Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Autumn 1992 No. 30 automne 1992

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

Visualizing and Understanding Diagnoses
Suhayya Abu-Hakima
Visualiser et Comprendre les Diagnostics

Report on the 4th International Symposium on Artificial Intelligence
Compte rendu sur 4e le Symposium International sur l'Intelligence Artificielle

Pitch Expert: Canadian AI Success Story # 1
Pitch Expert: Récite #1 sur les Succès d'IA au Canada

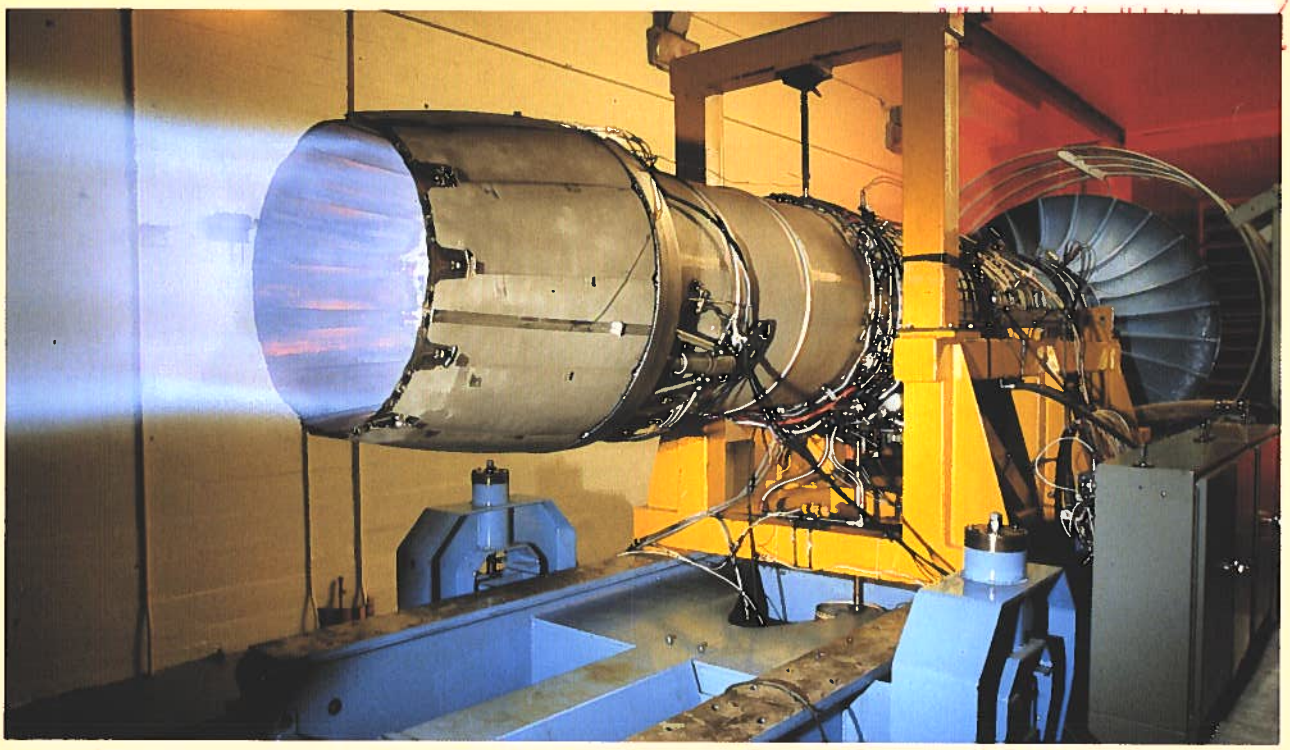

General Electric F404 Aircraft Engine being diagnosed in test cell at the National Research Council of Canada, Ottawa.

Advanced AI Tools for Applications Development

NEURON DATA

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

Soft Warehouse Inc

Soft Warehouse Inc – muLISP-90: Lisp language programming environment. Turns a PC into an integrated workstation for developing AI software. With compiler, $490. DERIVE: Mathematical assistant for the PC. Solves large and complex problems. $310.

Talarian Corporation

RTworks™ – The tool for distributed AI applications development. Originally developed for advanced aerospace applications, RTworks combines realtime data acquisition (RTdaq), realtime graphical user interface for data display and input (RThci), an extremely high performance inference engine (RTie), full data archiving (RTachive) and retrieval (RTplayback). Its powerful distributed architecture allows multiple copies of RTworks processes to cooperate transparently over mixed networks of several vendors's workstations. Available on most workstations including Sun, DEC, Silicon Graphics, Hewlett Packard, IBM, etc. From $43,000

VOICE CONNEXION

Micro IntroVoice – Complete voice input/output system providing voice recognition of 1,000 words and unlimited text-to-speech synthesis. Can operate in noisy environments in excess of 85 db for reliable performance in factory floor or other high-noise areas. Provides DOS compatibility for use with IBM XT, AT, 386 and PS/2 Micro Channel computers. $1580.

NeuralWare Incorporated

NeuralWorks Professional II Plus – Advanced, easy-to-use "graphic-intuitive" approach results in a seamless interface. User-defined neurodynamics, and loadable control strategies provide complete freedom in designing multi-paradigm networks or developing completely new network types. IBM PC AT or PS-2 and MAC Plus, MAC SE, MAC II $2,775. SUN/3 or SUN/4, $5,850.

DRAGON SYSTEMS, INC.

Dragon Systems, Inc. – DragonWriter and DragonDictate voice recognition/adaptive systems have been acclaimed the best voice recognizers for IBM XT, AT and 386 machines. They have constantly outperformed systems costing several times more. The design is based on a sound set of algorithms developed by two of the top researchers in the research community, Drs. Jim and Janet Baker. The devices are ideal for such applications as command/control, data entry/retrieval, form filling (hospital, government), documentation creation or dictation, front-ending of expert systems, and support systems for the handicapped. From $1,950 to $11,000.

AI

Applied AI Systems, Inc.

RTworks™ Extensions – Applied AI Systems, Inc. announces plans for the integration of RTworks with other subdisciplines of AI thus enabling RTworks users to apply the power of fuzzy set theory, neural networks, case-based reasoning and genetic algorithms in their real-time applications when needed. This will increase the flexibility and applicability of the tool for a wide variety of applications. Scheduled release dates for the integrated products include RTworks + Fuzzy Logic in December 1992; RTworks + Neural Networks in March 1993; RTworks + Genetic Algorithms in June 1993; and RTworks + Case-Based Reasoning in September 1993.

Applied AI Systems, Inc.

VIP (Voice Interactive Package) – Complete voice recognition and synthesis package for developing a multi-modal interface to applications developed using NEXPERT Object. Speaker dependent voice recognition of up to 1,000 commands and unlimited number of output texts. Especially valuable in "hands busy, eyes busy" situations and where access to the system is limited. $5,500. A version that interfaces MS-DOS directly under development.

IS Robotics — Exploring the Future

Intelligent Robots – Using Subsumption Architecture or behaviour-based approach developed at MIT's AI Research Labs, robots are becoming a practical and commercial reality in labs, industry, classrooms and the home. Their small size, intelligence and sensing capabilities make the robots ideal for research, prototyping, sensor placement and tool carrying in remote, cramped or hazardous environments. Atilla-II, a six-legged micro-robot, can walk as fast as a person and is equipped with 150 sensors of 12 different types. T-1, T-2 and T-3 are highly manoeuverable tracked robots suitable for mixed terrain applications. T-2 is equipped with a colour camera and T-3 with sophisticated sensors. R-1 and R-2 are wheeled robots designed to support research in cooperative behaviour studies and application-specific experiments. Pebbles is designed for large payload capacity, is sealed to dust and can be sealed to water, oil, and chemicals.

Togai InfraLogic, Inc.

Togai InfraLogic, Inc. – A family of hardware and software tools which allow seamless integration of fuzzy logic in a wide variety of applications including factory automation, motor control, robot arm control, pattern recognition and classification, realtime process control, sensor fusion and task scheduling. TILShell is an object based, graphical window development environment for fuzzy expert systems; TILGEN is an automated rulebase generation tool; FC110 Digital Fuzzy Processor is a single chip VLSI processor for embedded fuzzy logic applications, and Single Board Fuzzy Controller (SBFC) is designed for control applications. From $6,000.

Carnegie Group

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

ARITY

Arity Knowledge Systems Development Pack – A new generation AI prototyping tool for PCs. A prolog interpreter/compiler, a rule/frame-based Expert Systems shell with object hierarchy and inheritance, the external file and database access, a screen design kit, linkage to conventional languages, and an IBM SQL access. From $1,220. Now available with Arity Windows toolkit.

TRIMETRIX

Axum – Superior technical/scientific graphics and data analysis package for PC's. Compatible with WordPerfect™, Word™, Ventura™, PageMaker™; import data from Lotus™, dBase™, and ASCII. Intelligent defaults permit easy, fast and high-quality results. Diverse plot types are included with numerous fonts and symbols in addition to a sophisticated data editor. $585.

ENQUIRY AND ORDER: Call (613) 592-7729 or send Purchase Order to Applied AI Machines & Software, Inc. Shipping and Handling charges and all applicable sales taxes apply. Applied AI Machines & Software, Inc. is an authorized dealer of the above AI software/hardware products. Prices are in Canadian dollars and are subject to change without notice.

AI

Applied AI Machines & Software, Inc.
Suite 504, Gateway Business Park
340 March Road
KANATA, Ontario, Canada K2K 2E4
Tel: +1 613 592 7729
Fax: +1 613 592 9762

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Autumn 1992 No. 30 automne 1992

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 — masrani@skyler.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

APPLAI - 93

Second Workshop on the Commercialization of AI Technology in Canada

Announcement and Call for Papers

The second annual Canadian Applied AI workshop will be held at York University in Toronto on **May 18, 1993**. The objective of the workshop is to get together both representatives of Canada-based companies and institutions which work on the development and commercialization of AI software. The one day workshop will consist of a technical session with paper presentations, a "show and tell" session, and a forum on the current and future expectations of Applied AI in Canada. The workshop is held in conjunction with Vision/Graphics Interface 93 which is held from May 18-21.

People interested in attending are invited to submit a short description of their institution/company and its products. Those wishing to make a presentation are requested to submit an abstract of 2000 words max. We are particularly interested in hearing about Canadian success stories. Submissions on products under development or near completion, however, are also welcome. All submitted abstracts will be made available to the participants. A subset of abstracts will be selected for presentation in the workshop. Selection of presentations will be made in an attempt to reflect the diversity of centres/companies carrying out these kinds of research and at the same time to reflect the integration into the applied sectors.

Contributions are invited from (but not limited to) the following domains: space, manufacturing, agriculture, natural resources, transportation, utilities, human factors, communications, government operations, and finance.

Potential applications: advisory/support, training, control, planning/scheduling, visual inspection and guidance.

Submitted abstracts must be accompanied by a cover page with the following contents: Paper title, author(s), company/ institute, address, telephone, fax and e-mail address (if available). Abstracts are due on **February 8, 1993** and should be sent to the address given below.

For further information please contact:
Jan Mulder
Alberta Research Council
6815 -8 Street NE
Calgary AB, T2E 7H7
phone: (403) 297 - 7570, fax: (403) 275 - 3003, e-mail: mulder@skyler.arc.ab.ca

Organizing Committee:
Jan Mulder, Alberta Research Council
Russ Thomas, National Research Council

Letter to the Editor

Fuzzy Logic

In issue 29 (Summer'92) of **Canadian Artificial Intelligence** *Magazine you published a two page article entitled "Fuzzy Logic: A Basic Phenomena (sic)".*

Was this a leftover from a delayed April Fools' issue? Even under that assumption it wasn't very funny.

Although **Canadian Artificial Intelligence** *is not a refereed publication I think it is important to restrict publication to articles that are written in correct English or French, and are technically accurate. The article on fuzzy logic met neither condition.*

Sincerely,
Dick Peacocke

New Publishing Schedule

Canadian Artificial Intelligence Magazine is now published 3 times per year in February, June, and October.

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

AI NEWS
NOUVELLES DE L'IA

News Release

INFORMATION TECHNOLOGY STATISTICS PROMISING

OTTAWA, Ontario, May 22, 1992 - Michael Wilson, Minister of Industry, Science and Technology and Minister for International Trade, and Information Technology Association of Canada (ITAC) President Janice Moyer today released the 1991 Statistical Review of Information Technologies. The report indicates that the information technology sector generated over $40 billion in total revenue in Canada during 1990.

"The information technology sector, which is strategic to the Canadian economy, now employs more than 280,000 Canadians and accounts for 35 per cent of industrial research and development expenditures," said Minister Wilson. "With an average growth rate of 9 per cent predicted for this sector, information technology will be a major contributor to Canada's future prosperity," he added.

The information technology sector has grown at three times the rate of the national economy between 1986 and 1990. Expenditures on computers and peripheral equipment have more than doubled to $5.1 billion. Canada's imports and exports of information technology products are on the rise. As a proportion of Gross Domestic Product, the value Canadians add to products and services, the IT manufacturing sector now equals that of the pulp and paper industry, one of Canada's traditional economic engines.

The Information Technology Statistical Review was prepared by Industry Science and Technology Canada (ISTC) based on data from Statistics Canada, private research firms and ISTC estimates. Communications Canada and ITAC were consulted in the preparation of the document.

"We are pleased that ISTC proceeded with this collective effort", said ITAC President Janice Moyer. "The report clearly shows how important the information technology industry is to the Canadian economy and it will be a useful tool for helping the industry and governments devise strategies to keep Canada on the leading edge of this rapidly evolving business," she added.

