Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Winter 1993 No. 31 hiver 1993

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

Expert Systems: How to Develop One that Really Works
Oliver Vadas
Systèmes Expert: Comment Enréaliser un qui Fonctionne Vraiment

Un Module Question-Reponse pour la Conception
Anne Parent
A Question-Answer Module Used for Design

Artificial Intelligence Used to Predict Sports Events
Lynn Sutherland
L'Intelligence Artificielle est Utilisée pour Prédire les Evénements Sportifs

expert system (ek-spert sis-tem) *n*: a computer program which attempts to emulate the manner of problem solving and decision taking of an expert within a restricted context.

Definition from *Artificial Intelligence Vocabulary*

Canadian Artificial Intelligence Conference

Photo credit: Greg Klymchuk

Banff Springs Hotel and surrounding area, Banff National Park, Banf, Alberta, Canada

CSCSI '94 is the tenth biennial conference on Artificial Intelligence sponsored by the Canadian Society for the Computational Study of Intelligence. It will be held in conjunction with the Vision Interface and Graphics Interface conferences.

Reserve these dates on your calendar now. A detailed announcement will appear in the June 1993 issue of *Canadian Artificial Intelligence* magazine.

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Winter 1993 No. 31 hiver 1993

Canada's National AI magazine.

Editor · Editeur
Roy Masrani
Editor Emeritus · Rédacteur emeritus
Graeme Hirst
Managing Editor · Editeur en chef
Arlene Merling
Production Team · Equipe de production
Greg Klymchuk, Arlene Merling, Carol Tubman
Academia · Académiques
Vacant
Canadian AI Success Stories
Récite de Succès d' IA au Canada
Peter Turney
PRECARN Update · Nouvelles de PRECARN
Jean-Claude Gavrel
World Watch · Vue sur le monde
Russ Thomas
Book Reviews · Critiques de livres
Graeme Hirst
Translation · Traduction
Michel Addison
Advertising · Publicité
Arlene Merling

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

Copyright © 1993, Canadian Society for Computational Studies of Intelligence. All rights reserved; *Canadian Artificial Intelligence* may not be reproduced in any form without the written permission of the editors. Printed in Canada by Tri-ad Press Ltd. *Canadian Artificial Intelligence* is published with the assistance of the Alberta Research Council. The opinions expressed herein are those of their respective authors and are not necessarily those of their employers, CSCSI, *Canadian Artificial Intelligence*, the editors, CIPS, or the Alberta Research Council.

Copyright © 1993, Société canadienne pour l'étude de l'intelligence par ordinateur. Tout droit réservé; *Intelligence artificielle au Canada* ne doit être reproduite par quelque moyen que ce soit sans le consentement écrit des éditeurs. Imprimée au Canada par Tri-ad Press Ltd. *Intelligence artificielle au Canada* est publiée avec l'aide du Conseil de Recherche de l'Alberta. Les opinions exprimées dans ce magazine sont celles de leurs auteurs respectifs et non pas nécessairement celles de leurs employeurs, de la SCEIO, de l'*Intelligence artificielle au Canada*, des éditeurs, de l'Association canadienne informatique, ou du Conseil de Recherche de l'Alberta.

CSCSI Executive

President · Président: Janice Glasgow, Dept. of Computing & Information Science, Queen's U., Kingston, ON K7L 3N6; janice@qucis.queensu.ca
Past-President · Président Précédent: Ian Witten, Head of Computer Science, U. of Calgary, Calgary, AB T2N 1N4, 403-220-6780; ian@cpsc.UCalgary.ca
Vice-President · Vice-Président: Peter Patel-Schneider, AT&T Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974, USA, 1-908-582-3399; pfps@research.att.com
Secretary · Secrétaire: Stan Matwin, Computer Science Dept., University of Ottawa, Ottawa, ON K1N 6N5; stan@ csi.u Ottawa.ca
Treasurer · Trésorier: Eric Neufeld, Dept. of Comp. Sci, Univ. of Saskatchewan, Saskatoon, SK S7N 0W0; eric@spr.usask.ca
Editor · Editeur: Roy Masrani, Alberta Research Council, 3rd Floor, 6815, 8th St. N.E., Calgary, AB T2E 7H7; 403-297-2354; masrani@skyler.arc.ab.ca

Contents

Feature Articles **Gros Titres**

Academia **Académiques**

Canadian AI Success Stories **Récite de Succès d' IA au Canada**

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

A Note from Your Production Team:

Seated from left to right: Greg Klymchuk, Carol Tubman, Arlene Merling

The Autumn 1992 issue of *Canadian Artificial Intelligence* magazine saw the passing of the baton from one Production Manager to another. At the same time, outside production costs were too high, and more control over factors such as design and quality were needed. We were able to take advantage of existing leading-edge technologies available within the Advanced Computing & Engineering department of the Alberta Research Council to facilitate in-house production. Under the continuing direction of Roy Masrani, we now have a full production team comprised of Arlene Merling, Greg Klymchuk, and Carol Tubman.

Once Roy has reviewed contributions and determined the content, Arlene and Greg do the layout and design. Arlene coordinates all facets of production, oversees the creative aspects, and Greg provides technical support. Michel Addison continues to provide the invaluable service of translating material into French, to produce a truly Canadian publication. When it's all put together, Carol reviews the material for errors of omission and commission.

The challenges associated with producing each issue vary, and these challenges allow us to grow and improve. Our efforts are continually focussed on maintaining a high standard of quality for you, our readers. *Canadian Artificial Intelligence* magazine is **your** forum for showcasing current works in the AI community, providing an opportunity for you to submit your thoughts and ideas. We welcome all contributions; share your best with us.

Arlene, Greg, & Carol

Apologies are extended to Suhayya Abu-Hakima for inadvertently leaving out the following two footnotes from her article *Visualizing and Understanding Diagnoses*, which appeared in the Autumn 1992 issue of Canadian Artificial Intelligence.

1) NRC document 33220
2) Paper is based on article published in the AAAI '92 Workshop Proceedings on Communicating Scientific and Technical Knowledge.

AI NEWS
NOUVELLES DE L'IA

News Release

ISTC Invests $2.8 Million in CRIM for Creation of a Software Engineering Services Division

MONTREAL, Quebec, January 11, 1993 - The Honourable Michael Wilson, Minister of Industry, Science and Technology and Minister for International Trade, today announced a federal investment of $2.8 million in the Centre de Recherché Informatique de Montréal Inc. (CRIM) for the creation of a division to provide software engineering services to Canadian industry. This project supports the "Investing in Growth" initiative within the Economic and Fiscal Statement introduced in the House of Commons December 2, 1992.

Noting that the project will initially create 10 engineering jobs, Mr. Wilson said, "The federal government's investment in this initiative underlines our commitment to support the partnerships that play an increasingly crucial role in our efforts to respond to the challenges and opportunities in the global economy."

The investment will assist a CRIM-led alliance of industry partners in the creation of a new division, Applied Software Engineering Centre (ASEC), to provide access to the best managerial and technical solutions, to help the Canadian software community raise its competence in software design and management. It will offer services in areas such as software engineering process (assessment, evaluation, certification), training, awareness and conferences for special interest groups.

"The world demand for increasingly complex software to enhance industry productivity is growing rapidly, and it is crucial for Canadian industry to take the steps needed to secure a competitive position in this strategic sector," said Mr. Wilson. "ASEC will bring together several Canadian high technology leaders to promote the development of Canadian software engineering capabilities that will benefit a wide range of industries, including aerospace, energy, banking and transport."

The founding members of ASEC are six Canadian companies, specializing in aerospace, microelectronics and computer software, including CAE Electronics, Canadian, Kéops Informatique, Oerlikon Aérospatiale, Systems Paramax Canada and Spar Aerospace. ASEC will develop close associations with other Canadian organizations involved in software engineering to cooperate in training and the exchange of information.

Founded in 1985, CRIM is a non-profit private organization, with more than 60 members focusing on excellence in research and development and on the transfer and application of information technologies. The centre mobilizes resources in universities, companies and other agencies, encouraging their participation in its efforts to promote and conduct research and development, assist in the dissemination and transfer of knowledge, and help train a highly-qualified work force in the computer science field.

The federal government's investment in the project is made available through Industry, Science and Technology Canada's (ISTC), Technology Outreach Program (TOP) which is committed to improving the productivity and competitiveness of Canadian industry, through the support of start-up costs and sustaining support for technology centres.

For further information contact: Alain Letendre, ISTC, (613) 954-3468

Intuitive Computing

The Japanese government will launch a project in October to develop a next-generation computer with abilities similar to a human's sense of intuition.

The 10-year, 70 billion yen ($688 million) project will be carried out by a research co-operative established in July by 11 Japanese computer and consumer electronics manufacturers and a steel industry association.

About 30 scientists from those companies will work at a new research laboratory scheduled to open Oct. 1 in Tsukuba, a city near Tokyo.

Eight foreign countries including the United States will also participate in the project.

The new computer is expected to be able to perform such jobs as global environmental assessments, analysis of the safety of nuclear devices, automatic translations over the telephone, the development of intelligent robots, and early-stage cancer diagnosis.

Asahi

EXPERT SYSTEMS:

How to develop one that really works

Oliver Vadas
Senior Technical Specialist - Expert Systems
Pulp and Paper Research Institute of Canada

Background

Starting in early 1988 I was involved in the development of a large scale expert system (ES). This particular application is a diagnostic system, helping to solve dirt deposit related quality problems in wood pulp production. It is called the Pitch Expert.

In order to make the Pitch Expert truly useful, the knowledge base became quite large. It contains about 1200 rules and 2800 schemata. In any given session the system may ask 80-100 questions. Upon the request of the user, it can explain the question, help in finding the answers and can justify the question. After each session it can provide full justification for its recommendations. These features increase user confidence and make it an effective training tool. The development took three and a half years. The system is centrally located, hence minimizing site installation costs and helping system maintenance. It is running on a work station with 32 megabytes of memory and a 1000 megabyte disk. The users may log in through modem.

After a six-month trial with the involvement of 12 users, the realized cost savings already exceeded the cost of developments. At present there are over twenty-five mills who have access to the system, representing about 55% of potential Canadian users.

For the benefit of those readers who are not familiar with paper making, I am providing an introduction to the domain.

All wood species contain a sticky material called resin. During the pulping process this material is released, and has the tendency to deposit on surfaces. It is often referred to as pitch deposit. Later it breaks off and appears as dirt in the final product. This presents a very serious problem costing Canadian mills about 100 million dollars annually. Years of research at Paprican have led to the development of a method to best deal with the problem. This method was successfully applied by a human expert.

Since this expert system is now complete, and in use by a significant number of mills, I am in the fortunate position of being able to assess the real, practical value of this technology.

Discussion

In the following, I will discuss ES technology, a subset of Artificial Intelligence (AI), and the important issues you need to deal with to be successful in developing a truly working Expert System.

SUCCESS? How do we recognise and measure success? The measure of success of any undertaking is the extent to which you meet your objectives. So you must have a clearly defined set of objectives with which to start.

Most first ES projects begin with the objective to develop a practical Expert System, evaluate the technology and gain experience in applying it. However, in many cases, by the time the project is completed (or put on hold), the objectives change to: evaluate the technology, gain experience and, if possible, develop a prototype to be useable as a demonstrator.

My first important point is that the main objective should remain the development of a working and practical ES.

As for evaluating a technology? You cannot really meet this objective if you did not develop a practical, truly working system. This applies to the remaining objective as well; if what you put together does not really work, you may have learned very little, or came to incorrect conclusions.

As you can tell by the objectives, the fact remains that the developers are looking for an opportunity to assess the technology. No one does or should deny that. To be truly successful in evaluating the technology, the end result should not be limited to a demonstration of possibilities, but be truly a working application. Otherwise it will enforce the growing view that AI is good only for developing rudimentary prototypes.

What Can We Realistically Expect of ES

What ES can be expected to deliver is not as bright as some hoped for say 10 years ago, but not as dark as many sceptics try to paint. It is somewhere in the middle. In my opinion Expert Systems represent an important area of information technology and can deliver practical benefits, if we properly select where and how to apply them.

The key to success is selecting the area of application by verifying that: (1) a method of solution exists, (2) the scope of the project is realistic in relation to available resources, and (3) the end result will be truly useful.

By definition, Expert Systems substitute human experts. But this is part of the problem. This definition covers a large area. In reality, not all the potential applications are realistic, practical, possible and justified. If the scope of the project is too small, the end result will be trivial, consequently useless. Excessively large applications, however, may prove to be beyond the limitations of available resources (too complex). So you should select one that is realistic AND practical AND possible AND justified. I do not recommend choosing an application based only on its importance. It is better to solve a lower priority problem than tackle and fail with an important one. Expert Systems are not panacea; not a solution for everything. Take the time and look at many possible applications and select the most appropriate one.

To help make the right selection there are three questions to ask:

1. Has it been done? Did someone develop and apply a method that works? Does a solution for the selected problem exist?

Expert Systems do not create solutions, only convert the existing knowledge and experience of the individual (the expert) into a computer code, making it more available to those who can benefit from it.

To make the selected application a truly realistic one, there must be someone (the domain expert) who developed a method to solve the problem of the selected domain. That method must have a proven track record. There should be no or minimal need for further development in the domain knowledge. It is also important that the method is,or at least can be structured. People think, machines do not! With the appropriate structure they may appear to.

2. Can it be done? Are the available tools and other resources sufficient to meet the needs of the application? Is it feasible?

ES development support tools are improving rapidly. However, they still have a limited range of practical use. Other resource limitations, like the availability of trained knowledge engineers and allocated funds, may determine whether what is needed can or cannot be done. In some cases the building of a prototype is necessary to determine the feasibility. If the evaluation of the prototype indicates that the project does not appear feasible, minimum objectives cannot be met. Accept this conclusion and look at other possible applications.

3. Should it be done? When the limit of what can be done is established, evaluate whether it will still be useful and practical.

The first step is basically evaluating the justification for the development. Besides conventional methods like return on investment, overall cost and non-tangible benefits, other factors also come into the picture. These may include the value of the ES as a training tool or its use as a method of documenting knowledge. As prevailing limitations lead to reducing the scope of the project, the final system may only deliver relatively simple results. Determining the targeteduser group in advance is vitally important, because that is the only way to judge whether the system results for that group are useful or trivial. In the latter case the system will be clearly useless and should not be built.

The road to success is the proper design - which is not a trivial task. A well-designed structure, and inclusion of domain and background knowledge, can make the system function almost as well as a human expert does. Continuously maintaining the knowledge base will ensure long-term utility.

Suppliers of development tools (shells) often claim designing an ES is an easy job requiring minimal experience (as long as you use their product, of course). And in a limited number of cases this may be true. For relatively simple applications a basic understanding of ES technology may be sufficient to design operational systems. However, what initially may appear to be a simple case, the nature of the application, or the "simple minded design approach", may push the value of the final product down and will deliver only trivial results.

There is very limited range for applications that are useful (not trivial) and can be properly designed by someone without formal training. For most practical applications experienced Knowledge Engineers are needed.

You may already be contemplating purchasing a development tool and assign one of your enthusiastic, young engineers to develop an Expert System. Not impossible, but ....

The objective of the development is to create a structure that will simulate and function as the Human Expert. If you want to be above the trivial range, that structure is not going to be simple. The multiple interactions and overlapping of different areas of knowledge represents a major task. The designer must know the "tricks of the trade" to successfully deal with that task.

