we need to look at various levels of knowledge representation, going from refined to coarse data and vice versa. The methodology of rough sets formally recognizes that an exact description of data patterns may not always be possible due to the incomplete or imprecise nature of available information. Consequently, the rough sets characterization of information is based on the idea of approximately defined sets or rough sets. By applying the theorems of rough sets it is possible to:

- deal with uncertainty in data
- analyze hidden facts in data
- find a minimal representation of knowledge
- create classification\predictive models

Rough sets are used in DataLogic/R to focus on knowledge discovery and expert advice processes. DataLogic/R analysis and pattern discovery involves a number of computational procedures which can be summarized as follows:

1. elimination of redundant attributes while performing the dependency analysis
2. elimination of redundant data while performing calculations of approximation space
3. uncovering data patterns while generating classification rules
4. auditing and validating rules by means of cross-validation

After the rules are generated, they can be used to predict the outcome values for new data. All system decisions are traceable to rules and, through the Rule Strength Report, to cases supporting the rules. The user can inspect all decisions for their relevance to previous experience (data cases). We demonstrate below how DataLogic/R works.

Datalogic/R at Work

We have written DataLogic/R in C so that it will be easy to implement on a variety of computers. The packaged version of DataLogic/R runs on the PC AT. It works on a maximum of 2000 variables per case (in DataLogic/R+ version). The customized DataLogic/R systems' requirements depend on the size of the database and the required response time.

An Example

Here is an example of DataLogic/R+ at work. We have selected an example of "mining" a marketing database to demonstrate the knowledge acquisition capabilities of DataLogic/R+. The problem was quite challenging. It had a large number of variables and cases, and relatively weak and noisy patterns. The database selected was a consumer product survey database containing information on the use of over a thousand products in several hundred thousand households. In total, over 1,300 variables, including demographic information relevant to consumer behavior patterns, were available in the database. In our example, in addition to information on the use and purchase of different products, the data contained information on various interests or hobbies actively pursued by the households surveyed. Our objective was to identify profiles, in terms of the type of interests pursued and products purchased of the households, in which Fitness/Exercise was an interest or hobby.

The analysis of the problem started with the analysis of the patterns and key variables. First we looked at exact patterns (100 percent probable) and the relevant attributes. The rules discovered for the exact patterns (see Table 1) were supported by only a limited number of cases. It was evident that the patterns were also weak and noisy and that the discriminating power of the available attributes was poor.

Next, we looked at inexact rules. This feature has been provided in DataLogic/R for the discovery of knowledge from imprecise and/or noisy data. The inexact or imprecise

Table 1.

EXACT RULES FOR PHYSICAL FITNESS PROFILES

Rule #	Rule Description	Rule Prob.%	Support Cases,%
1	other hobby actively pursued is travel, does not use sugar substitutes, uses pain relievers and colds, not interested in gourmet cooking	100	1.1
2	employed female, other hobby actively pursued is travel, not interested in gourmet cooking, does not have upset stomach symptoms	100	1.5

Table 2.

PROBABILISTIC RULES FOR PHYSICAL FITNESS PROFILES

Rule #	Rule Description	Rule Prob.%	Support Cases ,%
1	other hobby actively pursued is travel, interested in eating more natural foods	69.7	14.3
2	female, wants to maintain or lose weight, interested in art/cultural events	67.7	5.3

patterns describe logical patterns combined with probabilistic information. The rules generated this time were strong but not exact, i.e. the rule decisions were met only at certain predetermined probabilities. Table 2 presents examples of the probabilistic rules. These rules (profiles) offered better understanding of the relationships between consumption of various products as well as descriptions of consumer behavior. The rules focus on "who buys what" which is important in micro-markets research.

The results obtained with DataLogic/R+ are generally superior to traditional statistical techniques which describe consumer profiles but not what to do about them. The rules are also accurate. By testing with new test data, we could correctly identify over 84 percent of the households with an interest in physical fitness and exercise.

Other Applications

DataLogic/R has been used for knowledge discovery and database mining in areas such as:

- agricultural research
- toxicology
- drug design
- medical diagnosis
- materials research
- spectral analysis

Also, DataLogic/R has been used for its machine learning capabilities in areas such as:

- process control
- fault diagnosis and detection
- adaptive control robotics
- business credit evaluation
- stock predictions
- information processing

to name a few examples. Since its introduction at the AAAI-92 Conference in San Jose, in July 1992, over 100 copies of DataLogic/R have been used in a broad range of applications.

Future Plans

When designing DataLogic/R we decided to incorporate a wide range of capabilities of rough sets - as much as a general purpose software package would allow us to do. DataLogic/R is designed for decision makers, analysts, researchers and engineers looking for new ways to squeeze knowledge from their data.

Future products will be tailored to meet more specific applications and vertical markets such as fault diagnosis in electronic circuitry, etc. We will also keep developing new DataLogic/R capabilities. We see a broad range of opportunities for DataLogic/R.

Suggested Reading

Pawlak, Z., *Rough Sets: Theoretical Aspects of Reasoning About Data.* Dordrecht, The Netherlands: Kluwer Academic Publishers, 1991.

Slowinski, R (Ed), *Intelligent Decision Support: Handbook of Applications and Advances of the Rough Sets Theory.* Dordrecht, The Netherlands: Kluwer Academic Publishers, 1992.

Szladow, A., "Datalogic/R: Mining of Knowledge in Databases, "PC AI. vol 7, no. 1, January/February 1993

Szladow, A. and W. Ziarko, "Rough Sets: Working with Imperfect Data," AI Expert. vol 8 no.7, July 1993.

Acknowledgement: NRC assistance under the IRAP-M program for parts of the development is gratefully acknowledged.

***Adam J. Szladow** is President of REDUCT Systems Inc. You can reach him at (306) 586-9408*

Edited by Peter Turney

The DPN Expert Advisor:

Canadian AI Success Story #3

Sue Abu-Hakima
Knowledge Systems Laboratory
Institute for Information Technology
National Research Council of Canada
(613) 991-1231

Sue Abu-Hakima

Résumé

Le Conseiller Expert DPN est un outil générique, non spécifique à une topologie particulière, utilisé pour la surveillance de réseaux et spécifiquement créé pour surveiller et diagnostiquer les réseaux à transmission par paquet (DPN) de Northern Telecom. Le Conseiller Expert travaille conjointement avec le Conseiller DPN qui est le système de gestion de réseau de Northern pour les réseaux privés, qui est inclus avec les produits DPN. DPN fournit aux clients, des connexions de communication point-à-point Les connexions DPN permettent aux terminaux ASCII à basse vitesse de communiquer avec des terminaux haute vitesse en utilisant des paquets de données synchrone. Le Conseiller Expert DPN fonctionne en temps réel et permet à l'opérateur de surveiller et diagnostiquer des événements réseau et modifier la banque de connaissance en ajoutant des descriptions de problèmes lorsque des incidents se produisent. Les incidents réseaux comprennent des alarmes mineures, critiques, avertissements aussi bien que des messages qui sont alors interprétés par l'opérateur. Les descriptions de problèmes utilisent un format bien spécifique, dans un langage naturel et qui est la propriété de Northern Telecom. Le Conseiller Expert DPN utilise une interface graphique pour entrer les descriptions de problèmes selon le format de Northern Telecom, pour ainsi s'intégrer complètement aux écrans utilisés fréquemment par l'opérateur.

Abstract

The Data Packet Networks (DPN) Expert Advisor is a generic, non-topology specific tool for network surveillance, used to monitor and diagnose Northern Telecom's Data Packet Networks. The Expert Advisor works in conjunction with the DPN Advisor, Northern's Network Management System (NMS) for private data networks within organizations who buy Northern's DPN products. Data packet networks provide customers with end-to-end communication connections. DPN connections allow low speed ASCII terminals to communicate with high speed terminals using synchronous packets. The DPN Expert Advisor performs in real-time and allows an operator to actively monitor and diagnose network events and modify the knowledge-base by adding Problem Descriptions as network events arise. Network events include major, minor, and threshold alarms as well as status messages that the operator interprets. The Problem Descriptions are input using an English-like language, proprietary to Northern Telecom. The DPN Expert Advisor has a window and menu-based user interface that has been fully integrated with the displays and browsers the network operators currently use.

History

The DPN Expert Advisor emerged from an earlier project, the Datapac ADvisor (DAD), which was prototyped for monitoring Datapac network alarms and advising Bell Canada operators at the National Data Network Center (NDNC). The work on DAD was started in 1987 at the Bell-Northern Research (BNR) Knowledge Technology group headed by Dick Peacocke. The project manager for the DAD prototype - and subsequently the DPN Expert Advisor product - was Sameh Rabie. The originators of the DAD prototype were Andrew Rau-Chaplin and Taro Shibahara. The project champions at Bell Canada were Jim Kennedy and Norm Robertson.

The first DAD prototype, implemented in the KEE shell from Intellicorp, was entered into a one-month field trial at Bell Canada's NDNC in 1988. DAD used a network event simulation and alarm files to advise the operator. The field trial was successful and the prototype was ported to Pascal running on an XMS, a proprietary workstation from Northern

Telecom by Chris Baird and Tony White at BNR. The main difference between the first and second prototypes was that the second ran in real-time and accepted alarms directly from the network. As is typical in many expert system developments, the KEE prototype was abandoned due to its inability to deliver real-time performance. The second field trial was completed in six weeks and was very successful; the operators were quite pleased with a tool that could be integrated with the network and provide expert advice on network events. Richard Lavictoire, who worked at NDNC at the time of the DAD trials, believes that they provided the developers with a day-to-day view of the problems operators were trying to solve. This lead to a system that was built to satisfy a real need. A simple mathematical cost-benefit model indicated that NDNC could save $700,000 in three years in operator time and network outage costs. This provided additional incentive to develop a product based on the prototypes.

A decision was taken early in 1989 to integrate the expert system functionality of the DAD prototype with the DPN Advisor. The DPN Advisor is a network management tool sold as a product to Northern Telecom's private data network customers. The expert system had to adopt the DPN platform and programming environment. As a result, the DPN Expert Advisor was re-implemented at BNR in C/C++ in the NT Signature windowing environment, which is now being evolved to X/Motif under the X Windows environment. The implementation was completed in early 1991 by Tony White and Andrzej Bieszczad and appeared in release 3 of the DPN Network Management System. A critical technical difference between DAD and the DPN Expert Advisor is that the latter included message passing between Problem Descriptions so that relational information could now be used for diagnosis. The main concepts in the knowledge representation of the DAD formalism were carried over. Another important extension was the implementation of a customization environment that included a simulator and a debugger which allows a DPN operator to modify or add knowledge.