Contact: Ed Theriault, ISTC
(613) 954-3317

To obtain copies of the report, contact:

ISTC
Tel: (613) 954-3317
Fax: (613) 952-8419

ITAC
Tel: (416) 493-8710
Fax: (416) 493-8758

Visualizing and Understanding Diagnoses

Suhayya Abu-Hakima

Knowledge Systems Laboratory
Institute for Information Technology
National Research Council, Ottawa, Canada
email: suhayya@ai.iit.nrc.ca tel: (613) 991-1231

Sommaire

Le diagnostic de systèmes physiques tels que les automobiles et les moteurs d'avion est une activité complexe. Les techniciens allient les manuels textuels (qui incluent des schêmas) avec l'analyse des données mesurées pour diagnostiquer et réparer les moteurs. Les concepteurs de système à base de connaissance ont ajoutés l'heuristique, pour lier le texte avec les graphiques (hypermedia) pour simplifier la tâche des techniciens, tel qu'implanté dans le projet JETA (Halasz92). Des outils pour examiner la connaissance sont utilisés par les concepteurs de tels systèmes pour structurer et injecter la base de connaissance et sont aussi utilisés pour aider l'expert à visualiser la connaissance et les différentes relations possibles. De tels outils ont été conçus pour afficher et éditer la base de connaissance du projet JETA. Dans RATIONALE, un système de diagnostics qui raisonne en fournissant des explications, les explications sont utilisées pour comprendre le raisonnement du système [Abu-Hakima90]. Cet article argumente que même si ces approches à base de connaissances aident à la visualisation et la compréhension des systèmes physiques, ils ont besoin d'être améliorés et mieux intégrés.

Summary

Diagnosis of physical systems such as car or aircraft engines is a complex activity. Technicians combine textual manuals with schematics and some analysis of measured data to diagnose and repair engines. Knowledge-based system designers have added heuristics to link text and graphic (hypermedia) representations of manuals to simplify the tasks of the technicians as implemented in JETA [Halasz92]. Knowledge browsers are used by the developers of such systems to structure and input the knowledge base and are also used in a limited capacity to help the domain experts visualize the knowledge and the various possible relations. Such a browser has been implemented to view and edit JETA's knowledge. In RATIONALE, a diagnostic system that reasons by explaining, explanation is used to understand system reasoning [Abu-Hakima90]. This paper argues that although these knowledge-based approaches help in the visualization and understanding of diagnoses in physical systems, they need to be improved and better integrated.

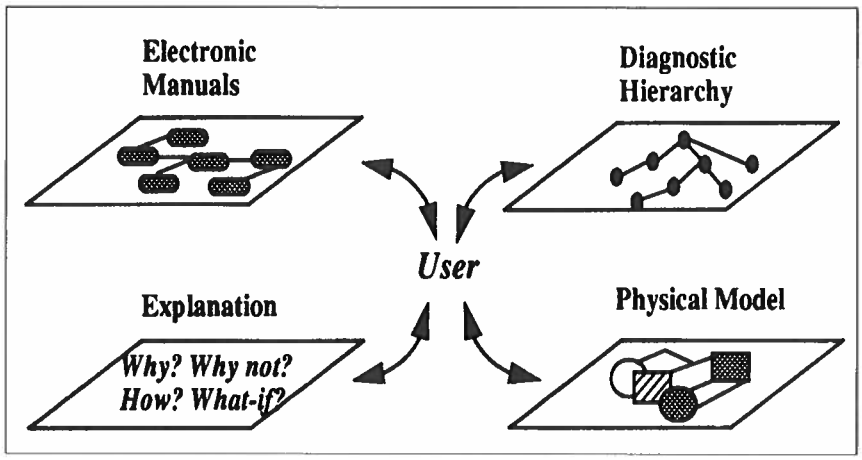

Figure 1: World of user in knowledge-based system

Introduction

Diagnosis of physical systems such as car or aircraft engines is a complex activity. Technicians combine paper manuals with schematics and some analysis of measured data to diagnose and repair engines. Knowledge-based systems provide the technicians with electronic manuals organized using hypermedia techniques as well as diagnostic hierarchies that represent the failure, test and repair actions of the diagnostic cycle. Such an approach has been followed for JETA, the Jet Engine Troubleshooting Assistant [Halasz92]. Other systems have followed modelling and simulation techniques that represent the actual physical system and attempt to diagnose it on the basis of the expected behaviour of the model [MBR91]. Some diagnostic systems

support the user with sophisticated explanations of diagnoses that help justify system behaviour and clarify the reasoning to the user. RATIONALE is a diagnostic system which was developed to provide the user with such explanations [Abu-Hakima90]. Figure 1 illustrates the information flow with respect to the user.

All these systems fall short in helping the user visualize as well as understand system diagnoses. The primary reason for this shortfall is the lack of integration and mapping of the information represented to the user. Two representations can be thought of in diagnosing a system. The first is a physical real-world representation of the various components and their interaction. This is best visualized using a multi-dimensional simulation true to the actual physical system. Supporting the real-world physical system is a paper manual describing its operation, failure testing and repairs. The second representation used to diagnose a physical system is the abstract one. This takes the form of representing the information known about the workings and failures of a physical system, derived from manuals and the experience of technicians, and encoding it into a diagnostic system.

This position paper elaborates on the various information sources that the world of a diagnostic system user holds. A diagnostic system user has hypermedia to index electronic manuals, knowledge browsing to understand diagnostic hierarchies and explanation to justify system reasoning. The final section of the paper describes how these information sources could be better integrated so that a user can better visualize and understand a diagnosis.

Navigating Electronic Manuals - Hypermedia

Hypermedia, the linking of textual and graphical information (and in some cases voice information) has become essential in the use of electronic manuals. The electronic manual, much like its paper counterpart, has indices and cross references. The difference lies in the ability to search the electronic manual in seconds versus minutes. Electronic manuals make use of hypertext links that display to the user definitions or entire procedures based on the selection of a keyword in context. These same links allow the user to access any cross referenced schematics or drawings. This capability has been implemented in JETA, the Jet Engine Troubleshooting Assistant [Halasz92]. The implementation of hypermedia in JETA has been facilitated by the existence of powerful user interface tools such as HyperNeWS.

HyperNeWS is a window interface tool [vanHoff91] written in PostScript which uses NeWS graphics primitives. It is similar in some ways to HyperCard™, but it can have client processes in LISP, Quintus Prolog and C. It is an object-oriented tool which runs under Unix. Its stacks can take any shape the developer creates using the drawing tool and can have text objects, check boxes, user created iconic buttons, etc. The primary advantage of this tool is its use of meaningful images, graphics and icons to relay information to the user.

JETA integrates heuristics with its hypermedia manuals. These heuristics form the basis of its frame-based knowledge representation. The knowledge representation was achieved by working with a domain expert in the structuring of the manuals. Its knowledge is organized into a diagnostic network based on a hierarchy of general diagnoses that are focused to more specific ones as described in the next section. These frames contain procedures and markers to schematics that are made available to the user as advice. JETA also has a help system with hypermedia links that a user can access at any time during the session.

Understanding Diagnostic Hierarchies - Knowledge Browsing

As a knowledge base is developed, it is more difficult to track relations and to maintain the consistency of the information. Knowledge acquisition becomes more difficult since gaps in the knowledge base are not easily found. In response to these issues, knowledge browsers are developed. A knowledge browser graphically represents the information in the knowledge base. Several researchers have been using knowledge browsers to graphically illustrate the relations between rules. Several such systems have attempted to address the problem of visually relating medical knowledge used in the diagnosis and treatment of patients [Banks88]. Another has been in the area of viewing and finding relations in statistical information [Stephenson88].

The objectives for browsers in all domains are similar: to produce an observable domain structure (such as a diagnostic hierarchy); to allow the user unrestricted movement in the structure; to allow for alternative views of the same information; and to make use of a mouse with menus to reduce keyboard entry. As a result of these objectives the user needs: a guide or map for the structure; the ability to graphically modify the knowledge structure; markers that support audit trails; line type and colour to indicate relations; filters to reduce the amount of information; keyword search mechanisms; consistency checking mechanisms for the knowledge structure; and user tailoring.

JETA is a frame-based rather than a rule-based system [Halasz92]. Its knowledge base is structured with defined relations represented by slots in the node frame. Thus, its knowledge browser has to be flexible enough to visually communicate the relations between the nodes. As a result, a diagnostic tree with colours on node and line features has been chosen to communicate relational information.

The knowledge structure is a hierarchy of nodes (as illustrated in Figure 2)with a root node at the top level. At the next level is a split between the phases of operation of the engine and the list of major symptoms. This level is followed by several levels of functional or procedural test and general nodes. The lowest level in the hierarchy includes

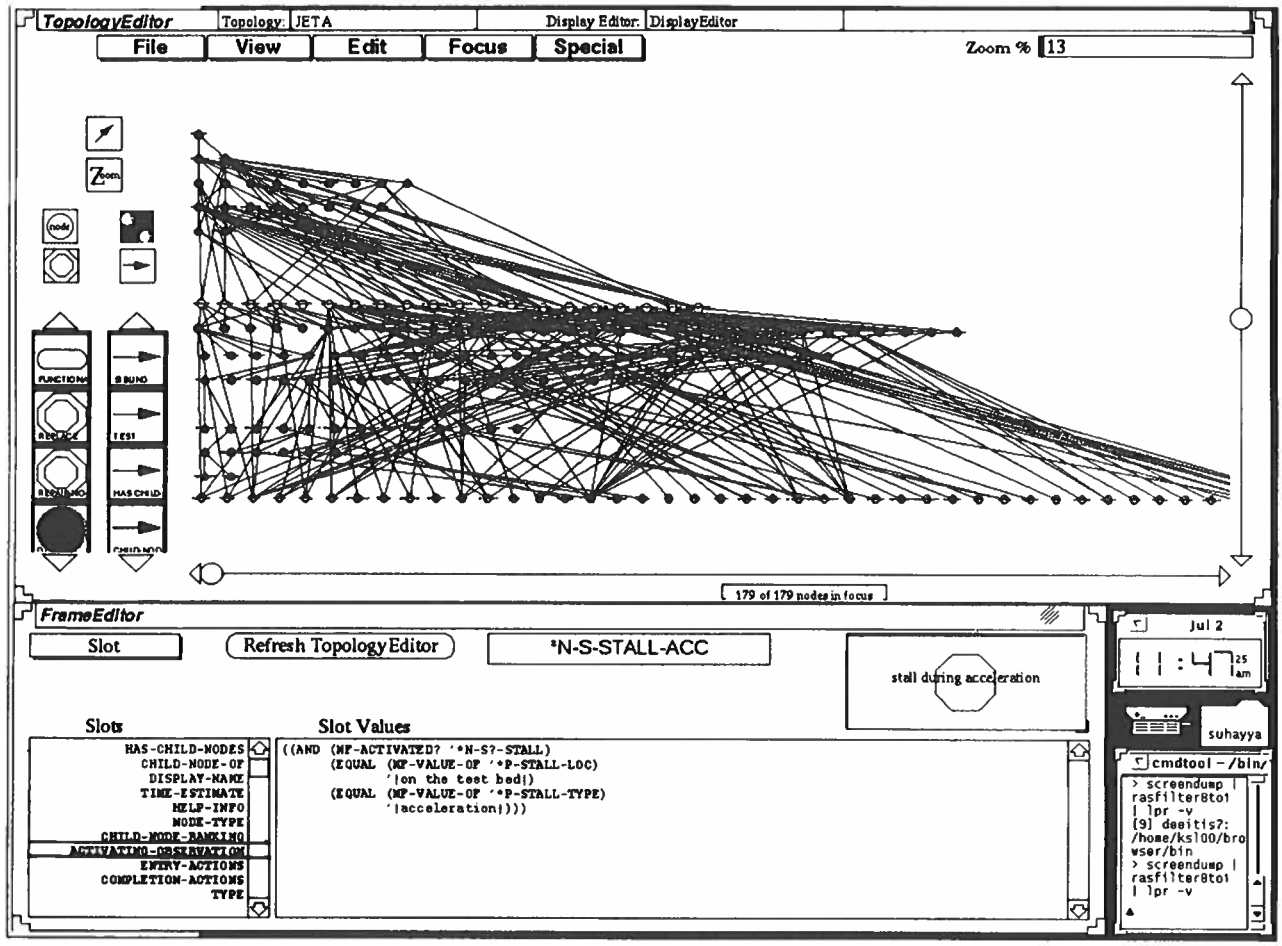

Figure 2: Knowledge Structure

replace, repair and adjust nodes. It is important to note that the hierarchy, although it appears as such, is not a true one and is a directed network since some of the terminal nodes may require further tests which cycle back to earlier nodes in the network. In addition, some nodes may have links to other nodes that are not directly connected.

An important feature of the browser is the ability to zoom in on the network. The zoom mode allows the user to 'rubber band' or put a box around an area they wish to zoom in on. The selected area is then expanded to the size of the current browser window. The browser allows the developer to edit the network graphically. Nodes and links can be added and deleted. Nodes can also be moved with their links shortened or lengthened accordingly.

An important difference between traditional knowledge browsers and JETA's is the ability to browse relations in a frame. In the browser the lines connecting the nodes can represent the child-node-of relation which is a slot in each frame representing a node. The arrow at the end (visible on zooming) indicates that the node pointed to is a subproblem of the previous node. At a glance the developer viewing and editing the knowledge base can view the structure of the knowledge base and find what key relations have been defined for a set of nodes.

The browser allows one to focus on a node. Focusing allows the user to view a node and its parent and sibling links without viewing any other nodes in the tree. Focusing is recursive in nature and allows the developer to repeatedly focus on the parents of a node, the parents of the parents of the node, etc.

Explanation

We and others have long argued for the need of explicit knowledge representation and reasoning strategies for the generation of concise, coherent contextual explanations [Abu-Hakima 88& 90; Neches 85]. To maintain the context of explanation in diagnostic systems, we have argued that the knowledge base be organized in a hierarchical structure. The most general knowledge should be represented near the root of the hierarchy and the more specific knowledge should be represented near the end nodes. Since some of the explanatory knowledge is represented locally in the frame that is reasoned about, explanatory contexts would allow

relations between nodes such as refinements, siblings or alternates to be explained [Abu-Hakima 88& 90]. Templates are used to structure generated text responses.