An expert knows his/her domain, but also has a very extensive background, a sort of common sense knowledge (the real basics like; Up-Down, Light-Dark, Liquid-Solid, Gravity, and so on). Furthermore, the expert learns something new every day.

The final ES has to meet all these requirements. It must have a knowledge base that includes what the expert knows about the domain. It also should have access to the part of the background knowledge on which the expert relies. It may not learn without help, but you should be able to add to its knowledge base to keep up with changes and with new knowledge. This continued effort is often referred to as knowledge base maintenance. The Expert Systems must have the structure to accommodate these requirements. To know how to design the proper structure you probably need a KNOWLEDGE ENGINEER!

The real value is in regular use. Informing the potential users is the first step. The ease of use and the benefits will ensure regular use.

No matter how carefully the application was selected and how well the system was designed, if it is not getting into the hands of the user and is not used, it will be no more than an academic exercise.

To be put into use, it has to be sold to the users. This selling job should start even during the development cycle. Getting the end users involved at the early stages can benefit the system by including features with which the users will be most comfortable. They also will be more receptive if they feel they were involved.

Often Expert Systems are developed for the internal use of an organization. They do not need to be sold in a conventional sense. Nevertheless, users must be made aware of the system and its benefits. Documentation and user help should make learning how to operate it an easy exercise.

Summary

The Pitch Expert is just one of many successful ES applications. However, there are just as many, or possibly even more, that are never successfully completed. This should not be interpreted that ES developments have a low probability of success. Careful selection of the application domain, the design approach and the project team will significantly increase the success rate.

Probably the most valuable asset of any organization is the knowledge and experience of its people. If, through successful Expert Systems, the utilization of this valuable asset can be increased by making the knowledge of individuals the working knowledge of many, the benefits can be substantial. While potential payback justifies the risk, careful planning will reduce it. Important lessons can be learned from evaluating the successes and failures. Applying the newly-acquired understanding of the capabilities of the technology will lead to reducing the risk of failure.

Acknowledgement

For the Pitch Expert, Dr. Larry Allen provided the domain expertise. The development was a cooperative effort with CRIM (a Montreal-based research organization specializing in information technologies). The system architect was Allan Kowalski of CRIM. Several other staff members at Paprican and CRIM made important contributions to the project.

Oliver M. Vadas *received an M.Sc. from the University of Technology of Budapest, Hungary, in 1965. From 1968 to 1972 he served as Product Manager with A.E.S. Data, Montreal. He then joined the Pulp and Paper Research Institute of Canada where he studies potential use of computer technology for industrial applications. He is currently involved in evaluating and applying new technologies for the pulp and paper industry, focusing on potential applications of Artificial Intelligence and Expert Systems and is the Manager of a major scale practical Expert System project. He is a member of the CPPA (Technical Section), IEEE Computer Society, Robotics International of SME, and Machine Vision Group of R.I. He is also a member of the CPPA Electrical Engineering Committee and the R&D Committee of the Machine Vision Association of SME.*

Food For Thought

An engineer, an educator, and an AI researcher were discussing the greatest achievements of human intellect. The engineer argued that the wheel was the greatest invention of all time. The educator argued in turn for the printing press. Finally, the AI researcher argued that the thermos was man's greatest achievement. His colleagues were taken aback. "What is so great about a thermos?" asked the engineer. "Well," the AI researcher replied, "you put hot things in and they stay hot, and if you put cold things in they stay cold." "So?" queried the educator. "So?!" replied the AI researcher in obvious shock, "Well, how does it know?"

Posted in Newsgroup comp.ai by Mike Dawson, Department of Psychology, University of Alberta.
Original source unknown.

Un Module Question-Réponse pour la Conception

Anne Parent

Laboratoire des systèmes de connaissances
Institut de technologie de l'information
Conseil national de recherches,
Ottawa, Ontario
Canada K1A 0R6
e-mail: anne@ai.iit.nrc.ca

Abstract

Although several question taxonomies have been elaborated to satisfy the needs of knowledge system users, few if any lend support to the generation of advice related to design tasks. In order to respond to this need, a methodology for analysing design-oriented dialogues was put together and a taxonomy of questions and responses was defined. This work contributed to the development of the question answering facility of an Intelligent Advisor System for Conceptual Data Modeling. The paper summarizes the results of our work, discusses its limits and suggests other research avenues. This work lies within the framework of a general programme of dialogue management for design.

Résumé

Quoique plusieurs taxinomies de questions aient été développées dans le but de répondre aux besoins des utilisateurs de systèmes de connaissances, aucune d'entre elles n'offre de support à la réalisation de tâches de conceptions. Afin de répondre à ce besoin, nous avons adopté une méthodologie permettant l'analyse de dialogues reliés à la conception et développé une taxinomie de questions et de réponses. Ces travaux ont conduit à l'élaboration du module question-réponse d'un système conseiller pour la conception appliqué au domaine de la modélisation conceptuelle des données. Cet article résume les résultats que nous avons obtenus jusqu'à maintenant, en discute les limites et suggère d'autres avenues de recherche. Nos travaux s'insèrent dans le cadre du développement d'un programme plus large de gestion de dialogues pour la conception.

Introduction

Cet article résume les résultats des travaux qui ont contribué à la définition du module question-réponse d'un système conseiller pour la conception. Le système conseiller vise à assister l'utilisateur d'une méthodologie de conception complexe afin d'optimiser l'efficacité avec laquelle il ou elle effectue une tâche de création (Brahan et al., 1992). Il est conçu dans le but d'offrir trois types de service. Le système peut assister l'utilisateur sous forme de ressource passive, active et tutorielle. Sous forme de ressource passive, le conseiller répond aux interrogations qui lui sont adressées par l'intermédiaire du module question-réponse. Sous forme de ressource active, il supervise l'évolution de la tâche et interrompt l'utilisateur afin de l'avertir, le cas échéant, de l'infraction d'une contrainte de la méthode de conception. Sous forme de tuteur, il utilise une approche fondée sur l'apprentissage guidé pour enseigner les concepts et les tâches du domaine. Le système est présentement appliqué au domaine de la modélisation conceptuelle des données. Des études récentes dans ce domaine soulignent l'intérêt d'outils intelligents capables de conseiller les créateurs de modèles conceptuels, d'offrir diverses solutions et de permettre l'examen des conséquences des décisions envisagées (Tauzovich, 1989).

Afin de déterminer la nature des communications requises du système, des protocoles contenant l'interaction entre un expert du domaine de la modélisation conceptuelle et des utilisateurs cibles ont été recueillis et analysés. Le procédé expérimental adopté pour recueillir les données est une version modifiée de la technique du 'Magicien d'Oz' (Bierman & Kamsteeg,1987). Selon cette technique, un expert du domaine démontre la performance attendue du système en communiquant avec un 'client' par l'intermédiaire d'un terminal. Les informations qui ne peuvent être observées

CATEGORIE de QUESTION	FORMULE GENERIQUE	EXEMPLE
Validation-modèle-future (simple)	Est-ce que je devrais <tâche>?	Est-ce que je devrais créer l'entité COMMANDEMENT?
(comparaison)	Est-ce que je devrais <tâche> ou <tâche>?	Est-ce que je devrais assigner les cardinalités 1,n ou 0,n au lien entre la relation 'est-envoyé' et l'entité RAPPORT?
Validation-modèle-courant	Vérifie la validité de <objet>.	Vérifie la validité de la paire de cardinalités 0,n entre l'entité COMMANDEMENT et la relation 'reçoit'.
Habilitation (action)	Comment puis-je représenter le triplet <sujet><verbe><complément>?	Comment puis-je représenter le triplet <unité><envoie><rapport>?
(assignation)	Comment puis-je décider si je devrais assigner <attribut>?	Comment puis-je décider si je devrais assigner une paire de cardinalités 1,n au lien entre l'entité COMMANDEMENT et la relation 'reçoit?
(décision-action)	Comment puis-je décider si je devrais <tâche>?	Comment puis-je décider si je devrais créer une entité REPONSE?
(décision-comparaison)	Comment puis-je décider si je devrais <tâche> ou <tâche>?	Comment puis-je décider si je devrais créer un attribut ou une relation 'partie-de'?
(décision-de-temps)	A quel moment devrais-je <tâche>?	A quel moment devrais-je placer les cardinalités?
(décision-de-valeur)	Comment puis-je décider de <valeur-attribut>?	Comment puis-je décider du nom d'une entité?
Elaboration	Quelle est la définition de <concept>?	Quelle est la définition de cardinalité?
Clarification	Que veux-tu dire?	Veux-tu dire que je devrais retirer l'entité <REPONSE>?

Figure 1: Nom, formule générique et exemple de chacune des catégories de question identifiées.

par le système, issues du comportement non-verbal et de l'expression faciale de l'utilisateur, sont ainsi rendus inaccessibles à l'expert. L'analyse des données a consisté en la description du contenu et de l'organisation des échanges entre les partenaires. Ce travail a résulté en l'identification d'une taxinomie de réponse et de six catégories de questions, Validation-modèle-future, Validation-modèle-courant, Habilitation, Elaboration, Clarification, et six sous-catégories (e.g. comparaison).

L'étude de protocoles était nécessaire afin de déterminer les besoins des utilisateurs, c'est-à-dire les questions posées, et les procédés de description et d'explication pertinents à une tâche de conception, c'est-à-dire les réponses. Des études montrent que différents types de texte (e.g. une narration) suscitent différents types de questions qui nécessitent différentes formes de réponse (McKeown, 1985). Quoique plusieurs taxinomies de questions aient été développées dans le but de répondre aux besoins des utilisateurs de systèmes de connaissances (e.g. Lehnert, 1978; Hartley & Pilkington, 1987; Graesser et al., 1988), aucune d'entre elles n'offre de support à la réalisation de tâches de conception.

La Taxinomie de Questions

L'analyse des dialogues a donné lieu à l'identification de cinq catégories et six sous-catégories de questions. Les cinq catégories se nomment Validation-modèle-future, Validation-modèle-courant, Habilitation, Elaboration et Clarification. Elles sont extraites de protocoles portant sur la modélisation d'un système de gestion de rapports sur la condition insatisfaisante de matériel (en anglais, Unsatisfactory Condition Report). Ce système est utilisé par la division génie terrestre et maintenance du ministère de la Défense nationale. La figure 1 illustre chacune des catégories et sous-catégories et en présente un exemple.

Les questions de type Validation vérifient le bien-fondé d'une action passée ou future telle que la création d'une entité. Elles portent sur le modèle courant de l'utilisateur (Validation-modèle-courant) ou sur une version future de son modèle (Validation-modèle-future), avec ou sans comparaison (e.g. Est-ce que cette partie de mon modèle est valide? Est-ce que je devrais créer...? Est-ce que je devrais retirer ... ou ...?)

Les questions d'Habilitation ont pour but d'habiliter l'utilisateur à prendre une décision ou à poser un geste (e.g. Comment puis-je décider...? A quel moment devrais-je ...?) Cette catégorie comprend six sous-catégories de questions, soit action, assignation, décision-action, décision-comparaison, décision-de-temps et décision-de-valeur.

Les questions d'Elaboration ont pour but d'obtenir la définition d'un concept (e.g. Quelle est la définition de...?).

Les questions de Clarification recherchent une plus grande spécificité de la part de l'expert (e.g. Est-ce que je devrais ...?)

La Fréquence des Questions

Six sujets ont participé à cette étude. Trois d'entre eux n'avaient aucun formation théorique et peu d'expérience pratique. Un sujet avait un peu de formation et possédait un peu d'expérience et deux étaient des modélisateurs expérimentés.

Les résultats montrent que 55 des 85 questions posées, soit plus de la moitié, portaient sur le besoin de valider le modèle actuel ou envisagé. Le plus grand besoin de l'utilisateur cible du système semble être la validation. Trente-quatre questions impliquaient la validation d'une action envisagée. Vingt et une questions portaient sur la validation du modèle à l'écran. Seize questions cherchaient à habiliter l'utilisateur à effectuer une tâche, quatre recherchaient une définition et onze visaient à obtenir la clarification d'un message précédent.

Tel que supposé, les utilisateurs moins familiers avec la tâche (sujets #1, 2, 5) ont posé plus de questions que les utilisateurs plus familiers (sujet #3) et expérimentés (sujets #4, 6). Tous les sujets ont posé plus de questions de validation que toute autre catégorie de question sauf pour le sujet #3. Cet individu a posé plus de questions de type Habilitation. Ce résultat peut suggérer la présence d'un type d'utilisateur aussi intéressé à acquérir des habiletés de conception qu'à effectuer sa tâche rapidement. Nous présumons que l'utilisateur de ce système cherche avant tout à maximiser sa performance présente. Deux sujets plus avancés étaient inclus dans l'étude afin de vérifier la possibilité que des modélisateurs d'expérience tirent aussi profit (e.g. exploration) de cette ressource. Graesser & Franklin (1990) suggèrent que les utilisateurs qui ont des connaissances avancées profitent d'informations inhabituelles. Toutefois, les résultats de cette étude montrent que peu de questions sont posées par ces derniers.

Discussion

Cette étude a permis d'identifier cinq catégories de questions, Validation-modèle-future, Validation-modèle-courant, Habilitation, Elaboration, Clarification et six sous-catégories de questions reliées à la réalisation d'une tâche de conception, plus spécifiquement à la réalisation d'un modèle conceptuel. Elle a aussi résulté en l'élaboration d'une taxinomie de réponse permettant de décrire les échanges entre un expert dans une méthodologie de conception et des utilisateurs de cette méthodologie. Chaque catégorie de question est présentement associée à un schème de réponse contenant des unités d'information que nous espérons généralisable à diverses applications reliées à la conception. Le mécanisme question-réponse comprend aussi des règles qui modifient le contenu et l'organisation des réponses en fonction des connaissances particulières de l'utilisateur et de ses échanges précédents avec le système. L'approche que nous avons adoptée s'appuie sur une vision de la gestion du dialogue basée sur un plan partiel (Gilbert, 1987). Cette vision suggère que les élocutions d'un individu reposent sur une planification souple qui permet une adaptation facile aux caractères changeants des événements.

Nous présumons que le contenu et l'organisation des réponses du prototype actuel sont d'une part, pertinents à la résolution de problèmes dans le domaine de la modélisation conceptuelle des données et d'autre part, adaptés aux mécanismes d'apprentissage des utilisateurs. Le choix du contenu et de l'organisation des réponses du système s'appuient à la fois sur une étude de données empiriques et sur des connaissances théoriques contemporaines. Le contenu des réponses a en effet été déterminée à partir de protocoles de dialogue et de lectures complémentaires dans le domaine de la modélisation conceptuelle. La structure des informations a été déterminé à partir de l'ordre de présentation employé par un expert du domaine, et selon certains principes d'apprentissage humain (e.g. Piaget, 1970). La séquence des informations présentées dans une définition, par exemple, suit ce que nous présumons être l'ordre du processus de conception, soit l'identification de la classe à laquelle appartient le concept, la reconnaissance de sa fonction, puis sa forme et représentation symbolique (Brunet, 1985). Le module question-réponse tente aussi d'assurer la cohérence des informations communiquées par le système (e.g. Cawsey, 1990). Le rapport entre plusieurs réponses est établi par l'intermédiaire de références aux communications précédentes. Lorsque, par exemple, une question d'élaboration fait suite à une question de validation-modèle-courant, les règles de validation du concept sont ajoutées à la définition donnée. Au cours d'une session de travail, le système tient compte des échanges qui précèdent une question quel que soit le mode dans lequel celles-ci ont eu lieu.