The development of the DAD prototypes and the DPN Expert Advisor consumed a total of 10 person-years in its implementation between 1987 and 1991. It is currently supported by BNR with the efforts of 2.5 people per year, mainly Pat Paley, Junming Gao and their manager Bill Brocklebank. It is now in use as an integrated DPN Expert Advisor product. There are approximately 100 Network Management Customers worldwide of whom approximately 50 have also purchased the Expert Advisor capability.

Choice of Problem

The problem addressed by the DPN Expert Advisor is typical of those faced in many telecommunications systems where streams of real-time data need to be interpreted and used for system management. Finding solutions to these problems is of strategic importance to BNR and Northern Telecom, and that is why Peacocke and Rabie started the work in this area.

Advising operators on the events that a data packet network generates is essential due to the hundreds of events that are generated with network component problems. For an operator to interpret these events without automated tools is a tedious task which can be very tiring. The DPN Expert Advisor performs quick searches through event databases which also helps an operator isolate a network problem. According to one of its developers, Tony White, the Expert Advisor has the capability to filter data alarms by a ratio of almost 10:1.

An advanced surveillance system such as this increases the confidence the users have in the reliability of the transmission of their data. The Expert Advisor reduces network outage times significantly. The system can also be customized to issue operator commands to the network, so obvious problems can be fixed faster. The system can also query a network about a possible problem; for example, a query can be sent to a port to find out whether the service characteristics of a line are normal. Such queries can be used for automated diagnostic functions or in support of operator troubleshooting. Also, the advisor has the capability of providing the operator with audio (synthesized voice) warnings about problems.

Description of the System

The DPN Expert Advisor is a generic non-topology specific tool designed to monitor and diagnose Northern Telecom's Data Packet Networks. The DPN Expert Advisor identifies problems from a stream of network events. Network events include those generated from devices and network connections. The Expert Advisor's main objective is to collect in a single object - a Problem Description - all the events relating to a common fault from the network with support from other data sources such as event archives. The Problem Description (PD) is then presented in a graphical user interface to the operator charged with monitoring the network. The PD information is presented to the user as a problem hierarchy, which can be browsed through the use of successive windows that provide table-driven, textual summaries. The hierarchy relates components through names and functional relations on the network. According to Andrzej Bieszczad, the operators find these summaries very convenient for diagnosing network alarms. The Expert Advisor assists the operator in isolating network faults.

The knowledge is represented in a Problem Description frame with slots containing event filters that define network events that signal the start of the problem, events that are part of the problem once activated, and a database query used to extract relevant events from the event archive. There are five types of events: network alarms, network status messages, Expert Advisor internal alarms, Expert Advisor expectation alarms, and Expert Advisor state messages. The first two events are network events while the other three are internal events.

PDs can be considered classes of problems which are instantiated by network events or deduced facts within the Advisor. Each network component - a logical or physical entity that may have events associated with it - has Problem Descriptions. The Problem Descriptions themselves may also be related across components. Problem Descriptions have three slots which contain rules. The first rule slot is evaluated on the creation of a PD instance, the second is evaluated when an event associated with the instance is detected and the third is evaluated when the instance of the PD is deleted.

The DPN Expert Advisor consists of approximately 200,000 lines of code. There are 12,000 lines of Problem Description Language statements and 6,000 lines of problem help in the knowledge-base. There are approximately 70 Problem Descriptions to date.

A major aspect of the system is the customization tool which allows technical support personnel, such as Richard Lavictoire in Northern Telecom's DPN product group, to add new problem descriptions and verify their integrity and relation to the remainder of the knowledge-base. The customization tool consists of the Debugger, the Problem Description Language Compiler, the Simulator, and the Problem Viewer. An Application Program Interface was developed to allow the sharing of information in ASCII text between the various modules. The customization tool is used in an off-line mode since the entire knowledge base must be included in the verification process. This tool is essential for the evolving network which has new components added to it continuously. The customization tool is used on demand to create new objects in the knowledge-base to maintain the integrity of the Expert Advisor. The user makes the changes to the knowledge-base and then verifies these changes by running a simulation of the stream of network alarms and the system's responses to check the validity of the new knowledge and its relation to the already existing knowledge. The results are displayed in the Problem Viewer. Knowledge is never eliminated due to the requirement of backward and forward compatibility of network software since at any one time it can never be guaranteed that all components are operating with the latest releases that trigger events.

Managerial Decisions

Both the prototype and the product were managed by Sameh Rabie of BNR. Stentor, formerly Telecom Canada, funded the prototype development while Northern Telecom funded the product development.

Making the prototype a real-time system was a key factor in bringing the Bell Canada network operators onboard. According to Tony White the operators were provided with a tool that was both "useful and useable". Thus the operators wanted to make use of it immediately to filter alarms and browse Problem Descriptions.

A key decision in developing the DPN Expert Advisor was its alignment and integration with an emerging product, namely the DPN NMS Advisor. This permitted the project to survive as an expert system integrated with a mainstream product. It also ensured its continued funding by Northern Telecom.

A drawback to its tight coupling to the network management system product was the necessity of adopting the C/C++ development environment. This forced the development from scratch of tools such as the Debugger and Problem Viewer whose functionality was not available within the implementation platform of the Expert Advisor.

In addition, developing an expert system in a software production environment was a challenge, according to Tony White. This was due to the restrictions placed on the development team, who had less time for innovation due to their required participation in many design reviews, code inspections, multiple document writings, and many meetings to discuss the development and integration of the software. This is the standard software cycle at BNR within which product releases are thoroughly tested and delivered to Northern Telecom customers who demand high software standards.

Technical Decision

There were four key requirements for the development of the DPN Expert Advisor. These were:

1) A requirement for real-time operation that implies handling 20 network events per second. This eliminated the use of expert system shells due to their inability to provide the required operator response time of 1-2 seconds.

2) The advisor had to support network expansion and modification. It was essential to maintain the network monitoring function while allowing the knowledge base to be modified. It was also necessary to provide facilities to handle major network upgrades that occur regularly. This implied that a system customization tool had to be included in the system.

3) The system had to be integrated with existing network surveillance tools already in use by the operator. This implied that specialized interfaces had to be developed.

4) The system had to have a graphical user interface of a quality that the operators were already accustomed to with their traditional tools.

As a result of the real-time and integration requirements, the system was developed on a SparcStation under UNIX in C/C++ with an X Windows interface. To achieve the ability to edit the knowledge base while the monitoring function remained active, rules within Problem Descriptions were selectively activated based on network events. Each network event triggers less than 5% of the PD knowledge base. Chains of rules associated with PDs are created based on events, changes of variable values and states. As a result, only a partial rule set is executed. So as not to miss events whose PDs are not activated, a global category was created as a catchall for all incoming events. This global category is later checked for events which do not have problem descriptions; if any such events are found, new PDs are created. This catchall function was a product requirement by Northern Telecom customers, who did not want to miss important network events.

The DPN Expert Advisor was initially launched in 1991 in Release 3 of the Network Management System. Currently it is in Release 8. After each release, a two week trial period, intended to uncover software problems, is held.

People

As noted above, the prototype and product development was achieved by a two person development team at BNR over a five year period. There were no clear distinctions between roles. The team was always managed at BNR. The intent was to succeed in delivering the software to the operators and satisfying the funders of the project. The Bell Canada operators helped in driving the specifications for the prototype which in turn helped define some of the functionality of the final product.

The DPN Network Verification personnel headed by Mike Verilli, were critical in providing the knowledge for the Expert Advisor. They helped in providing scripts of network events that could immediately be input into the system.

The developers of the Expert Advisor believe that there is very little trust in the use of expert systems contrasted with traditional software systems. There remains substantial resistance to the adoption of their knowledge-based technology even though the Expert Advisor has solved problems that the Network Management System was not able to resolve. They believe that such problems will continue in environments where the focus is traditional software rather than knowledge-based products.

Three development teams helped launch the Expert Advisor. This was not perceived as a problem, since the first team developed the KEE prototype, the second team developed the real-time XMS prototype and the third team developed the Expert Advisor product. Continuity was also achieved through maintaining the same Project Manager.

The development of a knowledge-intensive application in a traditional software development environment where knowledge acquisition was not understood was quite difficult according to Tony White. The integration of the Expert Advisor into the DPN development environment was restrictive and lacked many of the advantages gained in programming in knowledge-based environments. In addition, the absence of a debugger and a simulator in the form of a customization tool with a problem viewer was a difficulty that was resolved by their development from scratch, according to Andrzej Bieszczad. Despite these difficulties the developers believe that traditional software engineering features such as the window-based alarm browser has helped attract users to their product. They believe that the integration of their expert system with a traditional software environment did provide them with some useful functionality such as the interface to the network control system that provides the alarms.

Marketing Decisions

The tool is marketed by Northern Telecom. It requires no special installation and can be downloaded from a tape to a local or remote workstation. Training in the customization and features of the Expert Advisor is provided by Northern Telecom. The Expert Advisor adds value to Northern Telecom's network management system and seems to have found a natural niche within the DPN product family. It has helped differentiate the system it is integrated with from competing network products on the market. It has many customers all over the world. Any organization that has a private data packet network management system sold as the DPN NMS product by Northern Telecom can buy the Expert Advisor as an option. Richard Lavictoire, currently a Senior Operations Consultant at Northern Telecom, has assisted in the delivery of the DPN NMS and Expert Advisor to several customers. He believes that it is difficult for customers to make full use of knowledge-based systems without operations support. For this reason he spends anywhere between twelve weeks to one year providing customer support. He uses a list of alarms specific to a particular customer network to tailor the Problem Descriptions. For some customers, some irrelevant Problem Descriptions are deleted while relevant ones are added. One such customer is SITA in Paris. SITA provides the telecommunications networks for global airline communications. Lavictoire spent eighteen months supporting their installation of a full data packet network. He introduced some of the NMS functionality to SITA through the Expert Advisor. He finds that the Problem Viewer is quite useful in providing English-like descriptions rather than hexadecimal codes for alarms. He also experimented with having some Problem Description actions translated into French to evaluate how easily the Expert Advisor supports another language. He discovered that

> ***Another customer who is using the Expert Advisor is American Airlines in Tulsa, Oklahoma. American runs its SABRE reservation and ticketing system over DPN. They were also impressed with the system's capabilities in alarm filtering and Problem Description display.***

providing the PD actions in another language is quite easy. He finds that the customization tool is not very easy for customers at first, but after a training period, customers are able to make full use of its capabilities. SITA, for example, uses it to modify their PDs. He also showed SITA how to customize the alarm displays to show more alarms and activate alarms. SITA also deleted some of the PDs for events that did not occur on their network. For example the PDs associated with SL-10 switch events were deleted since this device was not on their network. Such deletions help speed up the performance of the Expert Advisor's monitoring function.