The explanations generated for RATIONALE, an earlier implemented diagnostic system that reasons by explaining, are textual but are available in a hypermedia interface that has been described previously [Abu-Hakima 89]. We are currently specifying explanations for JETA which will be modelled after RATIONALE's. In troubleshooting jet engines, technicians make use of manuals, graphs and schematics showing how to dismantle and repair subsystems. For our user interface to truly support a technician, it is necessary for it to access and link relevant schematics, graphs and template-based text in context. Further, it is essential that explanation be directly tied into this explicit reasoning and hypermedia interface.

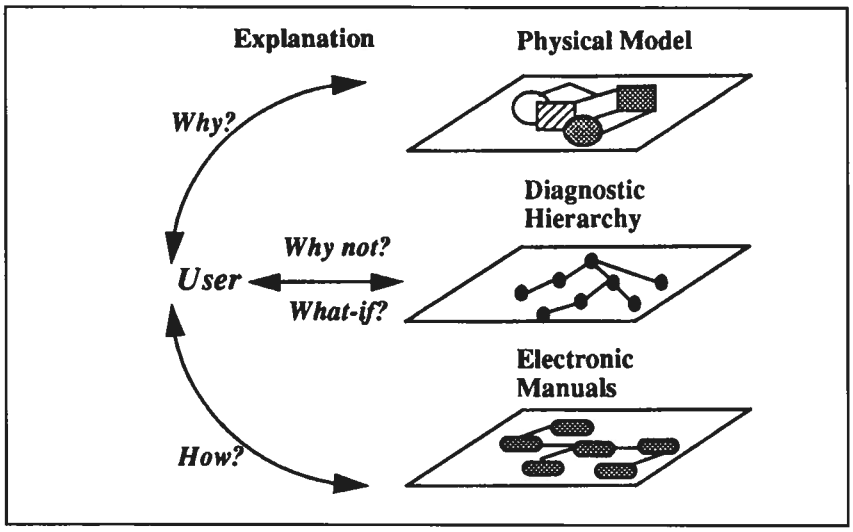

Figure 3: Integrated world of user in knowledge-based system

In RATIONALE, explanations are generated from traces of diagnostic reasoning. As described in [Abu-Hakima 90] they allow the user to ask questions about dynamic and static knowledge. This includes why, why-not and what-if questions about the dynamic trace. It also allows the user to ask about deductive and abductive methods of the system and how conclusions are drawn from both the static and dynamic knowledge.

Integration and Visualization

For a user to understand and visualize the workings of a knowledge-based system for diagnosis, several layers of information need to be better integrated as illustrated in Figure 3. There is a need for a visual model of the actual physical system. For example, in the case of a car engine, it would include the engine block, the carburetor, the pistons, the valves, the air and fuel intake valves, the fuel pump, the water pump, etc. All these components should be visually modelled so that the user can see the various components acting and reacting as they would in the real world. The user should then be able to zoom in and out on the subsystems of the model down to the component level.

A second layer directly related to this layer should be a knowledge representation and organization that allows the user to view the mapping from the real world physical system to the knowledge-based representation. This mapping could be in the form of a diagnostic hierarchy which includes the phases of operation of a system. In the case of the car engine rootnode, the next level could be start, accelerate, decelerate and stop. The next level in the hierarchy would relate the phases of operation to the symptoms that would cause components of the car engine subsystems to produce problems at each of the phases of operation. For example, in the case of an acceleration problem, we could have a fuel leak, a fuel pump malfunction, a distributor malfunction, etc. This level would then be made more specific down to the lowest component level of the engine that is causing the problem. The structure of this knowledge would be illustrated using the knowledge browser. Note that at the lowest level the user would see the individual components of the engine mapped from the real world physical model referred to above.

The third layer would be the hypermedia representation of the electronic manual. This would have hypertext links to procedures and tests included in the diagnostic hierarchy. In addition, any schematics or visual representation of procedures would make use of the real world physical model of layer one. Thus, if the user is required to dismantle a subsystem in a particular order, a simulation of the dismantling down to the individual components would be played back making use of the visualization of layer one. This is akin to playing back to the user a video of the system being dismantled into its individual parts. This approach communicates causal information quite easily and is related somewhat to the concepts discussed in [Feiner90].

The underlying link between the three layers comes in the form of explanation. Explanation of why the physical components in a subsystem have caused it to fail are much easier to generate with a visual aid such as those available in layers one and two.

An explanation of failure is much easier to visualize by showing the moving components of the physical system acting and reacting. In the car example, it is much easier to visually illustrate a piston seizing due to the lack of oil lubrication rather than textually describing it.

Another graphical enhancement to the content of an explanation would be the use of two to three level sub trees of the diagnostic hierarchy of the browser to illustrate relations. What if, for example, the user asks why a certain

conclusion was not made? If the conclusion was due to a refinement of the original hypothesis, the root of the subtree would be the original hypothesis and the subsequent refinements of that node would be illustrated using children of the root and how they failed to confirm a particular problem. In the case of a trace which required the traversal of many levels of nodes to reach a conclusion, a mechanism such as that used in RATIONALE's textual explanations [Abu-Hakima 88] could rank the nodes for explanation content in terms of relevance and complexity. In the cases where explaining the node is overly complex compared to its relevance to the current line of reasoning, it is omitted from the explanation. In the cases where it is relevant and not overly complex it is included.

Explanations should also be added to the browser itself. Explanations could be used to describe the various relations and how they are used in the diagnosis. A user who is not the developer, such as a domain expert, could benefit by viewing, commenting and possibly modifying the knowledge base.

A knowledge browser which uses visual cues such as colour and line shapes to provide the user with a structured view of the knowledge is essential for the average user of a system. With such a tool, the user could better visualize and understand the mappings between the real world physical layer, the electronic manuals and the diagnostic system.

Another level of integration would allow the user to generate hypermedia queries off the physical visualization. This could be offered in varying degrees of depth such as contextual help.

Conclusion

In conclusion, it is essential that diagnostic systems make use of existing visual aids such as multi-dimensional models and diagnostic network browsers to better serve users. The integration of such tools requires that the physical and abstract representations used in knowledge-based systems be better integrated. It is also important to recognize that there are many classes of users of diagnostic systems. There are developers, domain experts and end users (the technicians). Within each class of user there are varying levels of expertise which translate to specific needs that should be addressed. Finally, explanation is a powerful tool necessary for better understanding and visualizing diagnoses. It has thus far been poorly integrated and is essential for improving diagnostic systems for the end user.

Acknowledgements

The author would like to acknowledge a number of individuals who have contributed in various ways: Mike Halasz and Sieu Phan of the Knowledge Systems Laboratory who were key players in developing JETA; Dave Peloso of GasTops Limited and Tim Taylor of Phalanx Research who were contracted to implement JETA's HyperNeWS interface; and the Generalized Knowledge Browser/Editor respectively.

References

[Abu-Hakima88] Abu-Hakima S. RATIONALE: A Tool for Developing Knowledge-Based Systems that Explain by Reasoning Explicitly. Masters thesis, Carleton University, Ottawa, Canada. May 1988.

[Abu-Hakima89] Abu-Hakima S. and van Hoff A.A. An Intelligent Environment for Developing Knowledge-Based Applications. Proceedings of IJCAI '89 Workshop on A New Generation of Intelligent Interfaces. Detroit, MI. August 89.

[Abu-Hakima90] Abu-Hakima S. and Oppacher F. Improving explanations in knowledge-based systems: RATIONALE. Knowledge Acquisition journal (2). 301-343. December 1990.

[Banks88] Banks G., McLinden S., and Carlos G. Implementation of Medical Knowledge-Bases in Hypercard. Proceedings of the Twelfth Annual Symposium on Computer Applications in Medical Care. Washington,DC, November 6-9 1988. pp. 434-440

[Feiner90] Feiner S. and McKeowen K. Coordinating Text and Graphics in Explanation Generation. Proceedings of the Eighth National Conference on Artificial Intelligence, Boston, Massachusetts. July 29-August 3, 1990. pp. 442-449.

[Halasz92] Halasz M., Davidson P., Abu-Hakima S., and Phan S. JETA: A Knowledge-based Approach to Aircraft Gas Turbine Engine Maintenance. To appear in the Journal of Applied Intelligence: The International Journal of AI, Neural Networks and Complex Problem Solving Technology. Kluwer Academic Publishers. 1992.

[MBR91] Workshop Notes from the 9th National Conference on Artificial Intelligence (AAAI '91) Model-Based Reasoning. Anaheim, California. 1991.

[Neches85] Neches R., Swartout W.R., and Moore J.D. Enhanced Maintenance and Explanation of Expert Systems Through Explicit Models of their Development. IEEE Transactions on Software Engineering, Vol. SE-11, No. 11, 1337-1351 (November 1985).

[Stephenson88] Stephenson G.A. Knowledge Browsing - Front Ends to Statistical Databases. Statistical and Scientific Database Management. Fourth International Working Conference SSDBM Proceedings. June 27-30 1988. pp. 327-337.

[vanHoff91] van Hoff A.A. HyperNeWS 1.4 Manual. Turing Institute. Glasgow, Scotland. 1991

Sue Abu-Hakima *has been a researcher at the Knowledge Systems Laboratory since 1987. She is currently researching integrated diagnosis as applied to aircraft engines and is working on her Ph.D. thesis from Carleton University's Department of Systems and Computer Engineering.*

CALL FOR PAPERS

THE SECOND WORLD CONGRESS ON EXPERT SYSTEMS LISBON, PORTUGAL 10-14 January, 1994

Conference Chairman:
Professor Jay Liebowitz
Department of Management Science
George Washington University
Washington, DC 20052
Tel: 202-994-6969
Fax: 301-770-2978
jayl@aic.nrl.navy.mil

Honorary Conference Chair:
Professor Edward Feigenbaum
Stanford University

Regional Chairpersons:
Africa:
Denham Collett Grey
H.L. & H. Forest, S. Africa
Asia:
R. Sadananda
Asian Institute of Technology
Thailand
Australia:
John Debenham
Univ. of Technology, Sydney
Canada:
Ching Y. Suen, Concordia Univ.
I. Burhan Turksen, Univ. of Toronto
Caribbean:
Myron Chin, Univ. of West Indies
Eastern Europe:
Vladimir Milacic
Univ. of Belgrade, Yugoslavia
Egypt:
Ahmed Rafea
Central Lab for Agricultural ES
Europe:
Aleksandar Jovanovic
MPA Stuttgart, ***Germany***
Roy Rada, Univ. of Liverpool, ***England***
Juan Pazos Sierra
Politecnica de Madrid, ***Spain***
Gian Piero Zarri, CNRS, ***France***
Far East:
Jae Kyu Lee, KAIST, ***Korea***
Riichiro Mizoguchi, Osaka Univ., ***Japan***
Desai Narasimhalu
National Univ. of Singapore
Tu Xuyan
Beijing Univ. of Science & Tech. ***China***
Israel:
Martin Golumbic, IBM Israel
Mexico:
Francisco Cantu, ITESM
The Netherlands:
Johan den Biggelaar
Kennistechnologie, The Netherlands
Scandinavia:
Bernt Bremdal, GeoKnowledge, ***Norway***
Roar Fjellheim
Computas Expert Systems, ***Norway***
South America:
Jose Ramirez, Paradigma, ***Venezuela***
United States:
Yi-Tzuu Chien, *National Science Foundation*
Local Technical Liaison:
Joso Martins, *SISCOG, Portugal*
Technical Advisory Committee:
Ann Harrison, *JJM Systems*
Pat Harrison, *US Naval Academy*
Milton White, *IAKE*

In cooperation with:
American Association for AI
Australia SIG in AI
Canadian Society for Computational Studies of Intelligence
Caribbean AI Interest Group
European Coordinating Committee for AI
Egyptian ES Interest Group
IEEE Technical Committee on ES Applications
International Association of Knowledge Engineers
Japanese Society for AI
MPA Stuttgard
Portuguese Association for AI
Singapore Computer Systems SIG in AI
United Nations Industrial Development Organization
Venezuelan Association for AI

The second **WORLD CONGRESS ON EXPERT SYSTEMS** focuses on expert systems technology, applications and management with an international perspective. Papers from universities, industry, private and government institutions in all countries are sought in, but not limited to, the following areas:

ES Technology: knowledge acquisition; knowledge representation; developing ES shells and tools; verification, validation, testing and evaluation of ES; methods for handling uncertainty in ES; search and problem solving; neural networks; multimedia; object-oriented programming; genetic algorithms; embedded ES; case-based reasoning; fuzzy logic; learning; user interface design; integration of ES with conventional and other AI systems.

ES Applications: business and management; engineering; manufacturing; agriculture; medical; telecommunications; environmental and energy; social science; scientific; military; government; law.

ES Management: ES technology transfer; cost justification; project management; legal and societal issues.

Special themes: Linkage of ES to innovation efforts necessary for companies to survive/grow and for national economies to compete • Enabling technologies in the mid to late nineties • ES with global implications

Persons wishing to submit a paper should send five (5) copies written in English to:

The Second WORLD CONGRESS ON EXPERT SYSTEMS
Congress Secretariat, c/o Congrex (USA) Inc.
7315 Wisconsin Ave., Suite 606W
Bethesda, MD 20814
Tel: 1-301-469-3355 Fax: 1-301-469-3360

The paper should identify the area and sub-category to which it belongs (see above). Papers will be evaluated with respect to their originality, correctness, clarity and relevance. **Please limit papers to a maximum length of eight (8) pages, single-spaced.** A $50.00, non-refundable, per paper processing fee must accompany each submission. For accepted papers the fee will be applied toward registration. Papers will not be accepted by fax or any other electronic means. All submissions should include an Abstract, 200 words or less, and a listing of 5 Key Words. Please specify the author to whom all correspondence should be addressed.

Important Dates: Papers must be received by 15 March 1993. Notices of acceptance or rejection will be sent 31 May 1993. A final copy of accepted papers, camera-ready for inclusion in the Congress Proceedings, is due 30 July 1993.

The Fourth International Symposium on Artificial Intelligence

Randy Goebel and Francisco J. Cantu-Ortiz

The Fourth International Symposium on Artificial Intelligence (ISAI) was held in Cancun, Mexico, November 13-15, 1991. What? Another international AI conference you say? In Mexico? Yes.