Le prototype actuel ne permet toutefois pas à l'utilisateur d'obtenir la clarification des messages qui lui sont ambigus. La fréquence relativement élevée (11/85) de cette catégorie de question suggère l'importance d'offrir cette possibilité à l'utilisateur. De plus, selon Ringle & Bruce (1982), dans le cadre d'une conversation, les erreurs de compréhension sont la règle plutôt que l'exception. Afin de remédier à cette difficulté, le système pourrait modifier le vocabulaire d'une explication ou révéler de façon explicite l'intention derrière une communication. Paris (1988) suggère l'intérêt de varier à la fois la nature du contenu et le détail des informations présentées. En mode superviseur, le système communique présentement à l'utilisateur la règle qui a été transgressée mais n'offre pas de suggestions pour corriger la situation. Des questions à cet effet pourraient être ajouté au menu actuel. Par exemple, lorsque le système avertit l'utilisateur que les cardinalités qui décrivent la relation entre deux entités ne peuvent contenir les valeurs 1,1 de chaque côté de la relation, il pourrait aussi recommander, sur demande, la modification des cardinalités ou le retrait d'une entité redondante. En mode tutoriel, quoique le système tienne compte du contenu des questions précédentes, il ne peut toutefois faire référence au contenu des enseignements passés. L'efficacité du module tutoriel d'une part, et la richesse du module question-réponse d'autre part, seraient sans doute augmentées par la possibilité de mettre en relief un enseignement antérieur à l'aide de techniques telles que la comparaison et le contraste, par exemple. Enfin, nous suggérons l'intérêt de plus amples recherches de type 'Magicien d'Oz' dans le but de compléter l'inventaire de questions recueilli et d'améliorer la qualité des interventions du système en mode superviseur et en mode tutoriel.

Conclusion

Quoique la taxinomie de questions que nous proposons soit sans aucun doute incomplète, nous croyons qu'elle peut offir un point de départ utile à la satisfaction des besoins des utilisateurs du sytème conseiller.

Nous cherchons présentement des collaborateurs industriels oeuvrant dans le domaine de la modélisation conceptuelle qui participeraient à l'évaluation du prototype, ou oeuvrant dans un domaine de conception différent, participeraient au développement d'une autre application et à la commercialisation de la technologie. Pour plus d'informations, on peut communiquer avec M. J. W. Brahan au Laboratoire des systèmes de connaissances, Institut de technologie de l'information, Conseil national de recherches Canada, (613) 993-2484.

Références

Bierman, D.J. & Kansteeg, P.A. (1987). A methodology for the development of educational systems. In *Trends in European educational research.* Bramschweiger Studien Reihe, Dusseldorf.

Brahan, J.W., Farley, B., Orchard, R.A., Parent, A., Phan, C.S. (1992) A designer's consultant, *Research and Development in Expert Systems, IX Proceedings of Expert Systems 92, The Twelfth Annual Technical Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, U.K, Dec. 15-17. sous presse.

Brunet, L., Morin, S., Poulin H., Sasseville, M., & Thibaudeau, V. (1985). *Manuel du cours principe de la logique*, Presses de l'Université Laval, Québec.

Cawsey, A. (1990). Generating explanatory discourse. In *Current research in natural language generation*, London, Academic Press.

Gilbert, G.N. (1987). Advice, discourse and explanations. *The Third Alvey ExplanationWorkshop: Proceedings*, University of Surrey, Guilford, September.

Graesser, A.C., Lang, K. & Horgan, D. (1988). A taxonomy for question generation. *Questioning Exchange*, 2, 3-15.

Graesser, A.C. & Franklin, S.P. (1990). QUEST: a cognitive model of question answering, *Discourse Processes*, 13, 279-303.

Hartley, J.R. & Pilkington, R.M. (1988). Software tools for supporting learning: intelligent on-line help systems. In P. Ercoli and R. Lewis (Eds), *Artificial intelligence tools in education.* Amsterdam: Elsevier.

Lehnert, W.G. (1978). *The process of question answering*, Lawrence Erlbaum Associates, Hillsdale, New Jersey

McKeown, K. (1985). *Text generation*. Cambridge University Press, Cambridge.

Paris, C. (1988). Tailoring object descriptions to a user's level of expertise. *Computational Linguistics*, 14, 64-78.

Piaget, J. (1970). *L'épistémologie génétique*. Paris, Presses Universitaires de France.

Ringle, M. H. & Bruce, B.C. (1982). Conversation failure. In W.G. Lehnert and M.H.Ringle (Eds), *Strategies for natural language processing*, Hillsdale, N.J.: Lawrence ErlbaumAssociates.

Tauzovich, B. (1989). An expert system for conceptual data modelling. *Proceedings of the 8th International Conference on Entity-Relationship*, Saugus, Ca., USA, 329-344.

Anne Parent *détient un Doctorat en Psychologie de l'Université Laval à Québec. Elle est agent de recherches au Laboratoire des systemes de connaissances, Institut de technologie de l'information, Conseil national de recherches Canada dupuis 1987. Elle travaille en collaboration avec d'autres membres du Laboratoire, sur un projet visant le développement de la technologie logicielle fondamentale sous-jacente à la création et l'implantation de systemes conseillers. Ses recherches portent présentement sur l'analyse et la gestion de dialogues.*

Canadian
Artificial
Intelligence

Intelligence
Artificielle
au Canada

The Committed Assistant

Sometime in the not-so-distant future, you are having trouble with your new household robot. You say, “Willie, bring me a beer.” The robot replies, “OK, boss.” Twenty minutes later, you screech, “Willie, why didn't you bring that beer?” It answers, “Well, I intended to get you the beer, but I decided to do something else.” Miffed, you send the wise guy back to the manufacturer, complaining about a lack of commitment.

After retrofitting, Willie is returned, marked “Model C: The Committed Assistant.” Again, you ask Willie to bring a beer. Again, it accedes replying, “Sure thing.” Then you ask: “What kind did you buy?” It answers: “Genessee.” You say, “Never mind.” One minute later, Willie trundles over with a Genessee in its gripper. This time, you angrily return Willie for overcommitment.

After still more tinkering, the manufacturer sends Willie back, promising no more problems with its commitments. So, being a somewhat trusting customer, you accept the rascal back into your household, but as a test, you ask it to bring you your last beer. Willie again accedes, saying, “Yes, Sir.” (Its attitude problem seems to have been fixed.) The robot gets the beer and starts towards you. As it approaches, it lifts its arm, wheels around, deliberately smashes the bottle, and trundles off.

Back at the plant, when interrogated by customer service as to why it had abandoned its commitments, the robot replies that according to its specifications, it kept its commitments as long as required-commitments must be dropped when fulfilled or impossible to achieve. By smashing the last bottle, the commitment became unachievable. Despite the impeccable logic, and the correct implementation, Willie is dismantled.

Extracted from the paper “Intention Is Choice With Commitment” by P. R.Cohen and H.J. Levesque, as it was presented in ARTIFICIAL INTELLIGENCE 42 (1990). The same paper is included in the book “Intentions in Communication” published by MIT Press.

Artificial Intelligence Used to Predict Sports Events

Lynn Sutherland
Alberta Research Council

The Advanced Computing and Engineering Department of the Alberta Research Council was contracted by CBC Sports to predict the play-by-play action and the final score of the Canadian Football League 1992 Grey Cup football game. The predictions were televised during the Grey Cup pre-game special. ARC had previously been contracted by CBC to predict the 1991 National Hockey League playoffs series for use on a television special that was aired in April 1991. Both projects were used to explore the application of artificial intelligence techniques to prediction problems, to entertain, and to provide greater awareness of computing technology to the public. The projects generated a lot of public interest, as many sports fans were curious and felt personally challenged to beat the computer's predictions. This report briefly describes the techniques used and the results of the two projects. More details are available in [Suth91, Suth92].

The NHL hockey playoff predictions used traditional statistics (discriminant analysis), enhanced with the experimental use of neural networks, to predict the winning team in the eight semi-final, four division, two conference, and Stanley Cup games.

Back propagation neural networks were trained on the results of large subsets of the 840 games in the regular hockey season. The trained neural networks were tested on games in the season that had not been used in training. They were able to correctly predict winning teams 65% of the time. Upon inspection of the networks, it seemed that they learned to base the result on the win-loss-tie percentages of the two team's previous match-ups. Training with additional information, such as each team's win percentage in its previous ten games, did not improve the results. The results in this area are preliminary and should be investigated further.

The discriminant analysis of 23 common hockey variables (such as total goals for and against, power play percentage, and home/away win percentages) between teams within a division was used to select the variables that correlated to final regular season standing. The five or six most highly correlated variables, along with the win-loss-tie records between the teams, were then used to predict the outcome of the match-up of two teams in the playoffs.

Researchers Jeff White and Lynn Sutherland working on NFL prediction project.

The program correctly predicted 5 out of 8 of the semi-finalists and 2 out of 4 of the division winners. Because of the cascading nature of the playoff format, incorrect predictions early in the playoffs meant that correct predictions were impossible in later rounds. The program did not correctly predict the Stanley Cup finalists, but once it was given the finalists, correctly predicted the Stanley Cup winner. In comparison with human experts that year, most made the same errors as the computer predictions (because there were two "upsets" that year), and none performed any better than the computer.

Predicting the Grey Cup winner presented a different problem. In this case, we were asked to generate a typical play-by-play of the football game, knowing which teams were actually in the Grey Cup game. We decided to implement the program using object-oriented simulation and a simple form of case-based reasoning.

"Playbook" case histories for the quarterbacks, kickers, defences, and penalties were created from the play-by-play sheets from the CFL's regular season and semi-final games. The play-by-play sheets, which detail every individual play in a game by team, field position, downs, yards, type of play, actors in the play, and yards gained or lost on the play, were scanned into the computer using optical character recognition hardware and software. The play-by-plays were then sorted by actor, field position, down, yards to go, and plays remaining in the game, and stored in separate files for easy access.

The program runs a simulated football game by keeping track of game variables such as home team, visiting team, current offence and defence teams, score, field position, down, and yards to go. The simulation selects an offensive play from the offensive team's quarterback playbook, a defensive player for the play from the defence team's playbook, and a penalty on the play based on the offensive and defensive team's history of penalties. The plays of the game are selected by looking at what the quarterback or kicker has done in the same situation in the past, as stored in the playbook, and selecting a random play from that list of plays – no attempt at a more "fuzzy" match was implemented. This simple from of case-based lookup chooses a play that was used by the quarterback in a similar situation in the previous season.

The playbooks for the two playoff teams, Calgary and Winnipeg, were input into the simulator and 200 simulated games were run. Statistics were then determined from the simulated games. The statistics included, number of wins for each team, average score, average score differential, and score frequency distribution. The predicted winner of the Grey Cup game was the team that won the most simulated games. The specific simulated game that was chosen as a predicted play-by-play for the game was determined by the scores that were typical of the two teams, based on a frequency distribution of the results of the simulated games.

The CFL computer prediction correctly predicted the winning team, was close on many end of the game team statistics, and correctly forecast a number of specific plays during the game. The main difference between the computer prediction and the actual game was the lack of offence by Winnipeg during the game that resulted in a lower scoring game than predicted.

The two projects were performed under limited time and budget constraints. The hockey project involved three people over two weeks and the football project required three people over three weeks. Under these constraints, and the complexity of predicting a sporting event, the results were satisfactory.

The main results from these preliminary investigations are: that artificial intelligence techniques can be used to predict the outcome of sporting events when there is sufficient historic data; there is still a long way to go to harness the inherent predictive information in the data; and as in all predictive systems, even a good prediction of the most likely outcome of an event, doesn't mean that's what's going to happen.

References

[Suth91]Sutherland L., Schack B., Thomas R.; NHL Hockey Playoff Prediction 1991. Alberta Research Council OFR 91-18, 1991.

[Suth92] Sutherland L., White J.; Canadian Football League Grey Cup Predictions 1992. To be published in Alberta Research Council OFR Series.

***Lynn Sutherland** is an Associate Research Officer in the Pattern Analysis and Machine Intelligence group at the Alberta Research Council. She received a B.Sc. in Mathematics and Computer Science from Queen's University in 1982 and is a perpetual M.Sc. student at the University of Saskatchewan. Her interests include real-time knowledge-based systems and massive data analysis.*

Interested in

- **CSCSI/SCEIO Membership**

For more details, see inside back cover

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Interested in

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

Contributions for the Autumn 1993 issue of *Canadian Artificial Intelligence* magazine will be accepted until August 31, 1993. Please send your contribution, electronic preferred, with an abstract, photograph and short bio to:

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

THE GREATEST ANALOGIES IN THE HISTORY OF SCIENCE

Paul Thagard

Philosophy Department,
University of Waterloo,
Waterloo, Ontario, Canada N2L 3G1

Résume

Cet article présente une collection et l'analyse de plusieurs des plus importantes analogies utilisées lors de la découverte, conception et évaluation de théories scientifiques. Les mécanismes cognitifs utilisés dans plusieurs de ces analogies vont au delà deceux modelés dans d'actuelles théories informatique de pensée analogique, particulièrement au respect de la construction et de l'application de ces analogies.

Abstract

This paper presents a collection and analysis of many of the most important analogies that have been used for discovering, developing, and evaluating scientific theories. The cognitive mechanisms used in many of these analogies go beyond those modeled in current computational theories of analogical thinking, particularly with respect to the construction and application of the analogies.

Several years after his famous kite experiment was carried out, Benjamin Franklin responded to a query concerning how he came to propose it by quoting from a journal that he had kept at the time:

> *Nov. 7, 1749. Electrical fluid agrees with lightning in these particulars: 1. Giving light. 2. Colour of the light. 3. Crooked direction. 4. Swift motion. 5. Being conducted by metals. 6. Crack or noise in exploding. 7. Subsisting in water or ice. 8. Rending bodies it passes through. 9. Destroying animals. 10. Melting metals. 11. Firing inflammable substances. 12. Sulphureous smell. — The electric fluid is attracted by points. — We do not know whether this property is in lightning. —But since they agree in all the particulars wherein we can already compare them, is it not probable they agree likewise in this? Let the experiment be made (Franklin, 1941) p. 334).*

The kite experiment was thus inspired by the analogy Franklin noticed between lightning and electrical phenomena such as sparks.

Historians, philosophers, and psychologists of science have documented many instances of analogical thinking, but there has been little general analysis of the scientific uses of analogy and of the cognitive processes underlying those uses. This paper presents a collection of many of the most important analogies that scientists have used, and it provides an account of the main mechanisms required for analogical thinking. The various contributions of analogy to the discovery, development, and evaluation of scientific theories have involved a number of different ways of representing, constructing, and using analogies.

Great Scientific Analogies

It would be easy to compile a list of hundreds of analogies that have been used by scientists, but my concern has been to identify the analogies that qualify as most important according to two criteria. First, the analogy must have clearly contributed to some vital stage of a scientist's thinking, whether it was the discovery or development of an idea or the later argumentation in its defense. Second, the scientist's thinking that involved the analogy must have contributed to a major theoretical advance. The theory to which the analogy contributed need not be accepted today, but it must have been important in its own time and context.

Because of the possibility of additional defensible candidates, the following list is not absolute. But it suffices to provide a broad sample of highly significant analogies for analysis and generalization, presented here in chronological order. Recorded use of analogies goes back at least as far as the old testament parables (written before 1,000 B.C.), the poems of Homer (before 700 B.C.), and the fables of Aesop (around 600 B.C.). Rich analogies abound in the writings of Plato, Aristotle, and other Greek thinkers (Lloyd, 1966). But I have identified only one scientific analogy of great significance before the modern era.

