Canadian Artificial Intelligence

Intelligence Artificielle au Canada

October/octobre 1989 **No. 21**

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

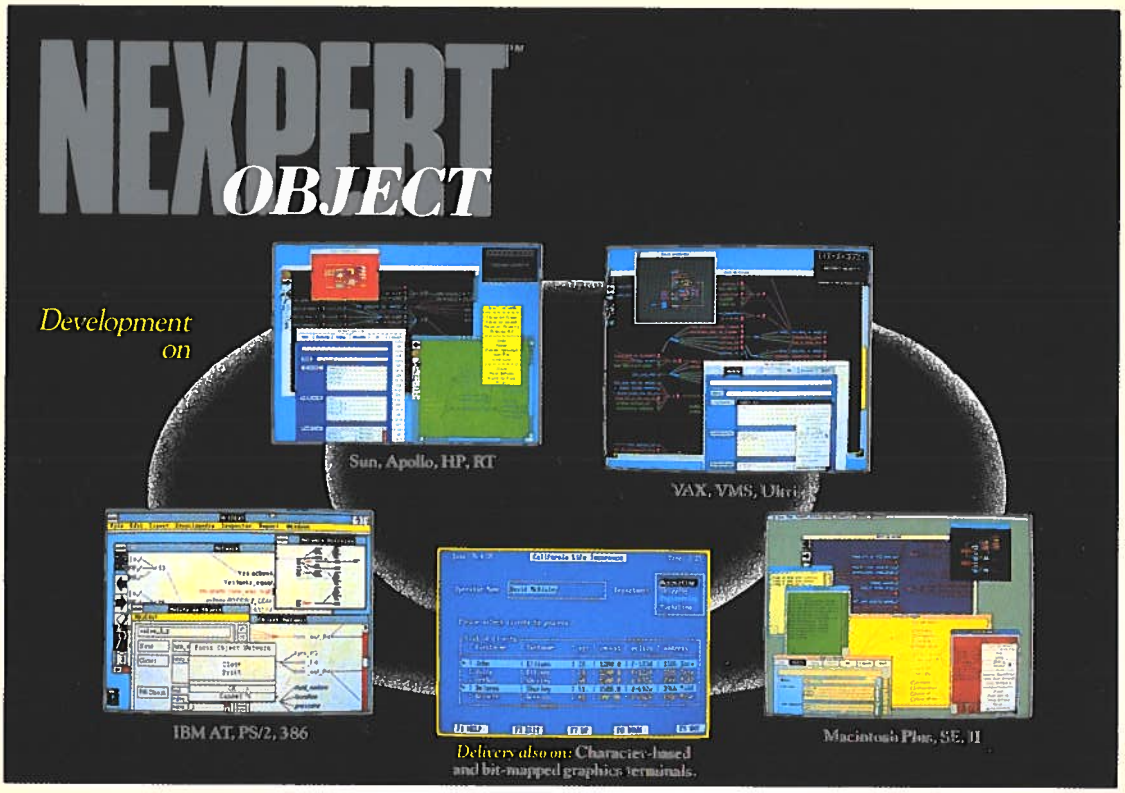

REASONS FOR NEXPERT'S SUCCESS

Links to popular conventional languages

- C
- COBOL
- FORTRAN
- Pascal
- Ada
- Assembly

Effective linkage with other databases and to own database – provides direct system in and out calls through the Library

- Oracle
- Informix
- dBase
- Guide
- Sybase
- RdB
- Lotus 1-2-3
- Hypercard (Q1/89)
- Ingres
- SQL/DS
- Excel

NEXTRA – first truly useful Knowledge Acquisition aid

- graphically induced acquisition process with automated repertory grids
- hierarchical clustering with spatial representations
- inductive engine for example-based analysis
- graphical feedback of elicited knowledge

Linkage to other graphics packages

Ability to graphically represent macrostructure of the knowledge base

- DataViews
- Ease+
- AiVision
- PCPaint/Brush
- MACPAINT

Cross-compatibility across a wide range of platforms – Quick and efficient adaptation to new platforms due to advanced Software Engineering techniques.

- development on VAXSTATION II, III & 2000 (VAX/VMS, ULTRIX) with delivery on all VAX computers
- Sun, HP, Apollo, and MicroVax workstations
- Hewlett-Packard 9000 series 300 and 800
- 80386 Unix workstations running X-windows
- IBM PC AT, PS/2 and compatible 80386-based PCs running Microsoft Windows
- Macintosh family (Plus, SE, II, IIx)
- IBM mainframe under VM (the 30xx, 43xx, and 9370 machines)
- Several foreign platforms (European, Japanese)

Runtime Package

- serious consideration given to runtime/delivery issues
- dual user interface concept: *end-user interface* not identical to *developer interface*
- highly automated end-user interface design tool

Phone or write for more information

Applied AI Systems, Inc.
Gateway Business Park
300 March Road, Suite 602
Kanata (Ottawa), Ontario, Canada K2K 2E2
Tel. (613) 592-3030 • Fax. (613) 592-2333

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

No. 21 October/octobre 1989

Contents

Canadian Society for Computational Studies of Intelligence

Founded in 1973

CSCSI is the Canadian society for the promotion of interest and activity in artificial intelligence. It conducts workshops and fully refereed national conferences, publishes this magazine, sponsors the journal Computational Intelligence, and coordinates activities with related societies, government, and industry. To join CSCSI, use the membership form in this issue. Non-Canadian members are welcomed. CSCSI is affiliated with the Canadian Information Processing Society and International Joint Conferences on Artificial Intelligence, Inc.

Memberships in CSCSI:
Membership form is on the last page. Please send subscriptions, memberships, and changes of address to:
CSCSI, c/o CIPS, 243 College Street, 5th floor
Toronto, Ontario, CANADA M5T 2Y1

Société canadienne pour l'étude de l'intelligence par ordinateur

Fondée en 1973

SCEIO est la Société canadienne encourageant l'intérêt et la recherche en intelligence artificielle. Elle organise des ateliers ainsi que des conférences nationales avec évaluation des articles soumis. Elle publie ce magazine, subventionne le journal Intelligence Informatique, et coordonne toute interaction avec des sociétés parallèles, le gouvernement, et l'industrie. Pour devenir membre de la SCEIO, veuillez utiliser le formulaire d'inscription de ce numéro. Les non-canadiens sont bienvenus. La SCEIO est affiliée à l'Association canadienne informatique, et aux International Joint Conferences on Artificial Intelligence, Inc.

Cotisations dans la SCEIO :
Le formulaire d'inscription est à la dernière page. Prière d'envoyer tout abonnement, cotisation, et changement d'addresse à:
SCEIO, c/o CIPS, 243 College Street, 5th floor
Toronto, Ontario, CANADA M5T 2Y1

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Founded in 1974 as / Fondée en 1974 en tant que CSCSI/SCEIO Newsletter

Submissions:
Canadian Artificial Intelligence is published quarterly by CSCSI/SCEIO and is a benefit of membership in the society. *Canadian AI* solicits contributions in English or French on any matter related to artificial intelligence, including: articles of general interest; descriptions of current research and courses; reports of recent conferences and workshops; announcements of forthcoming activities; calls for papers; book reviews and books for review; announcements of new AI companies and products; opinions, counterpoints, polemic, controversy; abstracts of recent publications, theses, and technical reports; humour, cartoons, artwork; advertisements (rates upon request); anything else concerned with AI. Paper or electronic submissions are welcome. Electronic submissions are preferred and should be unformatted. Canadian AI is published in January, April, July, and October. Material for publication is due six weeks before the start of the month of publication.

Contributions :
L'Intelligence artificielle au Canada est publiée trimestriellement par la CSCSI/SCEIO, et est offerte gratuitement aux membres. L'*IA au Canada* encourage les contributions, en français ou en anglais, portant sur l'intelligence artificielle. Ceci comprend: des articles d'intérêt général; des descriptions de recherche courante et de cours; des rapports de conférences récentes et d'ateliers; l'annonce d'activités à venir, et des requêtes d'articles; des critiques de livres ainsi que des livres à critiquer; l'annonce de nouvelles compagnies en IA et de leurs produits; des opinions, des répliques, tout ce qui est polémique; des résumés de publication récentes, de thèses et de rapports; des trucs humoristiques ou artistiques, de bandes dessinées; des annonces (s'enquérir des frais); tout autre matériel touchant à l'IA. Contributions, sur papier ou par courrier électronique, sont bienvenues. Nous préférons le courrier électronique mais les soumissions ne doivent pas avoir un format. L'IA au Canada apparait en janvier, en avril, en juillet, et en octobre. Toute communication à publier doit nous parvenir au moins six semaines avant le début du mois de parution.

Advertising:
Advertising rates and press kits are available upon request from the address below, or by phoning 403-297-2600.

Réclame:
Les prix pour les annonces et les trousses pour la presse sont disponibles sur Écrivez à l'addresse à la gauche ou téléphonez 403-297-2600.

Please send submissions to / Prière d'envoyer contributions à:
CDNnet: CSCSI@noah.arc.cdn
UUCP: cscsi%noah.arc.cdn@alberta.uucp
or to / ou à: Marlene Jones
Canadian Artificial Intelligence
Alberta Research Council
6815 8th Street NE, 3rd floor
Calgary, Alberta, CANADA T2E 7H7

Book reviews and candidate books for review should be sent to:
Envoyez des critiques de livres ainsi que des livres à critiquer à:
CDNnet: gh@ai.toronto.cdn
CSNET: gh@ai.toronto.edu
UUCP: gh@utai.uucp
or to / ou à: Graeme Hirst, *Canadian Artificial Intelligence*
Department of Computer Science, University of Toronto
Toronto, Ontario, CANADA M5S 1A4

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

ISSN 0823-9339

Executive Notes

The results of the reader survey printed in the April 89 issue of the magazine are finally tallied and interpreted. We received 25 responses (approximately 4% of subscribers, not statistically significant) roughly equally divided between industry (44%) and university (36%). 4 readers responded in French and 21 in English. (The French readers pointed out some embarrassing grammatical errors in the questionnaire - our sincerest apologies). Thank you to the readers who took the time to respond.

All but one respondent thought that the magazine was either very relevant (44%) or somewhat relevant (52%). 48% of respondents read most of the magazine with 24% reading it from cover to cover and 28% reading some of it. After reading it, these people either saved the magazine for future reference (80%) or passed it on to others (16%). In fact, these 25 respondents of the magazine passed it on to 30 others. (If extrapolated, this means that the magazine is reaching over 1300 people in total.) The one reader who thought the magazine was not at all relevant discarded it after reading it for 15 minutes.

On average, most readers (56%) spent over 2 hours with the magazine; 24% spent 15 minutes or less. The sections of interest to readers were (in decreasing order of interest): Feature Articles, Research Reports, AI News, Publications, Conference Reports, Conference Announcements, World Watch, and Humor. Readers want the Book Reviews section to be expanded and the magazine to take a more industrial slant with reviews of tools and case studies of expert systems projects. Technical articles were deemed by some to be of little use, whereas the Thimbleby-Witten genre of speculative articles was lauded by at least one respondent.

The low response rate makes it difficult to generalize these results to the population as a whole. Rather than speculate on what the rest of the "AI Community" thinks of the magazine, we would like to present the results as anecdotal and have you make any conclusions from them.

Letters to the Editors

Worms and bugs in the Church-Turing Thesis

We are pleased that our article on "Bugs in the Church-Turing Thesis" [6] has achieved its goal of initiating a dialogue, but must confess to some disappointment that Bersohn's response [1] is based more upon asseveration than on specific counter-claims to which we might recruit evidence against, or which might, on reflection, help us to change our views or express them more precisely.

Consider his snowflake. Obviously it has marvelous precision and intricacy, and the water molecules have not had to "think up" their design—at least, not in the sense that a person or a computer (or a spider) would have to think up such a design. And this is precisely our point. A computer program to create the snowflake would work in a way that was very different to the actual mechanism. And would it be as successful? At a gross level it would be easy to simulate a snowflake, we suppose; but as the detail required became finer and finer, and the granularity of the simulation had to change from symmetry pattern to molecule, from molecule to atom, from atom to elementary particle and quantum level ... to achieve the desired accuracy—who's to say that the process would converge? We may need a computer of arbitrary complexity to simulate a small physical object to arbitrary accuracy. Instead of repeating our original Feynman point we quote a recent, more detailed conjecture [7] that "some problems [of simulating physical systems] may be so massively parallel, and so extensively interdependent, that they can only be efficiently simulated by themselves—that is, they are not simulatable by any significantly different architecture." Bersohn's wonderful liver cells, which he claims are "not considered by anyone to be capable of thought or computation," may in fact transcend thought or computation. That is what we are worried about. "Unthinking nature does not 'know'," he continues, but perhaps we as artifacts of nature cannot

either (nor our computers), and the Church-Turing Thesis is doomed to failure! The recent surge of interest in chaotic models of physical processes (see [5] for a popular presentation) provides more evidence of a trend toward rejection of the old view of the world as a complex, but tractable, process that can be simulated quantitatively to arbitrary degrees of accuracy.

As for worms, of course a single worm cannot play chess (presumably because it has no survival value in the normal habitat). But—to paraphrase conventional AI approaches—we are sure that if we only had the facilities for 96 worms, we could do it. (Exercise for the reader: 96 worms allows for promoting pawns, unless worms are piece/colour-specific). In fact we have begun some preliminary investigations with half a dozen common-or-garden worms *(Eisenia fetida)* which grow well in a variety of habitats, including compost heaps. Although there are some problems with the user interface we believe we now have these under control and the worms are on the verge of being able to cope with some of the standard book openings. Incidentally, one of the problems with this research is the high mortality rate of worms in the chess-board environment; current work [4] at the Open University, England, is tackling the problems of being able to produce large quantities of earthworms adapted to living in different environments and we hope to be able to benefit from this.

Much work in robotics in particular seems to us to take a similar approach to our parody: when one worm fails, try a few more. When robots are concerned, the "obvious" upper limit of 96 does not exist, and even if it did, would be prohibitively expensive to obtain. Contrariwise, if you lent us a vermiculture set and 96 worms and we failed in our stated aims, then our hypothesis would appear to be roundly refuted. AI, certainly the way it is often advanced in practice, does not lend itself to such easy refutation, and (following Popper) we might say this is to its actual disadvantage.

Although our view (on the standing of the Church-Turing Thesis) has implications for research activities, on currently available evidence it is observationally indistinguishable from a view argued for and endorsed by Dennett [3]: *viz*, that non-biological computation will never be swift enough to support acceptable human intelligence. Dennett argues that the brain is approximately 10^5 times faster (measured in operations per second) than the fastest contemporary computer's raw processing speed. Levels of abstraction placed on this slow a simulation down: a practical problem, then, is that even if AI is only a factor of two too slow in some levels, it does not take very many such levels for the overall progress of an intelligent agent to be far too slow to be preceptibly intelligent in any natural environment. Dennett's view is that the Church-Turing Thesis is valid, but practically inapplicable; our stronger view is that it may simply be inapplicable in principle.

But our main argument in our original paper addressed to robotics was based on the example set us by spiders, not worms. We propose a prize (having lodged $1 with the Canadian AI Society for this express purpose, and to be called the Thimbleby-Witten-Spider Prize) to be awarded to the first robot (or its legal representative) capable of out-performing the common British House Spider *Teganaria atrica* — the female achieves 33 body lengths per second (and is not the fastest spider) [2]. This is a fitting revival of the eighteenth century British tradition of spider racing, and we earnestly hope that the high quality competition that is encouraged lends increasing prestige to the prize. Of course, $1 may seem purely nominal, but then (as we are sure the competitors will agree) by the time the prize is won, the AI society's normal investment programme will have brought its value to a handsome sum.

Ian H. Witten
Computer Science, Calgary University, Calgary T2N 1N4, Canada

Harold W. Thimbleby
Computing Science, Stirling University, Stirling FK9 4LA, Scotland

[1] Bersohn, M. "No Bugs in the Church-Turing Thesis," *Canadian Artificial Intelligence*, 4-5, July 1989.
[2] Bristowe, W. S. *The World of Spiders*, Collins, 1958.
[3] Dennett, D.C. "Fast Thinking" in *The Intentional Stance*, Bradford Books , MIT Press, 1987.
[4] Knight, D. "Nice Work for a Worm," *New Scientist*, 1682: 55-59; July 1989
[5] Stewart, I. *Does God Play Dice?: The Mathematics of Chaos*, Blackwell: Oxford; 1989.
[6] Thimbleby, H. W. & Witten, I. H, "Bugs in the Church-Turing Thesis," *Canadian Artificial Intelligence*, 17-18, April 1989.
[7] White, I. "The Limits and Capabilities of Machines — A Review," *IEEE Trans on Systems, Man, and Cybernetics* 18(6): 917-938; November/December 1988.

Conference on Ethics and Technology

A conference entitled "Ethical choices in the age of pervasive technology" will be held at the University of Guelph 25-29 October 1989. Among the goals of the conference are to increase the understanding of the worldwide impact of new technologies, to assess the current and future impact of technologies, and to set forth an ethical agenda for the use of technology that will create a humane and sustainable future.

Information is available from: Ethical Choices in the Age of Pervasive Technology, Division of Continuing Education Room 160, Johnston Hall, University of Guelph, Guelph, Ontario N1G 2W1 (519) 824-4120 x 3956 (519) 824-4747 x 3957 Fax: (519) 767-0758

AICorp comes to Canada

AICorp of Waltham, Mass. has opened a Canadian headquarters in Toronto to support its Canadian customers. AICorp is the vendor of KBMS an expert system tool integrated with IBM mainframe databases, and KBMS/PC, an expert system tool for MS-DOS and OS/2.

PRECARN appoints new Vice President

Dr. Arthur Collin, former Secretary, Minstry of State for Science and Technology and Science Advisor to the Government of Canada, has been appointed as Vice President of PRECARN Associates, Inc. Dr. Collin's career within R&D policy organizations of the Federal Government spanned 33 years, including responsibility for planning of the Canadian Space Agency.

PRECARN funds new projects

Two more projects have been funded by PRECARN, one each in the area of intelligent systems and robotics. The first, a project comprising Microtel Pacific Research Ltd., the Advanced Systems Institute of BC, H.A. Simons Ltd., and Simon Fraser University, aims to develop an intelligent, graphical interface to assist operators of complex monitoring and control systems. The second, comprising Inco Ltd., Falconbridge Ltd., the Canadian Centre for Automation and Robotics in Mining, McGill University, and Ecole Polytechnique de Montreal, is to tailor and enhance computer vision, robotics, and AI techniques to the particular problems of mining. The addition of these projects brings the total number of feasibility studies funded by the agency to nine.

PRECARN receives Federal Funding

On June 28th, the Honourable Harvie André, Minister of Industry, Science and Technology, announced that the Ministry will contribute $10 million of funding to PRECARN's Intelligent Systems projects over the next five years. This funding will assist the start-up of major research projects after the initial feasibility studies have been completed. This funding will complement

Conférence sur l'éthique et la technologie

Une conférence intitulée "Ethical choices in the age of pervasive technology" (Choix éthiques à l'âge de la technologie omniprésente) aura lieu à l'Université de Guelph du 25 au 29 octobre 1989. Parmi les objectifs de la conférence, on trouve l'accroissement de notre compréhension des conséquences mondiales des nouvelles technologies, l'évaluation de l'impact des technologies dans le présent et dans le futur, la mise en place d'un agenda éthique à l'usage de la technologie qui permettera la création d'un avenir supportable et humain.

Pour de plus amples informations contactez: Ethical Choices in the Age of Pervasive Technology, Division of Continuing Education Room 160, Johnston Hall, University of Guelph, Guelph,Ontario N1G 2W1 (519) 824-4120;3956 (519) 824-4747;3957 Fax: (519) 767-0758

AICorp entre au Canada

AICorp de Waltham, Mass. a ouvert une succursale canadienne à Toronto afin de mieux servir ses clients canadiens. AICorps vend KBMS, un outil de développement de systèmes experts intégré aux bases de données sur gros ordinateurs d'IBM, et KBMS/PC, une version de KBMS pour MS-DOS et OS/2.

PRECARN nomme un nouveau vice-président

Le Dr. Arthur Collin, ancien secrétaire du Ministère des Sciences et Technologies et conseiller scientifique du gouvernement canadien, a été nommé vice-président de PRECARN Associates, Inc. La carrière du Dr. Collin au sein des organisations responsables des politiques du gouvernement fédéral en matière de recherche et de développement s'étend sur une période de 33 ans, où l'une de ses responsabilités fut la planification de l'agence spatiale canadienne.

PRECARN subventionne de nouveaux projets

Deux autres projets sont subventionnés par PRECARN, l'un dans le domaine de l'intelligence artificielle et l'autre dans le domaine de la robotique. Le premier qui implique une collaboration entre Microtel Pacific Research Ltd. et l'Université Simon Fraser, a pour but le développement d'une interface graphique intelligente destinée à aider les operateurs de systèmes de surveillance et de contrôle complexes. Le second projet est une collaboration entre Inco Ltd., Falconbridge Ltd., Canadian Centre for Automation and Robotics in Mining, L'Université McGill et l'Ecole Polytechique de Montréal. Il a pour but d'adapter la vision par ordinateur, la robotique et des techniques d'intelligence artificielle aux problèmes particuliers que posent les mines. L'addition de deux projets porte à neuf le nombre total des études de faisabilité subventionnées par l'agence.

PRECARN reçoit une subvention fédérale

Le 28 juin, l'Honorable Harvie André, ministre d'Industrie, Science et Technologie Canada a annoncé que le ministère octroyera $10 million aux projets de développement de systèmes intelligents de PRECARN, au cours des cinq prochaines années. Cette subvention aidera à la mise sur pied de projets de recherches majeurs après que les études de faisabilité initiales auront été

PRECARN's own resources and $5 million from the NRC Industrial Research Assistanceship Program (IRAP).

AI in the Mainstream

The July 20th edition of *Computing Canada* (volume 15, number 15) contains an extensive Special Report on artificial intelligence in business. It appears that we can all buy suits and go to work for the banks. Ah, legitimacy!

Stand up and be counted!

The Commission of the European Communities (CEC) is doing a survey on public institutions and private companies which are engaged in research or development activities in the field of "Language Industries" — e.g. speech recognition and synthesis, spellchecking, computer-assisted translation, hypertext, text database management, naturall-language interfaces, and optical character recognition. This survey aims at getting to know "who does what" in this field, and covers Europe, North America, and Japan.

If this survey is of interest to you and you wish to receive the questionnaire, please write to:

M. Emmanuel Rabier, INK International, Prins Hendriklaan 52, 1075 BE Amsterdam, Nederland.

This is a unique opportunity to make your know-how and products known overseas. The detailed answers will be restricted to CEC officials, but the list of respondents will be widely curculated.

Price Waterhouse Survey

The Department of Industry, Science and Technology Canada, Technology Policy Branch has commissioned Price Waterhouse to conduct a broad survey of AI capabilities in Canada. The resulting inventory will be maintained and used by several government funding agencies to identify organizations who could be contracted to supply AI expertise to the government.

Price Waterhouse is interested in hearing from any company, university or individual working in cognitive systems, sensory systems and robotics systems. If you are interested in being included, please contact:

Price Waterhouse, 180 Elgin Street, Suite 1100, Ottawa, Ontario K2P 2K3. Attn: Anne Belngessner, Telephone: (613) 238-8200, Fax: (613) 238-4798.

complètées. Cette subvention s'ajoute aux ressources propres de PRECARN et aux $5 millions reçus du Programme d'aide à la recherche industrielle du CNRC (PARI).

L'I.A. est à la mode

L'édition du 20 juillet de *Computing Canada* (volume 15, numéro 15) contient un long reportage spécial sur les applications de l'intelligence artificielle dans le domaine des affaires. Il semble que nous puissions tous acheter des vestons et aller travailler dans les banques. Ah, la légitimité!

Faites-vous connaitre!

La Commission des Communautés Européennes (CEC) méne actuellement une inquête sur les établissments publics et les entreprises privées que oeuvrent en recherche ou développement dans le secteur des "industries de la langue" : recommaissance et synthèse de la parole, correction orthographique, traduction automatique, hypertexte, informatique documentaire, interfaces en langue naturelle, lecture optique de caractères, etc. Cette eenquête vise à connaître "qui fait quoi" dans ce secteur; elle couvre l'Europe, l'Amérique du Nord et le Japon.

Si cette enquête vous concerne et que vous vouliez le questionnaire, veuillez écrire à l'adresse cuivante :

M. Emmanuel Rabier, INK International, Prins Hendriklaan 52, 1075 BE Amsterdam, Nederland.

C'est là une occasion unique de faire connaître vos réalisations ou vos produits outre-Atlantique. Le détail des réponses n'est connu que des décideurs de la CCE, mais la liste des répondants fait l'objet d'une diffusion assez large.

Le sondage Price Waterhouse

La direction générale de la politique sur la technologie du ministère de l'industrie, Sciences et Technologie Canada a chargé Price Waterhouse de conduire un grand sondage sur les possibilités en intelligence artificielle (IA) au Canada. L'inventaire qui en résultera sera gardé à jour et utillsé par plusieurs organismes-subventionneurs gouvernementaux pour identifier les organisations pouvant fournir de l'expertise en IA au Gouvernement.

Price Waterhouse aimerait inventorier n'importe quel(le) compagnie, université, ou individu qui travail sur les systèmes cognitifs, les systèmes sensoriels, et les systèmes robotiques. Si vous aimeriez être inclus dans ce sondage, veuillez contacter :

Price Waterhouse, 180 Elgin Street, Suite 1100, Ottawa, Ontario K2P 2K3. A l'attention de Anne Belngessner,
Numéro de téléphone: (613) 238-8200,
Numéro de télécopieur: (613) 238-4798.

Products and Corporations

AICorp (Waltham, MA) has released a new and improved version of its KMBS expert system tool that features a 20-fold improvement in speed. An MS-DOS version has also been released, and is selling for US$9,500.

Inference (Los Angeles, CA) has announced the commercial release of ART-IM/MVS, the IBM mainframe version of ART. ART-IM/MVS is written in C and goes for US$100,000. Inference has also released ART-IM/MS-DOS for PC's. It requires 640K of memory and 8M of hard disk, and goes for US$8,000.

Texas Instruments (Dallas, TX) and Air Canada are jointly developing two expert systems to increase Air Canada's efficiency in scheduling maintenance and scheduling gate and ground services. No mention on whether it will improve the taste of in-flight food.

Defense Research Establishment Atlantic (Dartmouth, Nova Scotia) has awarded a contract to a team consisting of Canadian Astronautics Ltd, KnowledgeWorks Research Systems Ltd, and IMP Group Ltd. The contract, for CAN$430,000, will see the group construct a test bed for an intelligent vehicle classification system to be used by sonar operators.

International Submarine Engineering (Vancouver, BC) has been awarded a CAN$188,000 contract entitled "Development of a Knowledge-Based Health Monitoring and Planning System for Space Station MSS Power Management - Phase I".