The first ISAI was held in 1988; this fourth consecutive annual conference drew the participation of visitors from several international AI communities including the United States, Mexico, Canada, Germany, Japan, England, France, Italy, The Netherlands, Spain, China, Belgium, Australia and Singapore—an impressive breadth of participants for a conference that has existed for only four years.

ISAI was born in the summer of 1987, when Francisco Cantu-Ortiz, the director of the Centro de Inteligencia Artificial at ITESM, together with support from senior administrators at ITESM, decided that hosting an international AI conference would help build an awareness of the technological importance of AI, and would help create an opportunity for Mexican researchers and industrial developers to learn more about AI from the international community. ITESM was an appropriate birthplace for the conference—it is known as the MIT of Latin America. ITESM is a private educational institution consisting of 26 campuses, the main campus of which is at Monterrey.

The conference has not only enjoyed the generous support of the ITESM system, but is also supported by the Mexican national AI society, Sociedad Mexicana de Inteligencia Artificial (SMIA), and IBM of Mexico. Internationally, ITESM has managed to secure support from several major AI organizations, including the Canadian Society for Computational Studies of Intelligence (CSCSI), the International Joint Committee on AI (IJCAI), the American Association for AI (AAAI), and the International Association of Knowledge Engineers (IAKE).

As we all know, there seems to be some form of AI conference almost every month, so it is no strange question to wonder what this conference offers that others cannot? One advantage is size. Unlike AAAI and IJCAI conferences, the conference is small enough so that it is possible to interact with the participants in a meaningful way, rather than spending a majority of time searching for discussion opportunities or being overwhelmed by the volume of presentations. The intimate atmosphere combined with the relaxed Mexican attitude creates a situation where everyone can comfortably discuss their favourite topics without the tension that is typical of the larger, faster-paced meetings like IJCAI and AAAI.

Another advantage is location. This year, the conference moved from Monterrey, where ITESM headquarters are located, to Cancun; those interesting little impromtu discussions happen in the shade of a palm by the pool, or on the deck by the beach.

Third, though the conference adjudicates a submitted paper competition that is growing in popularity (140 papers submitted this year), the list of invited speakers has always been impressive. For example, this year, Saul Amarel, John McCarthy, Adolfo Guzman, and Ray Reiter provided the invited presentations. Past years' speakers have included Woody Bledsoe, Judea Pearl, Alan Mackworth, Raj Reddy, Mark Fox, Richard Stern, Masaru Tomita, Mitsuru Ishizuka, Robert Cartwight, Nick Cercone, David Poole, Romas Aleliunas, Gerhard Fischer, and Sarosh Talukdar.

Yet another advantage is that each year provides a focussing theme. Past themes have included Manufacturing and Engineering Design, Knowledge-based Systems, and AI applications. This year's theme was "Intelligent Informatics," and included sessions with titles like "intelligent decision support systems", "intelligent data bases", "computer networks", "software engineering", and "management information systems." And the yearly theme encourages significant industrial participation, so that there is often an opportunity to investigate the potential industrial impact of recent research ideas.

Finally, the organizers of this symposium intend it to become internationally recognized as the "AI technology transfer conference"—the place to go to find out about what's happening in terms of using AI laboratory techniques in real world applications and in various socio-economical environments. These include industrialized countries, developing countries, as well as the various economic blocks being created worldwide, such as the North American block (Canada, the USA and Mexico), the EEC, and the Pacific Rim.

Enough about the general idea of the ISAI. This year, there was plenty to stimulate the participants' thinking, including two general talks in Expert Systems, four invited plenary talks, a panel discussion on the pending tri-lateral USA/Mexico/Canada trade negotiations, and, of course, the resort diversions of Cancun.

In the first of the invited presentations, Saul Amarel confronted the audience with a challenge to consider the use of computational design problems to focus AI research. In a presentation of epic proportion, Professor Amarel introduced his thesis with strong economic and scientific motivation, and pointed out that many of the major problems in developing computer-aided design tools beg issues that have long drawn interest in AI circles. One example is the development of design specifications that automatically emphasize their most appropriate features under particular views. For example, the hydrodynamic properties of a boat hull might be the focus of one view of a boat design, while structural integrity in terms of stress analysis might be another. As other researchers working on design of digital circuits have discovered, the principled development of a history of design evolution, typically in terms of a series of design records, is a fundamental challenge in the development of general design tools. In presenting these fundamental problems, Saul Amarel emphasized the need for AI researchers to embrace all of computing science, in order to use the best ideas in the most appropriate way.

Ray Reiter's talk focused on a specification of database updates in terms of the first order situation calculus. In reviving the situation calculus for database applications, Reiter stressed the relationship between a history of transactions and the description of an evolving logic database in terms of the situation calculus' description of state change as a term denoting a nested sequence of actions. Reiter's thesis is that his reconstruction of database updates in terms of the situation calculus provides a stable framework for understanding and extending several existing ideas concerning modern database theory and logic programming. These include the notion of projection as question-answering, and the application of logic programming's notion of unfolding as a most general kind of regression planning.

John McCarthy apologized for not being sure about what "situation semantics" is, but suggested that formalizing the notion of context might provide us with a way of understanding the role of context in common sense reasoning. McCarthy sketched a formalization based on a binary "is true" predicate, which provided the basis for asserting truths relating named contexts and first order sentences. While admitting to only a partially worked out theory, McCarthy demonstrated a few of what he considered to be fundamental problems, and suggested that they seemed to be what motivated the more abstract work in situation semantics.

Adolpho Guzman rounded out the invited speakers with a presentation that explained the value of viewing interconnected software tools as black boxes. In this view, Guzman addressed the problem of how two (or more) tools can be made to interact, when the code of the tools can not be changed, but only augmented. With so many software productivity tools available today, there is a frequent need to integrate, interconnect, or force collaboration amongst two or more such tools whose design never anticipated such interaction. In addition, the tools most likely were developed by different manufacturers, with a resistance to share or publish source code. Guzman presented some revisions for existing methods for solving this problem, and proposed extensions based on design time knowledge of anticipated future connections.

In addition to these speakers, a special session on Expert Systems included invited talks from Jay Leibowitz and Robert Moore. Moore is currently the president of Gensym, a successful vendor of expert system software development tools. He provided a fast-paced introduction to the structure and successful applications of the G2 system. In what was one of the presentation highlights of the conference, Jay Leibowitz seized the opportunity provided by the conference venue to provide the audience with a "dual" talk that interwove the most recent ideas about software development methodologies for expert systems and an historical analysis of the development of ancient Mayan culture. A description would not do justice to the talk—you really did have to be there. Professor Liebowitz, from George Washington University, is the general chair of the highly successful, first ever World Congress on Expert Systems held last December in Orlando, Florida.

A lively discussion panel, chaired by Francisco J. Cantu-Ortiz, addressed the theme of AI technology transfer, and focused on issues related to the role of government policies, university programs, industry requirements, consulting firms, and international agencies in the transferring of AI. The emphasis was on the role of technology in the formation of new economic blocks such as the Free Trade Agreement being negotiated between Canada, the USA and Mexico. Panelists included Randy Goebel and Aldo Dagnino from Canada, Paul Harmon and Dick Simmons from the USA, and Antonio Sanchez and Sergio Delgado from Mexico.

Perhaps the best summary of this fourth of the ISAI meetings is to note that, unlike many of the big impersonal meetings, the thought about why one actually attended is not the first that comes to mind during the trip back home. If you understand that, perhaps we will see you in Cancun next December 7-11, 1992, for the Fifth ISAI whose theme is manufacturing and robotics. Hasta luego amigos.

Randy Goebel *is an associate professor in the Department of Computing Science at the University of Alberta in Edmonton, Canada.*

Francisco J. Cantu-Ortiz *is currently director of the Center for Artificial Intelligence and professor of computer science and artificial intelligence at the Instituto Tecnologico y Estudios Superlores de Monterrey, Mexico.*

INTERNATIONAL
JOINT CONFERENCE
ON ARTIFICIAL
INTELLIGENCE

Call for Nominations

IJCAI-93 Awards

The IJCAI Award for Research Excellence

The IJCAI Award for Research Excellence is given at an IJCAI to a scientist who has carried out a program of research of consistently high quality yielding several substantial results. If the research program has been carried out collaboratively, the Award may be made jointly to the research team. Past recipients of this award are John McCarthy (1985), Allen Newell (1989) and Marvin Minsky (1991). The Award carries with it a certificate and the sum of US$ 2,000 plus travel and living expenses for the IJCAI. The recipient will be invited to deliver an address on the nature and significance of the results achieved and write a paper for the conference proceedings. Primarily, however, the Award carries the honour of having one's work selected by one's peers as an exemplar of sustained research in Artificial Intelligence.

We hereby call for nominations for The IJCAI Award for Research Excellence, which will be presented at IJCAI-93 in Chambery, France, 29 August - 3 September 1993. The accompanying note on Selection Procedures provides the relevant details.

The Computers and Thought Award

The Computers and Thought Lectures are presented at IJCAI conferences by outstanding young scientists in the field of Artificial Intelligence. Past recipients of this honour have been Terry Winograd (1971), Patrick Winston (1973), Chuck Rieger (1975), Douglas Lenat (1977), David Marr (1979), Gerald Sussman (1981), Tom Mitchell (1983), Hector Levesque (1985), Johan de Kleer (1987), Henry Kautz (1989), Rodney Brooks (1991) and Martha Pollack (1991).

The Award carries with it a certificate and the sum of US$ 2,000 plus travel and living expenses for the IJCAI. The lecture is given one evening during the conference, and the public is invited to attend. The lecturer is encouraged to publish the lecture in the conference proceedings. The lectureship was established with royalties received from the book Computers and Thought, edited by Edward Feigenbaum and Julian Feldman; it is currently supported by income from IJCAI funds.

We hereby call for nominations for the Computers and Thought Award, which will be presented at IJCAI-93 in Chambery, France, August 29 - September 3, 1993. The accompanying note on Selection Procedures provides the relevant details.

Selection Procedures for IJCAI Awards

Nominations for the IJCAI Research Excellence Award and the Computers and Thought Award are invited from everyone in the international Artificial Intelligence community.

There should be a nominator and a seconder, at least one of whom should not be from the same institution as the nominee. Nominees for the Computers and Thought Award cannot be older than 35 at the start of the conference. They must agree to be nominated. There are no other restrictions on nominees, nominators or seconders. The nominators should prepare a short submission of less than 2,000 words, outlining the nominee's qualifications with respect to the particular award.

An IJCAI Awards Search Committee has been established to encourage high quality nominations for IJCAI Awards. It consists of the three most recent past IJCAI conference chairs, currently: Alan Bundy (convenor), Wolfgang Bibel and Barbara Grosz. Members of this committee can be contacted for advice on nominations.

The IJCAI Award Selection Committee is the union of the Board of Trustees and the former Trustees of IJCAI, the Advisory Committee of IJCAI-93, the program chairs of the last three IJCAI conferences, and the past recipients of the IJCAI Award for Research Excellence and the IJCAI Distinguished Service Award, with nominees excluded.

Nominations should be sent to the Conference Chair for IJCAI-93 at the address below. The deadline for nominations is 1 December 1992. To avoid duplication of effort, nominators are requested to submit the name of the person they are nominating by 1 November 1992 so that people who propose to nominate the same individual may be so informed and can coordinate their efforts.

Prof. Wolfgang Wahlster, Conference Chair, IJCAI-93
German Research Centre for AI (DFKI)
Stuhlsatzenhausweg 3
D-W-6600 Saarbruecken 11, Germany

Tel: (+49 681)302-5252
Fax: (+49 681)302-5341
E-mail: wahlster@dfki.uni-sb.de

CALL FOR PAPERS

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

The City Chambers, Edinburgh, June 1 - 4, 1993 IEA[93]AIE

Sponsored by:
International Society of Applied Intelligence

Organised in Cooperation with:
Artificial Intelligence Applications Institute
American Association for Artificial Intelligence
Canadian Society for Computational Studies of Intelligence
European Coordinating Committee for AI
IEEE Computer Society
International Association of Knowledge Engineers
International Neural Network Society
Japanese Society of Artificial Intelligence
Southwest Texas State University

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

- **Case-Based Reasoning**
- **Computer Aided Design/Manufacturing**
- **Distributed AI Architectures**
- **Expert & Diagnostic Systems**
- **Intelligent Databases**
- **Intelligent Interfaces**
- **Intelligent Tutoring**
- **KBS Methodologies**
- **Knowledge Acquisition**
- **Knowledge Representation**
- **Machine Learning**
- **Machine Vision**
- **Model-Based & Qualitative Reasoning**
- **Natural Language Processing**
- **Neural Networks**
- **Pattern Recognition**
- **Planning & Scheduling**
- **Practical Applications**
- **Reasoning Under Uncertainty**
- **Robotics**
- **Sensor Fusion**
- **System Dependability**

Authors are invited to submit papers for the conference. These may either be long papers (up to 10 pages, suitable for presentation in 25 minutes) or short papers (up to 4 pages, suitable for presentation in 10 minutes). Long papers should present the results of original research or innovative practical applications relevant to one or more of the listed areas of interest. Short papers need not make such a substantial contribution to the field, and in particular may report work in progress, or describe a practical application developed using well-tried techniques. Please indicate the area of submission as shown above, and submit by November 1, 1992 four copies of your paper (single spaced) to the Program Chair. Authors will be notified of the committee's decision by February 12, 1993. Final copies of papers will be due for inclusion in the conference proceedings by March 19, 1993. General information and registration materials may be obtained from the General Chair.

Dr. Paul Chung
Program Chair, IEA/AIE-93
Department of Chemical Engineering
Loughborough University of Technology
Loughborough, Leicestershire
England, UK, LE11 3TU
Email: p.w.h.chung@uk.ac.lut

Dr. Moonis Ali
General Chair, IEA/AIE-93
Department of Computer Science
Southwest Texas State University
San Marcos
TX 78666-4616, USA
Email: ma04@swtexas.earn

I am interested in attending or receiving information on The Sixth International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems, Edinburgh, June 1-4, 1993. Please put my name on the mailing list.