1. Sound/wave. The analogy between sound and water waves was first used to suggest the nature of sound by the Greek Stoic Chrysippus around the second century B.C., but our knowledge of his views is fragmentary. A fuller use of the analogy was provided in the first century A.D. by the Roman architect Vitruvius in the course of explaining the acoustic properties of Greek amphitheatres (Vitruvius, 1960). He explicitly compared the sound of voices to water waves that can flow out and bounce back when obstructed, just as sound spreads and echos. Here we have the ancient origins of the modern wave theory of sound.

2. Earth/magnet. In his landmark work *De Magnete*, published in 1600, William Gilbert described important experimental investigations of the nature of magnets, and he proposed for the first time that the planet earth is a giant magnet (Gilbert, 1958). The basis for his hypothesis was a systematic comparison between the properties of the earth such as how it affects compasses and the properties of the small, spherical magnets on which he had performed many experiments. The earth is like these objects in many respects, so according to Gilbert we should infer that the earth is a magnet too and engenders the magnetism of the objects that were part of it.

3. Earth/moon. Galileo's *Dialogue Concerning the Two Chief World Systems*, published in 1630, contained two analogies that make important contributions to his contention that the earth moves (Galileo, 1967). First, Galileo compared the earth to the moon, which is also spherical, dark, opaque, dense, and solid with similar expanses of light and dark and of land and sea. Similarly, since the moon is known to move in an orbit, it is reasonable to suppose that the earth does too.

4. Earth/ship. Galileo used another analogy to counter an argument that the earth does not move: if a rock is dropped from a tower, it lands at the base of the tower, suggesting that the tower and hence the earth is not in motion. Galileo compared the tower to the mast of a ship that is moving and pointed out that a rock dropped from the top of the mast will fall and land at the base of the mast even though the ship is moving.

5. Light/sound. In his 1678 *Treatise on Light*, Christiaan Huygens used an analogy between light and sound in support of his wave theory of light (Huygens, 1962). That theory was eclipsed for more than a century by Newton's particle theory, but was revived in the early nineteenth century by Thomas Young and Augustin Fresnel who also exploited the analogy between light and sound to develop and defend a wave theory of light.

6. Planet/projectile. Towards the end of his celebrated *Principia* (1687), Isaac Newton used an analogy to help bring planetary motion within the scope of his theory of gravitation (Newton, 1934). He compared a planet to stones thrown upwards from the earth with greater and greater velocity. He presented a diagram to show how with greater force the path of the stone goes over into the path of an object in orbit around the earth.

7. Lightning/electricity. The beginning of this paper presented Benjamin Franklin's analogy.

8. Respiration/combustion. During the 1770s when Antoine Lavoisier was developing his oxygen theory of combustion, he also developed a theory of the role of oxygen in animal respiration. Much of his thinking was guided by an analogy between respiration and combustion, both of which involve a change of oxygen into carbon dioxide and a provision of heat (Holmes, 1989; Lavoisier, 1862).

9. Heat/water. In 1824, Sadi Carnot provided a thorough discussion of the motive power of heat, drawing heavily on an analogy between heat and waterfalls (Carnot, 1977; Gentner & Jeziorski, 1989). Heat acts on substances just as water acts on waterfalls, with the power depending in the former case on the amount of caloric (heat substance) and on the latter on the height of the waterfall. The idea of heat as a fluid was already well established by this time, but Carnot put it to much more systematic use.

10. Animals and plants/human population growth. Charles Darwin reported that he arrived at the basic idea of natural selection in 1838 by fortuitous reading of Malthus' tract on human population growth (Darwin, 1958). Darwin had been searching for a mechanism that could produce the evolution of species, and he realized from Malthus that rapid population growth in the face of limited food and land could lead to a struggle for existence. Darwin noticed the analogy between potential human strife produced by population growth outstripping resources and competition among animals and plants.

11. Natural selection/artificial selection. A different analogy played a much greater role in the development and evaluation of Darwin's theory of evolution by natural selection. He often compared natural selection to the artificial selection performed by breeders who exploited the inherent variability in animals and plants to choose desired features. Such selection leads to different breeds just as natural selection leads to different species. Darwin used this analogy in the *Origin of Species* (1859) and elsewhere, both in developing explanations and in arguing for the acceptability of his overall theory (Darwin, 1859).

12. Electromagnetic forces/continuum mechanics. James Clerk Maxwell was explicit and enthusiastic about the use of mechanical and mathematical analogies. The most important application in his own thinking was the construction in the 1860s of a diagrammed mechanical model for electrical and mechanical forces consisting of a fluid medium with vortices and stresses (Nersessian, 1992). He was able to abstract from this mechanical analog a general mathematical description that could be directly applied to electromagnetism.

13. Benzene/snake. One of the key developments in the history of biochemistry was Kekulé's discovery in 1865 of the molecular structure of benzene. According to Kekulé, he was led to the hypothesis that the carbon atoms in benzene are arranged in a ring by a reverie in which he saw a snake biting its own tail (Boden, 1990).

14. Chromosome/beaded string. In 1915, Thomas Morgan and his colleagues explained complex phenomena of inheritance by comparing chromosomes to a string containing beads corresponding to the various factors leading to inheritance (Darden, 1991). Within a few years, those factors had come to be called "genes". The beaded string analogy was most useful for describing how novel linkages could arise from crossover of chromosomes, just as new patterns of beads could arise from breaking and recombining the string.

15. Mind/computer. Numerous analogies have been used over the centuries in attempts to understand the nature of mind and thinking. By far the most fertile has been the use since the 1950s by Turing and many others of comparisons between thinking and computation (Johnson-Laird, 1988). Computational ideas have suggested hypotheses about the nature of mind that have led to much psychological and computational experimentation. As I will discuss below, this analogy is very dynamic and complex, since ideas about computation have evolved rapidly along with ideas about mind.

Notably absent from this list are two well known analogies that have often been used in teaching. Molecules of gases are often compared to billiard balls in motion, but I have not been able to find any use of this analogy by the developers of the kinetic theory of gases. Another famous analogy is the comparison of the Rutherford/Bohr model of the atom with the solar system, but the analogy does not seem to have played a role in the thinking of Rutherford or Bohr (Wilson, 1983). For an analysis of the analogy, see (Gentner, 1983).

These examples of scientific analogies vary along several important dimensions. Let us look in particular at the uses of the analogies and the cognitive mechanisms involved in their construction and application.

Uses of Scientific Analogies

Scientific analogies have at least four distinguishable uses: discovery, development, evaluation, and exposition. The most exciting is discovery, when analogy contributes to the formation of a new hypothesis. After a hypothesis has been invented, analogy may contribute to its further theoretical or experimental development. In addition, analogy can play a role in the evaluation of a hypothesis as revealed in the arguments given for or against its acceptance. Finally, analogies are often used in the exposition of science, when new ideas are conveyed to other people by comparing them with old ones. Although analogies are far more frequently used in exposition and instruction than in discovery, development, or evaluation, I shall have little more to say about educational analogies here, since my concern is with the more direct contribution of analogy to the original development of scientific ideas rather than their transmission to generations of students. We shall see that an analogy can have more than one use.

Of the uses of analogy, discovery is the hardest to document, since records are far less frequently kept of the very beginnings of hypotheses than of their development and evaluation. Nevertheless, three of the above analogies can clearly be seen as contributing to discoveries: Darwin's animals and plants/human population growth, Maxwell's electro-magnetism/mechanics, and Kekulé's benzene/snake. The cognitive mechanisms for producing the discoveries were quite different in these cases as we will see in the next section, but all the analogies played a crucial role in forming the hypotheses that were developed.

We can conjecture that several other analogies may have played a role in discovery. Chrysippus may well have been inspired to conjecture that sound moves in waves by noticing water waves and forming the sound/water wave analogy. And perhaps Franklin derived not only the idea for his experiment but also the basic hypothesis that lightning is electricity by grasping the lightning/electricity analogy. The famous story about Newton's theory of gravitation being inspired by a falling apple is not known to be false, so it is possible that the planet/projectile analogy played a role in the discovery of his theory as well as in the later argument for it. Lavoisier may have first grasped the role of oxygen in respiration by developing the respiration/combustion analogy. Perhaps further historical research will bring to light evidence that these and additional analogies were crucial for scientific discoveries.

Even if an analogy does not produce the initial formation of a hypothesis, it can aid greatly in its development. Two kinds of development are relevant: theoretical, in which a hypothesis is refined and linked with other hypotheses; and experimental, in which the empirical consequences of a hypothesis are worked out and translated into performable experiments. The light/sound analogy contributed to development of both these kinds. The diffraction properties of light were suggested in part by the ability of sound and water waves to go around corners. And the analogy suggested Young's landmark experiments in which coherent light from two pinholes was shown to exhibit interference. Finally, the analogy contributed to Fresnel's counter-intuitive but confirmed prediction that the central point in a shadow may be bright. Franklin's electricity/lightning is another clear case how analogy can serve to develop experiments. Carnot's heat/water analogy led both to new hypotheses about the properties of heat and to experimental tests of these experiments.

On the more theoretical side, Darwin stressed how useful the natural selection/artificial selection analogy was to him in constructing explanations and dealing with objections to his theory. The respiration/combustion analogy also seems to have played a strong role in the development of Lavosier's respiration theory.

Analogy is well known to be a risky form of argument, often apt to lead to false conclusions. So we might want to restrict its use to discovery and development of hypotheses, keeping evaluation pure of analogical taint. This restriction would, however, contravene the practice of several scientists of unquestioned reputation. Darwin was explicit in listing the natural selection/artificial selection analogy as one of the grounds for belief in his theory. (On the role of the analogy in Darwin's argument, see (Thagard, 1992).) Gilbert intended his earth/magnet analogy to be of more than heuristic use; it is part of his argument that the earth *is* magnetic. Similarly, when Galileo made the earth/moon and earth/ship comparisons, the analogies were in the service of his conclusion that the earth does in fact move. In addition, Newton's planet/projectile analogy is part of an argument for extending his gravitational theory to planetary phenomena. In none of these cases is the argument *purely* analogical. Darwin primarily advocated the acceptance of evolution by natural selection on the basis of its ability to unite and explain a very broad range of facts, and Galileo, Gilbert and Newton had substantial non-analogical considerations in support of their views. But analogy nevertheless played a partial role in evaluation in these important cases of scientific thinking. Similarly, psychological theories of how the mind works derive some of their non-experimental force from computational analogies.

In sum, our great scientific analogies were used roughly equally for discovery, development, and evaluation. Many seem to have been used for more than one purpose: discovery and development or development and evaluation. I have not been able to document any analogy used for all three purposes. The key question now is *how* analogical thinking can have these diverse functions.

Cognitive Mechanisms

An analogy involves two analogs, one of which is used to help explain or make inferences about the other. Following common practice in cognitive science, let us call the analog that we want to answer questions about the *target*, and call the other analog that is intended to help the *source*. I have been using the convention of identifying analogies by the pair **target/source**, as in **sound/wave** where sound is the target and water waves are the source. The fundamental questions that need to be answered to understand scientific analogies are:

1. Given the target, how can a source be found or constructed to provide what the target needs?
2. Given a possible source, how can it be applied to provide a solution to the target?

The scientific analogies show that the answers to these questions need to be somewhat more complicated than previous work on analogy in philosophy, cognitive psychology, and artificial intelligence has recognized. See, for example: (Falkenhainer, Forbus, & Gentner, 1989; Gentner, 1989; Hall, 1989; Hammond, 1989; Kolodner & Simpson, 1989; Riesbeck & Schank, 1989; Thagard, 1988; Thagard, Holyoak, Nelson, & Gochfeld, 1990; Winston, 1980).

Origins of the Source

Here is a simple story about how analogy works. In trying to solve a target problem, a scientist notices or remembers a source problem, then simply transfers over to the target the relevant aspects of the source. Remembering is a matter of retrieving a plausible source from a store of previously solved problems, and transfer involves creating a mapping between the target and source that shows how to apply insights from the source to answer the relevant questions about the target. Examination of the great scientific analogies reveals, however, the need to distinguish at least four ways in which sources can originate: noticing, retrieving, compiling, and constructing.

What I call noticing occurs when serendipity provides a source to apply to a target currently under consideration. We can imagine, for example, Chrysippus pondering the phenomena of sound, absentmindedly throwing a pebble into a pond, and being struck by the motion of the waves produced. Similarly, Franklin's lightning/electricity analogy may have been produced during contemplation of the behaviour of lightning while he happened to be creating sparks. Noticing a source may also occur when the target is not ready to mind so that the *target* has to be retrieved from memory. This is different from the simple story of problem solving we started with, in which a source is retrieved to be applied to a target. Darwin's discovery of natural selection came about because the source (human population growth) that he noticed in Malthus reminded him of the target problem on which he had long been working; the target, not the source, was retrieved from memory.

There are undoubtedly cases where a source is retrieved from memory in order to deal with a target. Because of limitations in the historical record, these are difficult to identify. But some of the more straightforward analogies probably arose when a source was remembered to help with a target. When Huygens thought about light, for example, he must often have been reminded of similar sound phenomena. Carnot's investigations of heat may well have prompted him to think of waterfalls. The work of memory, however, is not always confined to retrieving a source as a whole, ready made complex that can be applied to the target. Memory may be called upon to *compile* a complex of information that was not previously connected in any tight way. Compiling a source amounts to retrieving not a single problem but rather remembering and putting together various pieces of information. Franklin's list of electrical properties presented above was surely not retrieved as a whole but rather compiled over time as a result of reflection and recollection. Similarly, Galileo compiled a list of respects in which the earth and moon are similar. Darwin devoted much study to gathering information about artificial selection to use in his deliberations about natural selection.

The most cognitively complex origin of a source comes when it is not noticed, retrieved, or compiled, but must be constructed. Construction may involve aspects of the other three processes, but goes beyond them in the extent to which the source is different from anything that was previously known to the scientist. For example, Kekulé's unconscious thought processes did not simply produce recollection of a snake biting its tail; he may never have encountered such an image. He nevertheless generated a complex source that drew upon the target problem (benzene structure) and much biochemical knowledge stored in his memory. The new source went beyond both the target and what he knew by providing a structure - the snake rendered circular - that could inspire the hypothesis of benzene ring structure. Similarly, Maxwell could not simply use an existing mechanical system to generate answers to his problems about electromagnetism. Rather, he used his deep knowledge of the source domain to *build* a new mechanical model that could be used to generate mathematical understanding of electromagnetism. Newton's analogy involved creating the thought experiment of throwing a stone harder and harder until it went into orbit. Obviously, he had no memory of such an occurrence but rather constructed the complex projectile source analog involving projectiles to serve the ends of the target, understanding planetary motion. When Morgan and his colleagues talked about beads on a string, they were not remembering any particular string they had encountered, but rather were constructing a new source involving a special kind of organization and transformation designed to help understand the mechanism of crossing over. In cognitive science, computational ideas are not simply taken over directly to produce psychological theories, but instead involve constructing complex processing systems that are then used as analogs for thought. Thus to understand the uses of analogy in science, it is important to realize that source analogs often involve very complex designs and constructions that go beyond mechanisms implemented in current computational models of analogical thinking.