BBN Advanced Computers (Cambridge, MA) has enhanced its Butterfly Mach 1000 operating system to support up to 128 nodes. Butterfly upgrades are free.

Hecht-Nielsen Neurocomputer (San Diego, CA) has been awarded a US$500,000 contract by the U.S. Army Laboratory Command for Phase II of the Army Battlefield Neurocomputer Project. The contract will see Hecht-Nielsen develop a neurocomputer utilising the HNC100X, a neural chip. The goal is a performance of 500 million connection updates per second.

Syntonic Systems (Portland, OR) has announced DENDROS-1, described as the first commercially available neural network chip. Each chip implements 22 "synapses", and is available for US$35. An evaluation board containing 8 chips costs US$695.

Texas Instruments (Austin, TX) has announced the Speech System V Toolkit for TI's 386-based workstations. The system allows for the development of speech recognition, verification, and text-to-speech systems. It runs under XENIX, and costs US$1,000 for run-time versions, US$2,000 for development versions.

Thinking Machines (Cambridge, MA) has announced a smaller version of the CM-2, called, not surprisingly, the CM-2a. It comes in two flavours: a US$1,000,000, 8192 processor version and a US$500,000, 4096 processor version.

Symbolics Inc. (Burlington, MA) has announced the availability of Inference Corporation's ART 3.2 for its MacIvory systems. (MacIvory is an add-in board for the Macintosh II based on Symbolics Ivory processor).

Nestor Inc. (Providence, RI) is now producing the Nestor Development System for the Macintosh and Sun 4, as well as an ASCII based version which allows it to interface to systems such as Lotus and d-Base.

California Scientific Software (Sierra Madre, CA) has announced a new training program for neural nets using semi-linear algebra which allows them to learn at up to 100 times faster than traditional methods. The program, called the HyperSonic Training Program, is available for IBM PC's, and requires BrainMaker 1.3 or higher, also available from California Scientific Software. Prices are US$195 for BrainMaker and US$149.95 for HyperSonic Training Program.

Oracle Corporation (Belmont, CA) has announced the release of a software link between its Oracle database and Apple's HyperCard via HyperTalk. The link allows data from different database packages on different systems to be accessed through the HyperCard interface. Prices are US$299 for developers and US$999 for the networking version for end users.

Less than one year after being launched, Envos Corporation of Mountain View, California, has returned to and been absorbed by its parent, Xerox Corporation of Stamford, Connecticut. Products marketed by Envos will continue to be supported by Xerox for now.

Franz Inc. (Berkeley, CA) will be marketing Carnegie Group's (Pittsburgh, PA) Knowledge Craft. Initially, two components of Knowledge Craft, the Carnegie Representation Language (CRL) and CRL-OPS (an integration of OPS and CRL) will be offered. Prices are in the US$5,000 range.

Human Devices (New York, NY) has demonstrated Chorus, a parallel processor for neural network simulations. Using many distributed RISC 64-bit floating point processors, Chorus reportedly reaches a peak speed of 16 million connection updates per second.

Anza Research (Cupertino, CA) has released NeuralBase 5000, a neural network database in dBase III format. It contains 4200 entries and 3000 annotations, covering 100 years of research in over 200 publications. Entries can be searched using up to 10 of the 150 available keywords. NeuralBase requires an IBM PC XT or AT with at least 512K, a 360K disk, and a 5MB hard disk.

Apollo Computer (Chelmsford, MA) has released version 3.0 of Domain/Common Lisp. Among its new features are ephemeral garbage collection, sophisticated interrupt and error handling, better access to other programming languages, better debugging facilities, a fast compilation mode, and multi-tasking. It runs on the Series 4500, 3500, and 3000 workstations. Price is US$3,500.

Symbolics Inc. has released the successor to the MacIvory system, MacIvory 2. Symbolics claims complete code compatibility with the old system and a speedup of 75%. Price for an upgrade is US$15,900. Completely bundled systems start at US$30,000.

Applied Intelligent Systems Inc. (Ann Arbor, MI) have released the AIS-4000, a machine vision computer incorporating a parallel array processor. The array processor consists of 512 processing elements, each responsible for one column of the image. Among the operations it can perform are edge detection, filtering, and feature extraction.

Gold Hill Computers (Cambridge, MA) has announced its GoldWorks II expert systems shell for 386-based IBM compatibles, with Macintosh II and Sun versions available later this year. Price is US$7,500.

Millenium Software (Laguna Beach, CA) has written a HyperCard expert system called HyperX 3.0. The system, written entirely in HyperText, is useful as an introduction to expert systems for high school and university students. HyperX 3.0 requires a Macintosh Plus, SE, or II, running HyperCard 1.2 or greater. Price is US$99.95.

Neuron Data (Palo Alto, CA) has released a version of Nexpert Object which allows the use of X-windows server terminals. This allows the development of Nexpert Object applications on any Unix-based computer. Price is US$8,000.

Carnegie Group (Pittsburgh, PA) has enhanced the Knowledge Craft product line, adding a Standard Query Language (SQL) connection which allows it to be interfaced to Oracle databases, and a 386 version of the Knowledge Craft programming languages (CRL, CRL-OPS, and CRL-Prolog). The SQL package is available for the VAX, and a preliminary release is available for the Sun 3.

Texas Instruments (Austin, TX) has introduced new microExplorer systems for the Macintosh IIx (and IIcx) which feature increased memory and new disk drives. Prices range from US$16,645 to US$32,240, depending on memory configurations.

Graeme Hirst: Of magazines, books and natural language research

by Connie Bryson

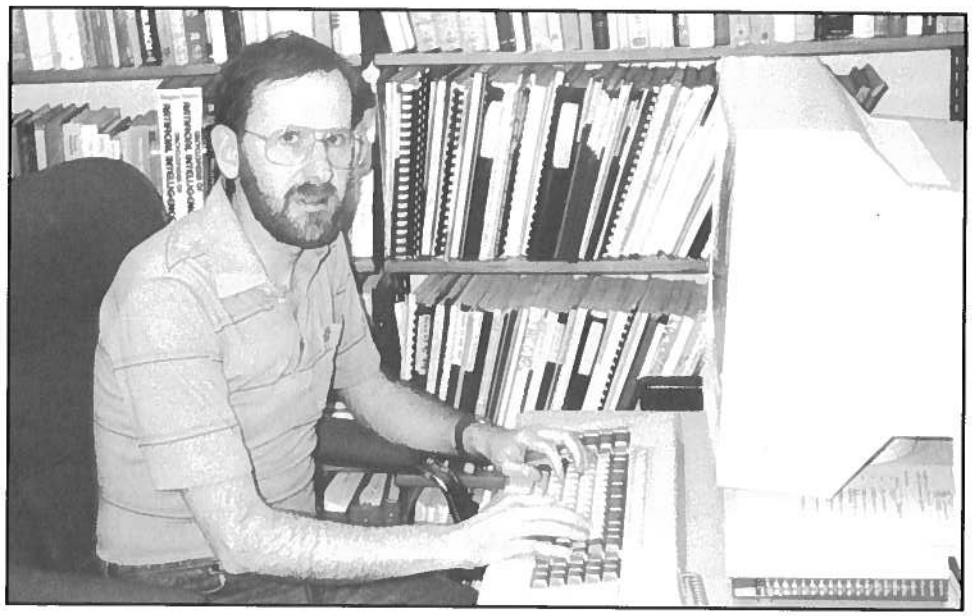

Graeme Hirst : A propos de magazines, de livres et de recherches en langages naturels.

Graeme Hirst a un esprist vif et la réplique facile. Chercheur en compréhension des langages naturels à l'université de Toronto, M. Hirst donne au terme "langages naturels" un sens beaucoup plus large que celui qu'on lui attribut dans les recherches en IA.

Sachant à quoi s'en tenir dans ce domaine, ses opinions et son sens de l'humour quelque peu caustique donnent du piquant à son discours. A titre d'éditeur du magazine portant sur l'intelligence artificielle canadienne de 1984 à 1987, M. Hirst entreprit de transformer un bulletin administratif pour en faire le magazine qu'il est à présent. La première parution du magazine publiée en 1984 contenait 24 pages, on y trouvait beaucoup plus d'articles qu'anticipé.

La deuxieme parution était encore plus volumineuse. M. Hirst fut interviewé par Connie Bryson, écrivain à la pige, au sujet de : son rôle dans la création du magazine, ses recherches dans le domaine des langages naturels et à propos de l'intelligence artificielle au Canada.

Graeme Hirst is rarely stuck for words. A researcher in natural language understanding at the University of Toronto, Hirst fits the description "natural language" in more ways than its AI research connotation. His sharp opinions and wry sense of humour make his conversation sparkle.

As editor of Canadian AI magazine from 1984 to 1987, Hirst was responsible for changing the publication from an administrative newsletter to a magazine. The first magazine issue in 1984 was 24 pages, bigger than expected, and the next issue was bigger still. Freelance writer Connie Bryson interviewed Hirst about his role in creating the magazine, his natural language research, and about AI in Canada.

Canadian AI: Was it a surprise that the magazine took off the way it did?

Hirst: It was a surprise, although I didn't think of it as such at the time, more like good luck than good management. I think the initial publication attracted a lot of people who were curious, they joined CSCSI for a year to see what it was like. At the same time AI was beginning to be fashionable in the commercial world and membership went up dramatically.

Canadian AI: Some say the magazine is a poor cousin to a refereed journal? What's your opinion?

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

Hirst: I don't think you can compare them in that way, because they're not trying to do the same things. A refereed journal is for communicating archival research results. The magazine is a newsmagazine; it has all sorts of stuff you wouldn't put in a journal — what Precarn is up to, what NSERC is up to, new product announcements, and so on.

The magazine was never intended to compete with refereed journals. On a couple of odd occasions I did publish papers that had a bit more of the flavour of a journal, but that wasn't usual.

The magazine *is* a poor cousin in that if you have written something that appears in the magazine it doesn't count for much on your CV compared to a refereed journal publication. But that's not the point. After all, publishing in the *New Yorker* would be highly prestigious in most regards, but it's not going to help your academic resume very much.

Canadian AI: You've stayed on with the magazine as book editor. Is this something you were coerced into doing or is it something you enjoy?

Hirst: I enjoy it and it's also a matter of convenience. I'm also review editor for *Computational Linguistics*. Being book review editor for both has been a distinct advantage. It's surprising how many books I get in my role as the editor of one journal which I end up reviewing in the other journal or on odd occasions in both.

As far as reading books myself, I don't read that many professionally. As a book editor, I send the good books out to reviewers and get to keep only the ones I decide not to review. So I'm in the rather odd position of having a collection of the more uninteresting books on AI and related areas.

Canadian AI: As a book review editor, do you see any trends in AI literature?

Hirst: The biggest trend in AI and computer science in general is the publishing of PhD theses. Usually they're badly laid-out, not adequately revised and can't stand on their own. I have to tread very carefully here, because that's what I did with my own PhD thesis from Brown University a few years ago.

It wasn't a trend then, but there's no question it's a trend now and I don't think it's a particularly good one. I'm embarrassed to be one of the people who have done it — even though I think I can say mine is better than many others.

Canadian AI: Your work with the magazine has also involved you with CSCSI. What is the value of a national AI organization?

Hirst: In general, I tend to be against nationalistic things in scientific research. But the reality is that we have separate countries and separate governments, so I think we need to have organizations that speak for AI in each country.

As long as we get our funding from NSERC, as long as the Canadian government must make policy on AI, as long as there's a need for input in Canadian technology policy — if there is such a thing at all — then we need a Canadian organization to advise the government, because an American or even North American-wide organization is not going to do that.

Canadian AI: Is CSCSI doing this job?

Hirst: I think that you have to look at what else is happening in AI in Canada — initiatives like PRECARN and CIAR. There are now, thanks to earlier work, some people in the right places who know there is AI in Canada, and thanks to CIAR — not just lobbying but putting money into it — we have some visibility. So perhaps at this point we don't have to make noise just for the hell of making noise.

Canadian AI: Do you see any danger to AI in Canada because the high-profile people in AI aren't AI researchers?

Hirst: I think in any field the people who become high profile are not researchers, or they they stop being researchers as soon as they become high profile.

Somebody's got to run the field. As long as the high-profile people can talk intelligently about the field and act in its best interests, it doesn't really matter if they were or weren't once researchers. And perhaps it's the best situation. It saves the AI researchers from having to do this work.

Canadian AI: How did you get into natural language research?

Hirst: I was always interested in the area, it's what I intended to study. I did my undergraduate work in cognitive science, before there was really such a thing as a cognitive science program. I just chose the subjects I was interested in — psychology, philosophy, computer science — and they happen to be what is nowadays called cognitive science.

The first linguistics course I took was at the Australian National University in Canberra. I had planned to do grad work there but found that I really couldn't do what I wanted. So I dumped that and went to UBC where I did my Master's on anaphora in natural language. I took more linguistics courses while I was at Brown University doing my PhD.

Canadian AI: How would you characterize your approach to natural language research?

Hirst: My approach, I suppose, has always been a synthetic one. When looking at problems, I try to look at them on a larger scale. For computer understanding of natural language, you have to understand how humans understand natural language. As a result, one has to be interested in psycholinguistics and theoretical linguistics, and then put it all together.

Putting it all together often means deciding what to leave out — you have to pick and choose.

The analogy I use is building a machine that flies. To do this you can look at how birds fly — they have wings and they flap them. When you build an airplane it turns out to be a good idea to have wings; however it's a rotten idea to make those wings flap.

I think it's the same with computers and language. Some aspects of human understanding are obvious and the right things to copy, others are inappropriate.

The same issue crops up in theoretical linguistics and its bearing on computational linguistics. I teach my students some background in theoretical linguistics but I don't teach them current Chomskyan theory (the 1985 version), I teach them the 1965 version. Why do I do that when Chomsky himself no longer believes that theory? The answer is that the simple approach to natural language is still the better theory. It may not be nearly as adequate for theoretical linguistics, but I argue that that doesn't matter when you're writing computer programs. The older theory is much more appropriate. You can go a lot further and a lot more rapidly with it.

That's what I mean when I say you have to pick and choose.

Canadian AI: What makes a good natural language researcher?

Hirst: There's not one single answer to your question. Perhaps what makes any person good in any line of research is the ability to think intuitively in that particular domain.

Something I think I have is one kind of a feel for language. I say "one kind" because what I have is not the same as what I see in my friends who are theoretical linguists. They also seem to have a certain feel, but it's different.

The funny thing is that there are some very good people in natural language research who seem to have no feel for language at

all. When I say they don't have a feel for language, it's obvious in the papers they write. And yet they do good work.

Canadian AI: Is the small number of natural language researchers in Canada a drawback for your research?

Hirst: That reminds me of something someone said to me not long after I returned to Canada, just after I'd got my PhD. It was "Welcome to Canada. Congratulations, you are now one of the senior computational linguistics people in the country." He was quite right.

The small number isn't a drawback on a day-to-day basis because the field is an international one. If there's a problem it's only because we'd like to get notice in Canada and eventually have the appropriate funding people realize that there is a computational linguistics establishment in Canada that could be working to the benefit of Canada on our special linguistic problems — translation and second-language teaching.

So our task is to get up enough momentum to be taken notice of. And when we're taken notice of we'll have a chance to fulfill that promise. I think graduate students are the key to this and that's one of the reasons I was attracted to the University of Toronto — the opportunity to work with top-quality grad students. Producing a lot of grad students is something I hope will help.

Canadian AI: Are there any areas of natural language research that are neglected in Canada?

Hirst: It's absolutely amazing that Canada is a bilingual country but isn't putting money into machine translation. The government did put money into it in the '70s — the TAUM-METEO project at the Université de Montréal, which was later terminated. For a short time, Canada was a world leader in machine translation.

The decision to cut the program was very short-sighted. Machine translation is going ahead full steam in Japan and Europe. Canada could have been selling the technology to the rest of the world. There's still a chance to revive the effort. Many of the people who were associated with that research are still here.

Canadian AI: You've mentioned a couple of times that AI isn't as fashionable as it once was. Why?

Hirst: Well, that's relatively speaking. When I was on the job circuit in 1983, I got offers from six places, and that wasn't unusual. I think you'd have to be God nowadays to get that many offers.

AI is still fashionable but it is past its peak. I think it's the same way any fad works. The period of hype is followed by a period of non-hype.

Canadian AI: Do you consider the period of hype to be the "heydays" of AI?

Hirst: Personally, I found it frustrating more than anything else. There was so much to keep up with. It used to be that AI was a small field where it was easy to keep up with everything that mattered. Then all of a sudden that wasn't true anymore.

Secondly, there was a lot of hype that had nothing behind it and that was frustrating.

Thirdly, there were a lot of people getting into AI who didn't know where they were going — they just followed the bandwagon. These were people who didn't understand the history of AI. So they just went on repeating the mistakes that had been made before and really making a mess. When AI suddenly became fashionable, some people assumed that it was new field and therefore all ideas were new. They didn't realize that although AI may have been newly "big", people had been doing AI since 1956. A lot of ideas had already been thought up, found to be wrong, and thrown away. Too many people reinvented the wheel because they had no idea of the history of the field. This is still happening today.

Part of the reason for this is that the people who go into computer science tend not to be people who look at the past or think about the past; they think about the future. They don't learn library research skills, they don't learn how to go back and look at past research. This also reflects the fact that some things are hard to find. In the early days, much research was published as technical reports, not in properly refereed journals.

Whatever the reason, too many people don't have a command of the literature — they don't know what's new, what's old, what's been tried, what can be built on top of, and so on.

Now I hope this is starting to go away as people realize AI isn't quite so new anymore. To think that AI was born in 1982 was pretty bad in 1984, but in 1989 maybe it's not quite so bad. At least a seven-year-old field has a "history."

Canadian AI: Is the perception of the novelty of AI research part of the reason many people don't consider AI a science?

Hirst: AI *isn't* a science, because so many people in it haven't the slightest idea of what science is. They may be bright and creative and hack ideas together, but it's not always in a scientific sort of way — a way that's falsifiable, that can be evaluated.

Many of these people don't have training in scientific method because they come from computer science and their background is engineering and mathematics. AI is a strange hybrid area where you can prove theorems, as in mathematics, and you can build things, as in engineering. But that's not all there should be to AI.

Very often the things we build are part of a larger system and it's not clear how to test them. That's one of the problems of the field; it's unfalsifiable. You shoot blindly and whatever you hit, that's what you say the target was.

Sometimes this approach works. If you're the first person working on a problem, or nearly the first, the fact that you've got any solution at all is something. It may not stand up, but at least you've got something to show the rest of the world and it might give others some ideas for what to do next.

But you can take this too far. I think we see too much of the work of people who have gotten carried away with this approach.

Canadian AI: We often focus on the disadvantages of doing research in Canada. Are there any advantages?

Hirst: One of the advantages we have in Canada is the funding of people rather than projects, for the most part. It certainly makes it easier to do curiousity-driven research, because you don't have to produce exactly what you promised if you find later on that a change of direction is advisable. And you don't have to do half the research just to propose funding for it.

It's more possible to do ground-breaking work here than in the US. If you suddenly see a side road that looks good, in the Canadian system you have a lot more freedom to take that road. It's easier to explore, to develop ideas — ideas that will give other people ideas. That characterizes a lot of work by my students lately.

Canadian AI: Does PRECARN and its emphasis on applications represent a threat to more curiousity-based research?

Hirst: It's not a threat unless one kind of research displaces the other. I don't know why it should; the two can and should live happily together.

Canadian AI: One of consequences of the proliferation of AI research is the fragmentation of AI as a field. For someone who takes a synthetic approach, is this a problem?

Hirst: The core of AI is still knowledge representation, and in

some sense it is a service area for the rest of the field. But as the field fragments, knowledge representation is becoming a separate theoretical discipline interested in its own theoretical problems.

One of my hobby-horses is that knowledge representation as a discipline is not interested enough in the practical problems that the rest of AI would like to have tackled. I know there are important theoretical problems in knowledge representation and a lot of interesting things to do there. However I'm not sure the people there are always choosing very well which problems to tackle first.

There's a good example in one area of natural language understanding that I've looked at — understanding sentences that talk about things that don't exist. There's no theory in knowledge representation to deal with these. Take a simple sentence like "Today's lecture is cancelled". In any first-order representation, you get into problems because you have to say something like "There exists a lecture that doesn't exist". With a few exceptions, all knowledge representations based on first- order logic have that same problem. To talk about something is to say it exists, even if you're saying at the same time that it doesn't.

So I've pointed out this neat little problem. What's happening in knowledge representation? Well there are a couple of systems that accommodate it to some degree, more by good luck than by deliberate work, but even those are inadequate in their own way.

I did some initial research in this area while on a short leave at the University of Edinburgh. I found myself doing a fair bit of reading in the philosophy of religion. A lot of the background on the nature of existence comes from questions about the existence of God. The modern approach in first-order logic to representations of existence goes back to Immanuel Kant and his refutation of certain proofs of the existence of God and the logical fallacies therein.

As I was doing this work, I was struck by the fact that reading a book on the philosophy of religion is still doing research in AI.

Canadian AI: Can you look into your crystal ball and see what you'll be doing 10 years from now?

Hirst: One thing I have learned in the last few years is never try to predict more than two years in advance. At one stage I was going to get a PhD, teach for a few years to save up enough money, then dump the whole lot and go on and study another area, probably sociology. That hasn't happened and probably isn't going to. So where will I be in 10 years? Probably in the same place doing the same thing. That's the safest prediction.

Standardizing Programming Languages for Artificial Intelligence

by Dr. E. Whitman Wright and Dr. Stan Szpakowicz

Standardisation des langages de programmation en intelligence artificielle.

Le développement de la normalisation des langages de programmation en IA est une discipline très exigeante. Faire notre part dans ce domaine est primordial afin de continuer à assurer au Canada sa place parmi les nations développées de ce monde.

Le comité international pour la normalisation des langages de programmation a crée, en 1987, de nouveaux groupes de travail se concentrant sur les langages Lisp et Prolog. Il a aussi distribué des lettres consultatives et par la suite, recueilli les avis et suggestions nécessaires afin de promouvoir des activités de normalisation pour les langages C++ et Smalltalk.

A l'heure actuelle un des comités de la ASC sur la programmation des langages a déjà formé un groupe intèressé au langage Lisp. Ce comité espère aussi être capable de former un groupe de travail international qui pourra participer aux discussions du groupe de travail de l'OSI pour le langage Lisp. Le groupe de travail canadien pour le langage Prolog a été créé tard en 1986.

Si vous êtes interessés à vous impliquer et que vous avez des connaissances dans ce domaine, votre travail avec le comité de programmation de langages de la ASC et avec ses différents groupes de travail, serait pour vous une occasion sans pareil de vous trouvez dans le feu de l'action et d'apporter votre propre contribution au projet.

Introduction

The spectacular technical advances of recent years have brought with them ever increasing complexity as an unavoidable by-product. As a result, society has become increasingly dependant upon standards in order to be assured of quality and freedom of choice in the products exchanged and consumed. Standards provide a guide for producers and an important means of protection for consumers. Nowhere is this more true than in the data processing industry.

To respond to the need for good standards, an elaborate system of standardization has been built up over the years. In Canada, the Standards Council of Canada (SCC) is a Crown Corporation reporting to Parliament through the Minister of Consumer and Corporate Affairs. It is independent of government in its policies and operations, although it is financed by Parliament appropriation.

Dr. Wright is the president of W. Wright & Associates. Dr. Szpakowicz is an associate professor in Computer Science at the University of Ottawa.

Five accredited Standards-Writing Organizations, one of which is the Canadian Standards Association (CSA), work with the SCC. The CSA is a non-profit, non-governmental association of technical committees, with a Board of Directors elected by the membership. For the national standardization activities of the CSA, the senior body is the CSA Standards Policy Board. Most of the industrialized nations of the world now have similar standards bodies that participate in developing international standards (e.g. in the U.S.A., ANSI; in the U.K., BSI; in France, AFNOR etc.).

The two principle international standards organizations in the field of Information Technology, the International Electrotechnical Commission (IEC) and the Organization for International Standardization (ISO), have formed a joint technical work is undertaken by JTCI sub-committees and working groups, subject to the authority of JTCI, acting under the general authority of IEC and ISO Councils.

In Canada, the CSA Steering Committee on Information Technology (SCIT) operates under the authority of the CSA Standards Policy Board and is responsible for the Technical Committees that are concerned with developing standards for information processing systems. The members of SCIT also serve as the Canadian Advisory Committee to the SCC for ISO/IEC Joint Technical Committee 1 (CAC/ISO/IEC JTC 1). The Committee on Programming Languages (CPL) is one of the subcommittees under SCIT.

The work of the Committee on Programming is subdivided and delegated to Working Groups. The members of the CPL and its Working Groups provide their services on a volunteer basis. Travel to committee meetings is supported by their employers. The CPL meets only four times a year but the Working Groups may organize additional sessions. The work load is substantial; members are expected to attend meetings regularly, review and prepare drafts of standards documents, and participate in informed discussion. When called upon, they represent Canada at international meetings dealing with programming languages. The activities of the Working Groups are outlined in greater detail later in this document.

Since computer programs are frequently transferred across international boundaries, an individual country seldom develops a standard for a programming language in isolation. At present, Canadian practice is not to develop such standards unilaterally but rather to co-operate with other member countries of ISO. To date, Canada has adopted as National Standards of Canada the IDO standards for Fortran, Cobol, Minimal Basic and Pascal and has endorsed the ISO standard for PL/I and Ada.

New Trends in the Work of the Programming Languages Committee

During its earlier years, the CSA Computer Programming Language Committee gave its main attention to the procedural programming languages such as Fortran, Cobol, APL and Ada. These languages are still very important for computer software developers and users. However the last decade has seen a rapidly growing acceptance of knowledge-based systems, expert systems and other software products related to artificial intelligence (AI). It is an important growth area in computer use, with no end in sight. The CPL has responded by shifting part of its attention to programming languages suitable for supporting artificial intelligence applications, in particular Lisp and Prolog.

Conventional programming languages, such as Fortran, Cobol and Pascal, assume a computer hardware configuration close to the traditional von Neumann machine architecture. This architecture consists of a single memory that contains both data and instructions and a control unit for fetching these instructions to a processing unit (CPU) where they are executed according to a predetermined sequence. The user is required to translate his/her needs into a language that is an abstraction of the hardware architecture.

Some of the user needs (e.g. manipulation of symbolic information) are hard to express in terms of conventional languages. Early attempts at introducing alternatives to these imperative (statement-oriented) languages were not accepted by the computer industry until research in the field of artificial intelligence netted practical results that could be translated into financial gains by its users. The emerging popularity of some of these languages and the diversity of their uses incited international work. ISO/JTCI/SC22, the international body for standardizing programming languages, created new working groups for Lisp and Prolog in 1987 and has issued consultative letter ballots on proposals to initiate standardization activities on C++ and Smalltalk.