Name: _____
Full Postal Address: _____
Email Address: _____ Phone: _____
Please tick: I would like registration information information ☐ I would like to submit a paper ☐

Mail to the General Chairman, Dr. Moonis Ali, at the address given above.

Pitch Expert:

Canadian AI Success Story #1

Peter Turney

En tant que nouvel éditeur de* Intelligence Artificielle au Canada,** ***je me propose de faire une série d'articles sur les succès de l'intelligence artificielle au Canada. L'article suivant est le premier en lice. Chaque article présentera un système d'IA qui a connu un succès commercial. L'emphase sera mis sure les stratégies techniques, de mise en marché et de gestion qui ont fait de ces systèmes, des succès commerciaux. Si vous avez un tel récit à nous faire partager, n'hésitez pas à nous en faire part.

As a new contributing editor to* Canadian Artificial Intelligence,** ***I am starting a series of Canadian AI Success Stories. This article is the first in the series. Each article will present an AI system that has achieved commercial success. The emphasis in the articles will be on discovering the technical, managerial, and marketing strategies that made the system a success. If you have a success story, please let me know.

The following success story is largely based on discussions with Mr. Oliver Vadas of PAPRICAN. Mr. Vadas initiated and managed the Pitch Expert project. My thanks go to him for taking the time to tell me about Pitch Expert and for demonstrating the system to me.

Pitch Expert is an expert system that deals with pitch problems in kraft mills. The kraft process is a chemical method for the production of pulp from wood that employs chemicals at high temperatures to loosen the fibres. Pitch is a sticky resin found in wood. Pitch can cause problems by adhering to machinery and reducing the quality of the pulp.

Pitch Expert is the result of collaboration between the Pulp and Paper Research Institute of Canada (PAPRICAN) and the Computer Research Institute of Montreal (CRIM). PAPRICAN does applied research on problems that are relevant to the pulp and paper industry. Pulp and paper companies pay PAPRICAN for membership. The members of PAPRICAN represent about 90% of the Canadian pulp and paper industry. CRIM does computer research in a number of areas, including artificial intelligence [1].

History

In 1987, PAPRICAN approached CRIM with a proposal. The two institutes would work together to build an expert system for pitch problems. Both institutes would cover their own costs and both would have rights to the resulting software. The objective was to evaluate expert system technology and develop expertise in applying it, by the development of a practical, working expert system.

The project began in 1988. A year later, the prototype was completed. The prototype had two components, a model of a generic kraft mill and rules for solving some problems related to pitch. The production version of the system preserved most of the model, but the rules for pitch problems were significantly changed.

In 1991, Pitch Expert was ready for testing. Questionnaires were sent to thirteen kraft mills. The replies to the questionnaires were entered by an operator into Pitch Expert. The conclusions and recommendations generated by Pitch Expert were examined by a domain expert at PAPRICAN. Some of the output was edited by the domain expert and Pitch Expert was modified so that, in the future, given the same information, it would produce correct advice. The edited report was sent to the thirteen mills. The mills implemented the recommendations and observed the results. Three of the mills obtained results within the time frame of the testing exercise. As a result of the advice from Pitch Expert, these mills achieved combined savings of more than two million dollars per year. This was enough to pay for the development of Pitch Expert within a year. Note that these savings are only for the three mills that had results within the time frame of the test, but there are about 50 kraft mills in Canada.

In 1992, Pitch Expert became available to all members of PAPRICAN. The system can be used interactively through a modem. PAPRICAN worked closely with personnel at CRIM to ensure that PAPRICAN would be able to perform maintenance and minor revision of the knowledge base, after the completion of the development.

Description of the System

Pitch Expert [2,3,4] uses Inference Corporation's ART shell. Some enhancements to ART were written in LISP. The system runs on a Sun workstation. It consists of about 3,000 frames, 1,200 rules, and 150 customized LISP functions. The frames mainly encode the model of the generic

kraft mill, while the rules encode the expertise with pitch problems.

The user interacts with Pitch Expert through a text-based interface. The system asks a series of questions and the user supplies the answers. Questions are only asked when the information is actually required by the system. Most of the questions have four possible answers: "yes", "no", "later" (that is, "I will try to find the answer to that question at some later time"), and "unknown" (that is, "I don't know the answer and I probably never will know"). Some questions ask for numerical values: "What was the percentage resin in the dry deposit?" Other questions are more open: "What organic solvent was employed in the extraction?" To see what types of solvents are known to Pitch Expert, the user can type in "answers." Pitch Expert will respond with a list of acceptable answers to the question.

The goal of the system was to solve problems with pitch, but users report that this is only part of the value of the system. Pitch Expert also helps the users by teaching them. The questions focus the user's attention on important issues. An explanation facility supports the training ability of Pitch Expert. At any point in the dialogue, the user can ask for three types of explanation: "WHY did you ask that question?", "HOW can I find the answer to that question?", and "WHAT do you mean by that question?"

When enough questions have been answered, Pitch Expert produces a list of conclusions and recommendations. The conclusions are graded with subjective statements of the level of certainty and the recommendations are graded with subjective statements of the level of importance. The user can ask "why" a recommendation was made. When the provided information has elements that contradict each other, the system takes this into account when making recommendations. It also points out the presence of a conflict and recommends verifying the information.

Technical Decisions

There are several technical decisions that may partially account for the success of Pitch Expert. Perhaps the most important technical decision was the careful selection of an appropriate class of problems for applying expert system technology [5]. PAPRICAN has extensive experience with pitch problems. One PAPRICAN project, headed by Dr. Larry Allen, systematically studied problems with pitch deposits for more than ten years. Dr. Allen agreed to be the domain expert for Pitch Expert.

In 1988, ART was a good choice for an expert system shell, since it is still supported today, unlike many of the 1988 shells. A complex expert system will inevitably require some function that was not anticipated by the designers of the shell. A minimum requirement for a shell is that it should be possible to modify it. ART allows the expert system builder to incorporate new LISP routines.

Conventional wisdom says that about half of the development effort for a successful software system will be spent on the interface. Pitch Expert goes against this belief, since it uses a simple text-based interface. Since the main objective was to evaluate expert system technology, the decision was made that the project should focus on reasoning and knowledge engineering, not on the interface.

The system is currently running on a single Sun workstation at PAPRICAN. The system can easily be used interactively with a modem, since the interface is text-based. The hardware required at the mill sites is therefore simple and inexpensive. There are several advantages to a single, centrally located system. A prime advantage is the ease of maintenance. There is no need to distribute each new revision to all the users of the system. It is easy to monitor the usage of the system, to see what problems the users are having, what aspects of the system are popular, what needs revision, and what new features might be useful. The decision to include an explanation facility has met with enthusiasm from the users. Explanations help to train the users and increase their confidence in the system.

People

A small project does not require as many person-years (PY's) as a larger project. It seems plausible, however, that a commercial AI project must have a certain minimum size to be a success. If the size is too small, the project is trivial and cannot count as a note-worthy success. This minimum size implies a minimum number of PY's for success. Thus it is worthwhile to estimate the number of PY's involved in a successful project.

The Pitch Expert project involved about half a dozen people at PAPRICAN and another half a dozen at CRIM, although most of these people did not spend all of their time on this one project. PAPRICAN provided the initiative and the domain expertise, while CRIM provided the AI expertise. Now that Pitch Expert is available, most of the ongoing maintenance and improvement of the system is done by three people at PAPRICAN. (See below for a list of the people involved.)

The development of Pitch Expert spanned about four years (1988 to 1991) and required about five PY's each year, for a total of roughly 20 PY's. Ongoing operation, maintenance, improvement, and marketing will probably require around two PY's per year for the next few years. This is relatively small, given that Pitch Expert has more than 100 potential users, world-wide.

Managerial Decisions

One of the first managerial decisions was the selection of CRIM. The people at PAPRICAN recognized that they needed to go outside for AI expertise. PAPRICAN could hire new employees with expertise in AI, but one of the

objectives of the project was to evaluate AI technology. If the evaluation were negative, the new employees might have found themselves unemployed. An alternative would be to hire a consulting company, but this would have been an expensive way to evaluate a new technology. CRIM represented a relatively low-risk way to explore AI. CRIM brought extensive experience in applied AI to the project. CRIM was willing to cover its own costs, in exchange for shared rights to the technology. CRIM and PAPRICAN were well matched partners for the Pitch Expert project.

Now that Pitch Expert has shown the practical value of AI technology, PAPRICAN has established an Expert System Laboratory. Several staff members are involved full-time in the application of AI to problems in the pulp and paper industry.

Marketing Decisions

In a sense, PAPRICAN has a captive market for Pitch Expert. The companies that are members of PAPRICAN have paid for their membership and have a vested interest in applying the fruits of the research at PAPRICAN. This is perhaps an ideal situation for developing a successful AI system. However, there is still some effort in getting the companies interested and involved in the project.

To involve the companies in the project, early versions of Pitch Expert were demonstrated at annual PAPRICAN seminars, given to industry managers. The testing of the system, by thirteen kraft mills, can also be viewed as part of the marketing of the system. Even the explanation facility can help promote the system, since explanations increase the user's confidence in the system.

PAPRICAN and CRIM are jointly considering marketing Pitch Expert to the Canadian pulp and paper companies that are not members of PAPRICAN and to foreign pulp and paper companies. There has not yet been a decision on this matter.

Difficulties Encountered

It is often reported that there is a tendency to underestimate the time, money, and resources required to build an expert system. The building of Pitch Expert was not an exception. A number of difficulties were encountered that were not initially obvious and may be relevant to other expert system projects.

For example, one source of difficulty was the diverse nomenclature in the industry. As an illustration, there are about 15 major locations in a kraft mill and there are over 700 ways of talking about these locations. To cope with this diversity, Pitch Expert has tables of synonyms. Compiling these tables is, by itself, a time-consuming exercise in knowledge engineering. Expert system shells invariably need to be customized for each application. It is very difficult to

The People behind Pitch Expert

Larry Allen, PAPRICAN
Director of Research
(Domain Expert for Pitch Expert)

Diana Bouchard, PAPRICAN
Computer Scientist

Ron Crotogino, PAPRICAN
Director of Research

Roger Gauthie, CRIM
Director of Projects

Daniel Gauvin, CRIM
Knowledge Engineer

Allan Kowalski, CRIM
Senior Knowledge Engineer
(System Architect for Pitch Expert)

Monique Lefebvre, CRIM
President

Christine Lapointe, PAPRICAN
Associate Technical Specialist, Pitch Domain

Yves Larin, PAPRICAN
Knowledge Engineer

Jennifer Muise, CRIM
Knowledge Engineer (no longer with CRIM)

John Opala, CRIM
Knowledge Engineer (no longer with CRIM)

Louise Quesnel , CRIM
Vice-President Corporate Development

Jim Rogers, PAPRICAN
Senior Vice-President

Oliver Vadas, PAPRICAN
Senior Technical Specialist,
Expert Systems (Project Manager for Pitch Expert)

Peter Wrist, PAPRICAN
President

predict in advance the amount of customization that will be required. Pitch Expert required additional LISP routines to enhance the reasoning process and some C code for the link between the expert system and the user.

Warning

Mr. Vadas is now involved in the evaluation and development of other expert systems. He has found that the reusability of tools and techniques developed for the Pitch Expert is quite limited. He believes that the design of an expert system must reflect the methods that are used by the domain expert. No two experts work the same way, so no two expert systems should be designed the same way. What can be transferred to future developments, however, is the experience of the managers and the developers.

Conclusions

Many factors contribute to the success of Pitch Expert. Here is a summary of some of the points discussed above. On the technical side:

- the careful selection of a domain that was suitable for expert system technology and was an area of expertise for PAPRICAN
- the choice of a suitable shell: a shell that can be modified by the developer and that has a stable parent company
- the choice of a suitable interface: in this case, a simple text-based interface
- the decision to have a single running system, which can be accessed by modem
- the inclusion of an explanation facility

On the managerial side:

- the partnership of PAPRICAN and CRIM
- the selection of the project team

On the marketing side:

- the relation between PAPRICAN and its member companies, which helps to ensure that Pitch Expert has users
- the close involvement of the end users in the project

There is some overlap in these categories. For example, the explanation facility helps to market the system. Since the system is centrally located, it can easily be revised to accommodate new technology. This also helps to market the system.

Many of these observations have been made by others for other projects. The most surprising point to me was the decision to avoid a fancy graphical interface. In retrospect, this decision makes good sense. It means the project can focus on AI issues and it makes the system easily accessible by modem.

Acknowledgements

Thanks to Aurora Diaz of the NRC for directing me to PAPRICAN. Thanks to Jack Brahan and Mike Halasz of the NRC for their comments on this article. Thanks again to Oliver Vadas of PAPRICAN.

References

[1] J. Muise & D. Goupil, "Artificial Intelligence Research at CRIM", Canadian Artificial Intelligence, No. 18, 1989.

[2] A. Kowalski & D. Gauvin, "Pitch Expert - An Engineered Collection of Specialized Knowledge Structures", 8th IEEE Conference on Artificial Intelligence Applications, 1991.

[3] A. Kowalski & D. Gauvin, "Pitch Expert: A Productivity Improvement Tool for Pulp and Paper Production", 8th International Conference on CAD/CAM, Robotics and Factories of the Future, 1992.

[4] L.H. Allen & A. Kowalski, "Pitch Expert System - A Case Study of a Knowledge-Based System for Pitch Control in the Pulp and Paper Industry", Forest Sector 2000, The Application of Advanced Informatics to the Forest Industry, 1992.

[5] O. Vadas, J.R. Amyot, A. Kowalski, & D.C. Bouchard, "Expert Systems: Application Selection and Implementation Issues", 78th Annual Meeting Technical Session Canadian Pulp and Paper Association, 1992.