Construction of productive source analogs sometimes uses visual representations that differ from the propositional ones typically used in cognitive accounts of analogy. Source representations can involve mental pictures or paper diagrams or both. Kekulés analogy seems to have been visual: the image of the coiled snake directly suggested a similar image of the structure of benzene. The texts of published works by Newton, Maxwell, and Morgan involved diagrams representing their analogical constructions of, respectively, projectile motion, a mechanical system, and beads on a string. For recent discussions of visual representations and visual analogies, see: (Glasgow & Papadias, 1992; Thagard, Gochfeld, & Hardy, 1992; Thagard & Hardy, 1992).

Application to the Target

Strikingly, none of the uses of scientific analogy conforms very well to the familiar schema for proportional analogy,

A : B :: C : ?,

or A is to B as C is to what? Conformity would arise in cases where something needs to be filled in for the target which is known to involve C, and the relation between A and B in the source is intended to provide the clue. Perhaps the closest example is Kekulé's encoded as

snake : circle :: benzene : ?.

Putting it this way is misleading however, since it suggests that he first constructed the proportion and then filled in the answer, whereas getting an answer was part and parcel of constructing the analogical image. Similarly, Darwin did not reason:

human population growth : conflict :: animals : ?

and then fill in the struggle for existence for the question mark. Rather, what mattered about the analogy was grasping that this struggle could lead to evolution of different species. The proportion had no point outside the explanatory context of Darwin's trying to figure out a mechanism for evolution. Some scientific analogies have a structure similar to:

1. Why does T have properties A, B, C, etc.? (target)
2. S is like T in having properties P, Q, R that are like A, B, and C. (source)
3. S has P, Q, R because of X.
4. So maybe T has A, B, C because of X, which is a modification of X.

Thus we explain why sound spreads, reflects, and diffracts by noting that it is like water which spreads, reflects and diffracts through wave action. In other analogies, such as earth/moon, earth/ship, and lightning/electricity, the aim of the analogy is not so directly a matter of explanation. Galileo wanted simply to infer that the earth moves, and Franklin wanted just to suggest that lightning might be attracted by points. The aim of Maxwell's mechanical analogy was much more complex, since its main use was to enable him to construct a mathematical framework for electromagnetism. As for Galileo, the use of the analogy was indirectly explanatory, since Maxwell's general aim was to explain the electromagnetic phenomena, but his immediate goal was to work out the mathematics, just as Galileo's local goal was to construct and rebut arguments concerning the motion of the earth. The diverse purposes of the use of analogy make implausible the existence of any simple general theory of how aspects of a source can be transferred to a target.

Nevertheless, we can notice several interesting features of the mechanisms of transfer in the great scientific analogies. We saw that sources can originate through processes of varying complexity, and target completion can similarly range from virtually automatic to very complex. Inferring that the earth moves like the moon and that lightning may have points are instances of the simpler sorts of transfer. Often, however, transfer is not so straightforward as respects in which the source deviates from the target unproductively must be weighed. Vitruvius did not assume, for example, that sound waves are *just like* water waves: sound waves spread in many planes, not just one. Similarly, Fresnel and other later proponents of a wave theory of light realized to explain polarization they had to assume that light waves differed from sound waves in being transverse. Darwin had to struggle to explain how natural selection produces new species when breeders practicing artificial selection succeed only in producing new breeds.

I have been writing as if transfer is always a matter of going from source to target, which assumes that analogies are always used unidirectionally. In many cases, however, similar phenomena can be used to shed light on each other. For example, in the use of the analogy between mind and computer, accessible aspects of mind have been used to suggest new ways of doing computation, just as computation has provided new ways of understanding thinking. Planets can help us think about projectiles just as projectiles can help us think about planets, and we can learn something about electricity by studying lightning. In these cases, we have the achievement or the prospect of achievement of a unifying theory that specifies why two phenomena once thought to be disparate are fundamentally similar. The best example is the unification of projectiles and planets by Newtonian mechanics. Perhaps someday we will have a unified theory of information processing systems that will establish the depth of the analogy between mind and computer. Conjecturally, such a theory will only arise through a process of co-evolution of theories of mind and computation, with analogies among mind, brain, and computers going in various possible directions depending on the state of advancement of knowledge of each. Many scientific analogies, however, are not bidirectional. We do not use benzene to help understand snakes, or chromosomes to help understand beads, or atomic structure to help understand the solar system.

I hope my collection and analysis of great scientific analogies has been useful for several reasons. First, we find analogy playing an important role in scientific development in many epochs and fields. Second, analogy contributes to various stages of science, from discovery to evaluation, and serves various explanatory and inferential ends. Finally, it is clear that scientific analogies have required complex representations and processes whose understanding will require new projects in cognitive science.

Acknowledgments

This research is supported by the Natural Sciences and Engineering Research Council of Canada. Thanks to Lindley Darden, Nancy Nersessian, and Gregory Nowak for numerous suggestions

References

Boden, M. (1990). *The creative mind: Myths and mechanisms.* New York: Basic Books.

Carnot, S. (1977). *Reflections on the motive power of fire.* (Thurston, R., Trans.). Gloucester, MA: Peter Smith.

Darden, L. (1991). *Theory change in science: Strategies from Mendelian genetics.* Oxford: Oxford University Press.

Darwin, C. (1859). *The origin of species.* London: Murray.

Darwin, C. (1958). *The autobiography of Charles Darwin and selected letters.* New York: Dover.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence, 41*, 1-63.

Franklin, B. (1941). *Benjamin Franklin's experiments.* Cambridge, MA: Harvard University Press.

Galileo (1967). *Dialogue concerning the two chief world systems* (Drake, S., Trans.). Berkeley: University of California Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, *7*, 155-170.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 199-241). Cambridge: Cambridge University Press.

Gentner, D., & Jeziorski, M. (1989). Historical shifts in the use of analogy in science. In B. Gholson, W. Shadish, R. Beimeyer, & A. Houts (Eds.), *The psychology of science: Contributions to metascience* (pp. 296-325). Cambridge: Cambridge University Press.

Gilbert, W. (1958). *De Magnete* (P. Fleury Mottelay, Trans.). New York: Dover.

Glasgow, J., & Papadias, D. (1992). Computational imagery. *Cognitive Science*, 355-394.

Hall, R. (1989). Computational approaches to analogical reasoning: A comparative analysis. *Artificial Intelligence* , *39*, 39-120.

Hammond, K. (1989). *Case-based planning: Viewing planning as a memory task.* New York: Academic Press.

Holmes, F. (1989). Antoine Lavoisier and Hans Krebs: Two styles of scientific creativity. In D. Wallace & H. Gruber (Eds.), *Creative people at work* (pp. 44-68). Oxford: Oxford University Press.

Huygens, C. (1962). *Treatise on light* (S. Thompson, Trans.). New York: Dover.

Johnson-Laird, P. (1988). *The computer and the mind.* Cambridge, MA: Harvard University Press.

Kolodner, J., & Simpson, R. (1989). The MEDIATOR: Analysis of an early case-based problem solver. *Cognitive Science, 13*, 507-549.

Lavoisier, A. (1862). *Oeuvres.* Paris: Imprimerie Impériale.

Lloyd, G. E. R. (1966). *Polarity and analogy: Two types of argumentation in early Greek thought.* Cambridge: Cambridge University Press.

Nersessian, N. (1992). How do scientists think? Capturing the dynamics of conceptual change in science. In R. Giere (Eds.), *Cognitive Models of Science.* (pp. 3-44). Minneapolis: University of Minnesota Press.

Newton, I. (1934). *Mathematical principles of natural philosophy* (A. Motte and F. Cojou,Trans.). Berkeley: University of California Press.

Riesbeck, C., & Schank, R. (1989). *Inside case-based reasoning*. Hillsdale, NJ: Erlbaum.

Thagard, P. (1988). Dimensions of analogy. In D. Helman (Eds.), *Analogical reasoning.* (pp. 105-124). Dordrecht: Kluwer.

Thagard, P. (1992). *Conceptual revolutions.* Princeton: Princeton University Press.

Thagard, P., Gochfeld, D., & Hardy, S. (1992). Visual analogical mapping. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 522-527). Hillsdale, NJ: Erlbaum.

Thagard, P., & Hardy, S. (1992). Visual thinking in the development of Dalton's atomic theory. In *Proceedings of the Ninth Canadian Conference on Artificial Intelligence,* (pp. 30-37). Vancouver.

Thagard, P., Holyoak, K., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence*, 46, 259-310. Vitruvius (1960). *The ten books on architecture* (M.H. Morgan, Trans.). New York:Dover.

Wilson, D. (1983). *Rutherford, simple genius*. Cambridge, MA: MIT Press.

Winston, P.H. (1980). Learning and reasoning by analogy, *Communications of the ACM, 23*, 689-703.

***Paul Thagard** is Professor of Philosophy and Adjunct Professor of Psychology and Computer Science at the University of Waterloo. He is a graduate of the Universities of Saskatchewan, Cambridge, Toronto (Ph. D. in philosophy, 1977) and Michigan (M.S. in computer science, 1985). He is the author of* Conceptual Revolutions *(Princeton University Press, 1992),* Computational Philosophy of Science *(MIT Press, 1988), and co-author of* Induction: Processes of Inference, Learning, and Discovery *(MIT Press, 1986).*

The Fourth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems

André Trudel

The Fourth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'92) was held in Mallorca, Spain, July 6 - 10, 1992. The previous three IPMU conferences were held in Paris (1986 and 1990) and Urbino, Italy (1988).

The major issues dealt with at the conference were the acquisition, representation, management and transmission of uncertain knowledge. The majority of the papers presented were in the following areas: belief networks, dynamic systems, default reasoning, evidence theory, neural networks, image processing, probabilistic methods, databases, logical operators, decision making, fuzzy logic, reasoning under uncertainty, information measures, learning methods, questionnaires, clustering and classification, and measures in social and behavioural sciences. This is only a partial list of the areas. With 175 papers being presented in such a wide variety of areas, it was easy to find interesting talks to attend.

One of the areas with the largest number of papers and attendance at presentations was fuzzy logic. There were 6 sessions dedicated to fuzzy topics. There were also papers presented in other sessions which used fuzzy logic. Lofti Zadeh was present to give an invited lecture titled "The solution of equations in interpolative reasoning" and receive the *Kampé de Fériet Lecturer Award.* The majority of conference attendees were from Europe. Only a small number of attendees were Canadians. There are only 8 Canadian authors listed in the proceedings. Two of these authors are mathematicians and one is a psychologist. This is not a Canadian anomaly. There were many non-AI researchers at the conference. This is a good indication that AI researchers are not working on uncertainty problems in isolation. It was also nice to see how people in other disciplines approach AI related problems.

The proceedings are not publicly available, but some of the papers are being considered for a special journal issue. If you would like a photocopy of the index of papers or a particular paper, send email to Trudel@AcadiaU.ca.

Since I am writing this review a few months after the conference, the details are vague. One I thing I do remember vividly is the conference location. Mallorca is a favorite European vacation island in the Mediterranean. The conference site was a seaside resort with two pools and unbelievable buffet meals. The resort staff was very friendly and most spoke 5 languages. The banquet was held in the open courtyard of a circular medieval castle which overlooked the capital of Mallorca. The Spanish certainly know how to pick conference sites!

***André Trudel** is an assistant professor in the Jodrey School of Computer Science at Acadia University in Wolfville, Nova Scotia.*

VARMINT:

Canadian AI Success Story #2

Peter Turney

Ceci est le second article dans une sèrie sur les rècits de succès en Intelligence Artificielle Canadienne. Chaque rècit présente un système d'IA qui a atteint le succès commercial. Le but de ces articles de cette série est de découvrir les stratégies techniiques, de gestion et de mise en marché qui ont permis à ces systèmes d'IA d'atteindre le succès au Canada. Si vous avez de tels récits, faites le moi savoir.

This is the second in a series of Canadian AI Success Stories. Each story presents an AI system that has achieved commercial success. The goal of the articles in this series is to discover the technical, managerial, and marketing strategies that make an AI system successful in Canada. If you have a success story, please let me know.

VARMINT (Analyse des Vibrations pour Composantes Internes de Machinerie Rotative) est un système expert qui prédit les pannes mécaniques d'équipement, pour que l'entretien préventif soit effectué orsque requis. Ce système fut conçu pour etre utilisé sur des bateaux où il analyse les données de vibrations mécaniques de centaines de machines, incluant pompes, compresseurs, éventails, trèuils et moteurs. Une des fonctions de VARMINT est justement de prédire quand l'entretien sera requis (parexemple, maintenant, dans une semaine ou dans un mois), basé sur les données de vibration et sur un modèle pouvant etre paramétrisé, sur les parties rotatives de l'équipement.

VARMINT est le résultat de travail entre la Garde-cote Canadienne (CCG), le Centre de Développement du Transport (TDC) de Transport Canada, MacDonald, Dettwiler and Associates Ltd. (MDA), et Design Maintenance Systems Inc. (DMSI). TDC a commandité et dirigé le projet pour CCG. MDA a fournis son expertise en IA et en ingénierie de système. DMSI a apporté son expérience en produits et services d'entretien préventif. VARMINT est présentement mis en marché comme produit commercial par DMSI.

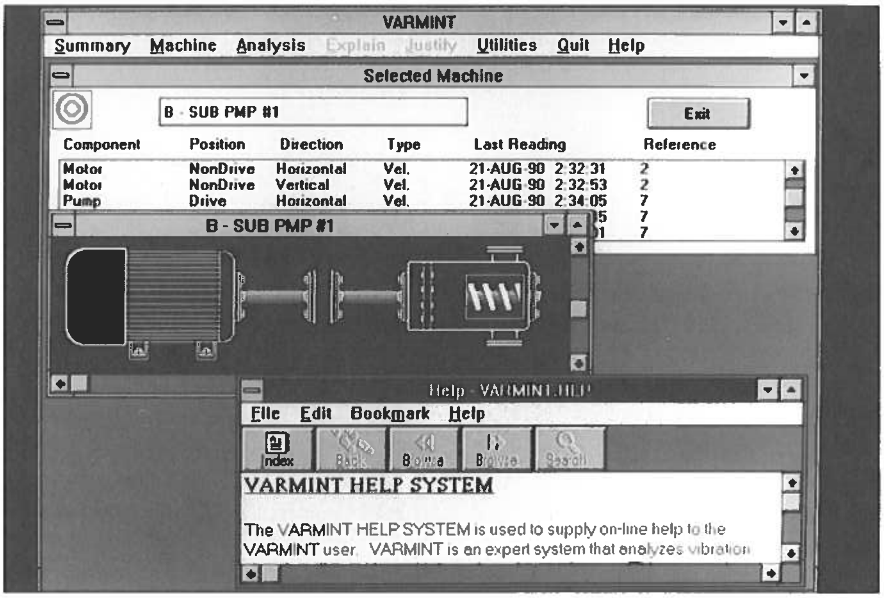

B - SUB PMP #1

Component	Position	Direction	Type	Last Reading	Reference
Motor	NonDrive	Horizontal	Vel.	21-AUG-90 2:32:31	2
Motor	NonDrive	Vertical	Vel.	21-AUG-90 2:32:53	2
Pump	Drive	Horizontal	Vel.	21-AUG-90 2:34:05	7
				5	7
				1	7

VARMINT HELP SYSTEM

The VARMINT HELP SYSTEM is used to supply on-line help to the VARMINT user. VARMINT is an expert system that analyzes vibration

Figure 1: On-Line Help: DMSI's VARMINT expert system includes on-line help, which makes the system easy to learn and use.