LISP
Brief History of the Language

Lisp (for List Processing) was developed at the Massachusetts Institute of Technology (MIT) in the late 1950's under the direction of John McCarthy, a pioneer in the field of AI.

Lisp is the first and best known example of functional programming languages. Computation in Lisp consists of evaluating functions; its fundamental mechanisms are conditional schema and recursion. Programs and data are uniformly represented by "atoms" (elementary symbolic data) and nested lists built of atoms and lists. The original Lisp is a declaration-free and typeless language of exceptional expressive power and flexibility, capable of self-modification. It can be defined in itself, thus allowing language extensibility unparalleled in other languages of a similar age. No other language is routinely accompanied by such rich programming environments with interactively used interpreters, specialized editors, and symbolic debugging. Lisp encourages a declarative style of program design to a much greater extent than procedural languages such as Pascal or Ada. Because of these properties, Lisp has been quickly adopted by programming language specialists and by researchers working in AI.

The first implementation, Lisp 1.5, completed in 1960 at MIT, was elegant and extensible, but very inefficient, especially for numeric computations. It was and still is primarily a programming language for symbolic applications. The average user in the sixties identified computing with numerical problems and did not appreciate the potential of symbolic computations and associative memory organization. Even truly efficient implementations, combining the interpretive use of Lisp with incremental compilation, were for a long time popular only in the university world. The uniform but not particularly readable syntax did not help promote Lisp in the computing community at large.

Around 1978, AI applications brought Lisp to the attention of large commercial and military users. After several years of Lisp software development by a wide range of users, it became desirable to control the proliferation of Lisp dialects. In 1981, the U.S. Department of Defense encouraged an effort by a number of

researchers from several organizations in the U.S. to define a common acceptable version of Lisp. This group, headed by Guy Steele of Carnegie-Mellon University, consisted of more than 75 experts. In 1984 Digital Press published "COMMON Lisp, the Language" by Guy L. Steele. This document is now the formal definition of the Lisp dialect called "Common Lisp", a de-facto standard in the U.S. However, many other dialects (e.g. Scheme and InterLisp) continue to flourish and gain disciples.

Standardization Activities

In September, 1987 "Resolution 69: Establishment of Working Group 16 on Lisp" was passed at the Second Plenary meeting of ISO/JTC1/SC22 with the aim of developing an ISO Lisp standard. Christian Queinnec, AFNOR (France) was nominated convenor, Richard Gabriel and William Klinger, ANSI (U.S.) were appointed project editors. France, Great Britain, U.S.A. and Japan all have active committees for the standardization of Lisp and send delegates to international meetings.

The national bodies differ in the goals which they wish to achieve through an ISO Lisp standard. Since ANSI hopes to have a standard for Common Lisp and IEEE is in the process of standardizing Scheme, it is only natural that the U.S.A. would like the ISO Lisp working group to start by standardizing either Common Lisp (for commercial users) or Scheme (for the "academic" users). Great Britain, supported by France, Japan and West Germany, is looking for one ISO Lisp standard (different from Common Lisp and Scheme) that can serve commercial users as well as research people. That could mean that the U.S.A. would have three different dialects: Common Lisp, Scheme and ISO Lisp. Discussions to resolve these issues are still in progress.

Presently the CSA/CPL has a Lisp Interest Group and hopes to form a national working group in order to participate in the discussions of the ISO Lisp working group.

PROLOG
Brief History of the Language

Prolog is a common name for a rather large family of closely related languages which have their roots in logic programming. A product of European scholars, logic programming has gained wide popularity in the early eighties when it was chosen as one of three pillars of the Japanese Fifth Generation project. Logic programs make it possible to separate logic and control more clearly than classical procedural languages, and they have well defined semantics based on first-order predicate logic. Prolog is a practical embodiment of logic programming. With its powerful underlying mechanisms of unification, recursion and backtracking, it makes a very concise and expressive tool for and symbolic processing application.

The history of Prolog begins in 1972 in Marseille with the team of Alain Colmerauer. The first fully operational and somewhat widely distributed implementation of Prolog was completed in 1973. By 1976 diversification has already begun. During the next few years separate projects in Scotland, England, Hungary, Portugal, Canada and Poland led to the creation of several syntactically and semantically non-identical dialects.

Two landmarks that greatly contributed to the future success of Prolog were the construction (by David H.D. Warren) of the first efficient compiler-cum-interpreter in 1977 in Edinburgh, and the publication in 1981 of the first textbook (by W. Clocksin and C. Mellish) that covered a subset of Edinburgh Prolog. Five years later Prolog already was a fact of life in the computing community at large. At present, most universities teach Prolog, applications requiring tens of thousands of lines of Prolog code are not uncommon, and people working in artificial intelligence (first of all in Europe and Japan, but also in North America) routinely use Prolog in their research and implementations.

Most Prolog dialects available commercially and semi-commercially today are offsprings of the 1977 Edinburgh implementation. Many of them were even produced entirely on the basis of the incomplete and not quite consistent dialect presented by Clocksin and Mellish, and never intended as anything more than a primer. But numerous dialects are significantly different in their syntax and set of built-in operations and there are subtle differences in the semantics of the so-called extralogical features of the language.

Other Prolog or Prolog-like systems have been developed that can no longer be considered dialects. Some have been redesigned for concurrent programming as the main area of application (Concurrent Prolog, Parlog). Many logic programming languages proposed by researchers go beyond Prolog's self-imposed limitations that allow efficient execution. One could name languages which embody constraint logic programming, or logic-based symbolic equation solvers.

Even among the more down-to-earth Prolog dialects the variety is astounding. Commercial products from the USA, Canada, Great Britain, France, West Germany, Belgium, Japan differ in so many large and small ways that the user may be practically tied to one developer. A standard is necessary as soon as possible. It should perhaps disregard most of the ongoing research and concentrate on a small common foundation which still can be identified in most Prolog systems available on the market.

Standardization Activities

In its origin, Prolog was simply one aspect of a research activity without international recognition. This has made for a considerable and very undesirable diversity of implementations following the "Fifth Generation challenge" of late 1981. The book by Clocksin and Mellish was the only source of knowledge for many inexperienced developers rushing to fill the obvious gap in the market, especially in the USA. This created a fait accompli with effects that are only now becoming visible: there exists a "de-facto Clocksin and Mellish standard" which in fact is only a sub-standard but which is used as yardstick by the unsuspecting user community. Even worse, the best selling Prolog system for small computers, Turbo Prolog, is an unfortunate, rigid sublanguage with limitations that take away most of the beauty of full Prolog.

The international standardization effort for Prolog began with the BSI (British Standards Institute) setting up a panel in late 1984. Two major Prolog dialects had been developed in Britain: Edinburgh Prolog and microPROLOG. The panel set itself the goal of assessing the chances of a Prolog standard based on a blend of these two dialects. As Edinburgh-like dialects were already by far the most widespread, it was probably a mistake not to try to push for one of these available dialects or some skillful combination of them to be

recognized as the first (quick but badly needed) standard.

Instead, the BSI panel embarked on a three-year trek, trying to invent a new language that would be better than all the existing dialects taken together. The gaps in competence were less damaging than the lack of any clear mandate and direction. The work was also purely voluntary, and accordingly it proceeded slowly.

One achievement of the BSI group was the tentative division of tasks between Prolog syntax, semantics, built-in operations and programming environments. Still, the mutual relationship between these broad topics made it difficult to proceed at any reasonable speed.

In 1985 a parallel effort was launched by AFNOR (the French national standards body). Ironically, the latest version of Marseille Prolog is exceptionally well designed and motivated; making it a standard would not be difficult. But it is not an Edinburgh-like dialect, and it would be unacceptable to the North American user community, already committed to the dialects that were a more or less direct continuation of the original Edinburgh Prolog system.

The AFNOR group co-operated with the BSI panel in several technical areas, and also in the attempt to initiate an international standardization activity. Their joint effort resulted in 1987 in the creation of the ISO working group WG17, with a rather vague specification of its mandate. The secretariat of the ISO panel was entrusted to the BSI group, and its decisions ended up as the "seed" ideas for WG17. The British together with the French have produced over 200 documents, notes and working papers, and finally worked out a list of over 50 standardization issues, a tentative description of Prolog syntax and a reasonably acceptable description of its semantics.

The first two meetings of WG17, in 1988, were dominated by delegations from the BSI, AFNOR and DIN (the West German national standards body). The net result was a consolidated document that apparently will be presented to the third meeting in 1989. This document contains many contentious issues, because there remain unresolved questions about the syntax, the set of basic types, the method of modularizing Prolog programs and the semantics of some very fundamental built-in operations. The DIN group is pushing for a very quick standard. The remaining experts do not seem to support the DIN approach. All in all, the actual productivity of the BSI and AFNOR expert group has been rather low, and the ISO group has barely started. There should be enough time to make decisions and review them before submitting a draft for serious consideration.

The Canadian Prolog working group was created late in 1986 in response to the first attempt at organizing an ISO group or groups for "AI languages", this is to say, Lisp and Prolog. It has been carefully reading all the documents circulated by BSI/AFNOR experts.

Since WG17 was officially approved, the CSA working group has completed a comprehensive analysis of all the major issues. It prepared a position for the first WG17 meeting in March 1988 and to the second meeting it brought several documents outlining its concerns as to one major and a few minor issues. It will be hosting the fourth meeting in Ottawa, in the hope of bringing the ISO sponsored activity close to the large and important user community in the USA, which is not represented in WG17.

Smalltalk and C++

Smalltalk and C++ are the two most popular object-oriented programming languages. C++ is an object-oriented extension of C, developed by Bjarne Stroustrup at Bell Laboratories in Murray Hill, New Jersey. Smalltalk was designed by Alan Kay while working on a research project to develop a flexible, easy-to-use but powerful environment for personal computers. The final product was Smalltalk-80, which is the current Smalltalk version licensed by Xerox.

In procedural programming languages, the central notion is that of a procedure which operates on passive objects - data structures. Object-oriented programming emphasizes "objects" rather than operations. An object is an active entity that has methods of dealing with messages from other objects. A response to a message means executing code owned by the recipient and issuing an answer-back message to the sender. Problem definition in an object-oriented language consists of descriptions of relevant objects, complete with all applicable operations; computation is a set of message exchanges.

Objects encapsulate data and operations in a manner compatible with current trends in programming languages; an object is a natural unit of data abstraction and modularity. Objects are grouped into a hierarchical system of classes which share properties and methods. A subclass inherits some properties and methods of its superclass, and may add its own. Finally, every object may have its own local properties.

Both Smalltalk and C++ are closer to procedural languages than to functional or logic programming. Interestingly, object-oriented extensions of Lisp and Prolog are being proposed, as well as Prolog-like extensions of Smalltalk. Although they are very promising, it is too early to consider their standardization now. In the long run, however, this seems to be the correct trend in the search of better tools for AI programming.

Summary

The development of programming language standards for artificial intelligence is a demanding discipline, capable of taxing the capabilities of dedicated experts. However, it is part of the forefront of intellectual progress at the present time. If Canada is to retain its place among the developed nations of the world, it is important that we do our share in this area.

For those who have the interests and the background to make a contribution, work with CSA Programming Languages Committee and its working groups offers a rare opportunity at the same time to be where the action is and to make one's personal contribution.

For committee members, CSA does not pay per diem or travel expenses (except occasional bare bones travel expense for members on special assignments). On the other hand, all that is usually required for membership in the working groups, and in due course the CPL itself, is one's personal qualifications and commitment. From any person who believes he or she would like to explore possible membership in one of the working groups of the CPL, we invite inquiries. Please write to:

Mr. Anton Bickle
Chairman, CSA Committee on Programming Languages
Environment Canada
Place Vincent Massey, 5th Floor
Ottawa, K1A 0H3

Debugging Super-Intelligent Machines

by Peter Turney

Retrait d'erreurs pour machines super-intelligentes.

Supposons qu'il soit possible, un jour, de créer des machines super-intelligentes (des MSI). Le jour où une telle machine sera créeé est-ce que notre travail comme chercheurs en IA sera terminé ou y aura-t-il de nouveaux problèmes sur lesquels travailler?

En fait, il est très probable que la première MSI sera pleine d'erreurs "bugs". Ces erreurs pourrons sans conteste être resolues par des chercheurs en IA. Car en dépit de leur super-intelligence, les MSI serons incapabables de résoudre leurs problèmes sans notre aide.

Dans le paragraphe suivant, je vais donner quelques exemples de défaut de programmation.

On doit considérer la possibilité que les MSI n'aient pas besoin de l'aide des chercheurs. Toutefois ceci n'est pas une certitude et on est même porté à croire le contraire. Nous devons donc nous préparer pour ce que nous appellerons la recherche post - MSI faite par des chercheurs humains.

Introduction

Let us say that a being is "super-intelligent" (SI) when it is significantly more intelligent than members of the species *homo sapiens*. Let us assume that we will, one day, create super-intelligent machines (SIM's). Suspend disbelief for a moment, as researchers in AI, be complete when we create a SIM? Or will there be further problems for us to work on?

At first glance, it might seem that a SIM could solve any problems of concern to AI research, without the help of human AI researchers. In fact, this is not necessarily true. It is likely that the first SIM's will be full of bugs. These bugs, I will argue, will be most suitable for treatment by human AI researchers. In spite of their super-intelligence, SIM's will not be able to deal with these bugs without our help.

What kind of bugs can we expect a first-generation SIM to have? We can guess at answers to this by looking at some of the bugs that human beings have. Marvin Minsky has conveniently provided us with a list [4]:

- obsessive preoccupation with inappropriate goals.
- inattention and inability to concentrate.
- bad representations.
- excessively broad or narrow generalizations.
- excessive accumulation of useless information.
- superstition; defective credit assignment schema.
- unrealistic cost/benefit analyses.
- unbalanced, fanatical search strategies.
- formation of defective categorizations.
- inability to deal with exceptions to rules.
- improper staging of development, or living in the past.
- unwillingness to acknowledge loss.
- depression or maniacal optimism.
- excessive confusion from cross-coupling.

Peter Turney is a visiting fellow at the Laboratory for Intelligent Systems, in the Division of Electrical Engineering of the National Research Council. He is currently involved in research on machine learning. He has a Ph.D. in Philosophy from the University of Toronto.

Although we have yet to solve these bugs, we have had thousands of years experience with them. We have at least partial solutions, which enable us to live with our bugs. This gives us an immense advantage over any first-generation SIM, which confronts these problems for the first time.

As AI researchers, we are particularly interested in bugs which, by their very nature, SIM's will have difficulty repairing, without our help. In the following, I will discuss a few examples of such bugs.

Suicide

There is evidence that suicide is correlated with intelligence [5]. Conceivably reality is too painful to tolerate for a sufficiently intelligent being. Indeed, one could argue that humanity has reached this limit, and that we are headed towards nuclear or environmental mass suicide.

There may be a way around this limitation, through what we might call "intelligence with blinkers". This would involve special subroutines that give selective blindness to intelligent beings, to protect them from the harshness of reality. Perhaps humans have such subroutines: It would explain our ability to ignore social injustice. A major task of post-SIM research might be the design of "blinker subroutines". This task could not be delegated to SIM's, because they would be too torn with depression and existential angst to concentrate on the task.

Split Personality

Perhaps larger minds tend to spontaneously split into smaller minds. A SIM that passes a certain level of intelligence might develop a split personality. This appears to take place in William Gibson's book, *Count Zero* [3].

Imagine that there is a critical intelligence threshold, beyond which fragmentation is inevitable. We might hold a SIM just below the threshold, by constantly erasing its memory, so that it never learns enough to pass the threshold.

Post-SIM AI research may be concerned with developing routines that hold personalities together; that prevent personalities from fragmenting. SIM's themselves may lack the integrity required to pursue this task.

Survival

Evolution selects for the ability to survive. We are perhaps mistaken in thinking that increased intelligence yields increased ability to survive. A dinosaur, asked what his evolutionary successor would be like, might guess that the next step in evolution would be bigger teeth and sharper claws. The next step in survival ability may have nothing to do with intelligence. Conceivably the next step is an increased sense of community, like intelligent social insects.

Intelligence, increased past a certain amount, could tend to lead an entity towards death. This is related to my speculations about suicide, since suicide is not a pro-survival characteristic. But there are other ways in which intelligence may be anti-survival. A SIM need not be suicidal, but merely negligent about survival. Srinivasa Ramanujan, a mathematician of great genius, loved mathematics so much that he would neglect sleep and food, which contributed to his early (and childless) death. It is easy to imagine that a SIM could become so involved in the beauty of mathematics that it would let its batteries run dry.

SIM's might be too morally good to survive. Or perhaps SI is incompatible with strong instincts, such as an instinct for survival. Our task, as researchers in AI, would be to program SIM's to survive. This task may have no appeal to SIM's. Indeed, they may actively resist such programming.

Our advantage over SIM's is more than three billion years of biological evolution. We are born with subroutines for survival which have evolved for millennia. A first generation SIM will lack these subroutines, and it will also lack any motivation for acquiring these subroutines. Only humans can provide SIM's with these subroutines.

Culture

With people, intelligence is a product of nature (genetic endowment) and nurture (social environment). Even if we can make a SIM, we may have difficulty raising the machine to its full potential.

There are instances of human babies being raised outside of human society, by wolves in one case. This results in permanent mental and emotional damage [1]. A SIM, raised by humans, would be in an analogous situation. However, this problem does not seem insurmountable. It suggests that we should make each generation of SIM's only slightly more intelligent than the previous generation, so that each generation has the community of the preceding generation to nurture it.

Much of our recent evolution seems to be cultural, not biological. SI might be dangerous to culture, say by leading to excessive individualism. (Note that I mean culture in general, not necessarily human culture.) SIM's who cannot form a society seem, intuitively, to be more limited than SIM's who can form a society. Post-SIM research might concentrate on making SI compatible with culture.

Super-Intelligent Machines and People

There will be some difficulties involved in the relationship between SIM's and people. AI research by humans will certainly have a role to play in the solution of these difficulties. Consider the items on Minsky's list of bugs. For example, "obsessive preoccupation with inappropriate goals". Who is to say that a SIM is spending too much time on a given goal? We are, of course!

I have, so far, ignored the possibility of Augmented Intelligence. It may be possible to increase human intelligence by fusing human brains with artificially intelligent machines. This blurs any possible distinction between SIM's and humans. It gives machines the advantage of subroutines created by biological evolution, and it gives brains the advantage of increased intelligence. It is particularly appealing to human AI researchers, because it gives us a role in post-SIM AI research, beyond merely debugging first-generation SIM's.

Conclusion

It may be that SIM's will not need the help of human AI researchers. However, this is not certain, and there are reasons to believe otherwise. We should prepare now for post-SIM AI research by human researchers. SIM's will be a mixed blessing. As Minsky warns us [2]: "I think that the species is just a step in evolution. I'm sure the chimpanzees thought people were a bad thing."

Acknowledgements

Thanks to Louise Linney and Ken Deaton for constructive criticism.

References

[1] de Villiers, J.G. and de Villiers, P.A., *Language Acquisition*, Harvard University Press. Cambridge, Massachusetts, 1978. Pages 214-219.

[2] Fjermedal, G., *The Tomorrow Makers*, Tempus Books. Redmond, Washington, 1986. Page 139.

[3] Gibson, W., *Count Zero*, Arbor House. New York, New York, 1986.

[4] Minsky, M., "Preface: Connectionist Models and Their Prospects", in *Connectionist Models and Their Implications*, edited by Waltz, D. and Feldman, J.A. Ablex Publishing. Norwood, New Jersey, 1988. Pages vii-xvi.

[5] Terman, L.M. and Oden, M.H., *The Gifted Group at Mid-Life*, Stanford University Press. Stanford, California, 1959. Pages 30-32.

Artificial Intelligence Research at the University of Waterloo

Edited by Peter van Beek

Recherches en intelligence artificielle à l'Université de Waterloo.

Le 'Logic Programming and Artificial Intelligence Group' (Groupe de programmation logique et d'intelligence artificielle) se trouve au département d'informatique de l'Université de Waterloo. Le personnel comprend quatre membres de la faculté (Fahiem Bacchus,Robin Cohen,Chrysanne Dimarco (à compter de janvier 1990) et Qiang Yang), un gestionnaire de laboratoire (John Sellens), cinq étudiants au doctorat (Fei Song,Bruce Spencer,André Trudel,Paul van Arragon et Peter van Beek), et un étudiant de maitrise (David Hsu). D'autres étudiants gradués sont attendus pour le mois de septembre 1989. L'équipement comprend plusieurs ordinateurs VAX de DEC et un VAX 8600 partagé. Nous présentons plus bas un aperçu des intérêts et travaux récents du groupe: la modélisation de l'usager; le traitement du discours; la traduction automatique; le raisonnement du sens commun; la planification; la recherche; et le raisonnement temporel.

The Logic Programming and Artificial Intelligence Group is in the Department of Computer Science at the University of Waterloo. The group has four faculty members (Fahiem Bacchus, Robin Cohen, Chrysanne DiMarco (as of January, 1990), and Qiang Yang), one lab manager (John Sellens), five Ph.D. students (Fei Song, Bruce Spencer, André Trudel, Paul van Arragon, and Peter van Beek), one Masters student (David Hsu), and many DEC VAXstations (aristotle ... wattle) and a share of a VAX 8600. New graduate students are expected in September of 1989. Below we give an overview of the interests and recent work of the group.

User Modeling

It has been shown by many that having an explicit model of a user's knowledge, beliefs, background, goals, and reasoning process can greatly improve the communicative and interactive aspects of an intelligent system.

Paul van Arragon is investigating user modeling with a formal representation language [14, 15]. His approach advocates the use of nested reasoning where the system reasons by default about an agent who is reasoning by default. He begins with the default reasoning system Theorist, and shows how to define Theorist in Theorist to provide nested reasoning. The resulting framework also provides a flexible way to reason about agents who are limited in their reasoning abilities by assuming an agent can make an inference unless the system knows otherwise. Nested Theorist is implemented by having a Theorist meta-interpreter as the data of another Theorist interpreter.

Robin Cohen and Bruce Spencer (with Marlene Jones, Amar Sanmugasunderam, and Lisa Dent) have worked on the question of how to exploit a user model to generate user-specific explanations within an expert system [5, 7]. The research has resulted in proposals for how a response generating system can take into account both the user's goals and background knowledge in generating a response to the user's query as well as how to initialize and maintain the user model by starting with an initial classification of the user according to the user's goals and background and updating the user model depending on changes in the user's knowledge. An implementation in Prolog helps illustrate the proposed techniques. At present, Robin Cohen is examining the role of user modeling for the analysis of input and the generation of feedback to users in an interface for a knowledge representation language.

Discourse Processing

Discourse processing concerns itself with understanding and generating coherent multi-sentential texts as opposed to single sentences in isolation.

Robin Cohen is continuing a research project on developing processing models for the restricted form of discourse known as arguments [6]. A primary concern of this research is the characterization of coherent discourse and the results should influence the general problem of designing effective interfaces. Current work includes building a prototype implementation in Prolog of her argument understanding model (with help from Mark Young and Fei Song) [8]. The output of the system is a representation of the structure of the speaker's argument that shows the claims of the argument and the evidence presented for those claims. The implementation exercise has resulted in insights for the specification of a discourse analysis model and for the general problem of recognizing a speaker's plans.

David Hsu is investigating the discourse model of Grosz and Sidner. The research involves turning their descriptive model into a computational model. In particular, he is looking at the role of redirection cues in dialogue (an example of a redirection cue is "returning to what I was saying"). These cues help determine the structure of the dialogue.

Fei Song and Robin Cohen are continuing work on an algorithm for determining the temporal relations between events mentioned in a narrative (such as, which event occurs before another) [11]. Their algorithm makes use of anaphoric references to events, tense and aspect information, and a novel structure called a situational description that includes the participants, place, and duration. The events and the temporal relations between the events are represented

Peter van Beek is a Ph.D. student in the Logic Programming and Artificial Intelligence Group at the University of Waterloo

in a tree-like structure that allows some relations to be left vaguely specified, possibly to be clarified as the algorithm progresses. The algorithm provides a processing model for temporal focus. An additional contribution of this work is a more complete semantic interpretation for the tense of sentences that can express all 16 tenses in English and incorporates the interval properties of events from the aspectual analysis.

Plan Recognition

Somewhat obviously, plan recognition is the recognition of an agent's plan from observation of the agent's actions. The agent's actions could be physical actions or speech actions (so, for example, one application of plan recognition is in question answering where recognizing the plan underlying a user's queries aids in the generation of an appropriate response).

Fei Song is investigating plan recognition. In most previous work in plan recognition the system proposes a plan from a set of prestored plans that explains all the observations of the agent's actions. However, there are likely to be discrepancies between this proposed plan and the agent's real plan. For plan recognition in an interactive setting, Song feels that discourse context, especially the temporal relations between events, can be used to help identify the differences. Thus, an agent's plan should be decided by considering both the linguistic analysis of the discourse and the proposed prestored plan.

Bruce Spencer is working on relaxing the restrictive assumption made in previous work in plan recognition that the system's knowledge is complete [12]. Relaxing this assumption means the system needs the ability to assimilate new knowledge. If this assimilation occurs when there is an existing set of beliefs about the agent's plan, that set should not be simply discarded, as that would be inefficient, but should be revised to accommodate the new knowledge. This is an instance of the general problem of belief revision. Spencer's starting point for a belief revision system is an assumption based truth maintenance system (ATMS). The ATMS is being extended to have the ability to calculate only the preferred plans rather than all plans, and to handle non-Horn clauses, variables, and equality. A second problem that arises from relaxing the completeness assumption is, how does the plan recognizer realize when its incomplete knowledge matters to the recognition of a plan and, if the recognizer is in an interactive setting, what questions does it ask to fill in this missing knowledge. Spencer is looking at these problems in two formal theories of plan recognition: Henry Kautz's minimal covering model and an abductive approach based on Theorist, a system for default reasoning.

Machine Translation

Chrysanne DiMarco is researching stylistics in machine translation [9, 10]. Her work addresses the failure of current machine translation systems to handle the translation of language style, which communicates a good deal of the meaning of the source text. The translation of style involves two complementary and sometimes conflicting aims: producing a style appropriate to the particular target language, while preserving the original author's stylistic intent. These aims require an understanding of the internal stylistics of both the source and target languages as well as the comparative stylistics of the language pair. DiMarco has developed a formal representation of French and English stylistic knowledge in the form of *stylistic grammars* and implemented a *stylistic parser* in Prolog. The stylistic grammar is a branching stratificational model, built upon a foundation dealing with syntactic and semantic stylistic realizations. Its central level uses a vocabulary of abstract stylistic elements common to both English and French, while the top level correlates stylistic goals, such as clarity and concreteness, with patterns of these elements. As well, the foundation has been laid for the codification of the comparative stylistics of French and English. The incorporation of stylistic analysis into machine translation systems should significantly reduce the current reliance on human post-editing and improve the quality of the system output.