Peter Turney *is a Research Associate at the Knowledge Systems Laboratory of the Institute for Information Technology of the National Research Council of Canada. Dr. Turney's current research is in Machine Learning.*

CALL FOR PARTICIPATION
AI-ED 93
World Conference On
Artificial Intelligence in Education

23rd - 27th August 1993
Edinburgh, Scotland

The World Conference on Artificial Intelligence in Education (AI-ED93) is one of a series of international conferences designed to report the best research in the field of AI and Education and to provide opportunities for the cross-fertilization of information and ideas on research and applications in this field. The conference is sponsored by the AI-ED Society of the Association for the Advancement of Computing in Education (AACE) and hosted by the Department of Artificial Intelligence (University of Edinburgh).

You are invited to submit proposals for papers, posters, tutorials/workshops, and panels. These will be reviewed for inclusion in the technical program. Papers should describe original and unpublished results of research work. Submissions should be no longer than 8 pages, including tables, figures, and references, with at least 1 inch margins all around and no smaller than 10 point font (12 point pica for typewriters), excluding cover sheet. Papers must be received by 12th Dec 1992.

Each submission must include 1 cover page and 6 copies of the complete manuscript. The cover page should include: the title of the paper with an abstract of no more than 500 words; keywords giving a clear indication of topic and subtopic; author names with affiliations, addresses, and phone numbers; and the E-mail address of the principal author. Electronic and fax submissions will not be accepted. All submissions should be in English.

For further details, and information about industrial and book exhibits, write to: AI-ED 93, c/o Helen Pain, Department of Artificial Intelligence, University of Edinburgh, 80 South Bridge, Edinburgh EH1 1HN SCOTLAND (E-mail: AIED93@aisb.ed.ac.uk; Fax: (UK): (31) 650-6516)

Areas of interest include but are not limited to: Intelligent tutoring systems; Learning environments and microworlds; Visual and graphical interfaces; Human factors and interface design; Non-standard and innovative interfaces; Intelligent multimedia systems; Authoring systems and tutoring shells; Collaboration tools; Training job skills; Principles for instructional design; Natural language interfaces; Knowledge representation; Knowledge and skill acquisition; Conceptual change; Metacognition; Motivation; Social and cultural aspects of learning; Cognitive development and errors; Student modelling; Cognitive diagnosis; Reading and writing; Computer-assisted language learning; Teaching higher-order thinking skills; Theories of teaching; Evaluation of computer systems; Assessment of learning outcomes; Educational robotics.

Schedule and Deadlines

Proposals - Submit: 12th Dec. 1992; Acceptance: 1st March 1993. Final Papers: 15th April 1993. Registration - Early: 1st April 1993; Late: 9th May 1993. Conference: 23rd-27th August 1993

For details of the AIED Society contact AACE, P.O. Box 2966, Charlottesville, VA 22902, USA (aace@virginia.edu).

ICO'93

4th Colloque international - 4th International Conference

Communication des Connaissances dans les Organisations
Communicating Knowledge in Organizations

May 4th to May 7th, 1993 Montreal will be the host of the 4th International Conference on Cognitive and Computer Sciences (CCS) for Organizations. The Conference is organized jointly by the Interinstitutional Research Group in Cognitive and Computer Sciences for Organizations (GIRICO) and Télé-université. It is sponsored by the following research centres: CRIM, CCRIT, ATO, CEFRIO, Alberta Research Council, INRIA and Neurope Lab, in co-operation with the following associations: AAAI, IEEE, CSCSI/SCEIO, AFCET, BCS, FIQ and CIEQ.

Over 300 participants are expected to attend this important scientific event. The program features main speakers from the United States and Europe, who will give thematic talks. Participants will be invited to attend presentations, videos and demonstrations, round-table meetings on CCS research and on North-South exchanges and forums dealing with integration of traditional computing and CCS, and information retrieval.

Conference themes

The themes are grouped under two categories: Knowledge Communication Support Systems (intelligent training; task performance support; scanning, analysis and management of texts and documents; decision-making process; systems integration) and, Tools and Techniques (knowledge modelling; natural language processing; intelligent interfaces and cognitive ergonomics; document structuring and media integration; knowledge telecommunication).

Pre-Conference tutorials

Pre-Conference tutorials, offered by experts, are scheduled for May 4th on: coordination science, intelligent training, integrating CCS in organizations, interfaces, knowledge telecommunication, and information retrieval and analysis.

The Call for Papers is in progress.
Deadline: November 10, 1992

For further information, please contact:

Claude Ricciardi Rigault
Programmation ICO'93
Télé-université
1001 Sherbrooke Est
Case postale 5250, Succursale C
Montréal (Québec), Canada
H2X 3M4

Du 4 au 7 mai prochain, se tiendra à importante manifestation scientifique, le Montréal le 4e colloque international en informatique cognitive des organisations, ICO'93. Les organisateurs de cette Groupe Interuniversitaire de Recherche en Informatique Cognitive des Organisations (GIRICO) et la Télé-université attendent quelque 300 participants. ICO'93 est parrainé par les centres de recherche: CRIM, CCRIT, ATO, CEFRIO, Alberta Research Council, INRIA et Neurope Lab, en collaboration avec les associations: AAAI, IEEE, CSCSI/SCEIO, AFCET, BCS, FIQ et CIEQ.

Une dizaine d'experts de réputation internationale ont déjà accepté d'y presénter une conférence thématique. Outre les sessions réservées aux communications, démonstrations et vidéos, les congressistes pourront participer à des tables rondes sur la recherche en ICO et, sur les perspectives d'échanges Nord-Sud. Ils seront, de plus, conviés à des forums sur l' intégration de l' informatique traditionnelle dans l' organisation ainsi que sur le repérage de l' information.

La thématique

La thématique du colloque été divisée en deux secteurs: les systèmes d' aide à la communication des connaissances (formation intelligemment assistée; assistance à la performance au travail; repérage, analyse et gestion de textes et de documents; soutien a la prise de décision; intégration des systèmes) et, les outils et techniques (modélisation des connaissances; traitement du langage naturel; interfaces intelligentes et ergonomie cognitive; structuration des documents et intégration des médias; communication des connaissances à distance).

Tutoriels pré-colloque

Le colloque sera précédé, le 4 mai 1993, de tutoriels offerts par des spécialistes et qui porteront sur les thèmes suivants: formation intelligemment assistée, "coordination science", intégration de l' ICO dans l' organisation, interfaces, communication des connaissances à distance et, repérage et analyse de l' information textuelle.

L'appel de communications
se termine le 10 novembre 1992.

Pour plus d'information, prière de s'adresser à:

Tél.: (514) 522-3540
Fax: (514) 522-3608
E-mail: ICO93@teluq.uquebec.ca

Jean-Claude Gavrel

August 1992

PRECARN Celebrates 5th Year

"An experiment to test the willingness of Canadian industry to work together on long-term, precompetitive research; an experiment to test industry's abilities to develop strong and trusting interfaces with the university community; and an experiment that will establish whether there is any real substance behind the Federal Government's call for research consortia such as those that already exist in other countries." Such was the description of PRECARN Associates Inc. by its President, Gordon MacNabb, in its infancy. Now, on the occasion of its fifth anniversary, it is rewarding to look back upon the early speeches to see that the experiment is indeed succeeding.

When PRECARN was first established, Gordon MacNabb was uncertain about the level of success the organization could achieve. Industry was being asked "to invest time and effort in research, an investment that is all too rare in Canada; to invest in long-term research, an act that is even rarer or non-existent in this country; and to make that long-term investment in cooperation with other industries, including perhaps, some of their competitors ."

As PRECARN's short history will attest, Canadian industry has responded to the challenge and is actively involved in helping this country stay competitive. As of July 1992, thirty-nine Canadian companies and agencies have committed to PRECARN's cause of promoting long-term, precompetitive research in artificial intelligence and robotics.

This industry commitment to long-term collaborative research has encouraged both the federal and provincial governments of Canada to support the PRECARN experiment. The first support came from the Federal Government through the National Research Council's IRAP program, quickly followed by a $10 million commitment by the Strategic Technologies Program of Industry, Science and Technology, Canada. Provincial Governments too have bought into the experiment, first British Columbia with support for the Intelligent Graphic Interface contribution to the Robotic Vision project (ARK) led by Ontario Hydro, and more recently the Quebec Government's decision to go ahead with a support of the Telerobotic Development Systems (TDS) project led by MPB Technologies. Another testimony to the success of the experiment is the recent decision by the Federal Government to extend its support of the PRECARN program with another $3 million.

PRECARN's five research projects and two feasibility studies are also an excellent measure of its success. PRECARN's first research project, APACS (Advanced Process Analysis and Control System), led by Ontario Hydro, has reached its first major milestone in December 1991, with the creation of its first prototype. The project team is now completing the functional specifications of phase 2, and is already starting the implementation of the second prototype. The phase 2 prototype, while not an "industrial strength" system, will be able to demonstrate the full functionality of an "APACS."

PRECARN's second project, IGI (Intelligent Graphic Interface) led by MPR Teltech Ltd., now involves eleven participating organizations. This project has been successful in carrying out its process control case studies and is now undertaking the implementation of its first prototype.

The ARK (Autonomous Robot for a Known Environment) project is also led by Ontario Hydro. The team is now conducting experiments at the AECL industrial bay where they expect the robot to be able to move point-to-point in the bay while recognizing objects as its moves.

Funding Agreements with the Province of Quebec for the TDS (Telerobotic Develop Development System) research project are being signed as this article goes to press. The Lead Contractor, MPB Technologies, has, however, already assigned resources and is setting up the research facilities.

The KAD (Knowledge-Aided Design) project was approved for funding in December 1991 and is in the process of finalizing its funding. It is expected to get underway very shortly.

PRECARN is also managing two feasibility studies. The first, MAP (Mining Automation Project), headed by Inco Limited, is now at mid-point. It will build on existing developments in underground communication systems to develop applications using artificial intelligence and robotics for fully automated mining.

The second feasibility study, CORFFA (Control for Robots for Future Applications), headed by Spar Aerospace, will conduct an R&D program aimed at demonstrating improved control of flexible dynamics and contact motion.

The projects that we have briefly summarized here are at the heart of the PRECARN program. They have been very difficult to select as PRECARN did not want to duplicate

research which is either done or that could be done effectively within other structures. For us it is rewarding to see that some of our projects are already confirming that the PRECARN approach is the "right one." In an article to be published in the next PRECARN Newsletter, Connie Bryson quotes the following:

Project Manager Bruce Nickerson doesn't think ARK would have gone ahead without the PRECARN mechanism behind it.

> "This is long-term research and although it is something that industry would want, I don't think it's something we would have paid the entire bill for. We get quite a bit of leverage working through PRECARN; I think everyone benefits."

So you have a snap-shot of where PRECARN stands after five years of operation. With the continued commitment by PRECARN's member companies and both federal and provincial governments, the experiment will undoubtedly succeed. The next five years will see the completion of research projects and beginnings of others. An exciting future - and one that is crucial to the competitiveness of our country.

For more information contact: Mrs. Lise McCourt, Manager of Corporate and Public Relations, or Mr. Jean-Claude Gavrel, Director of Research Programs PRECARN Associates 300-30 Colonnade Road, Nepean, Ontario, K2E 7J6 Tel: (613) 727-9576 / Fax: (613) 727-5672
Email: Gavrel@a1.atott2.nrc.ca

CALL FOR PAPERS

INTERNATIONAL CONFERENCE on CONCEPTUAL STRUCTURES

Theory and Applications

ICCS 93

August 4 - 7 1993
Quebec City, Canada

sponsored by PARAMAX, A Unisys Company (USA)
l'Université Laval (Canada) Butterworth Heinemann Ltd. (UK)

in cooperation with

AAAI	American Association for Artificial Intelligence
ACM	Association for Computing Machinery (requested)
CEFRIO	Centre francophone de recherche en informatisation des organisations
CRIM	Le Centre de recherche informatique de Montréal
IEEE	Computer Society (requested)
GIRICO	Le Groupe de recherche en informatisation des organisations
CSCSI	The Canadian Society for Computational Studies of Intelligence

General Chairman

John F. Sowa IBM Systems Research Institute (U.S.A.)

Program Committee co - Chairmen

Guy Mineau Université Laval (Canada)
Bernard Moulin Université Laval (Canada)

Program Committee

Jerrold Aronson	SUNY at Binghamton (USA)
Nick Cercone	Simon Fraser University (Canada)
Peter Creasy	University of Queensland (Australia)
Veronica Dahl	Simon Fraser University (Canada)
Peter Eklund	Adelaide University (Australia)
Gerard Ellis	University of Queensland (Australia)
John Esch	Paramax (USA)
Jean Fargues	IBM Paris (France)
Norman Foo	University of Sydney (Australia)
Carl Frederiksen	McGill University (Canada)
Brian Gaines	University of Calgary (Canada)
Roger Hartley	New Mexico State University (USA)
Martin Janta	CMR, St Jean (Canada)
Pavel Kocura	Loughborough Univ. (England)
Debbie Leishman	Hughes (Canada)
George Lendaris	Portland State University (USA)
Robert Levinson	Univ. of Calif. at Santa Cruz (USA)
Robert Meersman	Tilburg University (The Netherlands)
Sung Myaeng	Syracuse University (USA)
Tim Nagle	UNISYS Corp. (USA)
Heather Pfeiffer	New Mexico State University (USA)
Gérard Sabah	LIMSI-CNRS (France)
Doug Skuce	University of Ottawa (Canada)
James Slagle	University of Minnesota (USA)
Eric Tsui	Continuum Australia Ltd (Australia)
Paola Velardi	University of Ancona (Italy)
Eileen Way	SUNY at Binghamton (USA)
Yelena Yesha	Univ. of Maryland Baltimore (USA)
Michael Zock	LIMSI-CNRS (France)

Proceedings

Two types of papers will be considered: long papers (up to 5000 words) to be published by Morgan Kaufmann in a book which will be available at the conference; short papers (up to 2000 words) to be included in a supplementary notebook also available at the conference.