VARMINT (Vibration Analysis for Rotating Machinery Internals) is an expert system for predicting mechanical equipment failures, so that preventive maintenance can be performed as it is required. The system was designed to be used on board ships, where it analyses mechanical vibration data from hundreds of machines, including pumps, compressors, fans, turbochargers, engines, winches, and motors [1]. One of the functions of VARMINT is to predict when maintenance will be required (for example, now, in a week, or in a month), based on vibration data and a user-definable model of the rotating parts of the machine.

VARMINT is the result of work by the Canadian Coast Guard (CCG), the Transportation Development Centre

(TDC) of Transport Canada, MacDonald, Dettwiler, and Associates Ltd. (MDA), and Design Maintenance Systems Inc. (DMSI). TDC sponsored and directed the project for CCG. MDA provided their AI and systems engineering expertise [2]. DMSI supplied their experience with predictive maintenance products and services. VARMINT is currently marketed as a commercial product by DMSI.

History

The VARMINT project began with Ric Street of the fleet systems division of CCG. He believed that AI technology could be used to analyse vibration data on board CCG ships. CCG approached TDC with this idea, since TDC had experience with AI projects. In February of 1989, CCG and TDC sent a request for proposals to about 50 Canadian companies that were judged to have relevant experience in AI or vibration analysis. MDA and DMSI responded with a joint proposal to build an expert system.

MDA and DMSI began work on VARMINT in June of 1989 [3]. In December of 1989, they had documented the requirements for the system and they had a draft design. In August of 1990, VARMINT was installed on the CCG ship Sir William Alexander for a one-year field trial. The initial objective of the project was to build a demonstration prototype. By proceeding to a one-year field trial, the project actually exceeded the initial objective.

The story of the VARMINT prototype is only part of the story of the system's success. After the field trial, DMSI assumed responsibility for marketing and further developing the VARMINT system. This work is at least as important to the success of VARMINT. Of course, the complete story of VARMINT cannot be told, as long as the system continues to grow and develop.

I spoke with Pierre Jean of the CCG at Dartmouth, Nova Scotia, in November of 1992. He co-ordinates the CCG's use of the VARMINT system in the Maritimes region. He said that he is "quite satisfied" with the system, which has "taken a load off the shoulders" of the repair crew: "What used to take hours, now takes minutes." The CCG now has eight VARMINT systems in the Maritimes region, seven installed on their larger vessels and one at the regional office in Dartmouth.

Choice of Problem

The problem domain, as specified by the CCG, was the repair and maintenance of the approximately 200 pieces of rotating mechanical equipment on board CCG 1100-class icebreakers. Analysis of the costs of repair and maintenance showed that the area with the most room for improvement was preventive maintenance. As a precaution, machinery was often serviced when it was not necessary. Yet machinery sometimes failed, because it had not been serviced soon enough. The CCG believed that vibration analysis could help, by reducing the maintenance of healthy machinery and increasing the maintenance of machinery that is nearing failure. Vibration analysis was a good candidate for expert systems technology, because the 1100-class icebreakers were already equipped with portable vibration data collectors. This equipment, however, only performed an elementary analysis of the data. The CCG believed that a more sophisticated analysis would be highly beneficial.

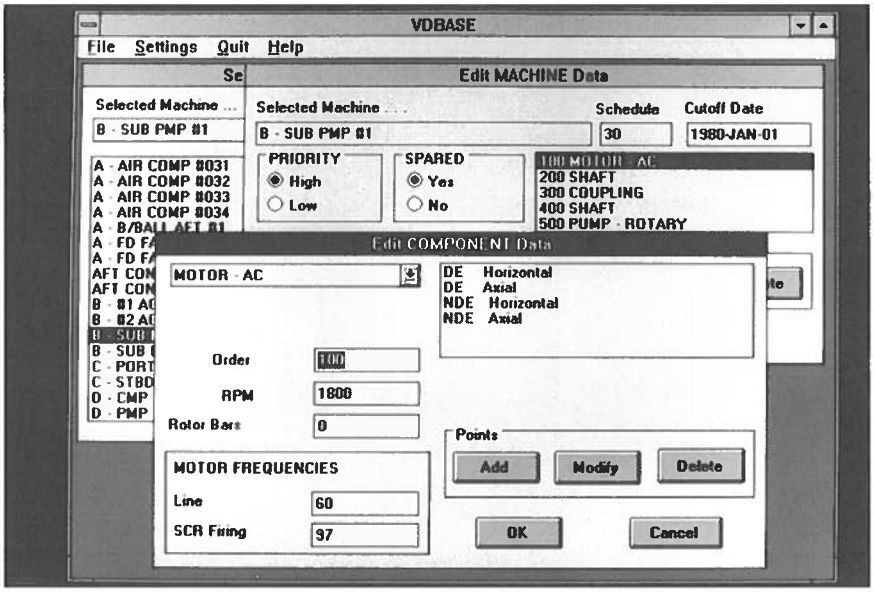

Selected Machine ...
B - SUB PMP #1

A - AIR COMP #031
A - AIR COMP #032
A - AIR COMP #033
A - AIR COMP #034
A - B/BALL AFT #1
A - FD F
A - FD F
AFT CON
AFT CON
B - #1 A
B - #2 A
B - SUB
B - SUB
C - PORT
C - STBD
D - CMP
D - PMP

Selected Machine ...
B - SUB PMP #1
Schedule 30
Cutoff Date 1980-JAN-01

PRIORITY
● High
○ Low

SPARED
● Yes
○ No

100 MOTOR - AC
200 SHAFT
300 COUPLING
400 SHAFT
500 PUMP - ROTARY

MOTOR - AC

DE Horizontal
DE Axial
NDE Horizontal
NDE Axial

Order 100
RPM 1800
Rotor Bars 0

MOTOR FREQUENCIES
Line 60
SCR Firing 97

Figure 2: Machine Train Model: The user "models" the machine train to tell DMSI's expert system as much as possible about the machine design and characteristics.

Description of the System

The VARMINT system uses the Nexpert expert system shell with Microsoft Windows 3.0. In addition to Nexpert, the system uses several subroutines written in C. The system runs on a PC-compatible, 386-based computer. It consists of about 20,000 lines of code, of which about 80% is conventional programming and 20% is expert system rules [1]. About half of the code is for the user interface.

The user interacts with VARMINT through a Windows graphical interface, using a mouse for most operations. VARMINT is a "non-consultative" expert system: It accepts data from portable vibration equipment, it automatically analyses the vibration spectra, then it sends the results of the analysis to the inference engine [3]. This approach minimizes the amount of information that comes from the user. The system is designed so that the user may query VARMINT, instead of having VARMINT query the user.

The VARMINT interface, in the standard Windows style, has a menu bar at the top, from which the user can select the main operations of VARMINT:

- prepare a "Summary" of all machines
- select a "Machine" to be analysed
- perform "Analysis" of the data for the selected machine
- "Explain" a conclusion in detail
- "Justify" the reasoning that leads to a conclusion
- get on-line "Help" on the use of VARMINT

When the user selects a machine, a C program analyses the vibration data for that machine. The C program extracts relevant features from the vibration data, such as peaks in the Fourier transform [3]. Part of the knowledge acquisition (KA) process was the derivation of these relevant features. These features are then fed into the Nexpert inference engine, which uses rules to form conclusions. A "conclusion" can be a symptom, a diagnosis, or a recommendation.

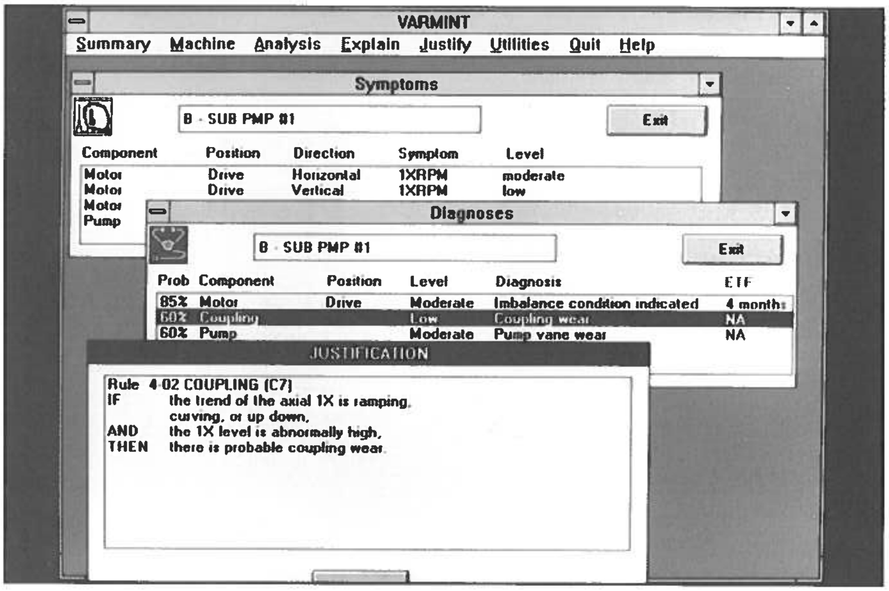

B - SUB PMP #1

Component	Position	Direction	Symptom	Level
Motor	Drive	Horizontal	1XRPM	moderate
Motor	Drive	Vertical	1XRPM	low
Motor				
Pump				

B - SUB PMP #1

Prob	Component	Position	Level	Diagnosis	ETF
85%	Motor	Drive	Moderate	Imbalance condition indicated	4 months
60%	Coupling		Low	Coupling wear	NA
60%	Pump		Moderate	Pump vane wear	NA

```
Rule  4-02 COUPLING (C7)
IF      the trend of the axial 1X is ramping,
        curving, or up down,
AND     the 1X level is abnormally high,
THEN    there is probable coupling wear.
```

Figure 3: Justification: At the users request, the DMSI expert system immediately makes available the justification behind the diagnosis. The justification includes rules utilized to produce the symptom, diagnosis, and recommended action.

Managerial Decisions

Timothy Bult of MDA emphasizes the importance of the engineering methodology that was used for the VARMINT project. He describes this methodology as a combination of ideas from the AI community and the systems engineering community. The methodology is discussed in detail elsewhere [1], but it is worthwhile to present it briefly here. The major elements are:

1. Project Plan: The plan is a sketch of the approach to solving the problem. It includes a description of the phases of the project, a schedule for the phases, and an estimate of the resources required for each phase.

2. Requirements Analysis: The requirements of the users must be analysed and documented. This process was guided by IEEE software standards [4]. Bult estimates that the requirements analysis took half of the scheduled time and 20% of the budget for the whole project. The requirements analysis document went through more than ten iterations [1]. This document should be relatively detailed. For example, the VARMINT requirements analysis document includes mock screens, to show how the user interface should look.

3. Knowledge Acquisition: The knowledge acquisition process begins before system design and extends through system validation. It is carried out in parallel with these other activities and it interacts with them. The VARMINT KA was highly systematic. Each interview of the domain experts was video-taped; the interviews followed an agenda with eleven items; several different formal and informal KA strategies were used. Seven domain experts were consulted, with expertise in different aspects of the problem domain. The knowledge engineers interview the domain experts and learn what they need to know in order to write a set of rules for the expert system and also a knowledge acquisition document. The KA document expresses the same information as the rules, but in a form that the domain experts can understand. This document is criticized by the domain experts. It is revised by the knowledge engineers, in parallel with the expert system rules, until the domain experts give their approval.

4. System Design: The system design involves the choice of hardware and software and the specification of the system architecture. This process was also guided by IEEE software standards [5]. Part of the system design is the choice of a scheme for knowledge representation, which will depend on the results of the knowledge acquisition.

5. Implementation: The development of the prototype began when the hardware and software, specified in the system design, became available. The initial work was done in Nexpert, but it soon became apparent that the system could not be implemented in Nexpert alone, since it was too slow for certain operations. Therefore several aspects of the system were implemented in C. This interaction between the design and implementation was important for the success of the system.

6. Validation: The validation process was performed using a "traceability matrix" [1]. This matrix traced items from one step of the project to the next step. Each item in the request for proposals was connected to the corresponding item in the requirements analysis document; each item in the requirements analysis document was connected to the

corresponding item in the system design document; and so on, up to the validation of the system. The system was tested with cases, some randomly selected and some deliberately selected. The results of the tests were connected back to the criteria for success, specified in the requirements analysis document.

For more information about this methodology, see [1] and [3].

Technical Decisions

Microsoft Windows is a typical window-based user environment. One feature which distinguishes Windows from some other products is its Dynamic Data Exchange (DDE) capability, which simplifies the exchange of data between application programs. This capability is not currently used by VARMINT, but it may be used in the future [3].

Nexpert satisfies two important criteria for an expert system shell: It is a popular shell, with a relatively large base of users, so it is likely to be strongly supported. Also, it is designed so that the expert system developer can customize it with code written in a standard programming language. Nexpert performed acceptably for most tasks, but there were a few operations that it did too slowly. These operations had to be coded in C.

Stephen Reilly of DMSI believes that the "non-consultative" design of VARMINT is very important to the system's commercial success. When it is feasible, it seems wise to minimize the amount of information that the user must supply interactively. Note that a "consultative" design could have been used: The system might have been designed to display various plots to the user, then ask the user a series of questions about the plots, such as, "Is there an unusual peak in the high frequency band?"

People

The partnership of MDA and DMSI was suitable for the development of the VARMINT prototype. DMSI wrote all of the software and did much of the design work for VARMINT. MDA was subcontracted for the KA and systems engineering. MDA has extensive experience in building AI systems, knowledge acquisition, conventional software, software engineering principles, and project management skills. DMSI has expertise in vibration analysis and software design. DMSI also has good contacts with domain experts. CCG was another source of much domain expertise. TDC has experience with sponsoring AI projects in the transportation sector.

Although about a dozen people were involved in the VARMINT project, most of them were only involved on a part-time basis. Timothy Bult estimates that the development of VARMINT required about six person-years (PYs), roughly five PYs from MDA and DMSI, plus one PY from CCG and TDC. TDC was charged about $300,000 by DMSI, of which about $100,000 went to MDA as a subcontract. CCG and DMSI also spent a significant amount of their own money on the project.

The project involved seven domain experts [3], but the KA meetings only took place about once a month, from the middle of the requirements analysis to the end of the validation phase [1]. Video-tape, a carefully planned agenda, and a KA document enabled the knowledge engineers to be efficient in their use of the domain experts' time.

Difficulties Encountered

The project presented a number of technical challenges [3]. One challenge was the format of the database in which the vibration data were stored. The project team discovered that it was difficult to access the data in the form that they required. It was necessary for them to write a special interface program for the database. This was done without any assistance from the company that designed the database.

Since there are many types of machinery on board a ship, VARMINT requires a model of the moving parts of each machine. In the initial design, it was decided that a fixed set of models would be adequate for the prototype, since it would be tested on a single ship, with a fixed set of machines on board. However, it soon became apparent that the supply of case histories for testing VARMINT was limited. There were no available case histories for the chosen ship. It became necessary to build a flexible, user-definable machine modelling system, so that VARMINT could be tested with the available case histories. This user-definable machine modelling system has turned out to be a very valuable feature of VARMINT.

Marketing Decisions

In November of 1992, I spoke with Roy Franco of DMSI in Vancouver about the marketing of VARMINT. He said that DMSI had sold 150 to 200 copies of VARMINT, world-wide, at $9,950 US each. In Canada, DMSI markets the system as VARMINT. Outside of Canada, SKF Condition Monitoring (based in California; the parent company, SKF, is based in Sweden) and Bruel Kjaer (a Danish company) market the system. They each have their own version of the system, tailored to work with their own proprietary vibration data collection equipment.