Another customer is DEC who have their network control center in Toronto. DEC automatically queries its data packet network for routing and if a trunk is down, communications are re-routed until the trunk is operational again. DEC bought the Expert Advisor two years ago but has only been actively using it in the last year. Richard Lavictoire points out that DEC required some training in realizing the full potential of the Expert Advisor. Key features that won the operators over included alarm filtering and the display of English descriptions rather than numeric codes. Alarm reduction in the field is 5:1 and is lower than the intended 10:1 quoted by the developers because not all the Expert Advisor functionality is fully utilized. In addition, the catchall facility is often deactivated by the customers since after the customization phase any events that occur should have Problem Descriptions.

Another customer who is using the Expert Advisor is American Airlines in Tulsa, Oklahoma. American runs its SABRE reservation and ticketing system over DPN. They were also impressed with the system's capabilities in alarm filtering and Problem Description display.

Richard Lavictoire finds each Expert Advisor customer unique and he tailors the system for each one. He believes that customer expectations are high because of the introduction of expert system technology. Thus far, he believes that these expectations are met because of the strong customization support provided by the Expert Advisor. One clarification he often makes to customers is that this tool does not replace operators, it simply makes them more efficient.

One drawback he points to is that there is a high maintenance requirement for the Expert Advisor. Every new alarm on the network requires an integration into the PDs that the BNR support group maintain. This support is necessary since private companies view data communications as an expense. They do not wish to add to this expense by training personnel to update PDs. Thus, from the customer's perspective a natural place for such maintenance to occur is at the Telco where communication products with alarms are developed. In an ideal situation, the customers would add new PDs to account for new equipment or links and their respective alarms on their private networks.

Further Development

The system is being adapted to other NT products such as the Passport Switch product under development at BNR for Northern Telecom. Passport will allow integrated voice, data, video and Local Area Network (LAN) communication using ATM packet technology. A Network Management System and the Expert Advisor will be integrated into the functionality of the Passport switch which is scheduled for release in March 1994 by Northern Telecom. The Expert Advisor's Inference Engine, Problem Browser and Problem Description Language Compiler were already ported to a Hewlett Packard workstation in a period of 4 months. The Knowledge Acquisition process which requires the generation of the Problem Descriptions from the events lists has been performed in parallel and the software is ready for release into the product cycle. The knowledge representation scheme has been completely applicable to the Passport application. In addition, it is estimated that 80% of the original Expert Advisor software has been re-used.

Conclusions

One of the most significant lessons learned in developing the DPN Expert Advisor is the need to integrate its functionality with a mainstream company product such as the network management system. This integration forced the developers to follow strict BNR software engineering guidelines which were restrictive in some ways but helped ensure that the Expert Advisor was well integrated with the operator's graphical user interface environment, operated in real-time and provided high quality alarm browsing and diagnosis functions. Another lesson the developers share is that the use of artificial intelligence or knowledge-based techniques was best downplayed; that is, the role of the Inference Engine was downplayed while the other traditional software features were emphasized. This helped gain the confidence of the operators and customers of the product.

The developers also believe that implementing a prototype is essential since the traditional product development binds one into a process within which experimentation is impossible. The use of industry standard tools for user interfaces such as X/Motif was also considered quite important. The exploitation of shared features within the development platform (such as alarm access functions) is a good approach to follow. The use of off-the-shelf software also simplifies the development cycle.

Finally, the developers believe that expert systems should be developed as open systems so that they can be enhanced to include tools such as those required for customization, knowledge input or modification.

Acknowledgements

I would like to thank Andrzej Bieszczad, Tony White and Richard Lavictoire for providing me with the information for this Canadian success story. I would also like to thank Innes Ferguson, Dick Peacocke and Taro Shibahara of BNR as well as Peter Turney of NRC for commenting on earlier versions of this article.

References

Baird C., White T., "A Real-Time Network Monitor", Proceedings of the 9th International Workshop on Expert Systems and their Applications 1989, pp. 35-41.

Bieszczad, A. and White, T., "Customization of Network Surveillance Expert Systems", Proceedings of GLOBCOM'92, Orlando, FL, 6-9 December 1992.

Peacocke D., Rabie S., "Knowledge-based maintenance in networks", Proceedings of Globecom 1987. Rabie S., Rau-Chaplin A., Shibahara T., "DAD: A Real-Time Expert for Monitoring Data Packet Networks, IEEE Networks Magazine, September 1988.

Peacocke D., Rabie S., "Network Management, and the DPN Expert Advisor — where's the AI?" Proceedings of Workshop on the Commercialization of AI, at AI-92 conference, University of British Columbia, 11th May 1992.

Rabie S., Rau-Chaplin A., Shibahara T., "DAD: A Real-Time Expert System for Monitoring Data Networks", IEEE Networks Magazine, September 1988.

White, T. and Bieszczad, A., "A customization Environment for the Expert Advisor Network Management System", Proceedings: Ninth Canadian Conference on Artificial Intelligence, 11-15 May, 1992. pp.82-89.

White, T. and Bieszczad, A., "The Expert Advisor: An Expert System for Real-Time Network Monitoring", European Conference on Artificial Intelligence, Proceeding of Workshop on Advances in Real-Time Expert System Technology, August 1992.

Sue Abu-Hakima *has been a researcher at the knowledge systems laboratory since 1987. She is currently coordinating projects that apply artificial intelligence techniques to telecommunications, cable, computing and wireless networks. She is also completing her PhD thesis in learning device models from diagnostic knowledge at Carleton University's Department of Systems and Computer Engineering.*

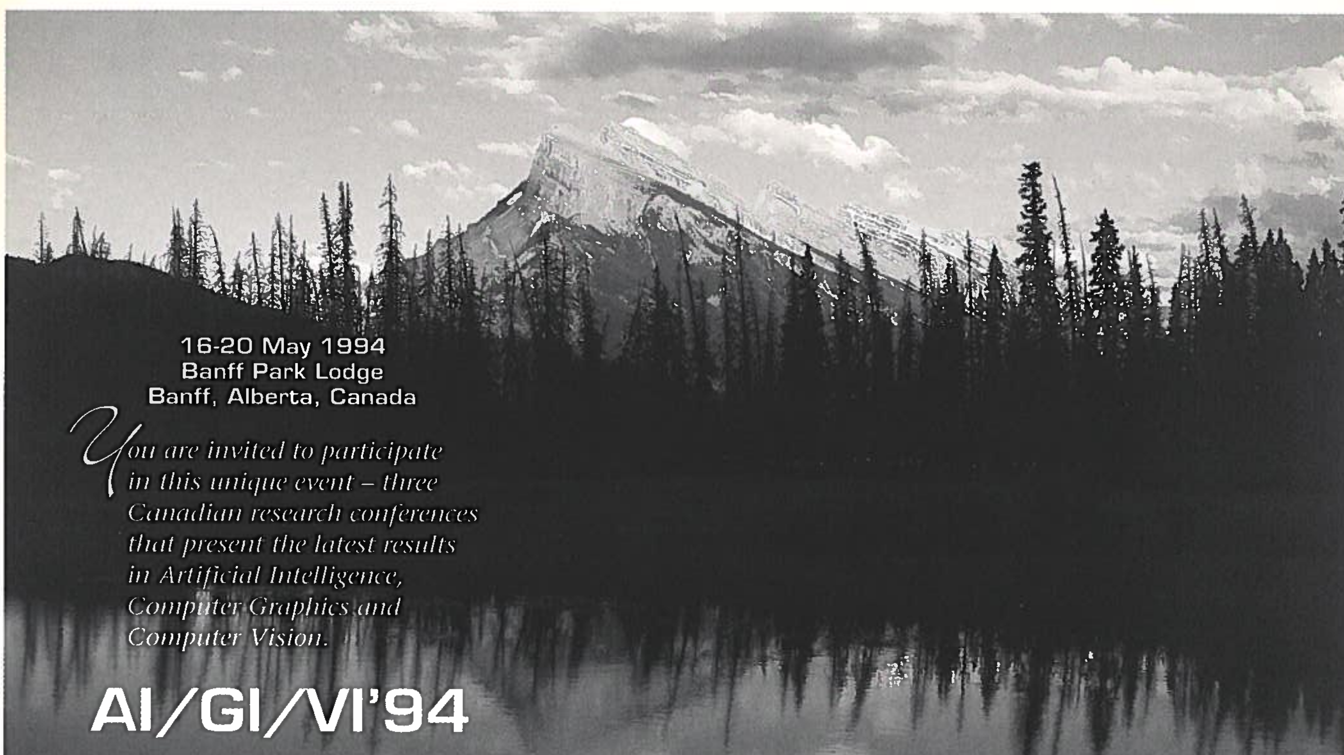

Photo credit: Greg Klymchuk

Mt. Rundle from Vermilion Lakes — Banff, Alberta, Canada

Registration form, page 20, this issue Canadian Artificial Intelligence magazine. Refer to Summer 1993 issue of Canadian Artificial Intelligence magazine for details of AI/GI/VI'94.

AI/GI/VI '94 - Registration Form

Mail Completed Form to: Box 1098, Summerland, BC, Canada V0H 1Z0

Please Type or Print in Block Letters

Family Name ____________________ Given Name ____________________

Affiliation ____________________ Business Phone ____________________

Mailing Address ____________________ Home Phone ____________________

City ____________________ Province ____________________ Fax ____________________

Country ____________________ Postal Code ____________________ e-mail ____________________

Registration Includes: 1- Proceedings, 1-Film Show Ticket & 1-Banquet Ticket

	Before/After 10 Apr 94	Totals
Member	$ 220/280	
Non-Member	$ 270/330	________
Student	$ 110/170	
Extra Tickets:	Film Show @ $ 10 ea.	________
	Banquet @ $ 40 ea.	________
Extra Proceedings @ $ 40 ea.	AI GI VI	________
Total Registration		________
Plus 7% GST (GST# R126170760)		________
Proceedings Selected ______	**Grand Total**	________

Society: CSCSI CHCCS CIPPRS *(circle appropriate society)*

Membership # ____________ *(must be included to qualify for the member rate)*

All prices are stated in CDN$. If paid by cheque in US$ deduct 15%. Credit Card charges will be in CDN$ and charged at the current exchange rate. Students ***must*** be full-time students and able to verify their status. To qualify for the membership rate, registrants ***must*** be paid up members of: CSCSI, CHCCS or CIPPRS.