Important Dates

author's submission	**December 1**	**1992**
notification of acceptance	**February 15**	**1992**
camera-ready final papers	**April 1**	**1993**

Two fields are devoted to the study and development of knowledge-based systems (KBS): artificial intelligence and cognitive science. Over the past 25 years, researchers have proposed several approaches for modelling knowledge in KBS, including several kinds of formalisms: semantic networks, frames , logics etc.

In the early eighties, John F. Sowa introduced the Conceptual Graph (CG) theory which provides a knowledge representation framework consisting of a form of logic with a graph notation and integrating several features from semantic net and frame representations. Since that time, several research teams over the world have been working on the application and on the extension of CG theory in various domains ranging from natural language processing to database modelling and machine learning. This international conference follows a series of seven annual workshops and aims at providing an active forum for researchers and practitioners to exchange ideas about the theory and application of conceptual graphs. It is also opened to researchers proposing alternative knowledge representation approaches, provided that they compare them to CG theory.

Subjects may include, but are not limited to the following topics.

Theory Foundations

- . Knowledge representation using CG
- . Logic systems using CG
- . Modality and truth maintenance
- . Operations on CG
- . Enhancements to CG theory
- . Reasoning using CG

Natural Language Processing (NLP)

- . Semantic representations based on CG
- . NLP systems using CG
- . Pragmatics
- . NLP systems using CG

Methods and Tools

- . Methodologies for knowledge acquisition
- . Database conceptual modelling using CG
- . Graph manipulation systems for CG
- . Learning approaches using CG
- . Expert systems and CG
- . Abstract machines for CG

Alternative Approaches

- . Comparing CG and approaches such as KL-ONE and other
- . CG compared with various logics
- . Alternative cognitive approaches for knowledge representation and manipulation

Standardization

- . Knowledge representation standardization efforts based on CG

Applications of all Kinds

Instructions to authors

Authors are invited to submit five copies of their papers, not exceeding 5000 words (long papers), or 2000 words (short papers), double spaced, written in English, including a page with: title of paper, author's name and address, phone and fax numbers, a 20 line abstract, a list of keywords.

Submitted papers should reach the program committee chairmen before December 1, 1992 at the following address :

Guy Mineau / Bernard Moulin
ICCS'93 Conference
Laval University, Computer Science Department
Pavillon Pouliot
Ste-Foy, Quebec, G1K 7P4 Canada
fax : 1 - 418 - 656 2324

BOOK REVIEWS

Probabilistic Reasoning in Expert Systems
Richard E. Neapolitan
(Northeastern Illinois University)
NY: John Wiley, 1990, xiii+433 pp
Hardbound, ISBN 0-471-61840-3,Cdn$65.50

Reviewed by
Fahiem Bacchus
University of Waterloo

Over the past 10 years, there have been tremendous strides in methods for reasoning probabilistically in expert system domains. Domains in which uncertainty is inherent, e.g., domains where cause and effect are not necessarily deterministic, can be modelled and reasoned about probabilistically by using a probability distribution to model the interrelationships in the domain. Inherent in the structure of the domain, there will typically be a large number of independencies between different variables. Modern techniques use these independencies to develop more efficient probabilistic models. This is accomplished by using the independencies to build up structured probability models reflecting those independencies. Structuring the probability models yields two immediate benefits:

(1) the structure can be used to better understand the probabilistic interrelationships and thus more easily elicit the probability model;
(2) the independencies allow a considerable improvement in computational efficiency when reasoning with the probabilistic model.

These structured probabilistic models are generally referred to as Bayesian, or causal, networks in the literature.

This new book by Neapolitan sets out to explain these recent developments to a non-specialist, and it does a credible job of it. These new techniques are important practical methods useful for builders of expert systems in commercial applications. In many domains, uncertainty is a fact of life, and older techniques, which usually involve attaching numeric grades to rules in traditional rule-based systems, are inadequate for dealing with uncertainty in a robust manner. Developing an expert system that has to deal with uncertainty using these outdated tools, which are typically the only ones available in current commercial expert system shells, is a very difficult task requiring a great deal of engineering effort to produce a functional system. The basic problem is that during computation, uncertainty propagates from cause to effect and backwards *simultaneously*, whereas forward- and backward-chaining in rule-based systems are unidirectional mechanisms.

The use of causal networks in expert systems is still mostly confined to research labs, but that situation is changing, and there are already a few commercial shells embodying these techniques on the market. It is clear that causal networks will become of increasing commercial importance as they become better known and as more supporting software is developed. In that respect, this book comes at a good time; it presents the new theory and algorithms at a level of detail that should be accessible to a reader with some undergraduate training in computer science, with a bit of effort.

The other recent book on causal networks is that of Pearl, *Probabilistic reasoning in intelligent systems*. Pearl and his colleagues have done more than anyone else to actually develop these new techniques, and he writes his book at a much deeper theoretical level. Pearl's book is essential reading for anyone who plans to do research in this area, but it is perhaps not best suited for someone who simply wants to use the techniques developed so far. For this purpose Neapolitan's book is a better choice. In fact, it is probably a good first choice even for a researcher, as a primer for Pearl's book.

An outline of the contents of Neapolitan's book follows. After an introduction, the second chapter addresses some issues in the foundations of probability. I found this chapter to be fairly weak. It does mention most of the major approaches to probability, but fails to do some of the positions justice. However, the foundations of probability are a very complex area, and I hardly expect someone to be able to do it justice in this amount of space.

Chapter 3 presents some results from graph theory that are useful in the sequel. The treatment is rather terse and might require some supplemental reading if this is the first time one is encountering graph theory. However, it is possible to understand the rest of the book without this background, if one is willing to take the graph theory results as givens.

Chapter 4 contains a useful discussion of the limitations of rule-based systems when it comes to dealing with uncertainty and the reasons why one needs a technique like causal networks.

Chapters 5, 6, and 7 contain the heart of the matter. These chapters present the basic theory and associated algorithms for causal networks. Even if one does not want to be bothered with the details of the theorems, Chapter 5 is important reading: it introduces the basic ideas that underlie causal

networks. In Chapter 6, algorithms for reasoning with singly connected networks are presented, and in Chapter 7, new techniques are developed for arbitrarily connected networks. Particularly useful in both of these chapters are detailed examples that one can use to step through the algorithms presented.

Chapter 8 presents techniques for reasoning abductively using causal networks. Abduction involves finding the best explanation for a set of evidence. In this situation, the best explanation is the set of variable instantiations that has the highest probability given the evidence. This is a hard problem for causal networks, and the algorithms presented seem impractically inefficient. This is clearly an area that needs more research, although there are probably no easy answers.

Chapter 9 presents, briefly, influence diagrams. These are causal networks in which decision nodes have been added. With these diagrams, one can compute decision-theoretic solutions to various problems, using graph-based algorithms similar to the causal network algorithms.

Chapter 10 closes the book with a discussion of the problems involved in obtaining the probabilistic information required by the causal networks. And it also discusses a method for representing one's uncertainty in this information via a second-order probability distribution, i.e., a distribution that gives the probability that a base level probability has a certain value.

In sum, the book is quite well executed. Its content is topical, and its presentation is at a level that should make it easily accessible, especially through detailed examples that help one in understanding the algorithms presented. There is a collection of problems after each of the major sections that might be useful if one is teaching a course from the book. It is to be hoped that this book and others like it will have the effect that certainty factors as a mechanism for dealing with uncertainty in expert systems will finally be laid to rest.

Fahiem Bacchus is the author of *Representing and reasoning with probabilistic knowledge* (The MIT Press, 1990).

Do the Right Thing: Studies in Limited Rationality
Stuart Russell and Eric Wefald
(University of California, Berkeley) Cambridge, MA. The MIT Press (Artificial intelligence series) 1991, xx+200 pp; hardbound, ISBN 0-262-18144-4)US$27.50

Reviewed by
Craig Boutilier
University of British Columbia

Anyone who has developed an AI program would probably claim that their program will "do the right thing" under whatever (limited) circumstances it finds itself. Indeed, that would appear to be a hallmark of intelligence. When pressed on just what "the right thing" is, the programmer will usually identify it with the program performing the appropriate external action, the one most likely to allow the program to achieve its goals. The "right thing" for a chess program is the move that maximizes its potential to win, while for a robot, it might be to turn and walk in the direction that allows it to get my coffee most quickly. Since such programs are meant to be intelligent, these actions often follow some deliberation by the program. Imagine now that my robot deliberates for several hours in an attempt to determine the "best" way to get my coffee.

Russell and Wefald maintain that the computational actions taken during the deliberation process are subject to the same test for appropriateness as external actions. Doing the right thing includes "doing the right thinking." Their book *Do the right thing* presents a theory of rational meta-reasoning and its relationship to limited rationality. Roughly, they attempt to characterize the trade-off between the cost of the complete, rational deliberation associated with discovering the best external action and the cost of executing some suboptimal action. The main thesis of this work is that the computation associated with selecting "better" actions is only worthwhile when the cost of the computation is offset by the expected gain in selecting a better action than the one currently "entertained." Hence, meta-reasoning lies at the heart of a theory of limited rationality. An agent must decide when it should reason further and when it should perform the "best" action it has been able to come up with so far.

Russell and Wefald's approach to meta-reasoning is decision-theoretic. The expected utility of the outcome of various computational actions is compared to the expected value of external actions. The theory they develop is very general (except for certain restrictions discussed below) and somewhat formal. In Chapter 1, some related work is presented, but their approach lacks a genuine predecessor. For these reasons, the theory is somewhat sketchy. (For example, it is hard to say much about the discovery of the expected value of various computations in a domain-independent manner.) However, it is given some technical teeth in Chapters 4 and 5, where they apply (and implement) the theory in the game-playing and general search domains.

The first chapter provides a very clear outline of their research strategy and the issues to be tackled in the rest of the book, although it does not provide a coherent description of the format of the book. Chapter 2 describes their view of reasoning, when it is appropriate, and how it should be accomplished. This includes a discussion of the problem of *infinite regress*; since meta-reasoning is itself reasoning, it can be subject to the same meta-level deliberations (meta-meta-reasoning) and so on. The core of Chapter 2 is a discussion of execution architectures for object-level and meta-level reasoning. Several categories of knowledge are identified: knowledge about states of the world, knowledge about the results of actions, and knowledge about the utility of world states. Classes of rules are also identified, corresponding to the "inference" of one type of knowledge

from another, and execution architectures are defined according to the types of rules on which they operate. This categorization captures (and distinguishes) a number of existing architectures, including decision-theoretic, goal-based, and production systems. The compilation of rules is discussed at length, and provides an interesting glimpse at how rules of different types can be combined and how deliberation can be partially eliminated. For instance, given certain decision-theoretic information, it may be that one action A is always best whenever a condition P holds. In such a case, all deliberation in such a situation can be replaced by a simple condition-action rule "If P then A." This, of course, is a simple production rule of the type familiar to all of us. Russell and Wefald suggest that such compilation be performed (perhaps incrementally) by the agent itself, and that discovering "approximately correct" rules may provide even larger computational gains. Finally, they discuss the extension of the various types of execution architectures and compilation techniques to the meta-level.

Chapter 3 forms the heart of the book, for it presents their view of rational meta-reasoning in its full generality. Admittedly, the theory does not provide a completely formal foundation for rational meta-reasoning, but under some simplifying assumptions allows a reasonable analysis of the trade-offs involved. The theory captures the meta-level deliberations of an agent as it decides whether to consider reasoning about its object-level decisions, or whether the expected utility of this reasoning is outweighed by the utility of immediate action. Naturally, one must model the *expected* outcome and value of the computation steps, for if the exact results were known, computation would be pointless. Three models of deliberation are presented, each refining its predecessor and more relevant to the construction of meta-level procedures.

The first is the *external model*, where an external observer determines objectively the "right thing" for an agent to do. Russell and Wefald make several assumptions to simplify the analysis. First, there should exist an objective utility function for states (or actions); that is, world states have intrinsic value. Second, the value of computation depends on the difference in utility between the default (current) action and the one suggested by the computation, and on the cost of such computation. If the gain in utility is outweighed by the cost, computation should give way to action. The final assumption is that the utility of the action suggested by further computation is equal to the utility of the state that results when that action is applied to the current state. Thus, the cost of deliberation depends only on the time it takes, not on the current state of the world. This is reasonable in the game-playing and search applications they develop later in the book; but, in general, this is a wholly unrealistic assumption. In certain states, an action must be taken immediately (for example, as a car heads toward me as I cross the street). In others the same action does not lose value as quickly. Thus, this model needs to be extended if it is to be applied to such tasks as planning and diagnosis.

The second model is the *estimated utility model*. Since an agent does not have full access to the results of deliberation at the meta-level (since the point is to decide whether to deliberate or not), a meta-level procedure must make use of *estimates* of the outcomes and utilities of various actions and computation steps. They provide an analysis of the value of computation that is meaningful for a meta-level decision procedure. Roughly, the idea is to treat "partial computations" as if they were complete (that is, the "default action" determined by a partial computation is taken to be the only possible action given the evidence brought to light by the computation), since extending partial computations completely is infeasible. A second approach, only briefly described, is to "learn" the expected utility of a partial computation by keeping track of its completions.

The final *concrete model* describes how the object-level decision procedure can influence the estimates for various outcomes and enable more accuracy in predicted utility.

Chapters 4 and 5 apply the theory to applications in game-playing and problem-solving search. The assumptions of Chapter 3 are made concrete and a thorough and very technical analysis is provided for the problems of estimating the probability distributions on the utilities associated with node expansion. Algorithms are also provided for the strategies suggested and an experimental analysis is provided comparing the results of these algorithms to the those of more standard AI techniques. Othello (though the rules of the game are never described!) and the eight- and fifteen-puzzles provide the experimental domains. Indeed, the experimental results are rather compelling and suggest that dynamic, rational meta-reasoning may be crucial in these domains.

These chapters are at their best when they describe examples of search trees and how the algorithms decide when to expand a node (deliberate) or when to make a move (act). There are also extremely brief descriptions of applications to robot path-planning and the parallel grouping problem in object recognition.