By licensing VARMINT to SKF Condition Monitoring and Bruel Kjaer, DMSI conveniently addresses the marketing of VARMINT. DMSI is now free to concentrate on the Canadian market. SKF CM and Bruel Kjaer facilitate the marketing of VARMINT by bundling the system with their hardware. Instead of selling an expert system, they sell vibration data collection equipment that has been enhanced by embedded AI technology.

Further Development

MDA was involved in the development of the first version only. DMSI now has full responsibility for support and further development. DMSI continues to improve the VARMINT system, using their expertise in vibration analysis and PC software, together with feedback from their customers. Most of their customers are not using VARMINT on board ships; they are using VARMINT in industrial applications, such as the pulp and paper industry.

I spoke with Stephen Reilly of DMSI in December of 1992 about the reasons for the success of VARMINT. He attributed much of the success to "revision after revision." Since the development of the prototype, VARMINT has undergone 18 revisions, 14 of which were released to the public. DMSI is now hard at work on version 2.0 of VARMINT, which will soon be commercially available. Reilly estimates that VARMINT 2.0 contains less than 10% of the source code of the original prototype.

Version 2.0 will no longer use Nexpert. It is written in C and C++, and it includes a custom inference engine that DMSI developed specially for VARMINT. Reilly cited three reasons for leaving Nexpert: (i) with Nexpert, DMSI must pay a royalty fee for each copy of VARMINT that they sell, (ii) DMSI has found that Nexpert does not work well in a multi-user network environment, and (iii) Nexpert is too slow.

The custom inference engine in VARMINT 2.0 performs resource-bounded computation. For example, the user can specify that the data must be analysed within 30 seconds. VARMINT 2.0 then limits its analysis to what it can do within the given time bound. The inference engine includes existential and universal quantification, parsing, objects, and backward chaining.

Conclusions

There are several factors that seem to have contributed to the success of VARMINT. Some of these factors are generally relevant to commercial AI projects:

- "revision after revision": DMSI continues to revise VARMINT, improving and expanding the system, based on their own experience and their customers' experience.

- the VARMINT project was user-initiated (as was Pitch Expert [6]); the CCG believed that expert systems technology could help in their problem domain.

- a carefully conceived software engineering methodology was applied to the project.

- the knowledge representation developed from the knowledge acquisition; there was no attempt to force a preconceived representation on the system (see Pitch Expert [6]).

Some other factors in the success of VARMINT seem to be dependent on the problem domain:

- VARMINT is a "non-consultative" expert system, which can increase user-acceptance, when it is appropriate.

- the VARMINT system is sold bundled with vibration analysis hardware; this seems to be a better approach than selling a stand-alone expert system, when it is appropriate.

It would be difficult to say which of these factors is most important. It is clear, however, that the work done after the completion of the prototype is at least as important as the development of the prototype.

Acknowledgements

Thanks to Timothy Bult of MDA for informing me about VARMINT; to Pierre Jean of CCG for answering my questions about CCG's experience with VARMINT; to Stephen Reilly and Roy Franco of DMSI for telling me about the development and marketing of VARMINT; and to Peter Clark, Mike Halasz, and Suhayya Abu-Hakima of NRC for their comments on this article.

References

[1] T. Bult, "A Mix of Software Engineering and Knowledge Engineering", *Operational Expert System Applications in Canada,* edited by C.Y. Suen and R. Shinghal, pp. 100-113, 1991.

[2] M. Krause, "Artificial Intelligence Work at MacDonald Dettwiler and Associates Ltd.", *Canadian Artificial Intelligence, Vol. 16,* pp. 23-27, 1988.

[3] S. Reilly and C. MacPhail, "VARMINT: An Expert System for Analysing Machine Vibrations", *Transport Canada Publication No. TP 10800E,* 1991.

[4] IEEE, "Guide to Software Requirements Specifications", *IEEE Standard 830,* 1984.

[5] IEEE, "Recommended Practice for Software Design Descriptions", *IEEE Standard 1016,* 1987.

[6] P. Turney, "Pitch Expert: Canadian AI Success Story #1", *Canadian Artificial Intelligence, Vol. 30,* pp. 14-17, 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. He is a regular contributing editor to* Canadian Artificial Intelligence *magazine.*

PRECARN UPDATE
NOUVELLES DE PRECARN

New PRECARN President Appointed

Mr. Mac Evans, PRECARN President and C.E.O.

PRECARN has seen many changes over its five year history, but none have caused such a flurry of emotion and excitement throughout its membership and staff. PRECARN has a new president and Chief Executive Officer, Mr. Mac Evans. The founding president, Mr. Gordon MacNabb, has decided to try retirement for a second time - the first time was about five years ago when he stepped down as president of NSERC.

Mr. Evans comes to PRECARN from the Canadian Space Agency, where he was Vice-President Operations. "Mac's experience is tailor-made for the task at hand, but even more important, Mac has demonstrated a remarkable ability to work harmoniously and effectively with others. In PRECARN, and in IRIS, that skill leads the list of prerequisites for the President's job, followed closely by the requirements for patience and perseverance!" states Gordon MacNabb. Mr. Evans has a B.Sc. in Electrical Engineering from Queen's University and an M.Sc. from the University of Birmingham. He began his career with the federal government in 1967 working on mobile communications for the Defence Research Board. Since then he has worked for the Department of Communications, the Ministry of State for Science and Technology, the Department of Industry, Science and Technology and the Canadian Space Agency.

Jean-Claude Gavrel

Gordon MacNabb won't quite make it to full retirement yet as he will continue to work with PRECARN on a part-time basis in 1993 as Director of the Institute for Robotics and Intelligent Systems (IRIS) with a primary responsibility of preparing the renewal of the Network, currently scheduled to terminate in 1994.

Other management changes at PRECARN/IRIS include the naming of Mr. Paul Johnston as Network Manager for IRIS, and yours truly as Vice-President for PRECARN. These changes reflect the maturing of PRECARN from an "experiment" to a operational research management organization. A number of challenges lie ahead for the Management and the Members of PRECARN, from delivering results of the research projects to ensuring continued long-term funding.

New Project Approved

PRECARN has recently approved funding for another Feasibility Study: it is in the area of computer supported collaborative work (CSCW) and is entitled "Team-Based Intelligent Productivity Systems (TIPS)". Participants in this project include MPR Teltech, the Alberta Research Council, Ernst & Young, Syncrude Canada and the University of British Columbia. The objective of the project is to develop advanced "intelligent" technology for computerized meeting support and will focus on same time/same place and same time/different place interactions. It will involve research in the areas of intelligent systems, multi-media and decision support systems. The project proposes to develop three automated "assistants": a Presentation Assistant, a Transcription Assistant and a Facilitation Assistant. The role of the Presentation Assistant will be to help in deciding the most appropriate form of presentation for a given set of data depending on the problem and based on individual needs. The Transcription Assistant will have the ability to capture salient points of the meeting proceedings - excerpts of video, audio, graphics and text. And the Facilitation Assistant will aim at increasing the consistency and effectiveness of meetings and to reduce the dependence on scarce, skilled humans. All three Assistants will be developed using the Distributed Agent Modelling (DAM) methodology. Broadly put, DAM is an instance of object-oriented programming, where the objects are agents with knowledge, goals and needs. The Feasibility Study phase is expected to be completed by the end of the year and will include an evaluation of industrial requirements and the definition of the technical framework for the research phase.

Technical Workshops

One of PRECARN's main objectives is the development and transfer of technology for exploitation by Canadian Industry. Technology Transfer is sometimes described as "a body contact sport" and an good mechanism for fostering these "body contacts" between the researchers and the users is workshops and conferences. In 1992, PRECARN and IRIS started to offer a number of project specific or technology oriented workshops. These were very well received by both our industry members and the university researchers. For 1993, we are already planning an extensive list of such activities, including two PRECARN project Workshops (ARK and IGI) in January and April respectively, IRIS project workshop workshops on Space Manipulators in March, on Bayesian Networks in April, on Human-Computer Interfaces also in April, on Speech Recognition in May, and on Cognitive Robotics in June. In addition we will be hosting again this year approximately 300 delegates from Industry, Government and University to our third annual IRIS~PRECARN Conference, to be held in Ottawa in June.

For more information contact:
Mrs. Lise McCourt, Manager of Corporate and Public Relations, or

Mr. Jean-Claude Gavrel, Vice President
PRECARN Associates
300-30 Colonnade Road, Nepean, Ontario, K2E 7J6
Tel: (613) 727-9576 / Fax (613) 727-5672
Email: Gavrel@a1.atott2.nrc.ca

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.

Become a member of CSCSI and receive *Canadian Artificial Intelligence* magazine. For more details, see inside back cover.

Canadian
Artificial
Intelligence

Intelligence
Artificielle
au Canada

CSCSI '94

Details to be announced in the June 1993 issue of *Canadian Artificial Intelligence* magazine.

Banff Park Lodge
16 - 20 May, 1994
Banff, Alberta CANADA

PACLING '93

First Pacific Association for Computational Linguistics Conference

April 21-24 (Wed-Sat) 1993

The Harbour Centre, Vancouver, British Columbia, Canada

Organizing Committee

Members:
Naoyuki Okada, Chair
Kyushu Institute of Technology, Japan
Nick Cercone
Simon Fraser University, Canada
Christian Matthiessen
University of Sydney, Australia
Yorick Wilks
New Mexico State University, USA

Local Members:
Dan Fass
Simon Fraser University, Canada
Paul McFetridge
Simon Fraser University, Canada
Fred Popowich
Simon Fraser University, Canada
Roland Sussex
Queensland University, Australia

Advisors:
Graeme Hirst
University of Toronto, Canada

Observers:
Minako O'Hagan
New Zealand Translation Center
New Zealand

Guest speakers: (tentative talk titles)
Dr. Takao Gunji
Osaka University, Japan
"An Overview of JPSG — A Constraint-Based Grammar for Japanese"
Dr. George Heidorn
Microsoft Corporation, USA
"Industrial Strength NLP: The Challenge of Broad Coverage"
Dr. Kathleen McKeown
Columbia University, USA
"Language Generation as Part of Multimedia Explanation"

Sponsors:
Natural Language Understanding and Models of Communication interest group of the Institute of Electronics, Information and Communication Engineers of Japan

The Australian Computer Science Society; Institute of Robotics and Intelligent Systems of Canada

The Advanced Systems Institute of British Columbia, Centre for Systems Science

Conference Aims

PACLING '93 will be a workshop-oriented meeting whose aim is to promote friendly scientific relations among Pacific Rim countries, with emphasis on interdisciplinary scientific exchange showing openness towards good research falling outside current dominant "schools of thought," and on technological transfer within the Pacific region. Papers have been received from many countries around the Pacific Rim (plus some from Europe) on a wide variety of topics within computational linguistics. Approximately 30 papers will be accepted for the conference. A list of accepted papers will be circulated early February.

Location of Conference and Hotels

The conference will take place at the Harbour Centre, the recently opened extension of Simon Fraser University at 515 West Hastings Street in downtown Vancouver. PACLING has secured special rates with three hotels, each only a few minutes walk from the Harbour Centre and from downtown shops, restaurants and nightlife.

Conference Registration Information

Full registration fees for the conference, besides attendance at conference sessions and use of guest e-mail facilities, include: copy of the conference proceedings; reception; banquet; and day trip to Whistler Village, home to two of the finest skiing areas in North America, Whistler and Blackcomb Mountains. The village and surroundings are very picturesque and have many shops and restaurants. Skiing is still good in April, weather permitting (ski passes not included in registration fee).

	Until March 15	**After March 15**
Full registration, reduced rate (full time student or unemployed)	CDN $105 US $ 88	CDN $105 US $ 88
Full registration, standard rate (everyone else)	CDN $210 US $ 17	CDN $270 US $230
Partial registration (partner of conference attendee reception, banquet and day trip only)	CDN $ 75 US $ 63	CDN $ 95 US $ 80

The registration fees include all taxes. We would prefer Canadian funds, but US funds are acceptable. Please pay by one of the following methods:

1. Bankers draft or cheque in Canadian dollars drawn on a Canadian bank
2. Bankers draft or cheque in US dollars drawn on an American bank
3. VISA card (supply VISA card number, full name and expiry date)

Please make bankers drafts and cheques payable to Simon Fraser University. Send your payment, complete with your name, address, phone number, and e-mail address (if applicable) to:

Fred Popowich: email: popowich@cs.sfu.ca
PACLING '93 Registration
Centre for Systems Science, Simon Fraser University
Burnaby, British Columbia Canada V5A 1S6
tel: (604) 291-4193 fax: (604) 291-3045

Attendees with Visa Problems

We are aware that some individuals may experience difficulties obtaining a visa. If people have paid for registration and then are not able to obtain a visa, their registration fee will be refunded. Unfortunately, hotel deposits probably cannot be refunded.

BOOK REVIEWS

Philosophy and AI: Essays at the Interface
Robert Cummins and John Pollock (editors)
(University of Arizona) Cambridge, MA: The MIT Press, 1991, xi+304 pp; Hardbound, ISBN 0-262-03180-9

Reviewed by
Peter Turney
National Research Council

If you are interested in both philosophy and AI (as I am), then you will enjoy this book. As the title indicates, the book is a collection of essays by philosophers and AI researchers, on topics of mutual concern. The range and depth of such topics may surprise some readers. It is not surprising that several of the essays deal with formal logic — mostly alternatives to classical deductive logic. However, other essays involve planning, learning, metapsychology, cognitive science, problem solving, probability theory, knowledge representation, natural-language competence (defined as both NL generation and NL understanding), extensions to object-oriented programming, and connectionist models of scientific explanation. This list covers a large fraction of the research that is done in AI. If an area is omitted from the list, it is only because the book is not an exhaustive treatment of the interface between philosophy and AI (of course, it does not claim to be exhaustive); it is not because there is no room for fruitful collaboration between philosophy and AI in that area.

The book begins with a concise introduction by the editors. They give eloquent explanations of "how philosophers drift into artificial intelligence" and "how artificial intelligence researchers drift into philosophy." As a "drifter" myself, I can support the aptness of their observations.

The book contains a dozen papers. My favourite was "Artificial intelligence and hard problems", by Glymour, Kelly, and Sprites. This paper examines the average-case complexity of problem solving and it arrives at some remarkable conclusions. For many of the problems that appear in the AI literature, the fastest algorithms that have been proposed have a worst-case complexity that is exponential in the size of the problem. A worst-case complexity bound is difficult to use in practice. In a real-world application, we may not care that an algorithm takes exponential time to solve 1% of our problems, if it can solve the remaining 99% in linear time. A more practical concern is the average-case (technically, the expected-case) complexity of the algorithm. With a few plausible assumptions, the authors show that the expected-case complexity of problem solving, for a large class of problems, is bounded by a constant, for all sizes of the problem. There seem to be two possible conclusions: Either problem solving is not as hard as we believed it to be, or something is wrong with the assumptions that the authors make. Either way, this is a stimulating paper.

Most of the papers propose ideas that were, to me, novel. For example, in "Memory, reason, and time," Elgot-Drapkin, Miller, and Perlis describe a formal logic that is inspired by cognitive psychology. This "step-logic" includes a formal model of different types of memory, such as episodic memory, semantic memory, and short-term memory. Unlike most formal logics, step-logic does not focus on final results; it does not model a chain of reasoning that ends with a conclusion. Instead, it models reasoning as a continuous process, occurring over time. It models fallible, limited, common sense reasoning, where contradictions and inconsistency are not catastrophic. In some ways, step-logic resembles default logics, but — unlike default logics — it emphasizes that reasoning is resource-bounded. Step-logic even models introspection; however, one element that is currently missing from step-logic is a representation of the reasoner's desires or goals.