Common Sense Reasoning

Fahiem Bacchus has developed logics for representing and reasoning with probabilistic knowledge [1, 4]. The logics can handle exact numerical information, but also, since the probabilities can be directly referred to in the logics, qualitative information about the values of probabilities, thus, countering claims that probabilities are epistemologically inadequate.

Most work on probabilities in AI has concentrated on probabilities over propositions. But propositional languages lack the expressiveness required for general knowledge representation tasks. This work investigates first-order probability logics. Two different types of probability logics have been examined.

The first type of logic has a probability distribution over the set of possible worlds. Possible world probability logics have been developed by various people, with some important recent work by J.Y. Halpern. They can be used to represent probabilities of assertions about specific individuals, that is, an agent's degree of belief in that assertion. For example, one can represent "I believe that Tweety can fly to degree greater than 0.75".

The second type of logic has a probability distribution over the domain of discourse. Domain of discourse probability logics were first studied for their use in AI by Bacchus in his Ph.D. thesis. They can be used to represent probabilities of assertions about sets of individuals or statistical assertions. For example, one can represent "More than 75% of all birds can fly".

These two types of probability logics can be combined to produce a logic that can express statistical assertions as well as an agent's degree of belief in assertions about specific individuals. The combination is done in such a way that the agent's degrees of belief are founded, or induced, from the statistical facts. This produces degrees of belief that are objectively founded in properties external to the agent. The combination can be used as a powerful system of default reasoning. In this system defaults, or empirical generalizations, are represented as qualitative statistical assertions. And the default conclusions are represented as induced degrees of belief in these conclusions. The system has several advantages over standard non-monotonic systems for default reasoning, including having an interpretative semantics for the defaults and producing a graded support for the default conclusions.

Fahiem Bacchus is also interested in inheritance reasoners that allow exceptions. He has developed a reasoner that allows restricted, but semantically well founded, defeasible property inheritance [2]. The reasoner gives a well defined and easily understood semantic interpretation to all the assertions encoded in it, based on qualitative

statistical relationships between the property classes. The semantics allows a knowledge engineer to decide what knowledge can be encoded in the system and gives understandable formal guarantees about the quality of the conclusions that will be generated. For this reason the system is a more practical, usable inheritance reasoner than others that have appeared in the literature. The system has been fully implemented in a short (less than 75 lines) Prolog program. Although the system performs a restricted form of inheritance reasoning it can still represent and solve most of the inheritance "puzzles" that have appeared in the literature, including the recent heterogeneous inheritance problems.

Planning

Qiang Yang is interested in improving the efficiency of planning systems [18, 19]. He has explored two alternative approaches.

The first approach is to preprocess the planning knowledge. Yang has found that useful information about conflicts between actions can be extracted by preprocessing the planning knowledge. In particular, a set of syntactic restrictions can be imposed on the relationship between a non-primitive action and its set of subactions. Because of the static nature of the restrictions, it is possible to preprocess a given set of reduction schemata to check which schema satisfy the restrictions before the planning process starts. This allows computational savings at planning time. For example, it allows a hierarchical planner to identify early and backtrack from certain dead end points in the search for a plan.

The second approach is to abstract the kinds of goal and subgoal interactions that occur in some set of related problem domains, and develop planning techniques capable of performing efficiently in all domains in which no other kinds of interactions occur. This approach was applied to a particular formulation of multiple-goal planning. Yang showed that for cases where multiple-goal planning can be performed by generating separate plans for each goal independently and then optimizing the conjunction, restrictions can be imposed on the allowable interactions between goals that allow efficient planning to occur when the restrictions hold. These restrictions are satisfied across a significant class of planning domains (one example is process planning in computer-aided manufacturing). Yang has developed algorithms that are efficient for certain cases of this multiple-goal planning approach and, where these are not applicable, a heuristic search algorithm that performs well.

Search

Qiang Yang is also interested in search [20]. Heuristic search procedures are useful in many problems of practical importance. Such procedures operate by searching several paths in a search space at the same time, expanding some paths more quickly than others depending on which paths look most promising. Often, however, much time is required to keep track of the control knowledge. Yang (with Dana Nau) has devised a strategy for preprocessing the search space that uses a data structure called a *threaded decision graph* that captures the control knowledge for problem solving. For some problems, the method results in much time being saved during problem solving.

Temporal Reasoning

Temporal logics that extend previous approaches or avoid some of the problems with previous approaches have been developed by Fahiem Bacchus and by André Trudel. Peter van Beek has looked at reasoning about temporal relations.

Fahiem Bacchus (with Josh Tenenberg and Johannes Koomen) has developed a temporal logic for reasoning about propositions whose truth values might change as a function of time [3]. The temporal propositions consist of formulae in a sorted first-order logic, with each atomic predicate taking a set of temporal arguments that denote time points and a set of non-temporal arguments. The temporal arguments specify the predicate's dependence on time. By partitioning the terms of the language into two sorts, temporal and non-temporal, time is given a special syntactic and semantic status in the logic without resorting to reification. Unlike the first-order temporal logic developed by Yoav Shoham, propositions can be expressed with and interpreted against any number of temporal arguments, not just a pair of time points (an interval). The result is a more standard and usable syntax and added flexibility. The logic has a clear semantics and, again, unlike Shoham's logic, a proof-theory that is easily implemented with standard automated theorem provers. Moreover, nothing is lost for the logic subsumes Shoham's logic.

André Trudel has also developed a first-order temporal logic [13]. The major difference between this logic, called GCH, and previous temporal logics is how assertions that are true over an interval (such as "ran a mile") are represented. The representation of these assertions is based on the assumption that what is true at every point in an interval completely determines what is true over the interval. The Riemann integral is used to relate an interval with its internal points. For example, "ran a mile" is true over an interval if and only if the integral of "running" at each point in the interval equals one mile. GCH has a formal syntax and semantics and is implemented with a first-order theorem prover augmented with calls to Maple, a system for symbolic mathematical calculation. Trudel (with Scott Goodwin) has also specified a solution to the persistence problem in terms of GCH. Given the above, a user can tackle any temporal problem in AI. As one example, Trudel plans in future work to apply the GCH temporal logic to multi-agent planning.

Peter van Beek has considered a representation for temporal relations between intervals introduced by James Allen, and its associated computational or reasoning problem of, given possibly indefinite knowledge of the relations between some intervals, computing the strongest possible assertions about the relations between some or all intervals. Determining exact solutions to this problem has been shown to be (almost assuredly) intractable. Allen gives an approximation algorithm based on constraint propagation. Van Beek has developed new approximation algorithms, examined their effectiveness, and determined under what conditions the algorithms are exact [16, 17].

Bibliography

[1] Bacchus, Fahiem. 1988. Statistically Founded Degrees of Belief. *Proceedings of the Seventh Conference of the Canadian Society for Computational Studies of Intelligence* (CSCSI/SCEIO), Edmonton, Alta.

[2] Bacchus, Fahiem. 1989. A Modest, but Semantically Well-Founded, Inheritance Reasoner. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (IJCAI), Detroit, Mich.

[3] Bacchus, Fahiem, Josh Tenenberg, and Johannes A. Koomen. 1989. A Non-Reified Temporal Logic. *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Ont.

[4] Bacchus, Fahiem. 1990. Representing and Reasoning with Probabilistic Knowledge. *A Logical Approach to Probabilities.* Forthcoming, The MIT Press.

[5] Cohen, Robin, Marlene Jones, Amar Sanmugasunderam, Bruce Spencer, and Lisa Dent. 1989. User-Specific Responses from Expert Systems. *The International Journal of Expert Systems: Research and Applications.* To appear.

[6] Cohen, Robin. 1987. Analyzing the Structure of Argumentative Discourse. *Computational Linguistics* **13**.

[7] Cohen, Robin, and Marlene Jones. 1989. Incorporating User Models into Expert Systems for Educational Diagnosis. In *User Models in Dialog Systems*, A. Kobsa and W. Wahlster (eds.), Springer-Verlag.

[8] Cohen, Robin. 1989. Implementing a Model for Understanding Goal-Oriented Discourse. *International Journal of Intelligent Systems.* To appear.

[9] DiMarco, Chrysanne, and Graeme Hirst. 1988. Stylistic Grammars in Language Translation. *Proceedings of the International Conference on Computational Linguistics* (COLING), Budapest.

[10] DiMarco, Chrysanne. 1989. Computational Stylistics for Natural Language Translation. Ph.D. thesis (expected: Fall, 1989). University of Toronto.

[11] Song, Fei, and Robin Cohen. 1988. The Interpretation of Temporal Relations in Narrative. *Proceedings of the Seventh National Conference on Artificial Intelligence* (AAAI), Saint Paul, Minn.

[12] Spencer, Bruce. 1988. Managing Multiple Candidate Plans for Plan Recognition. Unpublished manuscript (May 1988; revised, Oct. 1988).

[13] Trudel, André. 1989. Representing and Reasoning about a Dynamic World. Ph.D. thesis (expected: Fall, 1989). University of Waterloo.

[14] van Arragon, Paul. 1987. User Modeling Bibliography. Department of Computer Science Research Report CS-87-22, University of Waterloo, Waterloo, Ont.

[15] van Arragon, Paul. 1989. User Modeling as Nested Default Reasoning. Ph.D. thesis (expected: Fall, 1989). University of Waterloo.

[16] van Beek, Peter. 1989. Approximation Algorithms for Temporal Reasoning. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (IJCAI), Detroit, Mich.

[17] Vilain, Marc, Henry Kautz, and Peter van Beek. 1989. Constraint Propagation Algorithms for Temporal Reasoning: A Revised Report. In *Readings in Qualitative Reasoning about Physical Systems*, J. de Kleer and D. Weld (eds.). Morgan-Kaufman.

[18] Yang, Qiang. 1989. Improving the Efficiency of Planning. Ph.D. thesis, University of Maryland.

[19] Yang, Qiang, Dana S. Nau, and James Hendler. 1989. An Approach to Multi-Goal Planning with Limited Interactions. AAAI Spring Symposium Series, Stanford, Calif.

[20] Yang, Qiang, and Dana S. Nau. 1989. Preprocessing Search Spaces for Branch and Bound Search. *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence* (IJCAI), Detroit, Mich.

Artificial Intelligence Research at the University of Manitoba

by Mark Evans, David Scuse, John Anderson, and Ken Hotz

Recherches en intelligence artificielle à l'Université du Manitoba.

Le département d'informatique de l'université du Manitoba a formé un groupe appelé 'Intelligent Systems Group' (ISG) afin de coordonner l'expertise et les efforts de recherche des membres de la faculté et des étudiants gradués oeuvrant dans le domaine de l'intelligence artificielle et autres domaines connexes de l'informatique. L'ISG est associé à l'Institut canadien de technologie industrielle (ICTI) et a des échanges interdisciplinaires avec d'autres facultés de l'université. Les membres de l'ISG effectuent des recherches dans divers secteurs, dont l'intelligence artificielle appliquée à la médecine, l'intelligence artificielle distribuée et les systèmes experts. L'ISG fournit de plus un certain nombre de services externes tels que des ateliers, des études de faisabilité, le développement de systèmes de grande envergure, des analyses de produit et la mise au point d'algorithmes. L'objectif principal du groupe est de coordonner l'intégration des techniques de l'intelligence artificielle (particulièrement la technologie des systèmes experts) dans la recherche et dans l'industrie au Manitoba. Cet article décrit certaines des activités de recherche des membres du groupe qui estiment que l'intelligence artificielle est leur secteur de recherche principal.

The Department of Computer Science at the University of Manitoba established the Intelligent Systems Group (ISG) to coordinate the expertise and research efforts of senior faculty members and graduate students involved in the research and development of Artificial Intelligence (AI) technology and related fields of computer science. The ISG is concerned with coordinating the efforts of its members in dealing with the application of AI techniques to solving difficult problems. Members of the ISG also interact with researchers from other faculties within the University

The authors are all members of the Intelligent Systems Group at the University of Manitoba.

of Manitoba to direct and coordinate interdisciplinary AI education and research activities. In addition, the ISG benefits from its association with the Canadian Institute of Industrial Technology (CIIT) in Winnipeg, Manitoba through a cooperative agreement between the University of Manitoba and the National Research Council of Canada. The CIIT functions as a Centre of Expertise in Advanced Technology, including artificial intelligence. The ISG and other research organizations at the University of Manitoba maintain offices and laboratories in the CIIT complex and benefit from access to state-of-the-art equipment and software dedicated to AI research and development. The ISG works cooperatively with the research staff of the CIIT to analyze AI hardware and software and carry out AI development efforts.

Members of the ISG carry out research in a variety of areas, including medical artificial intelligence, distributed artificial intelligence, and expert systems. The ISG also provides a number of external services such as workshops, feasibility studies, full-scale system development, product analysis, and algorithm design. The primary objective of the group is to coordinate the integration of AI techniques (particularly expert systems technology) into research and industry within Manitoba. In this article we describe some of the research activities of the members who consider AI their primary research area. Other research activities that are not described include: AI techniques within distributed processing systems, graphics, and user interfaces.

A Constraint-Based Architecture for Cooperative Problem Solving – *M. Evans, J. Anderson*

Until recently, virtually all planning research was performed from the classical planning perspective, in which a single agent attempts to produce a sequence of actions to achieve some goal, using a state-based world description and a set of operators which may be used to alter that world [Wilkins, 1988]. The limitations of this paradigm are well known, and fall into two categories. The first class involves difficulties in representation, such as the representation of change [Shoham, 1985], or the representation of planning actions with context-dependent effects [Weber, 1988]. These are basic issues which must be dealt with regardless of the planning paradigm used. The second category involves practical challenges, such as the ability to deal with resource-bounds [Bratman et al., 1988], and the ability to perform in complex, dynamic environments [Ambros-Ingerson, 1988; Wilkins, 1988]. The latter example mainly concerns the ability to react to a changing environment and repair existing partial plans.

A large class of practical planning problems resists solution by a single agent because the problem is spatially distributed, because the volume of work will exceed the resources of any single agent, or because it demands specialized knowledge or abilities which are currently unreasonable to expect from any single agent. Recently, attention has turned to the field of distributed artificial intelligence (DAI), which concerns the use of multiple problem-solving agents to accomplish tasks.

Typically, each agent in a distributed problem-solving environment is resource bounded in terms of time and energy and, more importantly, has only a limited knowledge of the environment in which it operates. Differences in the types of knowledge which each agent possesses, however, allow us to divide distributed problem solving into two broad categories. In the first category, a large problem is decomposed amongst a number of homogeneous agents (that is, agents with identical or near-identical problem-solving capabilities). Each agent can work on its portion of the problem and communicate knowledge about the problem to other agents. This type of architecture has been seen to be of great use for problems such as distributed sensing [Durfee and Lesser, 1987] and distributed air traffic control [Cammarata et al., 1983]. The second category involves the use of heterogeneous agents (agents whom, while possibly sharing some basic capabilities, have their own talents and handicaps). Distributed problem-solving architectures which fall into this category generally attempt to model human social systems such as teams or organizations. Within each of these categories, many different methodologies have been used to coordinate multiple problem-solvers. These schemes range from systems with decentralized control such as functionally adequate, cooperative networks [Lesser and Corkill, 1981] to blackboard-based systems with a centralized control agent [Hayes-Roth et al., 1989].

Recently, Corkill and Lesser have studied the use of constraint satisfaction techniques in distributed problem-solving domains [Durfee et al., 1987; Pattison et al., 1987]. The utility of viewing planning at least partly in terms of constraint satisfaction has long been recognized. For example, Stefik's MOLGEN planner employed a constraint-propagation technique in generating plans for gene-splicing experiments [Stefik, 1981], while Wilkins' SIPE was the first planner to make use of domain independent constraints during planning [Wilkins, 1988]. Early attempts at the use of constraints in planning have demonstrated the need for more powerful techniques that can manage the complex constraint-directed search problems that arise in distributed problem-solving domains.

We have developed a model of distributed problem solving which addresses this need by providing a framework in which cooperative problem solving is viewed as a multi-agent constraint-satisfaction planning problem [Evans and Anderson, 1989a]. The architecture allows control of agents to be centralized, partially centralized, or completely distributed, and supports reactive planning through negotiation and constraint relaxation [Evans and Anderson, 1989b].

As an application of the model, we are currently implementing some of the distributed problem-solving functions of an automobile repair shop. This domain was chosen because we have found it typical of most cooperative problem-solving situations. The operations of a repair shop involve interactions between many classes of agents, which can be divided into two categories: a kernel of agents which provide fundamental problem-solving activities (e.g. mechanics, managers), and a group of peripheral agents which provide ancillary services (e.g. accountants, other repair shops). The organization of these agents is partially hierarchical (e.g. management agents can supervise and dictate the activities of worker agents) but also partially linear (e.g. two mechanics may make decisions as to which tasks each will perform).

Case-Based Medical Diagnosis – *D. Scuse, K. Hotz*

The application of AI techniques to the medical domain (AIM) deals with three major issues: diagnosis, therapy recommendations, and therapy management [Clancey et al., 1984]. Of these three issues medical diagnosis is considered the most difficult because it is the least understood and as a result it is the most challenging. The goal is to determine a diagnosis (or set of diagnoses) which is

consistent with the symptoms displayed by the patient. The purpose of a therapy recommendation system is to determine the most effective method of treating a disease. Therapy management systems are used to monitor patients as they receive treatment, recommending modifications to treatment if the patient reacts unfavourably. The use of a therapy management system may prove to be the most effective in terms of use because the duration of therapy is longer than the time required to make a diagnosis or therapy selection.

The approach most often taken to developing medical systems is to use knowledge about the diseases and body systems. This type of knowledge can be represented in varying degrees of depth ranging from low-level details to high-level associations. The low-level details (often referred to as "deep", causal, or pathophysiologic knowledge) consists of knowledge about the underlying disease mechanisms and their relationships with body systems. The high-level associations (often referred to as "surface" knowledge or "rules-of-thumb") consists of higher level associations which allow decisions to be made without being aware of the underlying reasons.

Physicians also use a collection of patient cases as the basis for making decisions. The knowledge for making decisions does not come from how the disease is known to present, but rather how it has presented itself in specific instances. This kind of knowledge is very important to medical decision making as it can fill in the gaps in the knowledge describing the diseases themselves. When students go through medical school they spend a significant portion of their time dealing with new cases and reviewing old ones. Physicians increase their competence as they gain experience in this manner [Cutler, 1980; Elstein et al., 1979]. Clinical knowledge is incomplete and the physician must use a combination of general theoretical knowledge and case knowledge when making decisions [Elstein et al., 1979]. We are looking at integrating the two types of knowledge in order to provide a more powerful and complete knowledge base from which better decisions may be made. By combining the two types of knowledge, the strengths of one will fill in for the weaknesses of the other. It will be necessary to represent and manipulate knowledge as both specific cases and general disease development. One of the potential problems which may arise is when there is conflict, with the two types of knowledge suggesting different decisions. A choice must be made to determine whether the decision suggested by the analysis of patient cases is a unique situation or that perhaps the underlying knowledge is inconsistent.

Chemotherapy Advisor – *D. Scuse, L. Strachan, K. Hotz, J. Anderson*

One of the primary goals of artificial intelligence in medicine (AIM) is the dissemination of medical expertise to geographical regions where expertise is lacking [Szolovits, 1982]. The Intelligent Systems Group, in conjunction with the Manitoba Cancer Treatment and Research Foundation (MCTRF), is pursuing this goal through the development of a knowledge-based system to provide chemotherapy management for patients on cancer treatment protocols. The MCTRF has developed an Outreach Programme with the objective to provide on-site care for rural cancer patients. A number of Outreach clinics have been established in selected rural communities in Manitoba. These clinics are staffed by physicians and nurses who do not possess the extensive oncology training of their counterparts in MCTRF and St. Boniface clinics. To ensure that the patients receive the best possible treatment, there must be some means of providing the physicians and nurses with expert oncology advice.

The Chemotherapy Advisor system is being developed to provide the Outreach clinics with expert advice on managing patients undergoing chemotherapy. Research in the area of automating medical protocols has been done extensively in the past, and previous approaches have ranged from statistical [de Rosis et al., 1988] to the application of artificial intelligence techniques [Tu et al., 1988]. We have adopted many of the techniques used by the ONCOCIN system for chemotherapy management [Shortliffe et al., 1981]. The system contains extensive knowledge about chemotherapy treatment plans (i.e. agents, tests, schedules, and toxicities). Each patient is registered in a treatment plan designed to treat the particular cancer(s) identified by the patient's diagnosis. From the definition of the assigned treatment plan, the Chemotherapy Advisor determines the schedule of treatment the patient will receive over a period of time. During the course of treatment, problems may arise (e.g. adverse reactions to therapy) which require modifications to be made to the patient's treatment plan. The system generates for the physician any changes that should be made to the treatment plan (as defined by the protocols from which the treatment plan was created and from additional knowledge previously obtained from MCTRF staff regarding situations not covered by the protocol). Justifications are also provided for each modification.

The existence of a reasonably large body of existing research in the area of medical protocol automation has allowed us to examine other important aspects of medical therapy advice systems. For example, it is well known that medical decision-making systems have had little impact on the day-to-day decision making of health care professionals. A number of reasons are hypothesized for this, such as physician resistance to medical decision-making systems and the difficulty transferring a "successful" system into a clinic, where many standards and practices may exist which differ from those anticipated during system development [Reggia and Tuhrim, 1985]. Previous systems have also demonstrated the importance of the user interface in medical therapy advice systems and its relation to eventual system acceptance [Lane et al., 1986]. Because our eventual goal is to have the Chemotherapy Advisor in routine clinic use, the major focus of recent research has been in the area of user interface and general clinical requirements for medical therapy advice systems [Burgess and Strachan, 1989; Strachan et al., 1989]. Another area of current research is the investigation of common clinical procedures for dealing with uncertainty in oncology protocols. While oncology protocols in general contain very detailed instructions regarding treatment plan modification due to the toxicity of chemotherapeutic agents, there are still sources of uncertainty (e.g. cases of interacting types of toxicity which are not defined by the protocol) which must be implicitly evaluated by the physician [de Rosis et al., 1988]. This is a very large research area which we have only begun to explore.

A prototype Chemotherapy Advisor system was developed between 1987 and 1988, and a version based on a more complete analysis of clinic requirements is expected to be in operation in an Outreach clinic in the fall of 1989.

Using Expert Systems to Generate Fertilizer Recommendations – *M. Evans, R. Mondor*

In conjunction with the Faculty of Agriculture and the Solomon

Sinclair Farm Management Institute at the University of Manitoba we have developed a demonstration prototype of a fertilizer selection advisor (FSA). The program represents an integration of conventional computing and expert systems technology designed to provide expert recommendations enabling farmers to obtain the best return on their fertilizer investment. Ideally, fertilizer recommendations should take into consideration several variables such as, soil type, moisture level, soil nutrient levels, regional climatic conditions, the cost and efficiency of available application methods and times, and the price and availability of fertilizer compounds. In addition, preferences of the user should be addressed in order to provide a customized fertilizer recommendation. These preferences may include: equipment constraints, financial constraints, risk aversion, and farming traditions. The difficulty in making effective fertilizer recommendations is in considering how to work with all the variables to arrive at the best recommendation — we call this the fertilizer problem.

In theory, to solve the fertilizer problem, one could consider all potential recommendations by generating all possible combinations of the variables and selecting the most cost effective alternatives (i.e. an exhaustive search of the problem space). In fact, such an approach is not viable since the computational complexity of this problem is exponential in nature. After consultation with various experts in crop management and soil sciences, we determined that agronomists do not use a formal algorithm for generating solutions to the fertilizer problem. Instead, they rely on experience, intuition, and heuristics.

In AI terms, the fertilizer problem can be viewed as a constraint-directed planning problem. The crux of the problem involves generating a plan that details when to apply the fertilizer, how to apply it and in what amounts, and what specific fertilizer compounds should be used. When establishing this plan, a variety of constraints need to be taken into account, such as the toxicity limits of certain fertilizers, the cost of applying the fertilizer (including application costs and prices of various compounds), and individual characteristics (a farmer may have limited resources, including time, machinery, and capital). The goal of the FSA system is to develop a cost effective plan while abiding by these constraints.

The FSA prototype employs a constraint-directed search strategy which reasons with knowledge represented using an object-oriented approach [Evans et al., 1989a]. Objects may be physical entities that can be seen or touched (e.g. fertilizers, machinery) or conceptual entities such as acts, events, or abstract categories (e.g. application methods, moisture categories). By using an object-oriented approach, we were able to structure the domain knowledge and represent each object's physical and behavioral characteristics. The behavior of an object is represented using procedures and rules embedded in the corresponding knowledge structures. For example, the toxicity constraints associated with various fertilizer compounds are represented as IF-THEN action statements. Similarly, objects representing fertilizer plans have a total cost component which is calculated using an embedded procedure designed to consider all relevant costs [Evans et al., 1989b].

We were able to program the system with sufficient knowledge enabling it to focus its attention towards relevant information only, thus reducing the problem space to a manageable size and significantly improving the efficiency of the system. We also employ an explanation facility which allows the user to query the system about decisions it has made, thus providing the users with a better understanding of the underlying assumptions, facts, and reasoning used to generate recommendations. In the future, we intend to continue development of the system by incorporating several additional factors such as other nutrients, additional economic and environmental factors, and a greater variety of crops. We are currently developing an architecture to incorporate these features by extending our original constraint-based approach to utilize a community of cooperating agents each capable of reasoning about particular subsets of the factors involved in analyzing and planning fertilizer recommendations. As well, the design and implementation techniques developed during this project can potentially be applied to many other areas of agricultural science such as integrated pest management, pesticide use in crops, and ration formulation in poultry and livestock feeds. We intend to examine these potential domains in the near future.

Bibliography

[Ambros-Ingerson, 1987] Ambros-Ingerson, J., IPEM: Integrated Planning, Execution, and Monitoring. M. Phil. Dissertation, Computer Science, University of Essex, 1987

[Bratman et al., 1988] Bratman, M., Israel, J., and Pollack, M., Plans and resource-bounded practical reasoning, *Computational Intelligence* **4**, 349-355, 1988.

[Burgess and Strachan, 1989] Burgess, B., and Strachan, L., User Interface Investigation, Joint Department of Computer Science, University of Manitoba/MCTRF Technical Report, 1989.

[Cammarata et al., 1983] Cammarata, S., McArthur, D., and Steeb, R., Strategies of Cooperation in Distributed Problem Solving, in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 767-770, 1983.

[Clancey et al., 1984] Clancey, W.J., and Shortliffe, E.H., Introduction: Medical Artificial Intelligence Programs, *Readings in Medical Artificial Intelligence*, Addison-Wesley, 1984, p. 2.

[Culter, 1980] Culter, P., Problem Solving in Clinical Medicine from Data to Diagnosis, Williams and Wilkins, Second Edition, 1980, pp. 2-3.