Chapter 6 describes how a program may come to learn the value of computation. As the program solves various problems, it makes observations regarding the state it is in and the error in its estimate of the value of various computation steps (for example, expanding certain nodes to certain depths). They discuss how the expected outcomes and values may be learned, but the chapter is disappointing in its brevity, and concentrates on game-playing applications.

Finally, Chapter 7 briefly describes ways in which this work may be extended and some preliminary research in these directions. These include the treatment of evaluation functions in search as "probabilistic" guides to the true value of a state, and the application of these methods to planning systems. Situated decision-making is described, and the model discussed seems to capture the scale of systems ranging from fully-reactive "planners" to the more

traditional fully-deliberative planners, according to the time-pressure under which an agent finds itself. Finally, the authors suggest methods for composing "anytime" algorithms, thus providing the ability to create new decision procedures from various components. They draw a crucial distinction between *interruptible* algorithms (that provide meaningful results when interrupted unexpectedly) and *contract* algorithms (that will provide results of increased quality given increased time, but require all of the time allotted to provide meaningful results). Obviously, the former tends to compose more easily.

To quote Russell and Wefald, "the size of the research agenda is rather scary." Indeed, this work tackles a huge problem, so it cannot be seen as the answer to the meta-reasoning and limited rationality questions. Chapters 3 through 5 provide a thorough analysis of the problem in a limited context, and while Chapters 6 and 7 are rather underdeveloped, they suggest a number of areas for future research. The book is very convincing in its suggestion that the future of traditional symbolic AI lies in this direction. Indeed, probabilistic and decision-theoretic analyses have become quite common in the traditionally symbolic domains, the recent surge of interest in decision-theoretic planning being one example. Thus, this book has made a timely appearance, and is an invaluable tool to researchers in the area. It describes an excellent research methodology and examples of its application in several contexts.

Craig Boutilier is an Assistant Professor with the Department of Computer Science at the University of British Columbia. His current research interests are in the areas of knowledge representation (particularly logics of knowledge, default reasoning and belief revision), probabilistic inference, and philosophical logic.

Programming for Artificial Intelligence: Methods, Tools, and Applications
Wolfgang Kruetzer and Bruce
McKenzie(University of Canterbury) Sydney: Addison-Wesley, 1991, xiv+682 pp (International computer science series) Hardbound, ISBN 0-201-41621-2

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

This is not a book about artificial intelligence! Does this remind you of one of the well-known openers of a French philosopher: "This is not a pipe?"

Even though many examples of AI applications are used throughout this text, the primary theme of the book is artificial intelligence as a practical activity. AI uses numerous metaphors, paradigms, tools, and styles of program development. The authors have created a foundation for the development of knowledge-based systems by proposing approaches to problem solving that introduce appropriate frames of reference and as little bias as possible.

In order to use this book, no prior knowledge of artificial intelligence or AI programming is required. A general knowledge of programming is assumed. Nonetheless, it could serve as the backbone for a course on AI programming, in combination with a methodology-oriented book. The bibliography alone is a good reference to begin the discovery of AI history and different AI programming approaches.

The book contains three divisions:

- discussion of major programming tools and environments
- reviews of a number of paradigmatic programming metaphors, and
- discussion of significant issues associated with user interfaces.

The crucial components of computer-based tools from programming are identified: languages, metaphors, and environments. Because of the related interaction among these aspects of the programming culture, the book proposes to study them together, not in isolation.

In the first division, the problems of managing intellectual complexity are discussed, which is believed by the authors to be the heart of AI programming. Prototypical tools, such as Lisp and Prolog are presented to delineate procedural from declarative programming. Object-oriented programming is introduced, which is later enriched in the third division with a demonstration of the capabilities of Smalltalk. In addition, the styles of each language are critiqued to permit the reader an opportunity to understand their subtleties.

The second division outlines the different programming metaphors used within AI programming and an accompanying historical perspective for each:

- state space,
- pattern matching and associative databases
- production systems
- associative and semantic networks
- frames and scripts, and
- natural language processing (NLP) and augmented transition networks (ATNs).

Each metaphor is explained by means of associated toolboxes and AI programming examples. Many of the examples are available on disc directly from the author. The Scheme language is singled out for particular attention.

The third division is a detailed description of what has become one of the accepted standards for AI programming interface design: Smalltalk-80. The language is put into historical context and then defined in terms of its desktop metaphor; basic syntax, data structures, operations, and

control structures; classes, properties, and methods; and graphics.

The authors conclude that "interactive and exploratory styles of program development are very effective metaphors for complex situations and tasks". They feel that an empirical approach is most useful in a situation where there is insufficient understanding of the problem; AI programming generally deals with such problems.

Michael Sutton is the author of numerous popular and academic papers, with a recent emphasis on the Electronic Forms/Formware technology. His background includes academic studies in expert systems, knowledge engineering, theology, and philosophy.

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.

The core language engine
Hiyan Alshawi (editor) (SRI International, Cambridge) Cambridge, MA: The MIT Press (The ACL-MIT Press series in natural language processing), 1992, xvi+322 pp; hardbound, ISBN 0-262-01126-3, US$37.50

Proceedings of the second European conference on computer-supported cooperative work (ECSCW '91), Amsterdam, 24-27 September 1991
Liam Bannon, Mike Robinson, and Kjeld Schmidt (editors) (University of Amsterdam and Risø) National Laboratory) Dordrecht: Kluwer Academic Publishers, 1991, xii+349 pp; paperbound, ISBN 0-7923-1439-5, US$95.00

The language of first-order logic: including the Macintosh program Tarski's World; second edition, revised and expanded *Jon Barwise and John Etchemendy* (Indiana University and Stanford University) Stanford: Center for the Study of Language and Information, 1991, xiv+297 pp and 3.5-inch Macintosh disk, paperbound, ISBN 0-937073-74-1, US$34.95 (distributed by the University of Chicago Press)

Knowledge representation
Ronald J. Brachman, Hector J. Levesque, and Raymond Reiter (editors) (AT&T Bell Laboratories and University of Toronto) First published as a special issue of *Artificial Intelligence*, 49(1-3), 1991. Cambridge, MA: The MIT Press, 1992, 408 pp; paperbound, ISBN 0-262-52168-7, US$29.00

An artificial intelligence technique for information and fact retrieval: An application in medical knowledge processing *Nicholas V. Findler* (Arizona State University) Cambridge, MA: The MIT Press (Information systems series; research reports and notes), 1991, xi+155 pp; paperbound, ISBN 0-262-56060-7,US$25.00

Text understanding in LILOG
Otthein Herzog and Claus-Rainer Rollinger (editors) (IBM Deutschland and Universität Osnabrück) Berlin: Springer-Verlag (Lecture notes in computer science 546, edited by G. Goos and J. Hartmanis, subseries Lecture notes in artificial intelligence, edited by J. Sieckmann), 1991 xi+738 pp, paperbound, ISBN 3-540-54594-8 and 0-387-54594-8, no price listed

Semantic interpretation and the resolution of ambiguity
Graeme Hirst (University of Toronto) Cambridge, England: Cambridge University Press (Studies in natural language processing), 1987, reissued in paperback 1992, xiv+263 pp; paperbound, ISBN 0-521-42898-X, US$ 22.95

An introduction to machine translation
W. John Hutchins and Harold L. Somers (University of East Anglia and University of Manchester Institute of Science and Technology) London: Academic Press, 1992, xxi+362 pp; hardbound, ISBN 0-12-362830-X, US$42.50

Text-based intelligent systems: Current research and practice in information extraction and retrieval.
Paul S. Jacobs (editor) Hillsdale, NJ: Lawrence Erlbaum Associates, 1992, viii+281 pp; Hardbound, ISBN 0-8058-1188-5, US$59.95; paperbound, ISBN 0-8058-1189-3, US$27.50

Questions and information systems
Thomas W, Lauer, Eileen Peacock, and Arthur C. Graesser (editors) (Oakland University, Oakland University, and Memphis State University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1992, vii+374 pp; hardbound, ISBN 0-8058-1018-8, US$69.95; paperbound, ISBN 0-8058-1019-6, US$34.50

***HANDEY: A robot task planner**
Tomas Lozano-Pèrez, Joseph L. Jones, Emmanuel Mazer, and Patricia A. O'Donnell Cambridge, MA: The MIT Press (Artificial intelligence series), 1992, xvii+227 pp; hardbound, ISBN 0-262-12172-7, US$39.95

Text generation and systemic-functional linguistics: Experiences from English and Japanese
Christian M.I.M. Matthiessen and John A. Bateman (University of Sydney and Gesellschaft für Mathematik und Datenverarbeitung, Darmstadt) London: Pinter Publishers (Communication in artificial intelligence series), 1991, xxii+348 pp; hardbound, ISBN 0-86187-711-X

Connectionism and psychology: A psychological perspective on new connectionist research
Philip Quinlan (University of York) Chicago: The University of Chicago Press, 1991, xvii+293 pp; hardbound, ISBN 0-226-69960-9, US$49.95; paperbound, ISBN 0-226-69961-7, US$24.95

Connectionist natural language processing: Readings from *Connection Science*
Noel Sharkey (editor) (University of Essex) Dordrecht: Kluwer Academic Publishers, 1992, ix+375 pp; hardbound, ISBN 0-7923-1542-1, US$88.50

***Fuzzy, holographic, and parallel intelligence: The sixth-generation breakthrough**
Branko Soucek and The IRIS Group New York: John Wiley and Sons, 1992, xiii+350 pp; hardbound, ISBN 0-471-54772-7

Morphology and computation
Richard Sproat (AT&T Bell Laboratories) Cambridge, MA: The MIT Press (The ACL-MIT Press series in natural language processing), 1992, xv+295 pp; hardbound, ISBN 0-262-19314-0, US$35.00

Lexical acquisition: Exploiting on-line resources to build a lexicon
Uri Zernik (editor) (Research and Development Centre, General Electric Company) Hillsdale, NJ: Lawrence Erlbaum Associates, 1991, ix+429 pp; hardbound, ISBN 0-8058-0829-9, US$69.95; paperbound, ISBN 0-8058-1127-3, US$34.50

Interested in

- **Contributing an article to *Canadian Artificial Intelligence* magazine**
- **Contributing a news item or conference review**

The next issue will be published in February 1993. Please send your contribution, electronic preferred, with an abstract, photograph and short bio by December 18, 1992 to:

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

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 205, 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.

Intelligent Micro-Robots for ALife Experiments

IS Robotics — Exploring the Future

R-2™

Compact Robot System for ALife Experiments

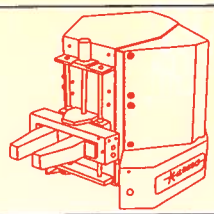

- Integrated, wheeled micro-robot
- Vertical direction parallel mechanical jaw gripper
- Ring of infrared proximity & bump sensors
- Cartesian manipulator for exploration
- Ideal for autonomous interaction, ALife, autonomous factory, cooperative construction experiments

T-1™

The Tracked Intelligent Robot Platform

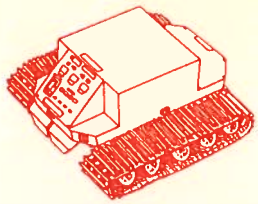

- Highly manoeuverable tracked micro-robot
- High power caterpillar tracks for rugged terrain
- Low or high gear ratio option
- Infrared transducers for collision avoidance
- Piezo-electric bump sensors
- Ideal for mixed terrain cooperative studies and prototyping autonomous Civil Engineering equipment

T-2™

The Tracked Exploration Robot System

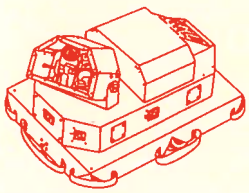

- Built on T-1 base (all features of T-1 apply)
- Color video camcorder and video transmitter (base receiving station included)
- Radio command/data link
- Optional infrared proximity sensors
- Optional vision processing module (future)
- Ideal for mixed terrain cooperative exploration
- Adapted by NASA for lunar surface exploration

T-3™

Tracked Human/Animal Interaction System

- Built on T-1 base (all features of T-1 apply)
- Pyro sensors for detecting human (animal) presence
- Light sensors to detect light intensity and gradient
- Optional microphones, speech recognition board, and digital speech output board
- Ideal for surveillance system, entertainment systems, intelligent toy development

Pebbles™

Sealed Tracked Robot for Hazardous Environments

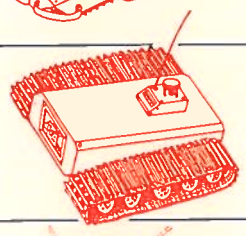

- Large payload capacity and expandability in the chassis
- Sealed to dust and can be sealed to water, oil, chemicals
- Can be hosed down and re-used if contaminated
- Black and white camera on front; video transmitter inside
- Ideal for exploration in areas unsafe for humans

Genghis-II™

The Six-Legged Walking Robot

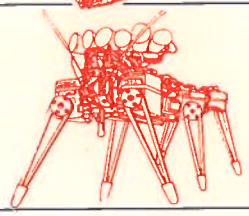

- Highly dexterous six-legged micro-robot
- Basic model detects forces, collisions, and people
- Extended package includes infrared proximity sensors, surface contact sensors, and pitch and roll inclinometers
- Expansion slots for future options
- Ideal for ALife behaviour learning and evolution studies

Attila-II™

The Legged Robot System

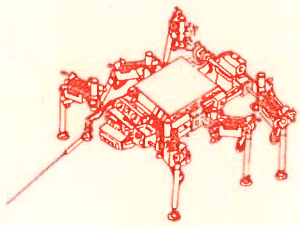

- Highly dexterous six-legged micro-robot
- Equipped with 150 high performance sensors
- Gyro stabilized CCD camera and rangefinder
- High performance multi-processor network
- Capable of determining environmental lay, texture, hardness and color
- Excels in "Find & Fetch" tasks
- Ideal for exploration, inspection and sample collection

For complete details on Intelligent Micro-Robot Technology, please contact:

AI Applied AI Machines & Software, Inc.
Suite 504, Gateway Business Park
340 March Road
KANATA, Ontario, Canada K2K 2E4
Tel: +1 613 592 7729
Fax: +1 613 592 9762

(Version 21OCT92)