Two other papers, "Plans and resource-bounded practical reasoning, " by Bratman, Israel, and Pollack, and "OSCAR: A general theory of rationality , " by Pollock, present work that is similar to step-logic. All three papers acknowledge the resource-bounded nature of reasoning. Like step-logic, OSCAR models introspection. Unlike step-logic, these two papers address "interest-driven" (OSCAR) or "means-end" (Bratman et al.) reasoning, which requires desires or goals. In other respects, these systems lack some of the psychological plausibility of step-logic. It would be interesting to attempt to combine elements of all three systems.

There are many other papers in this collection that are worth reading. Thagard's paper, "The dinosaur debate," presents a connectionist model of logic of scientific arguments. Loui's "Ampliative inference, computation, and dialectic" is an interesting treatment of an old question: What, exactly, is the difference between induction and deduction? There are several more, but this should indicate whether this collection matches with your interests.

One paper that I did not like was Doyle's "The foundations of psychology." The paper was written well, but the aim of the paper is (in my opinion) misguided. Doyle attempts to formulate a theory of mind that can encompass all possible thinking entities, whether they be the minds of strange, alien creatures; whether they can be realized in Turing-equivalent machines; or even whether they violate the laws of physics: "Many interesting psychologies may lie beyond the realm of what is realizable given the physics of our universe." This type of speculative meta-psychology is so abstract and so far removed from the real world that it has no appeal to me. I do not find it credible that anyone short of God could formulate such a theory.

The strength of this book is its diversity: The interface between philosophy and AI is evidently very broad. This strength, however, is also a weakness, because very few readers will have the breadth of background required to fully appreciate all of the papers in this collection (I am not one of those few). Nonetheless, it should be possible to follow all of the authors, even if it is not possible to constructively criticize them.

These minor criticisms aside, this is an excellent collection. If these kinds of problems are at all interesting to you, this book is worth reading.

***Peter Turney** is a Research Associate at the Knowledge Systems Laboratory of the Institute for Information Technology of the National Research Council of Canada. His current research is in Machine Learning. He obtained a Ph.D. in Philosophy from the University of Toronto, in 1988.*

BRIEFLY NOTED

SOAR: A cognitive architecture in perspective. A tribute to Allen Newell *John A. Michon and Aladin Akyürek (editors)* (University of Groningen) Dordrecht: Kluwer Academic Publishers (Studies in cognitive systems, edited by James H Fetzer, volume 10), 1992, xi+248 pp; hardbound, ISBN 0-7923-1660-6, US$89.00

Even before his death last year, several tributes had been planned for Allen Newell on the occasion of his 65th birthday. This book is one of them. The SOAR Research Group of the University of Groningen had been collaborating since 1987 with Newell's SOAR project on cognitive architectures and unified theories of cognition, and this book includes five papers presenting their results. In addition, there is a biography of Allen Newell, and a long paper by him in which he outlines the SOAR architecture and updates his ideas since the publication of *Unified theories of cognition*. Those not previously familiar with SOAR and Newell's later work will find this an easily accessible starting point.

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.

Logic primer
Colin Allen and Michael Hand (Texas A&M University) Cambridge, MA: The MIT Press, 1992, xiv+171 pp; paperbound, ISBN 0-262-51065-0, US$12.95

Logic programming: Proceedings of the Joint International Conference and Symposium on Logic Programming
Krzysztof Apt (editor) (CWI) Cambridge, MA: The MIT Press (Logic programming series, edited by Ehud Shapiro), 1992, xx+848 pp; paperbound,
ISBN 0-262-51064-2, US$75.00

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

Grammaires d'unification à traits et contrôlle des infinitives en français [Unification grammars with features and control of infinitives in French]
Karine Baschung (Université Paris X—Nanterre) Paris: 'Editions Adosa (Langues naturelles et traitement de l'information 2, edited by Gabriel G. Bés), 1991, 401 pp; paperbound, ISBN 2-86639-007-5, FF 360.

What robots can and can't be
Selmer Bringsjord (Rensselaer Polytechnic Institute) Dordrecht: Kluwer Academic Publishers (Studies in cognitive systems, edited by James H. Fetzer, volume 12), 1992, xiv+380 pp; hardbound, ISBN 0-7923-1662-2, US$111.00

The logic of typed feature structures
Bob Carpenter (Carnegie Mellon University) Cambridge, England: Cambridge University Press (Cambridge tracts in computer science 32, edited by C.J. van Rijsbergen), 1992, viii+270 pp; hardbound, ISBN 0-521-41932-8, US$34.95

Agency in action: The practical rational agency machine
S.C. Coval and P.G. Campbell (University of British Columbia) Dordrecht: Kluwer Academic Publishers (Studies

in cognitive systems, edited by James H. Fetzer, volume 11), 1992, xvii+206 pp; hardbound, ISBN 0-7923-1661-4, US$89.00

Generating referring expressions
Robert Dale (University of Edinburgh) Cambridge, MA: The MIT Press (The ACL—MIT Press series in natural language processing, edited by Aravind Joshi, Karen Sparck Jones, and Mark Y. Liberman), 1992, ix+277 pp; hardbound, ISBN 0-262-04128-6, US$39.95

Functional grammar in Prolog: An integrated implementation for English, French, and Dutch
Simon C. Dik (University of Amsterdam) Berlin: Mouton de Gruyter (Natural language processing 2), 1992, x+264 pp; hardbound, ISBN 3-11-012979-5

***Artificial intelligence applications in manufacturing**
A. (Fazel) Famili, Dana S. Nau, and Steven H. Kim (editors) (National Research Council of Canada, University of Maryland, and Massachusetts Institute of Technology) Menlo Park, CA: The AAAI Press and Cambridge, MA: The MIT Press, 1992, xiv+455 pp; paperbound, ISBN 0-262-56066-6, US$39.95

Logics of time and computation
(Second edition, revised and expanded) *Robert Goldblatt* (Victoria University of Wellington) Stanford: Center for the Study of Language and Information (CSLI lecture notes 7), 1992, ix+180 pp; distributed by the University of Chicago Press; hardbound, ISBN 0-937073-93-8, US$39.95; paperbound, ISBN 0-937073-94-6, US$19.95

Computers and writing: State of the art
Patrik O'Brian Holt and Noel Williams (editors) (Heriot-Watt University and Sheffield City Polytechnic) Oxford: Intellect Book (Great Britain only) and Dordrecht: Kluwer Academic Publishers (outside Great Britain), 1992, xi+387 pp; hardbound, ISBN 1-871516-20-X (Intellect) and 0-7923-1858-7 (Kluwer),$114.00, £67.00, Dfl~195.00

Metaphor and cognition
Bipin Indurkhya (Boston University) Dordrecht: Kluwer Academic Publishers (Studies in cognitive systems, edited by James H Fetzer, volume 13), 1992, xvii+456 pp; hardbound, ISBN 0-7923-1687-8, US$135.00

Languages of the mind: Essays on mental representation
Ray Jackendoff (Brandeis University) Cambridge, MA: The MIT Press, 1992, ix+200 pp; hardbound, ISBN 0-262-10047-9

Natural language processing: The PLNLP approach
Karen Jensen, George E. Heidorn, and Stephen D. Richardson (editors) (Microsoft Corporation) Boston: Kluwer Academic Publishers (The Kluwer international series in engineering and computer science; Natural language processing and machine translation, edited by Jaime Carbonell), 1992, xv+324 pp; hardbound, ISBN 0-7923-9279-5, US$80.00

The science of mind
Kenneth Klivington (Salk Institute for Biological Studies) Cambridge, MA: The MIT Press, 1989, reissued in paperback 1992, 239 pp; paperbound, ISBN 0-262-11141-1, US$19.95

***Genetic programming: On the programming of computers by means of natural selection**
John R. Koza (Stanford University) Cambridge, MA: The MIT Press (Complex adaptive systems series, edited by John H. Holland, Christopher Langton, and Stewart W. Wilson), 1992, xiv+819~pp; hardbound, ISBN 0-262-11170-5, $55.00 With accompanying video "Genetic programming: The movie" by John R. Koza and James P. Rice; 60 minutes, VHS format; NTSC, ISBN 0-262-61084-1, $34.95; PAL and SECAM, ISBNs not available, $44.95.Reason: Not CL; xfer to *Canadian Artificial Intelligence*

Towards high-precision machine translation: Based on contrastive textology
John Laffling Berlin: Foris Publications (Distributed language translation series, edited by Toon Witkam), 1991, xiii+178 pp; hardbound, ISBN 3-11-013388-1

Expressibility and the problem of efficient text planning
Marie Meteer (Rensselaer Polytechnic Institute) London: Pinter Publishers Ltd (Communication in artificial intelligence series, edited by Robin P. Fawcett and Erich Steiner), 1992, (distributed in North America by StMartins Press), xiii+188 pp; hardbound, ISBN 1-85567-022-4

A short introduction to modal logic
Grigori Mints (Stanford University) Stanford: Center for the Study of Language and Information (CSLI lecture notes 30), 1992, x+91 pp; distributed by the University of Chicago Press; hardbound, ISBN 0-937073-76-8, US$34.95; paperbound,
ISBN 0-937073-75-X, US$14.95

***Geometric invariance in computer vision**
Joseph L. Mundy and Andrew Zisserman (editors) (General Electric and Oxford University) Cambridge, MA: The MIT Press, 1992, xiii+540 pp; hardbound,
ISBN 0-262-13285-0

Enterprise integration modeling: Proceedings of the First International Conference
Charles J. Petrie, Jr (editor) (MCC) Cambridge, MA: The MIT Press (Scientific and engineering computation series, edited by Janusz Kowalik), 1992, xii+565 pp; paperbound, ISBN 0-262-66080-6, US$45.00

Grammatical competence and parsing performance
Bradley L. Pritchett (Carnegie Mellon University) Chicago: The University of Chicago Press, 1992, xii+192 pp; hardbound, ISBN 0-226-68441-5, US$55.00; paperbound, ISBN 226-68442-3, US$19.95

Connectionist approaches to natural language processing
Ronan G. Reilly and Noel E. Sharkey (editors) (University College Dublin and University of Exeter) Hove, UK: Lawrence Erlbaum Associates, 1992, xiv+472 pp; hardbound, ISBN 0-86377-179-3, US$89.95

The rediscovery of the mind
John R. Searle (University of California, Berkeley) Cambridge, MA: The MIT Press (Representation and mind series, edited by Hilary Putnam and Ned Block), 1992, xv+267 pp; hardbound, ISBN 0-262-19321-3, US$22.50; paperbound, ISBN 0262 69154-X, no price listed

Constraint-based grammar formalisms: Parsing and type inference for natural and computer languages
Stuart M. Shieber (Harvard University) Cambridge, MA: The MIT Press, 1992, xi+183 pp; hardbound, ISBN 0-262-19324-8, US$24.95

A functional perspective on language, action, and interpretation: An initial approach with a view to child modeling
Erich Steiner (Universität des Saarlandes) Berlin: Mouton de Gruyter (Natural language processing 1), 1992, x+289 pp; hardbound, ISBN 3-11-012379-0

Conference Announcement

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).

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)

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).

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

Applied AI	OBC
Knowledge Garden	Insert

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 Industrial Mobile Robot Prototyping and ALife Experiments

IS Robotics Exploring the Future

R-2™

Compact Robot System
for ALife Experiments

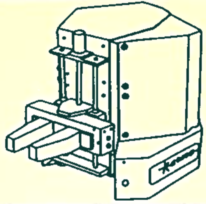

- Integrated, wheeled micro-robot
- Vertical direction parallel mechanical jaw gripper
- Ring of infrared proximity & bump sensors
- Cartesian manipulator for exploration
- Ideal for autonomous interaction, ALife, autonomous factory, cooperative work experiments

T-1™

Tracked Intelligent
Robot Platform

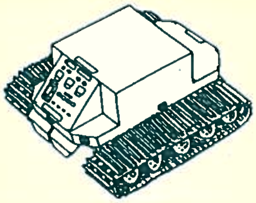

- Highly manoeuvrable tracked micro-robot
- High power caterpillar tracks for rugged terrain
- Low or high gear ratio option
- Infrared sensors for collision avoidance
- Piezo-electric bump sensors
- Ideal for mixed terrain cooperative studies and prototyping autonomous civil engineering equipment

T-2™

Tracked Exploration
Robot System

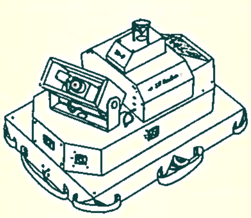

- Built on T-1 base (all features of T-1 apply)
- Color video camera and video transmitter (base receiving station included)
- Radio command/data link
- Optional infrared proximity sensors
- Optional vision processing module (future)
- Ideal for mixed terrain cooperative exploration
- Adapted by NASA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

T-3™

Tracked Human/Animal
Interaction System

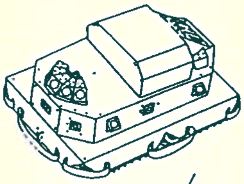

- Built on T-1 base (all features of T-1 apply)
- Pyro sensors for detecting human (animal) presence
- Light sensors to detect light intensity and gradient
- Optional microphones, speech recognition board, and digital speech output board
- Ideal for surveillance systems, entertainment systems, intelligent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

Pebbles™

Sealed Tracked Robot for
Hazardous Environments

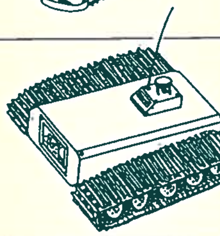

- Large payload capacity and expandability in the chassis
- Sealed to dust and, optionally, sealed to water, oil, chemicals
- Can be hosed down and re-used if contaminated
- Black and white camera on front; video transmitter inside
- Ideal for exploration and transportation in areas unsafe for humans

Genghis-II™

Six-Legged Walking Robot

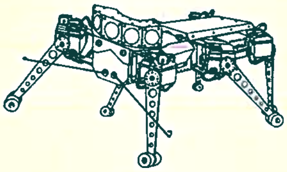

- Highly dexterous six-legged micro-robot
- Basic model detects forces, collisions, and people
- Extended package includes infrared proximity sensors, surface contact sensors, and pitch and roll inclinometers
- Expansion slots for future options
- Ideal for ALife behavior learning and evolution studies

Attila-II™

The Legged Robot System

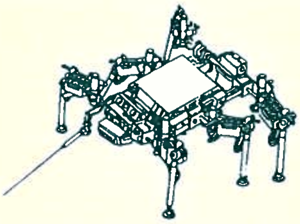

- Highly dexterous six-legged micro-robot
- Equipped with 150 high performance sensors
- Gyro stabilized CCD camera and rangefinder
- High performance multi-processor network
- Capable of determining environmental lay, texture, hardness and color
- Adapted by NASA for planetary exploration
- Ideal for exploration, inspection and sample collection

These robots based on Subsumption Architecture from MIT. For complete details, please contact:

Applied AI Systems, Inc.
Suite 500, Gateway Business Park, 340 March Road
KANATA, Ontario, Canada K2K 2E4
Tel: +1 613 592 3030 Fax: +1 613 592 2333
E mail: 71021.2755@compuserve.com

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 15FEB93)