For Further Information:
Phone: 604-494-9056
Fax: 604-494-4033
e-mail: davis@cs.usalberta.ca

Payment: *Cheque or Money Order (made out to: AI/GI/VI '94) or Credit Card: MasterCard or Visa*

Name on Credit Card ____________________ Signature ____________________

Credit Card Number ____________________ Expiry Date ____________________

Withdrawals and cancellations must be made in writing by 2 May 1994, and will be subject to a handling fee of $65, which includes one copy of the proceedings to be mailed after the conference. After this date no refunds will be made.

Housing Form - AI/GI/VI '94

Mail to: Banff Park Lodge
Box 2200, Banff, Alberta, Canada T0L 0C0
1-800-661-9266, 1-403-762-4433, Calgary 264-4440
15 - 21 May 1994

Please Reserve Accommodation For: # of Persons

Rate: $85.00 single or double, $10.00 for each extra person

Blocked Rooms will be held until 15 April 94

Name ____________________ Signature ____________________

Address ____________________

City ____________________ Province/State ____________________ Postal Code ____________

Phone: Office ____________________ Residence ____________________

Arrival Date ____________ Time ________ Departure Date ____________ Checkout 11:00 am

Credit Card ________ Expiry Date ________ Credit Card # ____________________

MasterCard not Accepted

If Accommodations are requested by telephone, request the **AI/GI/VI '94** Rate

INTELLIGENT ADVISOR SYSTEM[1]: A DEVELOPMENT LIFE CYCLE

Marie-Michèle Boulet
Computer Science Department
Université Laval
Québec, Canada G1K 7P4
Email (Internet): Boulet@vm1.ulaval.ca

Marie-Michèle Boulet

Abstract

Cet article traite d'un cycle de vie du developpement de systemes conseillers intelligents. L'originalite de l'approche suggeree est qu'elle met l'accent sur les besoins des personnes avant de considerer les aspects techniques. L'auteur insiste sur le fait que l'on doit definir une strategie d'intervention avant de penser a la methode de representation des connaissances ou a la selection des outils associes. Certaines metriques ayant trait a la qualite des connaissances sont egalement introduites.

Abstract

This paper presents a development lifecycle for intelligent advisor systems. Its main characteristic is to be user oriented. The author points out the importance of defining the strategy of intervention before thinking of the knowledge representation. Some measures of quality are proposed.

Introduction

An intelligent advisor is utterly linked to the use of what has been learned. It is an adaptive system aimed at intervening when the user of any application software performs transfer tasks[2]. Adaptive is a qualitative attribute that must be defined before any other decisions related to its function be taken. That means the kind of transaction of help[3] has to be detailed before beginning to think about the various logical and physical components of the advisor. A development life cycle that considers requirements for this particular instructional technology from a user standpoint is summarized here.

Software process

The primary function of a software process model is to determine the order of the steps involved in the development. Instead of focusing on how to navigate through each phase and how to represent phase products, which corresponds to a software method, a process model provides guidance on the order in which a project should carry out its major tasks. Reasons behind the order proposed in the intelligent advisor development life cycle are summarized in the following paragraphs.

At the beginning of the process of designing and developing any kind of intelligent computer-based instruction systems previously, we have used approaches either based on process and data modeling, object analysis, or rule-based formalism. Therefore, either we were looking for data and treatments of those data, objects, or rules used by experts to solve problems. In the development life cycle presented here, the system models are produced at the end of the third step. The main reason is human nature.

To design and develop an intelligent advisor system claiming to help a user performing a transfer task is a very complex undertaking. We believe that when the main concern is to identify entities and relations, objects, processes, or rules, while abiding by the estimates of a too tight deadline, analysts are inclined to only observe the superficial aspects of the communication, to minimize its importance. They identify questions, answers, explanations, feedback from the user; moreover, they will even propose a generic system which jumps to conclusions. But reality is not as simple. As stated by Leshin, Pollock, and Reigeluth (1992): different "types of learning require great differences in instruction" (p.130). Therefore, an explanation can be solely an example, solely a definition, or a fixed part being a definition with a variable part being examples. It can also be sometimes a definition with or without examples, sometimes a demonstration, sometimes merely a comment on errors, sometimes a comment on errors that reproduces the way the user made it, sometimes merely a message, sometimes a message linked with prerequisite knowledge, etc. Reality is not simple, but human nature causes individuals to try to simplify it. It is one of the greatest paradoxes that can be observed in the instructional technology industry: to propose the power of the technology to individualize, but to generalize the nature and content of transactions when designing. When the main concern is producing models, finding rules, or filling containers such as expert module, student diagnosis module, curriculum and instruction module, etc., analysts are not concerned with problems of defining what is an explanation, why an explanation is formulated in such a way at a certain moment, whether explanations are always formulated in a certain way, or are sometimes formulated in that way, etc. The content must fit in the container; if not, modify it regardless of the effects on the user.

The first three steps of the development life cycle proposed here emphasize the creative thinking process related to the identification of all the details of the strategy of intervention, the user being considered. The approach is user oriented; time and again, the focus is on how the intelligent advisor will help the user. The system analysis and models associated come after. To do so, existing approaches are used, and very useful.

Description of the steps

Step 1. Describing the strategy of intervention

1.1 Stating the aim of the advisor

To state the aim of the advisor is to describe what it will bring to the user; furthering the knowledge transfer, sharing rare expertise, instructing and transiting users from novice to expert behavior, providing flexible access to training material, or ensuring that all the members of an organization solve similarly problems, are examples.

1.2 Describing the target users

To describe the target users of the advisor is to characterize them. Examples are: weak knowledge of the domain, satisfactory knowledge with little to no experience transferring it, or satisfactory knowledge with experience transferring it, but having been appointed to other tasks for a certain period of time.

1.3 Describing the help transactions

There are three main types of help transactions documented in the literature: 1- Active, i.e., monitoring the user's progress to guide him or her based on the problem solving task the user is engaged; 2- Passive, i.e., responding only to the requests for help and questions initiated by the user; 3- Pedagogical, i.e., giving remedial teaching. Any combination can be made.

1.4 Describing the main characteristics of the course of interventions

For each type of help transaction selected, the interventions are described. The following is an example of passive: While the user composes, and in accordance with his or her request, the musical composition advisor answers questions, comments on the user's composition, creates its own compositions, and comments on them (Boulet, 1992).

1.5 Describing how the advisor adapts itself to the user

There are several possibilities. Examples are presented in the following. The advisor can engage the user in a dialogue; according to the user's feedback to advisor's explanation, it will modify the content of its next interventions.

Instead of engaging the user in a dialogue, it can decide to grade him or her. The advisor can propose levels to the user, or a grading test can be used.

Making inferences in regard to the user knowledge level is another possibility. The nature of user requests can be analyzed. This corresponds to the answer to the following question: Do the user requests refer to prerequisites corresponding to the lower level of a hierarchy of intellectual skills? The use of a model can facilitate the identification of prerequisites in regard of a specific task. Several models exist: Ausubel (1960) proposes advance organizers, Bruner (1966) proposes a spiral curriculum, Gagné (1977), a hierarchy of intellectual skills, Landa (1976), an approach based on algorithm, and Reigeluth (Reigeluth and Stein, 1983), the elaboration theory. Prerequisites being identified, it is possible to link requests to a level within a hierarchy; this level can be later used to determine the user one. To infer the user knowledge level, a list of systematic errors can be used. To produce such a list of errors with their causes, models exist. We refer to the field of measurement and evaluation in education (Scallon, 1988). One can decide to divide the domain in terms of systematic errors, and to link a question or a list to each error. One can decide to divide the domain in regard of prerequisites, and to link systematic errors to a performance objective. The user knowledge level can also be inferred in regard to the frequency of requests for help relatively to a same topic. The domain might be divided according to topics, or prerequisites. The main difficulty is to define the norm: how many requests are necessary to infer that it is enough. The user knowledge level can also be inferred in regard to the frequency of the errors for a same topic. Here again, the main problem is to determine the norm. Each error can be associated to a prerequisite, or to a division of the domain in terms of topics.

1.6 Overall characteristics of the subject of interventions

Taking into account the previous decisions, the overall characteristics of the intervention domain are described. Examples of questions to be answered are: Will the advisor do general comments on a problem solving process without considering the user's current task? Will the advisor comment on the problem solving process, or the results?

1.7 Overall characteristics of vocabulary of presentation

As an advisor aims at helping a user perform a task, the short term memory is the main concern; the means to allow the user to recall concepts, principles, and rules are planned. It is done for each type of help transactions selected. Examples are: using demonstrations, examples, generalities, index, maps, pictures, or procedures.

1.8 Knowledge depth

Having established the overall characteristics, the knowledge depth of each type of help transaction is characterized. This corresponds to the following questions: To what level will the advisor intervene (goal or step)? When will it give details related to a specific step of a problem solving process? When will it decide to explain all of the process (goal)? Will the advisor only describe what to do in regard of a goal or a step? Will the advisor only show, using examples, what to do in regard to a goal or a step? If the advisor uses examples, will these examples be general or related to the specific task currently being done by the user? Will the advisor both describe and offer examples of what to do in regard to a goal or a step? Will it be able to answer questions such as "What would happen if I did this"? or "What is the difference between these two ways of doing this"? Will the advisor consider links with prerequisites when it provides explanations, presents examples, or demonstrates?

1.9 Details relative to active resource

When the advisor is an active resource, many mechanisms allowing tasks recording, tasks analysis, and dialogue initiation must be installed. Which current tasks will be recorded must be determined. Examples are: each key pressed on, or a set of actions. How the advisor detects errors must be decided. Comparing the way an expert solves a problem with the user way is a possibility. If the advisor can take the initiative of establishing a dialogue with the user, its decision criteria are stated.

1.10 Details relative to passive resource

When the advisor is a passive resource, will it only answer to the user's request? Will mechanisms be installed in order to allow the advisor to divide its interventions in such a way that several user's feedback can be recorded, so the user controls the dialogue?

1.11 Details relative to pedagogical resource

When the advisor is a pedagogical resource, will it present general explanations, or explanations related to the field of the user? For example, with application software, the display of the content of the user guide after a user request in regard of a specific command is an example of general explanation. On the other hand, the explanation of a part of a conceptual database model representing a manufacturing enterprise to a user working on the elaboration of the model of such an enterprise is an example of explanation related to the field.

1.12 Interface

Several decisions must be taken in regard to the interface. Will the interface be based on natural language? Menu driven? Buttons? Keywords? A mixed one? Will the user be allowed to switch from one resource to another? How and when? Who will take the decision? The advisor, the user, both? If the advisor can decide, what are its decision criteria?

Step 2. Describing the strategy and means proposed for knowledge acquisition

2.1 To delimit the boundaries of the domain

The means that will be used to delimit the domain are described. The means corresponds to the analysis and characterization of the transfer task.