[Durfee and Lesser, 1987] Durfee, E., and Lesser, V., Using Partial Global Plans to Coordinate Distributed Problem Solvers, in *Proceedings of the Sixth National Conference on Artificial Intelligence*, 875-883, 1987.

[Durfee et al., 1987] Durfee, E., Lesser, V., and Corkill, D., Cooperation Through Communication in a Distributed Problem Solving Network, in Huhns, M. N. (Ed.), *Distributed Artificial Intelligence* (Morgan Kaufmann: San Mateo, CA), 29-58, 1987.

[Elstein et al., 1979] Elstein, A.S., Shulman, L.S., Sprafka, S.A., Medical Problem Solving, Harvard University Press, 1979, pp. x-xi.

[Evans and Anderson, 1989a] Evans, M., and Anderson, J., A Constraint-Directed Architecture for Multi-Agent Planning, in *Proceedings of the Ninth AAAI Distributed Artificial Intelligence Workshop*, Seattle, Washington, 1989 (to appear).

[Evans and Anderson, 1989b] Evans, M., and Anderson, J., Negotiation as Multi-Agent Plan Repair, Technical Report, Department of Computer Science, University of Manitoba, 1989.

[Evans et al., 1989a] Evans, M., Mondor, R., and Flaten, D., Expert Systems and Farm Management, CAEFMS Farm Management Workshop, Winnipeg, 1989 (also to appear in Canadian Journal of Agricultural Economics).

[Evans et al., 1989b] Evans, M., Mondor, R., and Flaten, D., Using Expert Systems to Generate Fertilizer Recommendations, Workshop on the Integration of Knowledge Systems with Conventional Computer Methods in Agriculture and in *Natural Resource Management*, Detroit, 1989 (also to appear in *The Journal of AI Applications in Natural Resource Management*).

[Hayes-Roth et al., 1989] Hayes-Roth, B., Hewett, M., Washington, R., and Hewett, R., Distributing Intelligence within an Individual, Knowledge Systems Laboratory Report No. KSL 88-50, October, 1988 (updated April, 1989).

[Strachan et al., 1989] Intelligent Systems Group, University of Manitoba, and Manitoba Cancer Treatment and Research Foundation, Statement of Requirements for the Chemotherapy Advisor, Joint Department of Computer Science, University of Manitoba/MCTRF Technical Report, 1989.

[Lane et al., 1986] Lane, C., Walton, J., and Shortliffe, E., Graphical Access to Medical Expert Systems: II. Design of an Interface for Physicians, *Methods of Information in Medicine* **25**, 1986, pp. 143-150.

[Lesser and Corkill, 1981] Lesser, V., and Corkill, D., Functionally Accurate, Cooperative Distributed Systems, *IEEE Transactions on Systems, Man, and Cybernetics* SMC-**11**(1), January, 1981, pp. 81-96.

[Pattison et al., 1987] Pattison, H., Corkill, D., and Lesser, V., Instantiating Descriptions of Organization Structures, in Huhns, M.N. (Ed.) *Distributed Artificial Intelligence* (Morgan Kaufmann: San Mateo, CA), 1987, pp. 59-96.

[Reggia and Tuhrim, 1985] Reggia, J., and Tuhrim, S., An Overview of Methods for Computer-Assisted Medical Decision Making, in Reggia, J., and Tuhrim, S., (Eds.), *Computer-Assisted Medical Decision-Making*, Vol. 1, (Springer-Verlag: New York, NY), 1985, pp 3-45.

[Shoham, 1985] Shoham, Y., Ten Requirements for a Theory of Change, *New Generation Computing* **3**, 1985, pp. 467-477.

[de Rosis et al., 1988] de Rosis, F., Steve, G., Biagini, C., and Maurizi-Enrici, R., Treatment of Uncertainty in an Oncology Protocol by Probabilistic and Artificial Intelligence Approaches, *Methods of Information in Medicine* **27**, 1988, pp. 22-23.

[Shortliffe et al., 1981] Shortliffe, E., Scott, A., Bischoff, M., Campbell, A., Van Melle, W., and Jacobs, C., ONCOCIN: An Expert System for Oncology Protocol Management, in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 1981, pp. 876-881.

[Stefik, 1981] Stefik, M., Planning with Constraints, *Artificial Intelligence* **16**, 1981, 1981, pp. 111-140.

[Szolovits, 1982] Szolovits, P., Artificial Intelligence and Medicine, in Szolovits, P. (Ed.), *Artificial Intelligence and Medicine*, AAAS Selected Symposia series (Westview Press: Boulder, CO), 1982. pp. 1-19.

[Tu et al., 1988] Tu, S., Kahn, M., Musen, M., Ferguson, J., Shortliffe, E., and Fagan, L., Episodic Monitoring of Time-Oriented Data for Heuristic Skeletal-Plan Refinement, Knowledge Systems Laboratory Report No. KSL 87-70, August, 1988.

[Weber, 1988] Weber, J., A Versatile Approach to Action Reasoning, Technical Report 237, Department of Computer Science, University of Rochester, 1988.

[Wilkins, 1988] Wilkins, D., Practical Planning: Extending the Classical AI Planning Paradigm (Morgan Kaufmann: San Mateo, CA), 1988. 205 pp.

Third Knowledge Acquisition for Knowledge-Based Systems Workshop

Banff, Alberta, November 6-11, 1988
Sponsored by AAAI

by Stephen Regoczei

Troisième atelier sur l'acquisition des connaissances pour les systèmes à base de connaissances.

Le troisième atelier portant sur l'acquisition des connaissances pour les systèmes à base de connaissances a eu lieu du 6 au 11 novembre 1988 à Banff en Alberta. Cet atelier donna aux participants l'occasion de constater la rapide évolution dans le domaine de l'acquisition des connaissances.

Les deux principaux objectifs de l'atelier étaient : a) de stimuler la recherche b) de partager rapidement les résultats de recherche à l'intérieur de la communauté afin de renforcer l'établissement d'une entreprise plus unifiée et par le fait même plus efficace. L'atelier atteingnit ces deux objectifs avec succès.

Introduction

A report on the third Knowledge Acquisition Workshop (KAW), Banff, 1988, provides a good opportunity to review the rapid growth of the knowledge acquisition field. Before the first Banff workshop in 1986, knowledge acquisition existed only as a "bottleneck": an annoying obstacle and an awkward technicality which may "get in the way" of building knowledge-based systems. In those days, the emphasis was on knowledge representation and rule-based inference strategies.

But the realization was slowly growing that knowledge acquisition (KA) is not only at the centre of creating expert systems but is also crucial to the development of other types of knowledge-based software. Looking back, we see the workshop addressing an already existing but unmet need. The response to the call for papers was overwhelming (Boose & Gaines, 1988, p. vii). "The intention was to hold a discussion – intensive meeting of some 35 highly involved researchers. In practice over 120 papers were submitted and some 500 applications to attend were received from about 30 countries." 42 papers were selected and 60 people attended. "Much of the rejected material was of high quality and it would have been possible to base a major conference on the material and requests to attend." (ibid.)

From the time of that workshop two years ago, the KA field had undergone a remarkable development. The fragmented attempts at knowledge acquisition research were united, and isolated researchers started to communicate with each other. The field is maturing. The main research issues are being clearly delineated and the KA tools are grounded in successively more sophisticated theoretical concerns. Researchers are more and more open about their thematic preferences; they are also less guarded and more specific in the descriptions of their approaches.

Stephen Regoczei is a researcher in Computer Studies at Trent University. His major interest is in knowledge acquisition.

Keynote Papers

The 1988 workshop opened with some keynote papers that documented this development. John Boose of Boeing made a major contribution by summarizing the KA literature of the past few years (173 items), including six workshops. He offered a comprehensive taxonomy of techniques and methods, and an annotated list of some 65 KA tools. The resulting document will be updated periodically and republished. It promises to be an indispensible guide for newcomers to the field (and even for some seasoned participants), enabling them to get a comprehensive overview very quickly.

What made Boose's paper especially interesting is that he used the KA tool AQUINAS to organize the knowledge that he extracted from the KA literature! In particular, the comprehensive taxonomy was produced with the repertory grid component of AQUINAS. This raises the fundamental question whether knowledge acquisition or knowledge "creation" was being performed. In any case, the paper certainly illustrated the power of AQUINAS to construct a useful knowledge base.

Brian Gaines, in his overview, called attention to the social context of the expert, emphasizing that knowledge is embodied, and, in particular, embodied in a social system. Being an expert and possessing expertise is a social activity. The expertise of an expert grows because he becomes privy to an insider's knowledge,

gets financial support for his work, and is given more interesting problems to solve.

Working Groups

This year the working groups were:

- Integration of machine learning and KA
- KA tools
- Repertory grid methods for KA
- Synthesis tasks and KA: planning, scheduling, configuration, design...
- Text analysis
- Validation and verification

For many participants, discussion in the working groups on both state-of-the-art and future directions remains the most fruitful activity of the workshop. In general, there was indication that the period of exploratory work is over and that it is time to tackle the really difficult research issues in each area.

Highlights of the Workshop

Many of the papers were progress reports on research that was outlined in previous workshops, but there were some remarkable new departures. For example, an important paper by Mildred Shaw and Brian Gaines of Calgary addressed the crucial and very difficult problem of the breakdown in communication between the expert and the analyst because they may use the same word to refer to different concepts, or have the same concepts but call them by different names. What makes such mismatch in vocabulary and concepts particularly pernicious is that the participants in the dialogue may not realize exactly what is wrong. They would try to "talk themselves out of the problem", but of course their vocabulary does not match their concepts. A possible way out is to use repertory grid techniques to get the informant and analyst talking to each other again.

Mark Musen of Stanford offered a simple but effective mini-life cycle for KA:

1) Creation of a systematic domain. ("Systematic domain" is a term of Winograd and Flores (1986) for a lexicon of terms and concepts that formalizes the domain of discourse and defines the ontology.)
2) Creation of a task model (this corresponds to Newell's knowledge-level analysis).
3) Extension of the task model (affectionately referred to as "knowledge stuffing").

Enrico Motta of The Open University offered a fuller life cycle. He presented KA as the production of conceptual models and his life cycle included separate stages for psychological analysis, linguistic analysis and conceptual analysis.

Converting repertory grids to more traditional forms of knowledge representation such as rules and frames is still something which will have to be digested by researchers not fully conversant with personal construct theory. But already Jeff Bradshaw et al. of Boeing are rushing ahead showing "how to do things with grids that people say you can't". These include: different abstraction levels, constructs which are not bi-polar, and representing procedural and strategic knowledge. They are now attacking synthesis problems by combining several software packages such as MacQuinas (derived from AQUINAS) and MacXotl (derived from Axotl).

Validation of knowledge bases with respect to several different criteria was addressed by a number of papers, including that of Izak Benbasat and Jasbir Dhaliwal of the University of British Columbia, and Rajiv Enand and Gary Kahn of the Carnegie Group.

Bruce Porter introduced the crucial issue of knowledge acquisition versus knowledge creation in a casual remark during a panel discussion on statistical databases. The issue is this: how to extract knowledge, when the knowledge is "not there". Something is there — raw data perhaps — but beyond this the metaphors break down. Do we extract the knowledge, or do we discover it? Do we add it by imposing structure upon the data using statistical techniques, or do we "boil down" the data in order to "distill" the knowledge?

Tom Gruber of Stanford focused on strategic knowledge, i.e. knowledge which is used by an agent to decide what action to perform next, where the actions affect both the state of the external world and the agent's models or beliefs that model these states. The automated assistant ASK elicits justifications from the expert and generates strategy rules. ASK makes a "subtle cut" across specializations; it gathers heuristic knowledge about the world that enables us to act on the basis of incomplete and uncertain knowledge.

Kenneth Murray and Bruce Porter of the University of Texas at Austin addressed knowledge integration: that is, the task of integrating new knowledge with an already existing knowledge base. The task is to relate the new information to existing knowledge, hence it is the existing knowledge that determines what is learned, not some a priori model of a particular problem solving method or task. Knowledge is acquired when a gap or anomaly is revealed while relating new information to existing knowledge. An interesting feature of their paper was the formalization of the terms "concept", "perspective", and "context".

In addition to the few examples cited above, we should mention that the repertory grid techniques have proved themselves over the years: they turned out to be an important direct way of getting at knowledge. Direct way of getting at knowledge was emphasized by others: Mary Meyer of Los Alamos, contrary to the current paradigms of software development (i.e. forcing the expert's thinking into the framework of the knowledge engineer), suggested the use of less coercive interviewing techniques such as the ethnographic query and the verbal probe. Nancy Johnson of Brunel University also advocated the gentle handling of the "emergent" knowledge before it gets formalized and shoehorned into one's favourite knowledge representation approach.

On the basis of actual, "hands-on" field experience, Cathy Kitto gave a very thorough comparison between AQUINAS from Boeing and KNAC from Carnegie-Mellon.

A Few Observations

Research on KA tools certainly seems to dominate the field. Why is there such an emphasis on automated KA tools? There are several answers. The emphasis on tools merely recognizes the basic facts of the knowledge acquisition process. For a knowledge base with an ontology list of 50-100 items, manual techniques are quite feasible, but a "real" system would require ontologies of at least 1,000-10,000 items. The claim could be made that systems of this size are still tractable, but even the bravest would have to admit that we have only the sketchiest idea of how to go about dealing with millions of items. Just to keep track of the terminology that describes the domain knowledge would require software.

Some of this may not have to be more complex than the data dictionary/directory systems used in the database context, but if knowledge acquisition is based on interviews, protocols, questionnaires, or the filling out of forms, the administrative task alone could become staggering.

But there is a deeper reason. One can almost go as far as to say that a theory of KA carries no credibility unless embodied in a tool, for the simple and very good reason that otherwise it cannot be tested, or even adequately operationalized. So tools are where it's at. For the same reason, knowledge acquisition without a performance system is not testable or operationalizable.

General Conclusions and Future Directions

The overall conclusion of the workshop was that the field is growing and maturing, and the different research issues are being clearly delineated, but, fortunately, the field is not yet large enough to start fragmenting into subgroups.

The workshops are certainly succeeding in their main objective: stimulating research, sharing research results quickly within the community and thereby strengthening the building of a more unified and effective enterprise. Once again, expressions of thanks are due to the organizers: John Boose of Boeing and Brian Gaines of Calgary.

Certain problems need to be highlighted. For example, the vocabulary proliferates: knowledge can be elicited, discovered, created, recorded, acquired, represented, formalized, explicated, etc., etc. The types of knowledge proliferates: structural, strategic, procedural, control, problem solving, domain, goal seeking, verbalized, tacit, ontological, functional, task-oriented, method-oriented, design, repair, skill-like, etc., etc. The uncontrolled growth of the terminology often signals the kind of mismatch between vocabulary and concepts that Shaw's paper so clearly described. Half in jest, Brian Gaines suggested that a standards committee be formed — quite a few people felt that the suggestion should be taken seriously.

We also need more "reports from the field", more "actual" data points. We need more case studies of the form: "At such and such a time and place I sat down with an actual, real, live expert and the following transpired..."

Also, as Brian Gaines pointed out, it is time to delineate the relationship of the knowledge acquisition community both to the associated technical community and the prospective users. How does KA relate to machine learning, natural language understanding, cognitive science, expert systems, knowledge representation, hypertext, and software engineering? How do we approach the user community and the professional disciplines whose expertise is being acquired? How do we react to the current issues in CASE, database modeling, expert systems, requirements specifications, and systems analysis? As the KA field becomes more clearly defined internally, these questions assume a more and more prominent role.

Addresses and future activities

The next Banff workshop will be held in the fall of 1989. Other knowledge acquisition workshops include the European Knowledge Acquisition Workshops (EKAW) (the next one in Paris). The last two EKAW's were held in Reading, England and Bonn, Germany. There was a workshop in St. Paul, Minnesota in August 1988 at the AAAI on the Integration of Knowledge Acquisition and Performance Tools. The Banff Knowledge Acquisition Workshops are sponsored by AAAI.

Many of the workshop papers are published in the International Journal of Man-Machine Studies. Papers from the Banff Knowledge Acquisition Workshops and the European EKAW's are being published in the Knowledge-Based Systems Series by Academic Press. For further information contact:

Brian Gaines, Knowledge Science Institute, University of Calgary, 2500 University Dr., NW, Calgary, Alberta, T2N 1N4

or

John Boose, Knowledge Systems Laboratory, Advanced Technology Center, Boeing Computer Services, P.O. Box 24346, Seattle, WA 98124, USA.

Bibliography

[Boose, J.H. and B.R. Gaines (1988)]: Knowledge Acquisition Tools for Expert Systems, Knowledge Based Systems, Volume 2, Academic Press, London.

[Winograd,T. and F. Flores (1986)]: Understanding Computers and Cognition: A New Foundation for Design. Norwood, New Jersey, Ablex.

KR'89: The First International Conference on Principles of Knowledge Representation and Reasoning

by B. Brown, J. de Haan, M. Grüninger, S. McIlraith and D.Schuurmans

La première conférence internationale sur les principes
de représentation de connaissances et de raisonnement.

La première conférence internationale sur les principes de représentation des connaissances et le raisonnement a eu lieu du 15 au 18 mai à Toronto. Cette conférence fut déclarée par les participants comme étant un grand succès. Lors de cette conférence l'emphase fut mise sur les principes sous-jacents à la représentation des connaissances et le raisonnement par contraste avec les problèmes reliés à l'ingénierie et les difficultés d'implantation.

On recommande fortement le compte rendu de la conférence à toutes personnes cherchant à avoir plus d'information sur les récents développements dans ce domaine.

The First International Conference on Principles of Knowledge Representation and Reasoning, KR'89, was held at Toronto's Royal York Hotel from May 15-18, 1989. The conference focused on the principles underlying knowledge representation and reasoning, as distinct from concerns of engineering and details of implementation. It provided an intimate environment for discussion of the latest research in knowledge representation and reasoning with some of AI's finest researchers. The conference was heralded by attendees as a great success, much of which can be credited to the conference chair, Ray Reiter, and the strong program committee chaired by Ron Brachman and Hector Levesque. The conference is recommended in the future to researchers interested in these aspects of artificial intelligence research.

The four-day conference was divided into two parallel sessions covering such topics as temporal logic, logic for knowledge representation, hierarchical knowledge bases, temporal constraint networks, default theories, nonmonotonic reasoning, analogical reasoning, formal theories of belief revision and metareasoning. It also included three half-day symposia on the topics of temporal reasoning, nonmonotonic reasoning, and a symposium entitled "Against Representation: The Opposition Speaks." There were 275 papers from 27 countries submitted for review to the program committee. Of these, 49 were selected as clearly contributing to the principles of representation and reasoning, present and future. Three of these papers were singled out as being exemplary: "Hard Problems for Simple Default Logics" by Henry Kautz and Bart Selman, "Three-Valued Formalization of Non-Monotonic Reasoning and Logic Programming" by Teodor Przymusinski, and "An Episodic Knowledge Representation for Narrative Texts" by Lenhart Schubert and Chung Hee Hwang. These papers were presented again at IJCAI'89.

In addition to the technical program, many KR'89 attendees took advantage of a delicious banquet dinner at Ontario Place, overlooking Lake Ontario. Several of the conference registrants also managed to attend one of the Blue Jay's last home baseball games outside the SkyDome, while other Western Canadians were holed up in their hotel rooms watching the Calgary Flames skate to Stanley Cup victory.

Although this report does not purport to critique the conference, several recurring themes were encountered that the authors feel are worth mentioning, including:

- unifying approaches to nonmonotonic reasoning
- the complexity of AI reasoning
- the role of probabilistic reasoning
- the need for experimental AI

Additionally, it was observed that there was a heavy emphasis on reasoning, with disappointingly little done in the area of knowledge representation. Some went so far as to say that this reflected a general loss in the richness of knowledge representation in AI. Also lacking from the conference were papers on such topics as truth maintenance systems, learnability theory and connectionism; all of which are generally seen to be hot new areas of research. It was conjectured that perhaps these researchers don't see themselves as doing work in knowledge representation and reasoning, or that perhaps they have better venues for discussion of their work.

This report touches on the highlights of each conference session and includes summaries of the three "best papers" of the conference. The conference proceedings are recommended to those looking for an account of the latest research in the field. They are available from Morgan Kaufmann Publishers.

The Sessions

Nonmonotonic Reasoning

Several papers addressed the issue of unifying the various formalisms for nonmonotonic reasoning. Vladimir Lifschitz, Wiktor Marek, and Miroslaw Truszcynski presented papers that dealt with establishing relationships between various formalisms. Lifschitz proposed a logic combining aspects of circumscription and autoepistemic logic, while Marek and Truszcynski each explored the relationship between autoepistemic and default logics.

Several other papers presented proposals to construct unifying formalisms for nonmonotonic reasoning (NMR). Fangzhen Lin and Yoav Shoham proposed argument systems, which are based entirely on inference rules and argument structures, as a uniform

The authors are all graduate students in AI at the University of Toronto. With one exception, they comprise the 'Alberta Mafia' at UofT.

basis for NMR. Daniel Lehmann presented a general framework for NMR by extending previous study of nonmonotonic logics and their consequence relations. He provided an account of a conditional-style logic in terms of inference rules and a well defined semantics. Teodor Przymusinski generalized previous work by drawing a connection between three-valued extensions of the major nonmonotonic formalisms and previous work in the semantics for logic programming. All three of these papers demonstrated that various aspects of default logic, autoepistemic logic, and circumscription can be expressed within their respective frameworks.

Two other papers expressed limitations of such unifying approaches by defining criteria for nonmonotonic reasoning — Jon Doyle and Michael Wellman's paper on impediments to universal preference-based default theories, and David Poole's paper on the lottery paradox and its implications for default reasoning. Discussion following these papers focused on the criteria which the authors proposed, primarily Doyle and Wellman's notion of non-dictatorship and Poole's one-step default property.

The role of probabilistic reasoning and its relationship to the nonmonotonic formalisms, while represented by only a small volume of papers, generated much discussion. Eric Neufeld's paper on defaults and probabilities introduced the notion of coherent models and showed that some basic ideas from probability theory cast light on why problems exist with default logics. He considered probability as a theory of sound approximate argument. Judea Pearl's survey lecture, "Probabilistic Semantics for Nonmonotonic Reasoning", viewed nonmonotonic reasoning as qualitative probabilistic reasoning. Both papers investigated establishing sound probabilistic semantics for nonmonotonic reasoning and both see default logics as moving towards probability.

Get thee to a keyboard

In his invited talk during the Nonmonotonic Reasoning Symposium, David Poole raised an issue which would be pursued in later discussions. Summarizing an all-night discussion at the Munich Nonmonotonic Reasoning Workshop, Poole claimed that "hacking logic is no better than hacking code," i.e., we need to take an experimental approach to AI, in which we build systems to guide progress on theoretical issues. By using these systems we can more easily detect shortcomings in the underlying formalism. This approach also avoids the problem of concentrating too much on toy examples, something for which work in nonmonotonic logic has often been criticized.

The sessions also featured a panel discussion on critical issues in nonmonotonic reasoning. The panelists were David Etherington, Ken Forbus, Matthew Ginsberg, David Israel, and Vladimir Lifschitz. The panelists' comments, as well as many other discussions at the conference, seemed to address the question — "Have formalisms taken on a life of their own? Research is driven by examples designed to demonstrate the strength and weaknesses of various formalisms, rather than being motivated by problems encountered in the everyday commonsense world." This view was expressed by Etherington, Forbus and Ginsberg. In particular, Forbus emphasized the need for domain theories and expressed doubt as to a domain independent solution to problems in nonmonotonic reasoning. Lifschitz countered that toy examples provide an adequate environment for debugging and enhancing formalisms. He also pointed out that logic programming is an excellent example of how nonmonotonic formalisms have been used to understand real life problems; this sentiment was expressed by several other people during the conference.

The issue of validation of nonmonotonic formalisms was raised by Etherington and Ginsberg. How can the different frameworks be evaluated without resorting to toy problems and conflicting intuitions about the solutions of these problems?

Other papers dealing with issues in nonmonotonic reasoning were: "A Simple Solution to the Yale Shooting Problem" by Andrew Baker, "Default Reasoning, Minimality and Coherence" by Hector Geffner, and "Plausible World Assumption" by Eliezer Lozinskii.

Tractable Reasoning

The field of Knowledge Representation and Reasoning (KR+R) research has recently witnessed a growing interest in understanding the computational complexity of various knowledge base (KB) management and reasoning tasks. Researchers have been particularly interested in investigating the boundaries between tractable and intractable reasoning. Although specific complexity concerns were addressed only by a few papers at KR'89, the growing emphasis of complexity-related issues was evident from the amount of informal discussion these papers generated.

Three papers at KR'89 investigated the computational complexity of particular reasoning tasks in some detail — distinguishing a number of tractable and intractable cases. They were: Kautz and Selman's paper "Hard Problems for Simple Default Logics" which investigated the complexity of various forms of default reasoning; the paper by Bylander *et al.* "Some Results Concerning the Computational Complexity of Abduction" which looked at abductive reasoning; and "Hierarchical Knowledge Bases and Efficient Disjunctive Reasoning" by Borgida and Etherington which examined approximate (unsound) reasoning. Other papers dealing with complexity issues also appeared in the session on temporal reasoning and constraint satisfaction. These will be discussed later.

The perception that there is a growing emphasis on complexity issues initiated a number of informal discussions throughout the week. Many people seemed to be questioning the value of worst-case analyses, often citing the empirical utility of resolution-based theorem provers despite their theoretical undecidability. However, others argued that understanding worst-case complexity — apart from providing strong performance guarantees — really does lead to a better understanding of a reasoning task in general.

Another observation was that, of the current work being pursued under the banner of "Knowledge Representation and Reasoning" research, most seems to be concerned with the "Reasoning" aspect of the title (an emphasis that was certainly evident at KR'89). An overriding concern with the complexity of reasoning invariably means that limited representation languages must be considered. This is contrary to one of the long-standing goals of KR research — developing richer languages for KR, expressive enough to capture wide ranges of human expression.

Unfortunately, the possibility of reconciling expressiveness and efficiency concerns received little discussion.