2.2 To identify problems

The means that will be used to identify problems an individual may have when he or she performs a complex transfer task X, while using an application software Y, are selected and described. A standardized test can be used. Values associated with its reliability and validity (such as level of confidence) can be used to estimate the proportion of target users not actually helped by the advisor. (Because some of their problems are not part of the test, it is not able to answer.) The number of errors that users could make (risk) despite the use of the technology can also be estimated with cost related factors (such as for the manufacturing sector, a lot of imperfect products).

It can also be decided to analyze how target users perform the transfer task. Will a representative sample of target users be taken? These users state their questions out load; then, it will be possible to list them. Because irrelevant variables can affect the environment, the process must be rigorously surrounded. Proportion of users the advisor will fail to help will be estimated, using values such as the proportion of users that may not be represented in the sample. Here again, these values help to evaluate the usefulness of the investment in regard to the need and the current situation of the organization.

It can be decided to set a simulation of the functioning of the advisor. An analysis of protocol will then be done. Data collected can be categorized. The means to ensure that the sample actually covers the variety of future users must be taken. Also, the means to ensure that the (subjective) process of linking a request for help to a category will be objectively done must be planned. To avoid a significant loss of information (simplification of the reality), the number of categories must be appropriate. Here again, statistics can be established: proportion of users that are not represented in the sample, reliability and validity of the case study with error associated, etc. The various values will then be used to estimate the proportion of future users that will not be helped (because their questions were not collected the advisor is not able to answer); the linked risk (i.e., they made mistakes) with the cost related factors, will be estimated.

2.3 To identify sources to write interventions

The means that will be used to write interventions the advisor does when an individual, performing a complex transfer task X while using an application software Y, makes a request for help or mistake, are selected and described. There are several possibilities: specialized manuals, course syllabi from academic or professional sectors, content of courses from the academic sector or professional sector, experts doing the task and expressing out loud how they perform the transfer task, or simulation of the advisor followed by an analysis of protocol allowing the identification of experts' answers facing a user request. Cautions mentioned in regard of the users' requests for help identification and categorization are relevant. Those mentioned in regard to calculating values in order to estimate the margin of error of the advisor (i.e., what proportion of users the advisor will fail to help) are also relevant. Analysis of these data must allow the identification of the useful explanations, i.e., those that actually helped the user to perform his or her complex transfer task. Values issued from statistics are used to estimate the proportion of users that will not be helped even if the advisor is able to answer their question (cost related).

To determine what will be the content of interventions, one can refer to a certain form recommended by researchers. Those researchers recommend it after having done several rigorously controlled experiments involving several samples; these experiments were done to identify how to facilitate the mastery of a particular capacity described by an action verb. For example, we mention instructional tactics proposed by Leshin, Pollock and Reigeluth (1992) in regard to each type of capacity they define. The Gagné and Briggs' (1979) principles of instructional design being stated after several years of experiments are another example. The use of those principles or tactics can contribute to minimize the margin of error. Moreover, the values necessary to estimate the utility and the risk of failure can easily be found in the literature related.

2.4 To set out the knowledge acquisition method

Having identified sources of knowledge and means, steps of the corresponding knowledge acquisition method are then detailed. Note that the method is not a generic one.

Step 3. Description of the content, the architecture, and the strategy of implementation in the environment

3.1 To analyze sources of knowledge

Sources of knowledge identified at the previous step are analyzed in accordance with the knowledge acquisition method.

3.2 To write the content of interventions

Each intervention is stated in accordance with the knowledge acquisition method. If there are persons that will write any part of the content (think about an automatic collecting of real cases, or an automatic collecting of exercises), its structure must be carefully detailed. Doing so, the style will be kept unvarying.

3.3 To produce the architecture of the advisor

The container being considered, the advisor structure is outlined, without taking any physical aspects into account. Global characteristics of the advisor are illustrated and documented. The main components are identified. This, in turn, brings about the identification of the main modules, and of links between them. To elaborate models of data and treatments, one can decide to use conventional techniques such as structured analysis for treatments (data flow diagrams), and conceptual database modeling for data (entity-relationship diagram); object oriented analysis can also be used.

3.4 To document the architecture

While producing the architecture, analysts have to document it.

3.5 To describe the strategy of implementation in the environment

Although the strategy of intervention of the advisor has been described at step one, the sources and means used to acquire the knowledge acquisition has been selected at step two, and the knowledge acquisition has been done at the beginning of the current step, the particularly complex problem to develop an advisor is not totally solved. Works done focus on one aspect of the problem, that being the transmission of the knowledge in a help context. Modeling tasks currently undertaken focus on another aspect, that being to find the best technological solution for the transmission of knowledge in regard of this particular context. Simultaneously, work related to the problem of implementing an instructional technology in the environment begin.

To be useful, an advisor must be used. To ensure that an event will occur, it must be managed. To be able to manage, all the events that can prevent a desired situation to occur must be identified. The desired situation is that the advisor be used at the beginning of its implementation, and after, even if some changes occur in the environment. The more the events are forecast, the more they can be managed.

First, the proper moment to introduce the instructional technology in the environment must be found. Then, the tactics are detailed: Will the target users be informed in advance? How long in advance? How many announcements will be done? No matter the tactic selected, the content of communication must be carefully detailed. When many communications are planned, content and sequence must be detailed. Mechanisms which periodically recall that an instructional technology is available can also be planned. The proper frequency can be found by analyzing data related to changes in staff.

The strategy of implementation of the advisor in its environment is defined to ensure that **all** the target users will **always** have **all** the information they need, regardless of changes that have occurred within the organization. An example is an employee that quits; will the information transmitted in order to prepare the environment to the installation of an instructional technology, i.e. aim, type of help, limits, etc., be available to the new person hired? To be sure that **everybody always** has **all** the information needed to use the advisor, or **always** has easy and quick access to this information, whatever changes occur in the organization, there must always be a detailed plan of communication for **all** advisors. It also means that someone must be **permanently** in charge of the management of the post implementation

Step 4. Description of the physical implementation

4.1 To convert logical models into physical models

Whatever methods and techniques used at the previous step, analysts did logical modeling of the advisor. They now have to convert it into a physical model. It can be decided to prototype. The following is an example of reason behind the decision to prototype a part of a musical composition advisor: "Domain experts elaborated at the previous step, a first version of explanations. ... The prototype being developed, content and form of explanations changed. At first, domain experts wrote a quite bookish content. ... using the prototype, they realized that this definition was not very helpful. In the same way, the prototype was used by a sample of future students ...: it was asked to each student to state in his or her own words what was the meaning of the explanation displayed." (Boulet, 1992).

4.2 To validate the content of interventions

Prototypes test parts with the target users before the whole advisor is programmed. Adjustments at this point are easier and cheaper than later.

4.3 To program

This is the point at which application programs are written or purchased.

4.4 To test the programs

Here methods proposed in the literature are used.

4.5 To implement the advisor

The programs are run, the different files are interfaced, and the human-interface is set.

Conclusion

An intelligent advisor system is an adaptive system aiming at intervening when the user of any application software performs transfer tasks. Examples of transfer tasks for which an advisor can be designed and developed are: 1. To compose a melody, while using a musical writing software such as CONCERTWARE™; 2. To prepare a research proposal, while using WORD™.

The life cycle presented here is user oriented. That means developers don't have in mind any a prior technical solution such as "To use Hypercard," "To develop a natural language interface," or "To use an interactive videodisc system;" that also means they don't make any a prior choice about encoding the knowledge in the advisor data structure. The first three steps of the development life cycle presented here emphasize the creative thinking process related to the identification of all the details of the strategy of intervention, the user being considered. Time and again, the focus is on how the intelligent advisor will help the user perform a complex transfer task. The system analysis and models associated come later.

References

Ausubel, D.P. (1960). The use of advance organizer in the learning and retention of meaningful verbal material. *Journal of Educational Psychology.* 51, 267-272.

Boulet, M.M. (1992). *Advisor systems: The development life cycle illustrated by two case studies.* Québec: Edition Dwayne.

Bruner, J. (1966). *Studies in cognitive growth.* New York: Wiley.

Gagné, R.M. (1977). *The conditions of learning.* New York: Holt, Rinehart and Winston.

Gagné, R.M., & Briggs, L.J. (1979). *Principles of instructional design.* New York: Holt, Rinehart and Winston.

Landa, L.N. (1976). *Instructional regulation and control.* Englewood Cliffs: Educational Technology Publications.

Leshin, C.B., Pollock, J., & Reigeluth, C.M. (1992). *Instructional design strategies and tactics.* Englewood Cliffs: Educational Technology Pub.

Merrill, M.D. (1988). Applying component display theory to the design of courseware. In D.H. Jonassen (Ed.). *Instructional designs for microcomputer courseware.* Hillsdale: Lawrence Erlbaum Ass.

Reigeluth, C.M., & Stein, F.S. (1983). The elaboration theory of instruction. In C.M. Reigeluth (Ed.), *Instructional design theories and models: An overview of their current status.* Hillsdale: Lawrence Erlbaum Ass.

Scallon, G. (1988). *L'évaluation formative des apprentissages: L'instrumentation.* Québec: Presses de l'Université Laval.

Notes

[1] Professor Boulet's researches on intelligent advisor systems are granted by: Canadian Workplace Automation Research Center (CWARC), Social Sciences and Humanities Research Council (SSHRCC 410-90-0578). Technical support was provided by IBM.

[2] "A transfer task consists of principles and decision rules that an expert uses to generate the appropriate performance for any given situation. It also consists of secondary content, such as concepts and information" (Leshin, Pollock, and Reigeluth, 1992, p.92)

[3] "An instructional transaction is a dynamic interaction between the program and the student in which there is an interchange of information" (Merrill, 1988, p.71)

***Marie-Michele Boulet** is a full professor in the Computer Science department at Universite Laval. Since May, 1993 she has been Vice-Dean of Administration, Faculty of Science and Engineering. She is responsible for the television course "Technologies de l'information"(distance education). Her main area of research is intelligent advisor systems.*

Call For Papers 1994 CSCHE Conference
October 1994, Calgary, Alberta
Session on Artificial Intelligence Applications

A session has been scheduled to bring academia scholars and industrial engineers together to discuss practical applications of artificial intelligence to chemical engineering. This session aims at promoting the research and development of AI technology in chemical engineering, and narrowing the gap between academic research and industrial applications.

Prospective attendees are invited to submit, by February 1, 1994, two copies of an extended abstract of up to 2 pages to:

Professor Ming Rao, Intelligence Engineering Laboratory
Department of Chemical Engineering, University of Alberta
Edmonton, Alberta T6H 2G6
Phone: (403) 492-8250 Fax: (403) 492-2881
E-Mail: usermrao@mts.ucs.ualberta.ca

PRECARN UPDATE
NOUVELLES DE PRECARN

Jean-Claude Gavrel

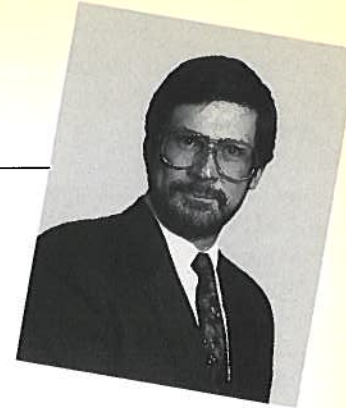

The summer of 1993 has seen a flurry of activity at the PRECARN office. Perhaps being located in Ottawa, with so much excitement surrounding the upcoming federal election, has given spark to our normally subdued summertime schedules. In this issue, we will highlight some of the results of our summer activities and present our plan for a vigorous Canadian Intelligent Systems research program that will surely lead to a more prosperous IS economy!

PRECARN Creates a PERK for its Members

With the first wave of research projects reaching the mid-point, PRECARN decided it was time to implement a new program, one that would facilitate and encourage industry to take the research knowledge or raw technology already existing in PRECARN or IRIS and take it one step further; that is, to develop application-oriented proof-of-concept prototypes.

This initiative, the Program for the Exploitation of Research Knowledge (PERK), was launched with a request for research proposals on July 9 with a deadline date of October 15, 1993.

The criteria for a successful PERK include: 1) shorter-term research proposals, with a duration of 18 to 24 months; 2) support from a minimum of three PRECARN members; and, 3) projects limited to the $500,000 range, with a maximum of $350,000 provided by PRECARN.

PERK proposals will be subjected to a rigorous evaluation by PRECARN's Technical Committee. The quality of the research, the qualifications of the team and the potential downstream for commercial applications will be measured. In December we expect to recommend a slate of proposals to the PRECARN Board of Directors, and hope to fund four proposals in this initial round.

Planning the Future of PRECARN

While the PERK initiative represents one component of the future direction for PRECARN, three research projects are already in the pipeline and ready to go.

The Team-based Intelligent Productivity Systems (TIPS) project assumed at increasing the efficiency and productivity of group activities such as meetings through the application of advanced intelligent systems, group decision support systems and multimedia technology.

The objective of the Knowledge-Aided Design (KAD) project is to use new knowledge-based technologies to enhance the degree of awareness, understanding, cooperation and coordination among engineering team members.

The Mining Automation Project (MAP) is aimed at developing the intelligent systems technologies necessary for an eventual Computer Integrated mine, one where machines would operate underground with humans controlling them from the surface.

Funding for these and other projects is a major challenge and we are busy preparing a submission to the federal government for funding of the next wave of research projects. At the same time, Industry, Science Canada has commissioned an independent evaluation of the PRECARN program. Results of that evaluation will provide valuable insight into the future directions for PRECARN.

PRECARN's Policy and Planning Committee, mainly composed of long-standing members of PRECARN, are looking at such issues as the Intellectual Property Policy, the membership structure and the overall research program. The Committee expects to report to the Board of Directors at its December meeting.

IRIS Passes a Grueling Test

A five-member Site Visit Committee, responsible for the evaluation of the IRIS (Institute for Robotics and Intelligent Systems) Network, attended meetings and interviewed industry and university representatives during the June IRIS/ PRECARN Conference,

The Committee was composed of Dr. Peter Will (University of Southern California), Dr. Brian Oakley (Logica Cambridge Ltd., United Kingdom), Dr. Ruzena Bajcsy (University of Pennsylvania), Dr. T.J. Tarn (Washington University) and Dr. Jack Minker (University of Maryland). Their report, which was submitted to the Natural Science and Engineering Research Council in early July, concluded that "...overall, IRIS represents an exceptional value as an investment by the Canadian government. The Committee strongly recommends continuance of funding."

Getting Ready for IRIS Phase 2

Preparation for the second phase of IRIS began more than a year ago, and included months of consultations with university and industry representatives across the country. Every effort was made to encourage those not presently involved in IRIS to submit proposals for Phase 2. From the enthusiasm of the participants, it quickly became apparent that IRIS was having a definite impact on the university community.

By the deadline date of June 11, 83 proposals were received, totaling over $120 million. IRIS' success would force some very difficult decisions.

Staff engaged the assistance of an external Review Committee, consisting of industrial and university representatives, to make an initial evaluation and selection of the best projects for the Network. The results were passed to the IRIS Research Committee and then to the IRIS Management Board. The thirty-one remaining projects were grouped into six research themes: Intelligent Computation, Human-Machine Interfaces, Machine Sensing, Control, New Devices and Integrated Systems for Dynamic Environments.

Although the consultation process and layers of evaluations has, at times, been exhausting, it will help ensure a successful proposal for the second phase of IRIS.

New Scholarship Fund Announced

The June 1993 IRIS/PRECARN Conference was the venue for the announcement of the establishment of the Gordon M. MacNabb Scholarship Foundation. The foundation was established in honour of PRECARN's founding President, Gordon M. MacNabb, who retired this fall.

During his years with NSERC, and more recently with PRECARN and IRIS, Mr. MacNabb has been recognized as championing the tremendous potential of Canada's youth. What better way to honour him than with the establishment of the Gordon M. MacNabb Scholarship Foundation.

The Scholarship, which will become available in the 1994/95 academic year, is being established for graduate students in the field of intelligent systems. It is anticipated that the Foundation will provide two scholarships of $5,000 each per academic year.

PRECARN Video

There are many ways to get your message across these days — and it seems that video is becoming a very popular means of communication for corporations. Well, we thought we'd try it too!

At its March meeting, the Board approved the PRECARN Video project; so, along with the film crew, PRECARN staff went on the road to talk to its members and researchers. The response and enthusiasm were overwhelming. Industry leaders such as Alan Winter, President of MPR Teltech Ltd., Monique Lefebvre, President of the Centre de recherché informatique de Montreal, Joe Wright, Vice-President (Research) of the Xerox Research of Canada, and university leaders including John Hollerbach and Ian Hunter of McGill University, John Mylopoulos of the University of Toronto and Tom Calvert of Simon Fraser University, as well as a host of others graciously donated their time to speak about the benefits that PRECARN and IRIS are bringing to their organizations.

We feel that these industrial and university leaders are sending a powerful message to government, and to potential members about the importance of models like PRECARN. We hope that government will continue to encourage collaborative efforts such as PRECARN's. Copies of the video are available from Mrs. Lise McCourt, Manager, Corporate and Public Relations.

For more information on any of the above, please contact Mr. Jean-Claude Gavrel, Vice President of PRECARN Associates.

PRECARN Associates Inc.
300-30 Colonnade Road, Nepean, Ontario K2E 7J6
Tel: (613) 727-9576 Fax: (613) 727-5672
E-mail: gavrel@a1.atott2.nrc.ca

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Are you interested in contributing an article, news item, or conference review to *Canadian Artificial Intelligence* magazine?

Contributions for the Spring 1994 issue of *Canadian Artificial Intelligence* magazine will be accepted until January 18, 1993. Please send your contribution, electronic preferred, with an abstract, photograph and short biography to:

Canadian Artificial Intelligence
c/o Alberta Research Council
3rd Floor, 6815 - 8 Street N.E.
Calgary, Alberta T2E 7H7
or - merling@.arc.ab.ca

BOOK REVIEWS

Understanding music with AI: Perspectives on music cognition *Mira Balaban, Kemal Ebcioglu, and Otto Laske (editors)* [Ben-Gurion University, IBM Thomas J. Watson Research Center, and Newcomp Inc) Menlo Park, CA: The AAAI Press and Cambridge, MA: The MIT Press, 1992, xxxviii+512 pp; paperbound, ISBN 0-262-52190-9 US$39.95

Reviewed by
Norma Welch
MASYCOM Consultants Inc.

Understanding Music With AI consists of a collection of 21 papers in the emerging interdisciplinary field of Cognitive Musicology. The papers arose from workshops in 1988 at the 1988 AAAI conference, the IJCAI, and the International Computer Music Conference. Musicians, especially those involved in music technology or theory, will find many of the papers interesting and challenging. AI researchers will also find interesting and novel applications of familiar and new techniques to the music domain. Alhough the papers introduce their terminology well, some familiarity with music, AI, and cognition is desirable.

The scene for this collection is set by a foreword containing a thoughtful and far-reaching dialogue between Otto Laske and Marvin Minsky. They discuss the question of "What is music?" and explore the possibilities of formalizing the musical activities of composing, performing, listening, analysing, etc. The issues discussed include: analogies between music and language, the desirability of modelling emotions, the relationship between composing and improvising, and the possibility of a musical common sense database. Marvin Minsky suggests that applying AI to music "would be making composers, or at least listeners."

An introduction by the editors brings us down to earth and the specifics of the questions addressed by current researchers and the papers in the collection. They discuss the inadequacies of AI methodologies for modelling musicological activities and provide overviews of the papers that follow. The papers are grouped in seven sections, each of which is introduced by the editors. The first ("Two views on cognitive musicology") presents the views of Otto E. Laske and Peter Kugel on the nature of AI and music; the second focuses on representations and modelling methods; the remaining five sections cover the music topics of composition, analysis, performance, perception, listening, and tutoring.

Otto Laske's paper, "Artificial intelligence and music: A cornerstone of cognitive musicology," provides an excellent overview of cognitive musicology from both a historical and practical point of view. Laske argues that cognitive musicology provides the possibility of integrating three dimensions: task environment (local knowledge), competence (musical knowledge), and performance (deployment of knowledge in task environments). He argues that traditional musicology fails to integrate these three dimensions and therefore sees a future for the study of AI and music. This paper includes a good overview of knowledge-based systems and a discussion of their potential to create structured computerized environments that will support a variety of musical tasks. In "Beyond computational musicology," Peter Kugel presents arguments to suggest that some musical thinking may exceed the power of normal computations. He presents powerful results from the theory of computation in easily understood form, with musical examples.

The section "General problems in modelling musical activities" is a technical section devoted to new methods of representing and manipulating musical objects. Stephen Smoliar focuses on listening behaviour; he provides an interesting discussion of musical memory and its interaction with mental states and concludes that propositional logic is insufficient to model musical behaviour. Two papers are devoted to the representation of musical items over time: Bernard Bel details a symbolic grammar-based representation, and Mira Balaban specifies a music structures language. Both papers describe the methodologies thoroughly and with many examples, and both refer to experimental computer systems. The final two papers in this section; "On designing a typed music language" by E.B. Blevis, M.A. Jenkins, and J.I. Glasgow, and "Logical representation and induction" by F. Courtot, describe two different languages designed to support compositional and analytical processes. The innovative approaches used by the authors of the papers in the section will interest the AI researcher as well as the more technically minded musician.