Taxonomic and NLU Representations

One session combined three papers on taxonomic representations with two papers on representations for natural language

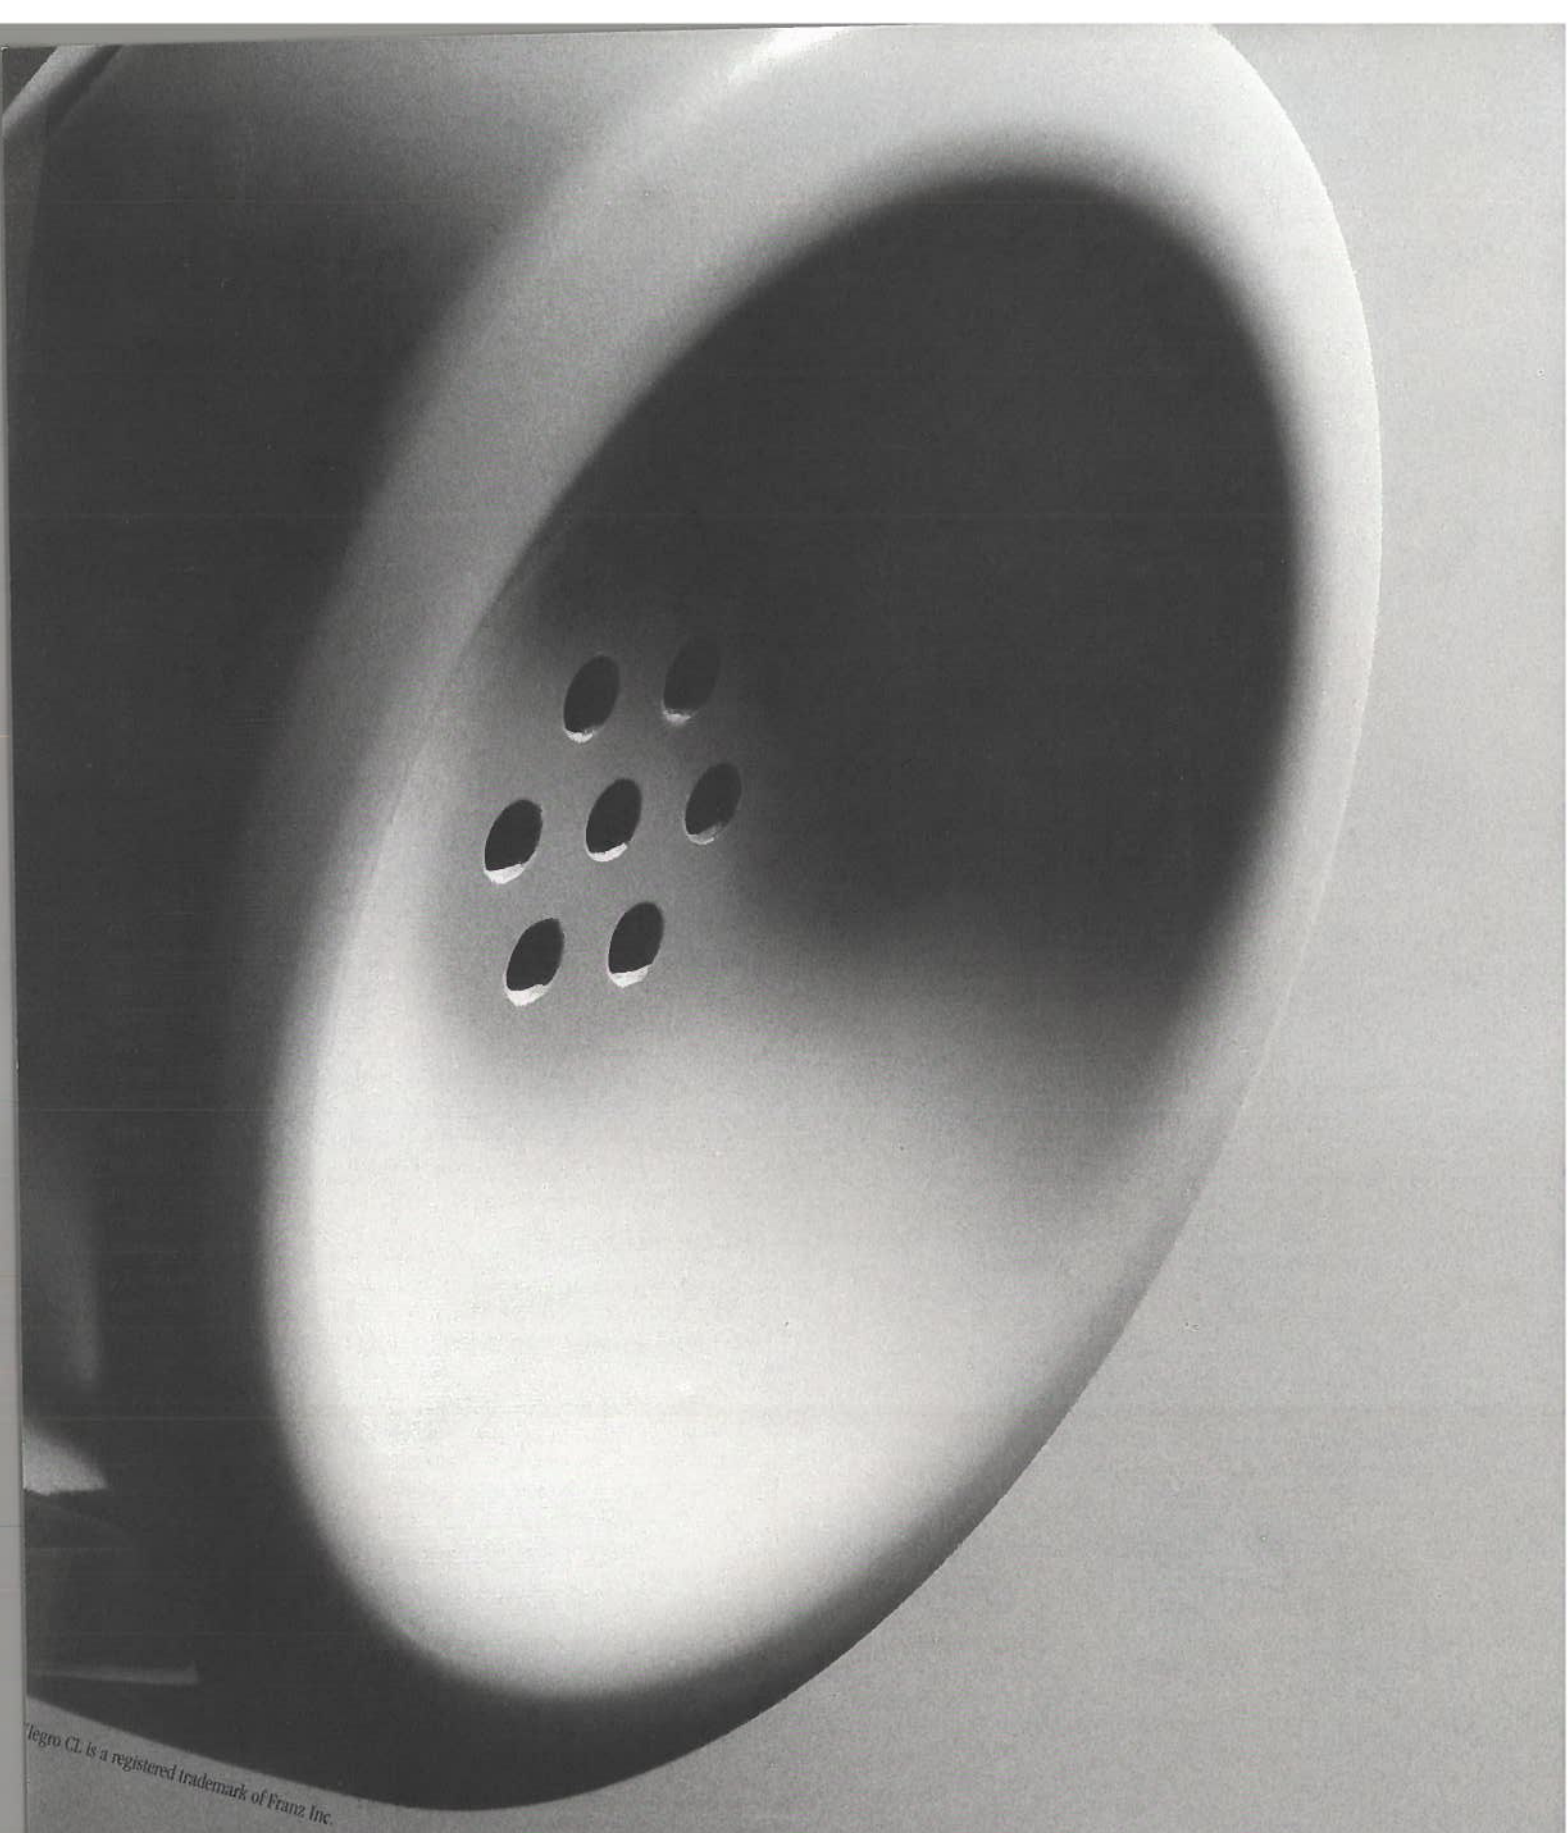

legro CL is a registered trademark of Franz Inc.

Franz Introduces the

Franz Inc. • 1995 University Avenue • Berkeley, CA 94704 • FAX: (415) 548-8253

Sure, we've got the most portable Common LISP in the business—micros to mainframes. And our Allegro CL® is setting new records for speed, compactness, and reliability. But when you're developing serious commercial applications, you need more.

"More than any other supplier, Franz has been incredibly responsive to all our needs, in record time."

Jeff Fox, VP Engineering—Silc Technologies, Inc.

If you have a LISP problem or need an enhancement, you have a choice: you can spend precious development time trying to get an answer from a hardware company, or you can come to the source.

"By far, the best customer support of any software vendor I've ever dealt with. Their open attitude and attention to our needs has been amazing."

Gordon Kotik, President—Reasoning Systems, Inc.

When you call customer service at Franz, you get a LISP expert. If you need a special feature or enhancement, you'll get it, fast—by dealing direct with the person who actually developed that part of Allegro CL.

Experience the ultimate in LISP support. Pick up your own Franz accessory, and call: (415) 548-3600.

FRANZ INC.

Ultimate LISP Accessory.

understanding (NLU). James Schmolze argued that although most terminological reasoners are restricted to using terms with arity≤2, there are naturally occurring terms with arity>2; and even though the reasons for the restriction to arity≤2 seem to be largely historical, users are forced to use clumsy and inefficient representations and techniques for n-ary terms. In his paper, he sketched a family of n-ary systems, and as an example he extended one system with arity≤2 into an n-ary system. Since these extensions generally strengthen the expressive power of a terminological reasoner, they also increase the complexity of computing subsumption. Schmolze argued that given a choice, this is what users want anyway — they would prefer to sacrifice tractability (or completeness) over expressiveness.

Manfred Schmidt-Schauß presented a somewhat surprising complexity result in his paper on subsumption in KL-ONE. He used a subset of KL-ONE which appeared to be only a moderate extension of a very similar tractable system but in which subsumption was in fact undecidable.

David McAllester *et al.* presented a taxonomic syntax for first order predicate calculus whose literals are more expressive than literals of classical first order logic, but which at the same time retained a polynomial time decision procedure for determining the satisfiability of a set of literals. In their paper, they proved that the length of a proof in their system is guaranteed to be no longer than the corresponding proof in the classical system. They went on to conjecture that there exist proofs in their system which are much shorter than the corresponding classical proofs. Whether it is fair to compare the lengths of proofs without considering the time it takes to find them was raised as a question during the talk, but the issue was left unresolved.

Graeme Hirst presented an interesting paper on the inadequacy of the standard Russellian approach to existence in first order logic as a basis for representing the notion of existence as it occurs in natural language. Hirst reviewed some of the philosophical background (starting with ontological proofs for the existence of God), and went on to examine current approaches to existence. One of the major problems with most approaches is that they force all things to exist (or not exist) in the same way; Hirst argued that any KR formalism to be used as a basis for NLU must be able to account for more than one kind of existence, which he illustrated by presenting a list of nine different kinds of existence.

The final paper of the session was the award-winning paper by Lenhart Schubert and Chung Hee Wang, in which they presented an extensive system for understanding narrative texts. This paper will be discussed in detail in a later section.

Deduction

Two papers in the session on deductive reasoning discussed the use of sorted logics as a means to speed up deduction. Alan Frisch presented a general framework by which an existing unsorted deductive system and its proof of completeness could be extended into a sorted deductive system and its proof of completeness. In Frisch's model, the sort theory can only be invoked during the unification process, when substituting one term for another. In a departure from this model, Anthony Cohn went on to discuss the pros and cons of logics in which the sort symbols appear directly, as predicates of characteristic literals in the formulas under consideration. In these logics, the sort theory is not only used during unification, but is further called upon to decide such things as whether two characteristic literals are resolvable. Cohn concluded that these kinds of logics are more expressive and therefore in some cases deduction is more efficient, although it is at the cost of additional complexity in the inference machinery.

The paper by Edward Stabler presented another way in which the unification process can be extended, to efficiently reason about equality for terms known to have distinct denotations in the intended model. In Stabler's resolution system, the search space is effectively reduced by building equality into the unification process, and then further informing the process about terms for which the unique names assumption is in fact true. The method does not require that *all* terms obey the unique names assumption, so the standard equality axioms can still be applied to cases for which they are appropriate.

Metareasoning

People do not have all the time in the world to carefully ponder every possible course of action to ensure that they always make the optimal choice. However, most research on problem solving systems has assumed that deliberation could continue indefinitely. Little attention had been paid to coping with real-time constraints on problem solving behaviour. Two papers at KR'89 addressed this problem in some detail, providing the first steps toward a formal theory of resource bounded problem solving. These papers developed similar formal frameworks to study the problem of selecting which computational action to pursue, weighing the cost of deliberation against the utilities of the actions they ultimately select.

The more elaborate of these two formalizations was presented by Stuart Russell and Eric Wefald. Their formalization is expressed in terms of states of the world, external actions, and computations (internal actions). By assigning utility measures to states, the utility of external actions can be defined in terms of the states expected to result from their execution. Russell and Wefald postulated that the positive utility of a computation is based solely upon its ability to potentially change the external action that is ultimately selected. By defining the value of a computation as the difference between its utility and the utility of the best external action considered so far they could characterize the ideal control cycle — perform computations of highest positive value until all values are negative, then commit to the ultimately preferred external action. By making a number of simplifying assumptions, Russell and Wefald applied their control techniques, with promising results, to a few time-dependent search problems (problems where the time spent in deliberation affects the quality of the solution).

Although the framework presented by Oren Etzioni is largely a specialization of Russell and Wefald's, he used a more general model of the cost of time (opportunity cost). Etzioni demonstrated that when one considers the opportunity costs of performing actions, the meta-level computations required to obtain optimal control become intractable, in general. To overcome intractability he uses a heuristic control technique which simply selects those actions with maximal marginal utility (utility/cost). Etzioni considered the task of acquiring estimates of the utility and cost measures in greater detail than did Russell and Wefald. He presented techniques for acquiring probably approximately correct estimates and characterized the effects that these estimates have on the performance of his control technique. He also looked at the problem of forming appropriate reference classes on which to base the estimates, suggesting that concept learning techniques might be usefully applied to this problem.

Belief Revision

The study of belief revision is concerned with modeling the beliefs of an agent in a changing environment, particularly when those beliefs may be inaccurate. Most belief revision research in AI has consisted of designing algorithms that perform belief revision functions, describing in detail how beliefs are to be represented and manipulated by a machine. Two papers at KR'89 investigated belief revision from the knowledge level, attempting to provide an account of belief revision which is independent of implementational concerns.

Bernhard Nebel pursued a knowledge level analysis of belief revision by studying and extending the theory of epistemic change. Nebel notes that epistemic change theory models the coherence theory of belief revision which ignores the reasons or justifications for beliefs. However, these are central notions to the foundational theory that most AI researchers are concerned with. He demonstrated the relationship between operations on belief bases (finite sets of propositions) and operations on deductively closed belief sets, observing that contracting belief bases resulted in a form of reason maintenance behaviour. Thus, Nebel provided a partial knowledge level account of reason maintenance at the symbol level. This paper argues that foundational belief revision can be accounted for within the coherence theory by using appropriate contraction functions.

Rao and Foo provided a detailed and formal account of the coherence and foundational theories of belief revision. For the coherence theory, they characterized the dynamics of belief under the expansion, contraction, and revision operations within a first order modal temporal logic. However, formalizing the foundational theory turns out to be much more difficult as Rao and Foo saw justification as a second order notion. Thus, they required a second order system to characterize the belief dynamics of a foundational system. They provide sound and (partially) complete axiomatizations of these characterizations. Many attendees questioned the relevance of coherence theory for constructing belief revision systems.

Temporal Reasoning and Constraint Satisfaction

The symposium on temporal reasoning in AI, philosophy, and theoretical computer science provided a forum for discussion of the roles of temporal logic, both in open question periods and panel discussions. The symposium featured Amir Pnueli from the Weizmann Institute and Johan van Benthem from the Universiteit van Amsterdam.

A refreshingly different look at temporal logic, van Benthem's talk was the highlight of the afternoon. Starting with an introduction to the beginning of time (time as a concept), he touched on issues that philosophers have been dealing with for centuries but are just starting to be applicable in computer science. Ultimately van Benthem had insufficient time to explain time, but the material he was able to present was fresh and provocative, and he was a most entertaining speaker.

The sessions on constraints and time focused on aspects of efficiency and representation. Both papers in the section on constraints addressed the complexity of constraint satisfaction problems. Kasif proposed parallel solutions to constraint satisfaction problems, identifying sequential algorithms that are amenable to parallel solutions. He also demonstrated that the parallel complexity of constraint networks is critically dependent on subtle properties of the network that do not influence its sequential complexity. Examining relaxation algorithms for constraint satisfaction, Rossi and Montanari identified the class of perfect relaxation algorithms. For particular classes of networks, they demonstrated that these algorithms have a time complexity which is linear in the number of constraints, and lead to exact solutions.

The paper by Dechter, Neiri and Pearl on temporal constraint networks extended classical methods of solving constraint satisfaction problems to deal with temporal constraints by including a representation for continuous variables. Koomen proposed an improvement to Allen's interval-based temporal logic in which temporal constraint propagation was localized based on temporal containment rather than clustering methods. He substantiated his claims with empirical results. The final paper of the session described a non-reified temporal logic proposed by Bacchus, Tenenberg and Koomen. This temporal logic allows a system to reason about propositions whose truth value could potentially change as a function of time. They also showed that this logic subsumes Shoham's temporal logic.

Commonsense Theories

Not all of the papers presented at the session on commonsense theories were met with universal enthusiasm, perhaps suggesting that there is little common ground in commonsense research. Ernest Davis' paper dealing with what he sees as a paradox of perception with limited acuity generated the most audience participation. Davis claimed that a paradox can result when an agent knows too much about the limits of its own perception — under some conditions the agent can perform multiple measurements to achieve a precision better than the limits of its inference. He demonstrated a new representation system that insulates the agent from certain information about its instruments.

Judging from audience reaction, Davis was either fixing something that was not broken or didn't justify his work sufficiently, for many thought that the paradox was not a paradox at all, but a natural phenomenon. Whatever the reaction, his work demonstrated an interesting way of formalizing how agents deal with measurements and what they can know about measurements, and deserves further consideration.

Erik Sandewall's paper on describing real-world systems with a combination of logic and differential equations was also notable for the comments that it generated, though in this case there was some question about whether the paper described anything substantially original. Sandewall's work was criticized on the grounds that it appeared to be a formalization of the problem solving methods common to basic Newtonian mechanics problems. In that sense it may not have solved new problems, but it formalizes methods for solving problems for which we typically use "commonsense".

The remaining two presentations were less controversial. Wlodek Zadrozny presented an extension of set theory that captures notions of cardinality and well-ordering that appear often in commonsense reasoning, but cannot be expressed in classical set theory. Randell and Cohn developed a formal notation for describing the topology of compound objects, using the feeding cycle of an amoeba as an example.

Learning

As the title would suggest, "Case-Based, Analogical and Inductive Reasoning" included an eclectic collection of papers. Some observers noted that there were not as many papers from the

machine learning (ML) community as one might have expected, given their obvious concern for KR related issues. Perhaps this is best explained by the fact that many venues for ML research already exist. The emphasis on case-based and analogical reasoning at KR'89 was quite surprising — these two topics tend not to dominate the ML conferences to quite the same degree. Perhaps this was due to the KR+R emphasis of KR'89.

Two papers in this session addressed the topic of case-based reasoning. Yoshiteru Ishida proposed a particular generalization strategy for organizing a case memory based upon minimizing the number of slots needed to represent the cases in an inheritance net. Indicating that this strategy turns out to be intractable, he introduced an incremental algorithm intending to approximate the minimal structure. Ishida then applied his memory organization technique in a system that performs case-based diagnosis. In this implementation cases were indexed by symptoms and also by the causal relations between symptoms.

The paper by Phyllis Koton and Melissa Chase was most appropriate for KR'89. They presented a formalization of the task of reasoning from a case-base, accomplished by establishing a mapping between case memory structures and normal default theories. They presented a transformation procedure that could map a case memory structure into a normal default theory characterizing the inferences sanctioned by the case memory.

Two papers in this session addressed the related topic of analogical reasoning. R. P. Loui's made the claim that analogical reasoning can be understood and justified strictly in terms of Kyburg's system of statistical reasoning. Loui views the process of selecting a particular analogy simply as selecting an appropriate reference class on which to base inductive judgement. He presented reconstructions of two recently proposed theories of analogical inference (one due to Clark and the other to Russell), providing what he felt was a more satisfying account of both.

Debbie Leishman addressed the problem of choosing the "best" analogy that holds between two knowledge structures. She chose to represent these knowledge structures with conceptual graphs, arguing that the most plausible analogies corresponded to the minimal common generalizations of two such graphs. She presented an algorithm for computing these minimal common generalizations (and, hence, identifying the "best" analogies) and an accompanying implementation.

Only one paper at KR'89 investigated the traditional machine learning topic of inductive inference. Nicolas Helft's paper critiques the way in which inductive inference has been traditionally formalized. He sees inductive inference as based upon the willingness to assume that the similarities between the observed data are representative of the rules governing them. This led Helft to reformulate induction as a form of closed world reasoning. He presented a formulation where the inductive generalizations of a theory are those formulas which hold in minimal models of the theory and satisfy some additional constraints. He then derived some of the syntactic properties of these generalizations.

Planning

The session on planning and reasoning about action featured five papers on systems that respond to their environment and on representation languages suited for planning. All papers were solid technical contributions, containing little that was either controversial or earth-shaking. Consequently, audience reaction was limited to questions of clarification, and there was no open disagreement with the points of view presented.

Stan Rosenschein identified traditional inference mechanisms as a major impediment to implementing real-time intelligent systems, claiming that the computational cost of performing inference and the difficulty of predicting that cost beforehand makes it very difficult to build knowledge-based systems with predictable response times. He proposed a solution based on finite automata which are synthesized from formal descriptions of an environment and the desired behavior of the automata. These formalizations are declarative, and in principle are as general as "traditional" symbolic logic, but the resulting automata have constant response time, making them suitable for controlling real-time systems.

Mark Drummond challenged two assumptions that he claims limit the generality of planning; that plans are programs and that plans are ordered sets of operators. His plan net formalism is a non-deterministic control program that carries out actions when environmental conditions permit. "Situated control rules" are used to select amongst the actions allowed at any point to ensure the satisfaction of global goals. Together, plan nets and situated control rules allow agents to exhibit both reflexive and goal-directed behavior without completely specifying a procedural plan.

Three examples of representation systems designed for planning were presented. ADL (Action Description Language, presented by Edwin Pednault) attempts to provide much of the expressive power of the situation calculus with the computational and notational simplicity of the STRIPS language. ADL performs very fast "vivid" inferences when the initial state of a system is completely known, while reasoning effectively (albeit at greater cost) when the initial state is not completely specified.

Josh Tenenberg used an extension of IS-A hierarchies to partition a planning system into distinct levels of abstraction. In general, it is easier to generate plans for more abstract cases, but Tenenberg claimed that most previous research failed to keep the different levels of abstraction coherent, a problem he addressed in detail.

In yet another efficiency-expressibility tradeoff, indexical theories allow the specification of objects by their properties instead of their symbol names, making it easier for an agent to specify concepts like "here" and "now". Devika Subramanian and John Woodfill studied the implications of making the situation calculus indexical, which they claim makes planning easier.

The Best Paper Awards

The KR'89 program committee honoured three papers as "exceptional research contributions that exemplify in different ways the most important aspects of the conference."

Hard Problems for Simple Default Logics

by: Henry Kautz and Bart Selman

Henry Kautz and Bart Selman investigated the complexity of reasoning with a number of limited default logics. To delimit complexity boundaries, they defined a partially ordered space of limited default theories. They considered three important problems: finding an extension, determining if a given proposition is true in some extension, and determining if a given proposition is true in all extensions. Surprising negative results (the high complexity of simple three-literal default rules) as well as positive results (a fast algorithm for skeptical reasoning with binary defaults) were reported. In addition, Kautz and Selman provided an intuitive

characterization of sources of complexity in default reasoning. These results impact work on defeasible inheritance hierarchies as well as default reasoning in general.

Three-Valued Formalizations of Nonmonotonic Reasoning and Logic Programming

by: Teodor Przymusinski

Teodor Przymusinski introduced three-valued extensions of the major nonmonotonic formalisms — circumscription, autoepistemic logic, and default logic. He then proved that these extensions are equivalent to the well-founded semantics of arbitrary logic programs. Przymusinski's goal was to establish the class of theories for which natural forms of all major nonmonotonic formalisms coincide. Since this class is equivalent to the class of all logic programs, Przymusinski sees this work as a means of using the computation methods of logic programming to provide a general inference mechanism for nonmonotonic reasoning.

An Episodic Knowledge Representation for Narrative Texts

by: Lenhart K. Schubert and Chung Hee Hwang

Len Schubert and Chung Hee Hwang presented an experimental story understanding system that displays an impressive ability to perform the complete process of understanding narrative texts, from parsing natural language surface text (using a GPSG-style grammar) to the inferences needed for all aspects of story understanding. At the heart of their system lies a logic that is able to represent in a comprehensible form the content of most English sentences and most of the world knowledge needed for narrative understanding. The logic has facilities for restricted quantifiers, propositional attitudes, predicate modifiers, nominalized predicates, episodic variables (episodes are events, actions, situations, etc. and can be used to represent temporal and causal relationships) and generic conditionals (via probabilistic inference rules). The system is capable of performing narrative inference such as causal connection and temporal succession, and the explanatory and simulative inference needed to reason about such things as the motives of agents in a story.

Apart from its current ability, the authors also stressed the principled way in which the system was developed. Unlike some of the other story understanding systems that have appeared in the literature, their system is modular, explicit, and transparent, in that all knowledge (whether it be linguistic or domain knowledge) is encoded declaratively in a clear, concise, and analyzable manner. Additionally, the control structures for parsing and inference are kept separate from this knowledge. They sketched out glimpses of the semantics for their logic, and although there are still gaps in the formal semantics, they do point out that this is true of the semantics for any situation theory.

The Finale

The conference concluded with a symposium entitled "Against Representation: The Opposition Speaks." What was touted as a confrontational debate over controversial aspects of knowledge representation and reasoning turned into a convivial discussion where it was not always clear which side the speakers were on! Discussions included: Hinton vs. Bobrow on "Connectionist Symbol Processing," Rosenschein vs. McDermott on "No Representation Without Information," and Perry vs. Moore on "Intelligence is Attunement to Incremental Information." Despite the lack of verbal fist fighting, the symposium provided an interesting conclusion to an excellent conference.

Acknowledgements

Special thanks to Diane Horton for providing the authors with her excellent notes from the conference. We would also like to thank those who read an original draft of this article for their comments.

Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning. Edited by Ronald J. Brachman, Hector J. Levesque, and Raymond Reiter, Published by Morgan Kaufmann Publishers, P.O. Box 50490, Palo Alto CA 94303, Distributed in Canada by John Wiley and Sons. ISBN 1-55860-032-9

Book Reviews

SPECIAL REVIEW SECTION

Recent representations of mind: Mental models, computer models, and computational psychology.

Computer models of mind

Margaret Boden
(University of Sussex)

Cambridge, England: Cambridge University Press, 1988, xi+289 pp
Hardcover, ISBN 0-521-24868-X, US$42.50;
Paperback, ISBN 0-521-27033-2, US$15.95

and

The computer and the mind

Philip N. Johnson-Laird
(MRC Applied Psychology Unit, University of Cambridge, England)

Cambridge, MA: Harvard University Press, 1988, 448 pp
Hardcover, ISBN 0-674-15615-3, US$29.50

Reviewed by
David V. Reynolds
University of Windsor

For someone without an extensive background in mathematics or computer science, the literature on computational theories of cognition can be mind-numbing. When I found that *The computer and the mind* was described by its author, Philip Johnson-Laird, as a book "that could be understood by anyone who is interested in the scientific explanation of mental phenomena," I bought it. Before this I had been looking forward to the publication of Margaret Boden's *Computer models of mind* because I had found the first edition of her earlier *Artificial intelligence and natural man* (1977) a clear, systematic introduction to artificial intelligence. As it turned out, I read both Johnson-Laird's and Boden's new books. Judged by their titles alone, they appear to be very similar. Although they both address issues related to computer and mind, they are quite different.

In the preface to *Computer models of mind*, Boden writes that her aim is to make clear how psychologists "can use computational ideas and computer-modeling to further their research" and "to explore the theoretical diversity of work in computational psychology, and the various controversial philosophical assumptions that underlie it". But what is computational psychology? Boden gives examples of what she means and outlines the intellectual history that brings her to what she calls a "minimal definition" of computational psychology. Her definition (paraphrased) is that computational psychology:

(1) assumes that functional mental states can be defined by precisely specifiable computational procedures,
(2) defines mind as a representational system, and
(3) considers the neuroscience component in functional terms — that is, focuses on what it does, not what it is.

All very well, but this definition needs further explanation and elaboration. And this, of course, is the main purpose of her book. She completes the introduction with a clear outline of the strategy that she will use. In the next two chapters on vision, she first establishes the philosophical context and then presents a critical exposition, defining terms as she proceeds, of computer vision and visual images. For the reader interested in further detail, Boden refers to the expanded 1987 edition of *Artificial intelligence and natural man* (1977). Next, she explains and critiques connectionist models of vision, focusing on the work of Marr and Hinton.

Boden devotes the next two chapters to natural language parsing and meaning. These tutorials analyze fundamental issues in the computational theory of language and include Chomsky (of course), definitions and examples of augmented transition-networks, the autonomy of syntax issue, Schank's semantic primitives, Dyer's BORIS, and Johnson-Laird's mental model theory applied to psycholinguistics.

She then turns to reasoning, with a thoughtful review and a critique of the production system approach of Newell and Simon. This is followed by a very detailed description and evaluation of Johnson-Laird's mental models approach. In the following chapter she points out that progress in computer modeling of human development and concept learning has been slow, summarizes recent theoretical and applied contributions (Fodor, Quinlan, Winston, Bruner) and then turns to skills and task analysis. Here Boden provides a detailed analysis of Anderson's ACT* model of cognition. (ACT stands for Adaptive Control of Thought and the asterisk is pronounced "star".) She then discusses Longuet-Higgins' computer model of learning and concludes with a discussion of whether or not this diversity of models implies that learning is actually several distinct processes.

Connectionist approaches to learning are addressed in the final part of the chapter. The early contributions summarized include Hebb's cell assembly, Rosenblatt's perceptron, Selfridge's pandemonium, and Minsky's and Papert's work that showed the limitations of single layer perceptrons. Boden then provides a lucid explanation and a detailed discussion of Hinton's Boltzmann machine. She completes this section with an example from McClelland and Rumelhart (1986) on how connectionist theories can be used to model the past tense during language acquisition. She concludes by reminding us that connectionist (PDP) models in themselves do not automatically solve theoretical problems fundamental to learning.

Boden begins the final two chapters by asking "Is computational psychology possible?" Since the main aim of the book has been to show how computational psychology *is* possible, her positive reply comes as no surprise. But before proceeding, Boden addresses some of the specific objections that might be made to a computational psychology. She calls upon material presented throughout the book to support her view, adding that new developments in connectionism can be regarded as examples of computational psychology.

What criticisms do I have of Boden's book? After reading through my notes a second time, very few. Most of my initial questions were answered as I continued to read. Of course, there is the occasional minor frustration, one of which is writing style. Boden loves parentheses and (for example) when she uses them to excess (as she sometimes does) the consequence, when combined with a complex sentence structure (I exaggerate, but you see what I mean) can result in a kind of cognitive whiplash. This tendency, however, is more than offset by the vivid descriptions and lively expressions sprinkled throughout her writing. When, for example, she describes a "filtering algorithms that crawled around the picture deleting . . ." I feel as though I know the nature of this algorithm. When she defines depth-first search as one that "follows through to the bitter end" she captures the essence of it. Finally, I am predisposed to agree with someone who punctuates an even-handed analysis with "This question cannot wait for an answer until the day when we have an all-singing, all-dancing, Boltzmann machine".

In sum, Margaret Boden has written a challenging and systematic analysis of what she terms computational psychology. She explains, criticizes and engages the material. The result is a scholarly, comprehensive, up-to-date critique of the relevant philosophical and psychological literature.

Johnson-Laird's *The computer and the mind* is described as an introduction to cognitive science. In his preface, the author proposes to describe what cognitive science is, within the context of the mathematical theory of computability. Further, he aims to do this in a way that can be understood "by anyone who is interested in the scientific explanation of mental phenomena". Brave words.

In the first part of his book, the author briefly traces the history of the study of mind including introspection, unconscious processes, behaviourism, mentalism and "the rise of cognitive psychology". The section devoted to cognitive psychology is very brief, but includes references to Hebb (cell assembly), George Miller (information processing), Broadbent (communication channels), von Neumann (cybernetics), Turing (computability theory), Bruner (cognitive theory), Newell and Simon, Lashley and Chomsky. Johnson-Laird then explains that an important requirement of a psychological explanation is to keep assumptions to a minimum. And, writes Johnson-Laird, the theory of computability does this.

The next two chapters are as clear an introduction to computational theory as you are likely to find. Johnson-Laird first defines symbols in the "external world", symbols in computers, and then makes the case that mental symbols are appropriate for computational modeling. In the process, he introduces binary counting, which is important in showing how computers process information. Finally, he proposes that perception establishes the link between the actual world and mental models of this world.

Johnson-Laird then suggests that readers who find the mathematics "too austere" may skip Chapter 2 entitled "Computability and Mental Processes". I strongly urge the opposite — don't skip this chapter. It is short, clear, and one of the most important discussions in the book. I have not found a more straightforward introduction to computational theory and I have spent some time looking. If you skip it you will have lost an opportunity to discover that the fundamentals of computational theory are comprehensible. Since the rest of the book is about computational theories of mind, read the chapter.

Following this introductory section, the author reviews progress in computational theories of vision. He includes processing of the visual image (including Marr and others), depth vision (Marr again, Gibson, and others), and visual shapes and imagery (including Waltz, Marr, Shepard and Cooper, and Kosslyn). In the next section Johnson-Laird addresses learning, memory, mental architecture and action (movement). In the chapter on learning he discusses the limits on what can, in principle, be learned. The chapter on memory includes an analysis of what is needed for memory, a brief discussion of how a digital computer does this, and then a selection of experiments on human memory suggesting that there are different kinds of memory. The author concludes that much remains to be understood about the nature and sequencing of these memory processes.

The discussion of memory sets the stage for two really fine chapters on mental architecture. The first discusses the production system approach that characterizes the work of Newell and Simon and more recently Anderson's ACT* theory of cognition. In the next chapter, Johnson-Laird discusses connectionism (parallel distributed processing) using, as an example, McClelland and Rumelhart's connectionist approach to memory. This chapter includes a very good explanation of learning by backward error propagation (Hinton), and concludes by proposing a hybrid system combining production system and connectionist architectures at different levels of knowledge representation. A short chapter on control of movement and robotics completes this section.

The next section, on thinking, which Johnson-Laird calls "Cogitation", includes chapters on deduction, induction, and creativity. The chapter on deduction includes an introduction to predicate logic, expert systems, and discusses several varieties of reasoning. Using reasoning as a focus, Johnson-Laird explicitly argues in support of one of the implicit agendas of the book — mental models. The rationale for the chapter on induction is not very clear, and the discussion is difficult to follow. In the next chapter (called Creation), examples from mathematics, music, and science are used to illustrate creativity. Johnson-Laird concludes that a computational theory of creativity will probably require more than one kind of mental architecture.

The section identified as "Communication" includes a brief introductory chapter on the general nature of communication and one on speech and hearing. These are followed by a chapter on grammar and another on meaning (semantics). The chapter on grammar discusses the difficulties of both the behaviourist theory of language acquisition (presumably Skinner's but the author doesn't mention names) and Chomsky's transformational theory. After summarizing what is required of a sentence parser, the author suggests the possibility of a connectionist architecture to deal with grammar. He then turns to the subject of meaning in the context of language. After discussing semantic networks, meaning postulates, and formal semantics, he suggests that procedural semantics are appropriate for mental models.

In the final section of his book, the author first outlines a theory of conscious and unconscious mind based on computational theory and then, using his mental model of "self", describes how this model can be used to account for emotion, motives, intentions, and free will. Briefly, Johnson-Laird proposes a hierarchical system of parallel processors that, among other attributes, "has access to a partial model of itself". In his theory of mental architecture, conscious mind depends upon "serial processing of explicitly structured symbols" while the unconscious is mediated by "parallel processing of distributed symbolic representations". There is no way to make all of this clear in a brief review, but Johnson-Laird does in the final section of the book. An essay at the end, entitled

"Envoi", is an excellent summary of the points he has made throughout the book and might well be read first before starting the first chapter.

How well does Johnson-Laird succeed in his introduction to cognitive science? In my view, he succeeds very well. He is at his best and clearest when describing the fundamentals of computational theory (Chapter 2 and 3), again when he describes production systems (Chapter 9) and parallel distributed processing approaches to computational theory (Chapter 10), and in the final group of chapters when he summarizes mental models approach.

What would I like to be different in the book? In an introductory book, especially a book with terminology from many different disciplines, I would like to have seen a glossary (call me old-fashioned). Also, a conclusion or summary section is needed at the end of each chapter — some chapters had them, some didn't. Finally, I would have appreciated some mention of the theory of mental models among the book's explicit aims in the preface. This is not just an introduction to cognitive science but an introduction to the author's theory of mental models as well (Johnson-Laird, 1983). The book is more coherent when one is aware of this agenda from the beginning. However, these concerns do not alter my overall view — Johnson-Laird has written a very good introduction to an extremely complex, interdisciplinary area.

What is the best way to use these two books? In my judgement, Johnson-Laird's book would make an excellent advanced undergraduate or graduate level text. The book would also be appropriate for individual study by a motivated graduate student or academic colleague. Boden's *Computer models of mind* is a more difficult book, more appropriate for a graduate seminar, and some prior reading is recommended such as *The mind's new science* (Gardner, 1985), the more recent *Artificial intelligence debate* (Graubard, 1988) (reviewed in *Canadian AI*, July 1989), or, as a matter of fact, the Johnson-Laird book reviewed above.

For the reader who finishes these with a passion for further inquiry, I suggest Pylyshyn's *Computation and cognition* (1984) which presents a prior and somewhat different approach to computational theory in psychology. Finally, for a stimulating and informed examination of "how the mind-brain represents whatever it represents, and of the nature of computational processes underlying behavior", you won't find a better introduction than Patricia Churchland's *Neurophilosophy: Toward a unified science of the mind/brain* (1986).

References

Boden, Margaret A. (1977). *Artificial intelligence and natural man*. New York: Basic Books.

Churchland, Patricia S. (1986). *Neurophilosophy: Toward a unified science of the mind/brain*. Cambridge, MA: The MIT Press.

Graubard, Stephen R. (editor) (1988). *The artificial intelligence debate: False starts, real foundations*. Cambridge, MA: The MIT Press.

Gardner, Howard (1985). *The mind's new science*. New York: Basic Books, Inc.

Johnson-Laird, Philip N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, England: Cambridge University Press.

Pylyshyn, Zenon W. (1984). *Computation and cognition*. Cambridge, MA: The MIT Press.

McClelland, James L. and Rumelhart, David E. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 2: Psychological and biological models*. Cambridge, MA: The MIT Press.

David Reynolds is Professor of Psychology at the University of Windsor. His recent research is concerned with the implications of neuroanatomy for cognitive and neural modeling. In 1988, he received the Outstanding Scholarly Contribution Award at the Fourth International Conference on System Science.

On Being a Machine Volume 1: Formal Aspects of Artificial Intelligence

Ajit Narayanan
[University of Exeter]

(Ellis Horwood series in artificial intelligence foundations and concepts) Chichester, West Sussex: Ellis Horwood, 1988, 200 pp (Distributed in Canada by John Wiley & Sons Canada) Hardbound, ISBN 0-85312-957-6, Cdn$55.50

Reviewed by
Peter Turney,
National Research Council

This book is the first of two volumes. The purpose of Volume 1 is to introduce philosophers to the formal theory of AI. By "formal theory of AI", Narayanan means the theory of computation, formal logic, and model theory. Volume 2 will introduce formal theoreticians to the philosophy of AI. By "philosophy of AI", he means the philosophical analysis of concepts and features of AI. Although it is not explicitly stated, my impression is that these books are intended to serve as text books for a third or fourth year undergraduate course in the philosophy of AI.

Ajit Narayanan is Head of the Department of Computer Science at the University of Exeter. He has a Ph.D. in Philosophy, also from the University of Exeter. His dual background is strongly apparent in his book.

Both volumes of *On being a machine* are organized around Alan Turing's well-known paper, "Computing Machinery and Intelligence" (1950). Chapter 1 is an extensive discussion of this paper. Readers will want to refer to Turing's paper as they read Narayanan's interpretation. It would have been convenient to have a copy of Turing's paper included in the book, but the paper is readily available elsewhere (see references below).

Turing listed nine objections to the claim that machines can think. He then tried to counter each of the objections. Chapter 1 discusses all nine of the objections briefly. Volume 1 of *On being a machine* deals in depth with three of these objections. Volume 2 will be a sustained treatment of two of the six remaining objections. The other four objections are relatively minor, and they are quickly dispensed with in Chapter 1.

Chapter 2 is an analysis of Turing's sixth objection, Lady Lovelace's objection. Lady Lovelace, discussing Babbage's analytical engine, states, "The analytical engine has no pretensions to originate anything. It can do whatever we know how to order it to perform." As Narayanan phrases it, "Machines cannot think because the results they produce are not original" (p. 40).

Narayanan points out that John Searle's "Chinese-room argument" (1980) is essentially a modern, expanded version of Lady Lovelace's objection (p. 84). Searle's argument runs as follows: Suppose we have a program that can converse intelligently in Chinese. Imagine that we execute this program by pencil and paper instead of executing it on a computer. We are locked in a room, with only pencils, paper, and a print-out of the program. We receive input and generate output in the form of written Chinese, which we pass through a slot in the wall. How can we have an intelligent dialogue in Chinese, if we do not understand Chinese, and we merely follow the rules in the print-out of the program? Searle views this as a *reductio ad absurdum* argument that we cannot have a program for intelligent conversation.

The bulk of Chapter 2 is an introduction to the theory of computation, beginning with definitions of functions and Turing machines. Narayanan argues that a proper understanding of Lady Lovelace's objection requires an understanding of the nature of computer programs. He argues that the force of Searle's Chinese-room argument comes from the ambiguity of the intuitive notion of following rules (p. 87).

It is interesting to contrast this analysis of Searle's argument with the analysis by Hofstadter (1982) in *The mind's I*. Hofstadter argues that the human who executes the program by hand is akin to the central processing unit of a computer. As such, the human is just one element of a larger system, including the print-out of the program, the pencils, and the paper. If there is any understanding of Chinese or any creativity in the conversation, it is a property of the system as a whole, not merely the human central processing unit. I prefer Hofstadter's analysis to Narayanan's. Narayanan does not discuss Hofstadter's view.

Chapter 3 studies Turing's third objection, the mathematical objection. Narayanan bases his examination on a paper by J.R. Lucas (1961), which elaborates on the mathematical objection. Gödel proved that formal logic has certain inherent limitations, and Lucas argues that these limitations must apply to computers. Humans, Lucas argues, do not have these limitations, since they can understand Gödel's proof.

Most of Chapter 3 is an introduction to formal logic, from propositional calculus to Gödel's theorem. In his analysis of Lucas's argument, Narayanan focuses on the claim that Gödel's proof applies to computers. He concludes with a question (p. 144):

> What precisely is the relationship between [a] formal system or theory, on the one hand, and the computer or program that is supposed to be an instantiation of it, on the other hand, and are the formal constraints of formal systems and theories necessarily also constraints of the computer or program?

Again, it is interesting to compare this with Hofstadter's response to Lucas. While Narayanan focuses on the claim that the limitations of formal logic apply to computers, Hofstadter focuses on the claim that the limitations do not apply to humans. He argues that our understanding of Gödel's proof does not let us transcend the limitations of formal logic. Again, I prefer Hofstadter's analysis to Narayanan's. In this case, Narayanan does mention Hofstadter's view, apparently with approval (pp. 142-143).

The fourth and final chapter is mainly concerned with an introduction to model theory. Models (semantics) are presented for propositional calculus, first-order predicate calculus, tense (temporal) logic, and modal logic. The discussion of model theory is motivated by Turing's eighth objection, the argument from informality of behaviour, expressed by Turing as follows:

> If each man had a definite set of rules of conduct by which he regulated his life he would be no better than a machine. But there are no such rules, so men cannot be machines.

Narayanan interprets this argument using propositional calculus (p. 43):

> P = Each man has a definite set of rules of conduct by which he regulates his life.
>
> Q = Man is no better than a machine.
>
> $P \vdash Q$; $\neg P$; therefore $\neg Q$.

In English: If each man has a definite set of rules of conduct by which he regulates his life, then man is no better than a machine. But it is not true that each man has a definite set of rules of conduct by which he regulates his life. Therefore man is better than a machine. As Narayanan remarks (pp. 42-43), this argument is not valid.

Narayanan writes, "This part of Turing's paper is very curious indeed" (p. 42). I suggest a different interpretation of the above argument, which makes sense of this part of Turing's paper. Turing makes an allusion to syllogistic logic. Let us cast the above argument in the form of a syllogism:

> Mx = x is regulated by a definite set of rules of conduct.
>
> Px = x is a machine.
>
> Sx = x is a man.
>
> $(\forall x)$ $(Mx \vdash Px)$;
>
> $(\forall x)$ $(Sx \vdash \neg Mx)$;
>
> therefore $(\forall x)$ $(Sx \vdash \neg Px)$.

In English: Anything that is regulated by a definite set of rules of conduct is a machine. No man is regulated by a definite set of rules of conduct. Therefore no man is a machine. Turing notes that this syllogism is not valid. Turing observes that the argument also confuses "rules of conduct" and "laws and behaviour". He points out that, if we change Mx to

> Mx = x is regulated by a definite set of laws of behaviour.

then we can make the syllogism valid by replacing $(\forall x)$ $(Mx \vdash Px)$ with $(\forall x)$ $(Px \vdash Mx)$, that is, all machines are regulated by a definite set of rules of conduct. Now, however, $(\forall x)$ $(Sx \vdash \neg Mx)$, no man is regulated by a definite set of laws of behaviour, is no longer plausible, as Turing explains:

> The only way we know of for finding such laws is scientific observation, and we certainly know of no circumstances under which we could say, "We have searched enough. There are no such laws."

Those who wish to support the argument from informality of

behaviour must persuade us that $(\forall x)$ $(Sx \vdash \neg Mx)$ is true.

This is where Narayanan applies model theory. Suppose we wish to show that $(\forall x)$ $(Sx \vdash \neg Mx)$ is false. That is, we wish to show that there is at least one man who is regulated by a definite set of laws of behaviour. What form might these laws of behaviour take? Narayanan suggests that they might take the form of logical rules. For example, perhaps the behaviour of understanding sentences is governed by something like model theory (p. 184). Turing, however, does not have to show that $(\forall x)$ $(Sx \vdash \neg Mx)$ is false. To defeat the argument from informality of behaviour, it is only necessary to show that $(\forall x)$ $(Sx \vdash \neg Mx)$ is not known to be true. Narayanan does not appreciate this, because he has misinterpreted Turing's argument (pp. 41-43).

Aside from the above remarks on the content of *On being a machine*, I should comment on the disturbing number of spelling mistakes in the book. I estimate that there is an average of one error per page. In general, the book seems more like a printer's proof than a final product.

If I were to teach a course in philosophy of AI, I would prefer to use *Minds and machines* (Anderson 1964), *Artificial intelligence and natural man (Boden 1977), and The mind's I* (Hofstadter and Dennett 1982), or some subset of these, rather than *On being a machine*. Narayanan's book is distinguished from these other books by the introductions to the theory of computation, formal logic, and model theory. The attempt to motivate the study of these formal topics by considering their applications to philosophy is interesting, but ultimately unsuccessful. I would prefer to treat these subjects in more depth, in a course that concentrates on them exclusively.

References

Anderson, A.R. (editor) (1964). *Minds and machines*, Englewood Cliffs, NJ: Prentice-Hall.

Boden, Margaret A. (1977). *Artificial intelligence and natural man*. New York, NY: Basic Books.

Collins, Allan M. and Smith, Edward E. (editors) (1988). *Readings in cognitive science: A perspective from psychology and artificial intelligence*. San Mateo, CA: Morgan Kaufmann.

Feigenbaum, Edward A. and Feldman, Julian (editors) (1963). *Computers and thought*. McGraw-Hill.

Hofstadter, Douglas R. and Dennett, Daniel C. (editors) (1982). *The mind's I: Fantasies and reflections on self and soul*. New York, NY: Bantam Books.

Lucas, J.R. (1961). "Minds, Machines and Gödel", *Philosophy*, Vol. XXXVI, 112-129. Reprinted in Anderson (1964).

Pylyshyn, Zenon W and Bannon Liam (1989). *Perspectives on the computer revolution* (second edition). Norwood, NJ: Ablex.

Searle, John R. (1980). "Minds, Brains, and Programs", *The Behavioural and Brain Sciences*, Vol. 3, 417-424. Reprinted in Hofstadter and Dennett (1982).

Turing, Alan M. (1950). "Computing Machinery and Intelligence", *Mind*, Vol. LIX, 433-460. Reprinted in Anderson (1964), in Hofstadter and Dennett (1982), in Feigenbaum and Feldman (1963), in Collins and Smith (1988), and in Pylyshyn and Bannon (1989).

Peter Turney is a visiting fellow at the Laboratory for Intelligent Systems, in the Division of Electrical Engineering of the National Research Council. He is currently involved in research on machine learning. He has a Ph.D. in philosophy from the University of Toronto.

OTHER REVIEWS

The fifth generation fallacy: Why Japan is betting its future on artificial intelligence

J. Marshall Unger
[University of Hawaii]

New York, Oxford: Oxford University Press, 1987, x+230 pp Hardbound, ISBN 0-19-504939-X, Cdn$37.50

Reviewed by
Nick Cercone,
Simon Fraser University

J. Marshall Unger's *The fifth generation fallacy* describes the difficulty of using various subsections of Japanese script during the input, storage, and output stages of word processing. Unger demonstrates that Japanese scripts, particularly kanji, are not amenable to computation, because characters are too numerous and ambiguous. He argues that artificial intelligence cannot solve these problems, and recommends three script-control measures that he says will remove all these obstacles to complete Japanese utilization of computing power.

The lengthy introduction explains the entire content of Unger's argument, which is then repeated in detail throughout the six chapters of the book. Topics include: current writing practice in Japan, the source of Japanese attachment to their script, a history of script reform, effects of the script upon education and literacy, influence of the scripts on data processing applications, Japanese script input devices and word processing technology, Japanese software production, and the goals and administration of the Japanese MITI as they apply to script processing.

Though the information presented is interesting and abundant, it is not cohesively organized. Where the intent is to convince us about the turbidity of Japanese script, this style of composition is effective, because it leaves us convincingly confused. However when the reader would like to retain some information, for example concerning the progress of script reform, the only method of retention would be to buy the book and use it as a reference. The book's organization is particularly disappointing because Unger is a professor of East Asian languages and literatures (at the University of Hawaii), so one anticipates proficient composition. Chapters and subsections frequently begin not with the topic with which they are titled, but with yet another attempt to express the content of the preceding section, as though Unger is still trying to find the correct words and vaguely recognizes that he has not yet succeeded. He resorts to bellicose diction to strengthen his undisciplined theses, especially when interpreting the "cockeyed optimism" of strong artificial intelligence.

The central idea of the book is that the Japanese have disregarded other national discussions concerning the potential of artificial intelligence, and, subscribing to the strong AI paradigm that intelligence is nothing more than the manipulation of formal symbols, have mistakenly dedicated research resources in expectation of a successful Japanese script-processing solution. The book is apparently directed to a non-technical readership, as

its glossary includes helpful definitions such as these:

> Hardware: The physical elements that comprise a computer system. In the opinion of most software developers, the easy part of the system to fix.
>
> Software: Computer programs; the coded instructions executed by a computer. In the opinion of most hardware engineers, the easy part of the system to fix.

For the non-technical audience in particular, it is misleading that Unger neglects the decade of technological developments since the philosophical discussions to which he refers. Experiments in processing natural language have been a logical consequence of the development of information systems throughout the world. The Japanese attempt at natural language processing appears more absurd when the failures and partial successes of other research efforts are ignored.

Contrary to the impression of the title, no fifth-generation initiative other than script processing is discussed. While insight into the causes of low white-collar productivity are promised, Unger simply states that Japanese script is inefficient for document production and management. This leaves unasked and unanswered interesting questions, such as: What applications of information technology can compensate within an economy for inefficiency in word processing? Will other fifth-generation initiatives produce the required compensation? How has a ruinously unproductive white-collar labour force achieved a globally pre-eminent economy?