The remaining five sections are more practical. In the 14 papers in these sections, we see the results of 20 years of research in music and AI. Every musician will be amazed at the breadth covered by these papers and the success that has been achieved. They are organized in the musical areas of music composition, analysis, performance, perception and learning/tutoring.

The authors of the four fascinating papers on music composition take an empirical approach to the theory of composition. A computer program that can compose music in a particular style embodies the theory of that style, and programs can be used to investigate the compositional process. In "The observer tradition of knowledge acquisition," Otto Laske describes programmed task environments that capture the actions of human composers at work, thereby bypassing the verbalization step of knowledge engineering. The empirical approach continues into the section on analysis, with the papers "An expert system for harmonizing chorales in the style of J.S. Bach" by K. Ebcioglu and "An expert system for analysis of tonal music" by John Maxwell. In the last paper of the analysis section, David Cope describes experiments using augmented transition networks to analyze a composer's work and synthesize compositions in the same style.

Included in the performance section is a fascinating paper, "A new approach to music through vision." In this, S. Ohteru and S. Hashimoto describe, though all too briefly, several projects developed in their "music and vision" laboratory: visual score recognition, two-way translation between printed music and Braille music, a music system that can follow a human conductor, and an automated dance accompanist. The group is aiming to integrate these tasks. Papers in the perception section focus on rhythm, meter, and quantization and describe a variety of knowledge-based and connectionist approaches to modelling the time dimension of music. The book concludes with two interesting applications of machine learning to the teaching of music.

The papers in this book provide wide coverage of the facets of both music and AI. Readers from both fields will not fail to find much of interest and can use the extensive references to lead them deeper into areas of particular interest.

Norma Welch is an AI consultant whose specialties include the application of AI techniques in railroad operations. She is also studying music at McGill University.

Case-based learning Janet L. Kolodner (editor) (Georgia Institute of Technology) Dordrecht: Kluwer Academic Publishers, 1993, 171 pp; originally published as a special issue of ***Machine learning*, *10*** (3), 1993, 195—363; hardbound, ISBN 0-7923-9343-0, US$105.00

Reviewed by
Michel Féret
Queen's University

This book is a reprint of an issue of the *Machine Learning* journal, initially published in the spring of 1993 by Kluwer Academic Publishers. It contains five articles concerned with ways of achieving learning for case-based reasoning (CBR) systems that are different and potentially more powerful than mere rote learning by accumulation of cases. The targeted audience of such a book is clearly researchers focused either on CBR or on machine learning (ML) issues. Given that the book version sells for US$105 and that most researchers have access to the journal, it is unlikely to show on the best-seller list, despite the quality of its content.

A brief introduction by the editor summarizes the issues relating CBR to ML. I will review these briefly, and then discuss each of the five papers.

Case-Based Reasoning and Machine Learning

CBR is based on a psychological theory of human cognition. It accounts for issues in memory, learning, planning and problem-solving. Its foundations lie in Schank's conceptual memory model, an integration of semantic networks and episodic memory.

CBR has traditionally been used as a stand-alone problem-solving method. A CBR system stores past experiences in the form of cases. When a new problem arises, the system retrieves the cases most relevant to the current problem, then combines and adapts them to derive and criticize a solution for a model of this process. If the solution is not satisfactory, new cases are retrieved to further adapt it in the light of additional constraints (expressed from the non-satisfactory parts of the proposed solution) until it is acceptable. After a problem is solved, a new case can be created and stored in the case base.

"Classical" CBR systems learn by accumulating new cases—a form of rote learning. This book is about other possible ways for CBR systems to learn. These alternative ways of learning are all linked to:

1. the difficulty of defining what it means for a past case to be *relevant* in the context of a new problem;
2. the indexing problem, i.e., how to store new cases and to re-organize and update the case base dynamically for maximum utility; and
3. the use of feedback information about usage made of retrieved cases by an external agent (e.g., a planner, a search algorithm).

The five papers that constitute this book address these problems with different perspectives and goals. Ram's and Hammond's papers belong to "true-faith" research in CBR, which is mostly interested in developing models for human cognition. The other papers belong to CBR "hard core" research, which focuses on techniques and methods to implement "true faith" ideas.

Individual Papers

Ashwin Ram's system, AQUA, aims at incrementally improving its understanding of incompletely understood domains by revising already-stored cases in the light of new information provided by the application of old cases to new problems, by re-indexing old (specific or general) cases, and by keeping track of unknowns that it has identified as potentially useful to investigate. Ram presents only a few, small examples that admittedly fall short of a proper scientific evaluation of his theory. Nevertheless, his paper provides valuable insights into issues related to explanation-based learning and memory management aiming at theory development.

Manuela Veloso and Jaime Carbonell present a CBR addition to PRODIGY used as a planner. The resulting system benefits from the association: The CBR component uses feedback from the planner that uses its cases for re-indexing old cases better. The planner becomes more efficient by avoiding previously made mistakes and by re-using successful past experiences. The association of a CBR system with another problem-solving system (e.g., a planner) is not new in itself. However, this paper's strength is its generality, inherited from PRODIGY itself. It provides a framework which integrates re-use and adaptation of past experience (through derivational analogy) with general problem-solving capabilities and illustrates the power of hybrid CBR systems.

Kris Hammond, Timothy Converse, Mitchell Marks, and Colleen Seifert focus on modelling opportunism as a means for learning. Their work stems from the contrast between the apparent complexity of the world and the overall simplicity of humans' methods for dealing with it. Their paper discusses, without ever defining it, a computational model for case-based planning that learns from its failures and from opportunities by using an appropriate opportunistic memory model. It also touches on "enforcement"—the idea that the planner's world is adapted to the planner's stored plans as much as the plans are to the world in which they were designed to perform. The TRUCKER system, presented as an illustration, provides a simplistic example of opportunism based on geographical closeness. This paper provides an interesting starting point for research on opportunism and raises a series of interesting questions. Like Ram's paper, it is mostly aimed at an audience of "CBR true believers."

Chris Owens's paper explores issues related to the concept of relevance. Most CBR systems rely on static similarity measures to define relevance. This creates problems: not all cases should be judged according to the same criteria, the current problem is often not described by appropriate features, some features are more expensive to use, others are more useful, etc. Owens's system, ANON, uses the case base to guide the extraction of features in the current problem, which is described in terms of what is known to the system. This allows a more focused, more flexible, hybrid top-down bottom-up, retrieval mechanism. ANON also measures the discriminating power (in the conceptual clustering sense) of features it uses as a predictor of their utility. Owens admittedly does not account for external utility measures and cites this as a path for future research. Although the paper does not provide any serious evaluation of the method, it represents a definite research advancement towards a better understanding of the notion of relevance.

Srinivas Krovvidy and William Wee use CBR to enhance heuristic search in the applied domain of waste water treatment systems. They use a representation of the domain that allows cases to be stored in a hierarchical manner that allows very effective retrieval as well as guided adaptation mechanism that is used to complement a traditional A* search algorithm. The main contributions of this paper are to present a working system in a real-world domain and to identify conditions under which CBR can effectively be applied to heuristic search problems. The latter is a precious contribution to research in heuristic search and in CBR.

Conclusions

This book reveals two trends in the CBR community:
1. the general need for proper evaluations of theories and systems, and 2. the tendency for applied systems to combine CBR with other problem-solving paradigms. The book provides a non-exhaustive introduction to case-based learning that focuses on means of learning for CBR systems other than rote accumulation of cases over time. Re-indexing cases, updating the content of old cases, labelling cases with utility measures, using feedback from another problem-solving paradigm or from a heuristic search algorithm are possible methods to achieve learning. Even though there are other possible ways of learning for CBR systems, the papers presented in this book are certainly a good starting point for more research in this area.

***Michel Féret** is a doctoral candidate in case-based learning at Queen's University.*

BOOKS RECEIVED

Reviewers are sought for books marked with a * in the list below. Readers who wish to review books for *Canadian Artificial Intelligence* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, or send electronic mail to gh@cs.toronto.edu or gh@cs.utoronto.ca. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

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

Active vision *Andrew Blake and Alan Yuille* (editors) (University of Oxford and Harvard University) Cambridge, MA: The MIT Press (Artificial intelligence series, edited by J. Michael Brady, Daniel Bobrow, and Randall Davis), 1992, xviii+368 pp; hardbound, ISBN 0-262-02351-2, US$50.00

The simulation of human intelligence *Donald Broadbent* (editor) (University of Oxford) Oxford: Blackwell Publishers (Wolfson College lectures series), 1993, vii+222 pp; hardbound, ISBN 0-631-18587-9, US$49.95; paperbound, ISBN 0-631-18733-2, US$19.95

Explanation and interaction: The computer generation of explanatory dialogues *Alison Cawsey* (University of Glasgow) Cambridge, MA: The MIT Press (The ACL—MIT Press series in natural language processing, edited by Aravind K. Joshi, Karen Sparck Jones, and Mark Y. Liberman), 1992, vii+232 pp; hardbound, ISBN 0-262-03202-3, US$29.95

Artificial experts: Social knowledge and intelligent machines *H.M. Collins* (University of Bath) Cambridge, MA: The MIT Press (Inside technology series, edited by Wiebe E. Bijker, W. Bernard Carlson and Trevor Pinch), 1990, reissued in paperback 1993, xiii+266 pp; paperbound, ISBN 0-262-53115-1, US$14.95

***Investigating explanation-based learning** *Gerald DeJong* (editor) (University of Illinois) Boston: Kluwer Academic Publishers (The Kluwer international series in engineering and computer science; knowledge representation, learning, and expert systems, edited by Tom Mitchell), 1993, ix+438 pp; hardbound, ISBN 0-7923-9125-X, US$117.50

Word Manager: A system for morphological dictionaries *Marc Domenig and Pius ten Hacken* (Universität Basel) Hildesheim, Germany: Georg Olms Verlag (Informatik und Sprache 1), 1992 ix+211 pp; paperbound, ISBN 3-487-09677-3, no price listed

Origins of the modern mind: Three stages in the evoluation of culture and cognition *Merlin Donald* (Queen's University) Cambridge, MA: Harvard University Press, 1991, reissued in paperback 1993, viii+413 pp; paperbound, ISBN 0-674-64484-0, US$14.95

Machine translation: A view from the lexicon Bonnie Jean Dorr (University of Maryland) Cambridge, MA: The MIT Press (Artificial intelligence series, edited by J. Michael Brady, Daniel G. Bobrow, and Randall Davis), 1993, xx+432pp; hardbound, ISBN 0-262-04138-3, US$45.00

Linguistic issues in machine translation *Frank Van Eynde* (editor) (University of Leuven) London: Pinter Publishers (Communication in artificial intelligence series, edited by Robin P. Fawcett and Erich Steiner), 1993, viii+239 pp; distributed in North America by St Martin's Press; hardbound, ISBN 1-85567-024-0, US$79.00

***Recent advances in qualitative physics** *Boi Faltings and Peter Struss* (editors) (Swiss Federal Institute of Technology and Siemens Corporate Research and Development) Cambridge, MA: The MIT Press (Artificial intelligence series, edited by J. Michael Brady, Daniel Bobrow, and Randall Davis), xii+449 pp; hardbound, ISBN 0-262-06142-2, US$39.95

A formalisation of design methods: A lambda-calculus approach to system design with an application to text editing *Loe Feijs* (Philips Research Laboratories, Eindhoven, The Netherlands) Chichester: Ellis Horwood (Ellis Horwood series in computers and their applications, edited by Ian Chivers), 1993, xv+502pp; hardbound, ISBN 0-13-106113-5, no price listed

The engineering of knowledge-based systems: Theory and practice *Avelino J. Gonzalez and Douglas D. Dankel* (University of Central Florida and University of Florida) Englewood Cliffs, NJ: Prentice-Hall (an Alan R. Apt book), 1993, xx+523pp and two 5.25" diskettes for MS-DOS; hardbound, ISBN 0-13-276940-9, no price listed

Neural network learning and expert systems *Stephen I. Gallant* (HNC Inc) Cambridge, MA: The MIT Press, 1993, xvi+365 pp; hardbound, ISBN 0-262-07145-2, US$42.50

Contributions to quantitative linguistics *Reinhard Köhler* and *Burghard B. Rieger* (editors) (University of Trier) Dordrecht: Kluwer Academic Publishers, 1993, xii+436pp; hardbound, ISBN 0-7923-2197-9, US$158.00

Techniques of Prolog programming, with implementation of logical negation and quantified goals *Tu Van Le* (University of Canberra) New York: John Wiley and Sons, 1993, xviii+601 pp and two MS-DOS diskettes; paperbound, ISBN 0-471-57175-X, no price listed

***Intelligent behavior in animals and robots** *David McFarland and Thomas Bösser* (University of Oxford and Westfalische Wilhelms-Universität) Cambridge, MA: The MIT Press (Complex adaptive systems series), 1993, xi+308pp; hardbound, ISBN 0-262-13293-1, $39.95

A new era in computation *N. Metropolis and Gian-Carlo Rota* (editors) (Los Alamos National Laboratory and MIT) Cambridge, MA: The MIT Press, 1993, xii+241 pp; based on the winter 1992 issue of *Daedalus*; paperbound, ISBN 0-262-63154-7, US$13.95

Conceptual structures: Current research and practice *Timothy E. Nagle, Janice A. Nagle, Laurie L. Gerholz, and Peter W. Eklund* (editors) (Unisys Corporation; School District 833, St. Paul Park, MN; Unisys Corporation; and University of Adelaide) Chichester: Ellis Horwood (Ellis Horwood series in workshops), xv+644pp; hardbound, ISBN 0-13-175878-0, no price listed

The alternative mathematical model of linguistic semantics and pragmatics *Vilém Novák* (Czechoslovak Academy of Sciences) New York: Plenum Press (International series on systems science and engineering, edited by George J. Klir, volume 8), 1992, xiii+204 pp; hardbound, ISBN 0-306-44269-8, US$65.00

The logic programming tutor Jocelyn Paine (Oxford University) Oxford: Intellect Books and Dordrecht: Kluwer Academic Publishers, 1992, xi+365 pp; paperbound, ISBN 1-871516-09-9 (Intellect) and 0-7923-1448-4 (Kluwer)

Language, music, and mind *Diana Raffman* (The Ohio State University) Cambridge, MA: The MIT Press, 1993, xi+169 pp; hardbound, ISBN 0-262-18150-9, US$22.50

The Soar papers *Paul S. Rosenbloom, John E. Laird, and Allen Newell* (editors) (Information Sciences Institute, University of Michigan, and Carnegie Mellon University) Cambridge, MA: The MIT Press (Artificial intelligence series), 1993, xli+1438 pp in two volumes; hardbound, ISBN 0-262-68071-8, US$95.00; paperbound, ISBN 0-262-68071-8, US$65.00

Computational linguistics and formal semantics *Michael Rosner and Roderick Johnson* (editors) (IDSIA) Cambridge, England: Cambridge University Press (Studies in natural language processing, edited by Branimir K. Boguraev), 1992, xix+321 pp; hardbound, ISBN 0-521-41959-X, US$74.95; paperbound, ISBN 0-521-42988-9, US$24.95

***Machine models of music** *Stephan M. Schwanauer and David A. Levitt* (editors) (Center for Computer Research in Music and Acoustics, Stanford University and HIP Software) Cambridge, MA: The MIT Press, 1993, xi+544 pp; hardbound, ISBN 0-262-19319-1, US$39.95

Computers and writing: Issues and implementations *Mike Sharples* (editor) (University of Sussex) Also published as a special issue of *Instructional Science*, 21(1-3), 1992. Dordrecht: Kluwer Academic Publishers, 1992, 224 pp; hardbound, ISBN 0-7923-1966-4, US$82.00

***Analogy and structure** *Royal Skousen* (Brigham Young University) Dordrecht: Kluwer Academic Publishers, 1992, xii+376 pp; hardbound, ISBN 0-7923-1935-4, US$124.00

The logical approach to syntax: Foundations, specifications, and implementations of theories of government and binding *Edward P. Stabler, Jr* (University of California, Los Angeles) Cambridge, MA: The MIT Press (The ACL—MIT Press series in natural language processing, edited by Aravind K. Joshi, Karen Sparck Jones, and Mark Y. Liberman), 1992, xiv+433 pp; hardbound, ISBN 0-262-19315-9, US$45.00

The phenomenon of common sense reasoning: Nonmonotonicity, action, and information *Dimitrios Thanassas* (Philips Research Laboratories, Redhill, Surrey) Chichester: Ellis Horwood (Ellis Horwood series in artificial intelligence, edited by John Campbell), 1992, x+254pp; hardbound, ISBN 0-13-663634-9, no price listed

***Artificial intelligence through search** *Christopher James Thornton and Benedict du Boulay* (University of Sussex) Dordrecht: Kluwer Academic Publishers and Oxford: Intellect Books, 1992, xvii+365pp; hardbound, ISBN 0-7923-1868-4, US$99.00.

What is cognitive science? *Barbara Von Eckhardt* (University of Nebraska, Lincoln) Cambridge, MA: The MIT Press, 1993, x+466 pp; hardbound, ISBN 0-262-22046-6, US$45.00

Computational models of American speech *M. Margaret Withgott and Francine R. Chen* (Interval Research Corporation and Xerox Palo Alto Research Center) Stanford: Center for the Study of Language and Information (CSLI lecture notes 32), 1993, viii+143 pp; distributed by the University of Chicago Press; hardbound ISBN 0-937073-97-0,US$39.95; paperbound, ISBN 0-937073-98-9, US$15.95

***Archon: An architecture for multi-agent systems** *Thies Wittig* (editor) (Atlas Elektronic GmbH, Germany) Chichester: Ellis Horwood (Ellis Horwood series in artificial intelligence, edited by John Campbell), 1992, 134 pp; hardbound, ISBN 0-13-044462-6, no price listed

Multilingual multimedia: Bridging the language barrier with intelligent systems *Masound Yazdani* (editor) (University of Exeter) Oxford: Intellect Books, 1993, x+210 pp; paperbound, ISBN 1-871516-30-7, no price listed

Intelligent systems for business: Expert systems with neural networks *Fatemeh Zahedi* (University of Massachusetts, Boston) Belmont, CA: Wadsworth Publishing Company (The Wadsworth series in management information systems), 1993, xxvii+658 pp; hardbound, ISBN 0-534-18888-5, no price listed. Distributed in Canada by Nelson Canada.

CSCSI/SCEIO Membership

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

☐ I am a student ($30.00* Cdn./yr.)
and/or

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

Name ______________________________

Mailing
Address ______________________________

____

____

____

Please mail your membership to:

CIPS
430 King Street West, Suite 106
Toronto, Ontario
M5V 1L5
Phone: 416-593-4040
Fax: 416-593-5184

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

Advertisers Index

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

Small Size

Modularity

High Computational Power

Easy to Control

Programmable Highly Intelligent Operations

Low Cost

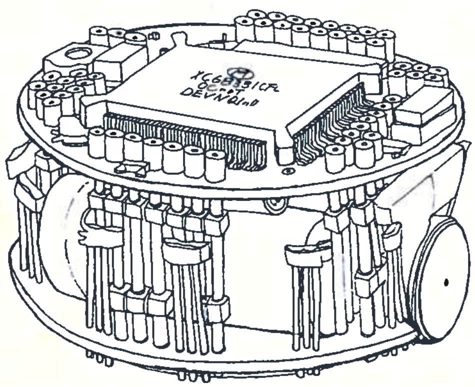

Khepera®

New possibilities in mobile robot experiments

-Experiments in a small area (e.g. a table top) with everything at hand
-Experiments involving many robots (e.g. Swarm Intelligence experiments in ALife research)
-Easy transportability (e.g. to other laboratories, conferences, demonstrations,...)

Possible applications

-Obstacle avoidance
-Wall following
-Target search
-Embodied learning algorithms
-Collective behavior (Swarm intelligence)
-Various Artificial Life (ALife) experiments

Operations

-Wired operation from host
-Downloaded, fully self contained operation

Specifications

Processor:	Motorola 68331
RAM:	256 Kbyte
ROM:	256 or 512 Kbyte
Motion:	DC Motors with incremental encoder (10 pulses per mm of advancement of the robot)
Sensors:	8 Infra-red proximity and light sensors
Power:	Rechargeable Ni-Cd batteries or external power supply
Autonomy:	20 min. max (basic configuration with maximum activity)
Extension bus:	The robot can be expanded by new modules added on the K-Extension bus
Size:	Diameter 55 mm. Height 30 mm.
Weight:	Approximately 70 g

Software environment (not included)

-Calm assembler (on PC, MAC & SMAKY)
-GNU C compiler (on all machines supported by GNU)
-LabVIEW® (on PC, MAC & SUN)
(LabVIEW is a trademark of National Instruments.)

Produced by: forelec
Jambe-Ducommun 6b
Switzerland

Developed by: lami
Laboratorie de Microinformatique
Lausanne, Switzerland

EPFL
Ecole Polytechnique Federale de Lausanne
Lausanne, Switerland