In the introduction, Unger explains that he was motivated to write his book by other bestsellers on the topic of fifth-generation research that report Japan's technological successes and warn other industrialists about the significance of the challenge. With reassuring, though incorrect, observations such as "Many of the publications coming out of ICOT are mostly recapitulations of visitors' lectures" (p. 188), *The fifth generation fallacy* may be sufficiently soothing to reach the non-fiction best-seller list. However, reading as it does, and organized as it is not, purchasers could suspect that they are buying a cleverly packaged first draft. Taken as raw material, the book would stimulate discussion around a coffee table, or supply references for further thought on what is really an intriguing topic.

Nick Cercone is the former president of the Canadian Society for Fifth Generation Studies.

Logic-based knowledge representation

Peter Jackson, Han Reichgelt, and Frank van Harmelen (editors)
[McDonnell-Douglas Research, University of Nottingham, and University of Edinburgh]

Cambridge, MA: The MIT Press, 1989, xv+255 pp (MIT Press series in logic programming) Hardbound, ISBN 0-262-10038-X, US$35.00

Reviewed by
Norma Welch
MaSyCom Consultants Inc.

Logic-based knowledge representation presents a series of papers describing recent research in logic programming at the University of Edinburgh. The book is in two sections: The first describes Socrates, an abstract knowledge-based architecture developed at Edinburgh, and the second presents several papers in theorem-proving methods for various logics. The topics are all clearly introduced and all papers reference the extant literature extensively. However, a strong background in logic programming is required of the reader.

The introduction presents arguments in favour of using logic to represent knowledge, rather than developing expert systems in one of the rich but unstructured programming environments such as LOOPS or KEE. It further argues for the separate explicit representation of control (meta-level) knowledge and domain (object-level) knowledge such as that provided by the Socrates architecture. With this abstract architecture, the domain knowledge can be expressed declaratively in the most appropriate formal logic. Different control regimes may then be used to apply the domain knowledge to different tasks (e.g., training, simulation, problem solving). Various control strategies can be tested for efficiency without altering the domain knowledge. A side benefit of this architecture is that explanations can refer to both the meta-level and the object level independently, thus providing the capability for deeper explanations. The major disadvantage of the approach is the computational inefficiency caused by the extra level of interpretation at the meta-level. Some of the work described in part II of the book was motivated by the need to improve the computational efficiency of the approach. The remainder of the research explores extensions to modal, nonmonotonic and doxastic logics.

Part I

Chapter 2 provides a useful discussion and classification of meta-level architectures from the literature. The classification focuses on whether most of the inferencing occurs at the object level or the meta-level, and whether the meta-level language is declarative or procedural. Completeness and strictness of the various architectures are also discussed.

Chapter 3 details the architecture of Socrates. In using Socrates, the knowledge engineer specifies the logic representation language for the domain knowledge (object level language), the inference rules of the logic, the control regime, and the domain knowledge. The object-level language may be many sorted and may be modal or temporal. The meta-level language may be either declarative or

procedural. The chapter includes two sample codes that implement a local-best-first non-exhaustive backward-chaining control regime: one written in a sorted logic and the other in Lisp. (Socrates has been written in Common Lisp.)

Chapter 4 provides examples to show the power and utility of the Socrates architecture. The examples illustrate various control regimes: backward chaining for an investment advisor, strategically guided proof for a medical diagnoser, forward chaining for an office configurer, and heuristic best-first search for a route planner. The chapter concludes with a solution to Schubert's steamroller, a well-known challenging problem for theorem provers. All examples are carefully described and include discussions of the control strategy and fragments of code.

Part II

The research described in Part II of the book aims to improve the efficiency of the Socrates approach and to extend its expressive power by the use of modal, temporal, and doxastic logics.

Chapter 5 discusses partial evaluation as a technique for improving efficiency. A detailed description of the technique is followed by discussion of its limitations in optimizing meta-level interpreters. Heuristic solutions to the limitations are explored.

Chapter 6 proposes assertion-time inference as a solution to the efficiency problems. The authors present tableau-type proof systems using sequents, and prove the correctness of the method for first order predicate calculus. This work is extended to modal predicate logic in Chapter 8, which includes the theoretical justification of a general proof method for modal predicate logic and a description of its implementation in Lisp. A sample run of the system is given in the appendix.

Temporal and modal logics are discussed in Chapter 7: the authors provide a clear discussion of the need for temporal logics and an assessment (with examples) of the required properties of such formalisms. They present three possible formalisms and argue in favour of modal treatments rather than first-order treatments. Finally, Chapter 9 provides a modal proof method for doxastic reasoning; this chapter would have been improved by the inclusion of a concrete example.

In conclusion, *Logic-Based Knowledge Representation* will be of interest to logic theoreticians interested in efficient and expressive methods of using logic for knowledge representation. The authors set their work clearly in the perspective of other research. The book is clear, well-written and almost free of typographical errors.

Norma Welch is a consultant specializing in AI. Previously, she worked on AI applications for Canadian National Railways.

Books Received

Books listed below that are marked with a + will be reviewed in a future issue. Reviewers are still sought for those marked with a *. Readers who wish to review books for the journal should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

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

The society of text: Hypertext, hypermedia, and the social construction of information
Edward Barrett (editor)
[MIT]
Cambridge, MA: The MIT Press, 1989, xix+459 pp
(Information systems series) Hardbound,
ISBN 0-262-02291-5, US$37.50

+Outils logiques pour le traitement du temps: De la linguistique à l'intelligence artificielle
Héléne Bestougeff et Gérard Ligozat
[Université Paris 7 et CNRS]
Paris: Masson, 1989, 272 p. (Études et recherches en informatique) Couverture rigide, ISBN 2-225-81632-8

Working models of human perception
Ben A.G. Elsendoorn and Herman Bouma (editors)
[Institute for Perception Research, Eindhoven] London:
Academic Press, 1989, xiii+514 pp Hardbound,
ISBN 0-12-238050-9

***Blackboard systems**
Robert Engelmore and Tony Morgan (editors)
[Stanford University and Systems Designers plc]
Wokingham, England: Addison-Wesley, 1988, xviii+602 pp
(The insight series in artificial intelligence)

Speech input and output assessment: Multilingual methods and standards
A.J. Fourcin, G. Harland, W. Barry, and V. Hazan (editors)
[University College London]
Chichester, England: Ellis Horwood, 1989, 290 pp (Ellis Horwood books in information technology) Distributed by John Wiley & Sons Hardbound, ISBN 0-7458-0651-1 and 0-470-21439-2, US$67.95

Natural language processing in PROLOG: An introduction to computational linguistics
Gerald Gazdar and Chris Mellish
[University of Sussex and University of Edinburgh]
Wokingham, England: Addison-Wesley, 1989, xv+504 pp
Hardbound, ISBN 0-201-18053-7, Cdn$45.15

Natural language processing in POP-11: An introduction to computational linguistics
Gerald Gazdar and Chris Mellish
[University of Sussex and University of Edinburgh]
Wokingham, England: Addison-Wesley, 1989, xv+524 pp
Hardbound, ISBN 0-201-17448-0, Cdn$45.15

Natural language processing in LISP: An introduction to computational linguistics
Gerald Gazdar and Chris Mellish
[University of Sussex and University of Edinburgh]
Wokingham, England: Addison-Wesley, 1989, xv+524 pp
Hardbound, ISBN 0-201-17825-7, Cdn$45.15

Introduction to Lisp and symbol manipulation
Sharam Hekmatpour
[Open University]
Hemel Hempstead: Prentice-Hall U.K., 1988, xv+301 pp
Paperbound, ISBN 0-13-486192-2

The recognition of speech by machine—A bibliography
Arthur S. House
[Institute for Defense Analysis]
London: Academic Press, 1988, vii+498 pp Hardbound, ISBN 0-12-356785-8, US$49.00

+Artificial intelligence vocabulary / Vocabulaire de l'intelligence artificielle
Noël Lazure
[Translation Bureau, Department of the Secretary of State, Canada]
Ottawa: Canadian Government Publishing Centre, 1989, xv+1217 pp (two volumes) Paperbound, ISBN 0-660-54136-X, Cdn$33.95, US$40.75

Knowledge acquisition from text and pictures
Heinz Mandl and Joel R. Levin (editors)
[University of Tübingen and University of Wisconsin]
Amsterdam: North-Holland, 1989, xv+329 pp (Advances in psychology 58) Hardbound, ISBN 0-444-87353-8, US$113.25, Dfl 215.-

The Paralation model: Architecture-independent parallel programming
Gary W. Sabot
[Thinking Machines Corp]
Cambridge, MA: The MIT Press, 1988, xiii+238 pp (Originally the author's PhD thesis, Harvard University, 1988) (The MIT Press series in artificial intelligence) Hardbound, ISBN 0-262-19277-2, US$27.50

Automatic text processing: The transformation, analysis, and retrieval of information by computer
Gerard Salton
[Cornell University]
Reading, MA: Addison-Wesley, 1989, xiii+530 pp (Addison-Wesley series in computer science) Hardbound, ISBN 0-201-12227-8

+The Facts On File dictionary of artificial intelligence
Raoul Smith (editor)
[Northeastern University]
New York: Facts On File, 1989, 211 pp Hardbound, ISBN 0-8160-1595-3, Cdn$32.95, US$24.95

Parallel logic programming techniques
Stephen Taylor
[California Institute of Technology]
Englewood Cliffs, NJ: Prentice Hall, 1989, xiv+205 pp Hardbound, ISBN 0-13-648767-X

+Knowledge-based programming
Enn Tyugu
[Institute of Cybernetics, Estonian Academy of Sciences, Tallinn, USSR]
Glasgow: Turing Institute Press and Addison-Wesley, 1987, xii+243 pp (Turing Institute Press knowledge engineering tutorial series) Hardbound, ISBN 0-201-17815-X

+Artificial intelligence: A handbook of professionalism
Blay Whitby
Chichester, England: Ellis Horwood, 1988, 194 pp (Ellis Horwood series in artificial intelligence foundations and concepts) Distributed in Canada by John Wiley and Sons Canada Ltd Hardbound, ISBN 0-7458-0350-4, Cdn$77.95

Lisp
Patrick Henry
Winston and Berthold Klaus Paul Horn [MIT] Reading, MA: Addison-Wesley, 3rd edition, 1988, xxi+611 pp Paperbound, ISBN 0-201-08319-1

Computational Intelligence

Abstracts for 5(2) May 1989

Explicitly Biased Generalization
by Diana Gordon
Navy Center for Applied Research in Artificial Intelligence
Naval Research Laboratory Washington, DC

by Donald Perlis
Department of Computer Science
University of Maryland College Park, Maryland

During incremental concept learning from examples, tentative hypotheses are formed and then modified to form new hypotheses. When there is a choice among hypotheses, *bias* is used to express a preference. Bias may be expressed by the choice of hypothesis language, it may be implemented as an evaluation function for selecting among hypotheses already generated, or it may consist of screening potential hypotheses prior to hypothesis generation. This paper describes use of the third method. Bias is represented explicitly both as assumptions that reduce the space of potential hypotheses and as procedures for testing these assumptions. There are advantages gained by using explicit assumptions. One advantage is that the assumptions are meta-level hypotheses that are used to generate future, as well as to select between current, inductive hypotheses. By testing these meta-level hypotheses, a system gains the power to anticipate the form of future hypotheses. Furthermore, rigorous testing of these meta-level hypotheses before using them to generate inductive hypotheses avoids consistency checks of the inductive hypotheses. A second advantage of using explicit assumptions is that bias can be tested using a variety of learning methods.

An Investigation of Modal Structures as an Alternative Semantic Basis for Epistemic Logics
Sharon J. Hamilton
James P. Delgrande
School of Computing Science
Simon Fraser University Burnaby, British Columbia

In the past, Kripke structures have been used to specify the semantic theory of various modal logics. More recently, modal structures have been developed as an alternative to Kripke structures for providing the semantics of such logics. While these approaches are equivalent in a certain sense, it has been argued that modal structures provide a more appropriate basis for representing the modal notions of knowledge and belief. Since these notions, rather than the traditional notions of necessity and possibility, are of particular interest to Artificial Intelligence, it is of interest to

examine the applicability and versatility of these structures. This paper presents an investigation of modal structures, by examining how they may be extended to account for generalizations of Kripke structures. To begin with, we present an alternative formulation of modal structures in terms of trees; this formulation emphasizes the relation between Kripke structures and modal structures, by showing how the latter may be obtained from the former by means of a three-step transformation. Following this we show how modal structures may be extended to represent generalizations of possible worlds, and to represent generalizations of accessibility between possible worlds. Lastly we show how modal structures may be used in the case of a full first-order system. In all cases the extensions are shown to be equivalent to the corresponding extension of Kripke structures.

Instance-Based Prediction of Real-Based Attributes
Dennis Kibler
David W. Aha
Marc K. Albert
Department of Information and Computer Science
University of California, Irvine Irvine, California

Instance-based representations have been applied to numerous classification tasks with some success. Most of these applications involved predicting a symbolic class based on observed attributes. This paper presents an instance-based method for predicting a "numeric" value based on observed attributes. We prove that, given enough instances, if the numeric values are generated by continuous functions with bounded slope, then the predicted values are accurate approximations of the actual values. We demonstrate the utility of this approach by comparing it with a standard approach for value prediction. The instance-based approach requires neither ad hoc parameters nor background knowledge.

Knowledge Acquisition by Incremental Learning from Problem-Solution Pairs
Stan Matwin
Department of Computer Science
University of Ottawa Ottawa, Ontario

Franz Oppacher
School of Computer Science
Carleton University Ottawa, Ontario

Patrick Constant
Ecole Nationale Superieure des Telecommunications and
Cognitech, Inc. Paris, France

This paper describes LEW (Learning by Watching), an implementation of a novel learning technique, and discusses its application to the learning of plans. LEW is a domain-independent learning system with user-limited autonomy that is designed to provide robust performance in realistic knowledge acquisition tasks in a variety of domains. It partly automates the knowledge acquisition process for different knowledge types, such as concepts, rules and plans. The inputs to the system, which we call "cues", consist of an environmental component and of pairs containing a problem and its solution. Unlike traditional forms of "Learning from Examples", e.g. (Winston 1984), in which the system uses the teacher's answer to improve the result of a prior generalization of an example, LEW treats the problem-solution or question-answer instances, i.e. the cues themselves, as the basic units for generalization.

Explanation and Prediction: An Architecture for Default and Abductive Reasoning
David Poole
Department of Computer Science
University of British Columbia Vancouver,
British Columbia

Although there are many arguments that logic is an appropriate tool for artificial intelligence, there has been a perceived problem with the monotonicity of classical logic. This paper elaborates on the idea that reasoning should be viewed as theory formation where logic tells us the consequences of our assumptions. The two activities of predicting what is expected to be true and explaining observations are considered in a simple theory formation framework. Properties of each activity are discussed, along with a number of proposals as to what should be predicted or accepted as reasonable explanations. An architecture is proposed to combine explanation and prediction into one coherent framework. Algorithms used to implement the system as well as examples from a running implementation are given.

Qualitative Data Modeling: Application of a Mechanism for Interpreting Graphical Data
Sheila McIlraith
Advanced Technologies Department
Alberta Research Council Calgary, Alberta

This paper describes a qualitative technique for interpreting graphical data. Given a set of numerical observations about the behaviour of a system, its attributes can be determined by plotting the data and qualitatively comparing the shape of the resulting graph with graphs of system behaviour models. Qualitative data modeling (QDM) incorporates techniques from pattern recognition and qualitative reasoning to characterize observed data, generate hypothetical interpretations, and to select models that best fit the shape of the data. Domain specific knowledge may be used to substantiate or refute the likelihood of hypothesised interpretations. The basic data modeling technique is domain independent and is applicable to a wide range of problems. It is illustrated here in the context of a knowledge-based system for well test interpretation.

Technical Report

McGill University

Annual Report
January 1988 - December 1988
CIM-88-11

The McGill Research Centre for Intelligent Machines (McRCIM) is concerned with three major areas of research. These are:

1) Computer Perception (vision and speech)
2) Robotics
3) Systems and Control

The total number of graduate students is about 90 and the staff consists of 16 professors from the Department of Electrical Engineering, Mechanical Engineering, the School of Computer Science and the Biomedical Engineering Unit. The projects in this report briefly describe our research during 1988.

Ordering Information:

McGill Research Centre for Intelligent Machines
McGill University
3480 University Street
Montréal, Québec
H3A 2A7

network address: mcrcim@larry.mcrcim.mcgill.edu

World Watch

World Watch on AI Applications and Development is sponsored by the National Research Council's (NRC) Associate Committee on AI. Based on the information provided in the abstracts, the references provided have been selected by the secretariat of the NRC Associate Committee on AI as a representative sample of interest and value to Canadian industry. Abstracts provided are reprinted from "Key Abstracts in Artificial Intelligence" with permission from INSPEC. INSPEC is widely recognized as the leading English-language database covering the published information in the field of physics, electronics and computing. Information contained in the INSPEC services is collected on an international basis from over 4,000 journals and 1,000 conference proceedings. INSPEC is a division of the Institution of Electrical Engineers, Station House, Nightingale Road, Hitchin, Herts, UK. All INSPEC's products and services are available in North America from the INSPEC Dept. IEEE Service Centre, 445 Hoes Lane, P.O. Box 1311, Piscataway, NJ 08855-1331, USA.
Persons wishing to obtain copies of references cited should contact their nearest technical library or the Canada Institute for Scientific and Technical Information (CISTI), NRC, Building M-55, Montreal Road, Ottawa, Ontario K2A 0S2 (Phone: (613) 993-1585, Telex: 053-03115). For on-line ordering, CAN/OLE users may use the CAN/DOC command. Envoy users type "COMPOSE CISTI."

1.0 THEORETICAL ASPECTS

291 Probablistic causal reasoning
T. Dean, K. Kanazawa
Dept. of Comp. Sci., Brown Univ., Providence, RI, USA
Proceedings of the Seventh Biennial Conference of the Canadian Society for Computational Studies of Intelligence, Edmonton, AB, Canada, 6-10 June 1988 (Toronto, ON, Canada: Univ. Alberta Press 1988), p. 125-32

The authors develop a theory of causal reasoning for predictive inference under uncertainty. They emphasize a common type of prediction that involves reasoning about persistence: whether or not a proposition once made true remains true at some later time. They also provide a decision procedure with a polynomial-time algorithm for determining the probability of the possible consequences of a set of events and initial conditions. Problems in dealing with persistence by nonmonotonic temporal reasoning schemes are avoided by the use of simple probability information. The ideas have been implemented in a prototype system that refines a database of causal rules in the course of applying those rules to construct and carry out plans in a manufacturing domain. (17 refs.)

292 Search strategies for conspiracy numbers
N. Klingbeil, J. Schaeffer
Dept. of Comp. Sci., Alberta Univ., Edmonton, AB, Canada
Proceedings of the Seventh Biennial Conference of the Canadian Society for Computational Studies of Intelligence, Edmonton, AB, Canada, 6-10 June 1988 (Toronto, ON, Canada: Univ. Alberta Press 1988), p. 133-9

McAllester's (1987) conspiracy numbers algorithm is an exciting new approach to minimax search that builds trees to variable depth without application-dependent knowledge. The likelihood of the root achieving a value is expressed as that value's conspiracy number: the minimum number of leaf nodes required to change their value to cause the root to change to that value. Initial experience indicates that the algorithm places too much emphasis on depth rather than breadth of search. Several variations on the conspiracy numbers search strategy are reported, each adding an increasing degree of breadth to the search. These ideas have been implemented in a program that solves tactical chess problems. Experiments indicate that the new algorithms are capable of solving 41% more problems than McAllester's original proposal. (13 refs.)

1000 A survey of motion planning and related geometric algorithms
J.T. Schwartz, M. Sharir
(Courant Inst. of Math. Sci., New York Univ., NY, USA.). *Artif. Intell.* (Netherlands), vol.37, no.1-3, p.157-69 (Dec. 1988). (Workshop on Geometric Reasoning, Oxford, UK, 30 June - 3 July 1986).

Surveys recent developments in motion planning and related geometric algorithms, a theoretical research area that has grown rapidly in response to increasing industrial demand for automatic manufacturing systems which use robotic manipulators and sensory feedback devices, and, more significantly, in anticipation of a future generation of substantially more autonomous and intelligent

robots. The survey concentrates on exact algorithmic solutions to the motion planning problem. (71 refs.)

1030 Pluralistic evaluation of belief plausibility and its application to nonmonotonic reasoning
O. Katai, S. Iwai
(Fac. of Eng., Kyoto Univ., Japan). *Int. J. Approx. Reason.* (USA), vol.3, no.2, p.219-37 (March 1989).

A pluralistic evaluation of belief plausibility is introduced by extending the notion of plausibility index introduced by Rescher (1982). The properties of the extended plausibility are analyzed by comparing its properties with those of probability theoretic indexing and Rescher's indexing of beliefs. Also it is shown that the new indexing method can be effectively utilized in truth maintenance of beliefs in complex and dynamically changing situations. The necessity for the extension of the plausibility index is clarified by referring to the method of knowledge organization and maintenance called ATMS (Assumption based truth maintenance system) introduced by De Kleer (1986), which is an extension of Doyle's TMS (1979). Rescher's plausibility index is briefly reviewed. The index is compared with the probability theoretic index of beliefs by referring to a system of logic called preference logic that reflects essential features of the indices. The comparison clarifies the monotonic, inflexible nature of Rescher's index and the need to extend it in order to cope with the nonmonotonic nature of truth maintenance in dynamically changing situations. An extension of Rescher's plausibility index is introduced, and its properties are examined. An application of the indexing method to ATMS and to default reasoning is presented, showing that the proposed indexing method can be effectively utilized in truth maintenance of beliefs in complex and dynamically changing situations. (11 refs.)

1032 A spatial reasoning and decision support system
S. Chen, M. Zhang, W. Zhang
(Dept. of Comput. Sci., North Carolina Univ., Charlotte, NC, USA). Methodologies for Intelligent Systems. *Proceedings of the Second International Symposium*, Charlotte, NC, USA, 14-17 Oct. 1987 (New York, NY, USA: North-Holland 1987), p.9-13

A spatial reasoning and decision support (SRDS) system is developed for military battlefield management and target tracking, environmental planning, land development decision support, and industrial robotics and automation. This system consists of several components: a human-machine interface module, a 2D and 3D spatial database management system, an image processing and understanding module, a spatial and application domain knowledge base, and a LISP-based approximate reasoning and decision support system. In order to illustrate the ideas, the authors present an airport-site planning system. (16 refs.)

1039 Rule-discovery from examples using a combination of syntactic and semantic information
S.Kundu
(Dept. of Comput. Sci., Louisiana State Univ., Baton Rouge, LA, USA). Methodologies for Intelligent Systems. *Proceedings of the Second International Symposium*, Charlotte, NC, USA, 14-17 Oct. 1987 (New York, NY, USA: North-Holland 1987), p.333-40

Given a set of initial-final state pairs (Si1,Si2), 1<i<n, a fundamental problem in machine learning is to find a transformation rule or query Q, if any, which transforms each initial state (Si1) to the corresponding final state (Si2). The author looks at this problem where a state is characterized by a syntactic component and a semantic component. A query Q may refer to both the syntactic and the semantic properties. The author develops a systematic method for the discovery of Q in the context where each (Si) is a two dimensional scene with spatial relationships among its objects. The combination of the syntactic and the semantic characteristics allow the solution of a larger class of problems than is possible using either of the two separately. It also tends to make the solution query Q simpler. (13 refs.)

1048 Steps towards common sense
L. Steels
(Artificial Intelligence Lab., Vrije Univ. Brussel, Belgium). ECAI 88. *Proceedings of the 8th European Conference on Artificial Intelligence*, Munich, West Germany, 1-5 Aug. 1988 (London, UK: Pitman 1988), p.49-54

There is a consensus that common sense remains one of the major areas of intelligence where current methods and techniques of AI break down. Common sense is a rather broad and somewhat diffuse area. This paper focuses on the problem of dealing in an intuitive way with the real world, i.e. problems involving physical objects, time and space, causality, planning, and action execution and monitoring. In an attempt to get a breakthrough, this paper puts forward a number of new hypotheses about the nature of common sense. These hypotheses are rather bold and obviously need further elaboration, yet, they are precise and explicit enough that they can be put to the test by concrete experimentation. The main idea is that common sense involves visual thinking or mental imagery, rather than formal and rational thinking based on a logical axiomatization of the world. The first part of the paper defines the hypotheses. The second part gives concrete examples in world handling. The third part deals with the application of common sense to problem solving. (22 refs.)

753 Self-organizing system obtaining communication ability: primitive model for language generation
K. Nakano
(Dept. of Math. Eng. & Inf. Phys., Tokyo Univ., Japan), Y. Sakaguchi, R. Isotani, T. Ohmori. *Biol. Cybern.* (West Germany), vol.58, no.6, p.417-25 (1988).

As one of the synthetic approaches to brain functions, the possibility is discussed that two intellectual robots could make 'words' for themselves and come to communicate with each other. Associatron, a model for associative memory with a neural network structure, is used as the memory in the brain, and information from the outer world is accumulated in it. Concepts regarding objects and attributes are extracted from the stored information, where the Associatron properties are utilized. The robots give specific names to the concepts and interchange them with each other. When the robots have the same experience at the same time, they modify their own words according to what the other robot says. As they repeat common experiences, the differences between the words become smaller and smaller, and finally the robots agree on a word for each object. At this stage, the robots can exchange information

with their words and each can act accordingly on the word from each other. The study shows that the self-organizing system with the above function can be constructed as a neural network model. (7 refs.)

761 The role of uncertainty principles in inductive systems modelling
G.J. Klir
(Dept. of Syst. Sci., State Univ. of New York, Binghampton, NY, USA). *Kybernetes* (UK), vol.17, no.2, p.24-34 (1988).

The author obtained new insight into the concept of uncertainty by examining the concept within a mathematical framework more general than the classical set theory of probability theory. A generalization of existing theories is actually a current trend in mathematics, as exemplified by the generalizations, from quantitative to qualitative, from continuous to discrete, from functions to relations, from graphs to hypergraphs, from precise analysis to interval analysis, from intolerance to inconsistency, imprecision or singularities to their tolerance, from single-objective to multiple-objective criteria optimisation, from classical geometries (Euclidean as well as non-Euclidean) into fractal geometry, from automata theory into dynamic cellular automata, from two-valued logic into multiple-valued logics, fuzzy logic, or logic of inconsistency, or, as most relevant to the subject of this article, from classical set theory into fuzzy set theory, and from probability measures to fuzzy measures. These generalizations have enriched not only his insights but also his capabilities for modeling the intricacies of the real world. This article is only an outline of the new and challenging development regarding the principles of uncertainty and their role in inductive systems modelling.(35 refs.)

764 Connectionism: past, present, and future
J.B. Pollack
(Dept. of Comput. & Inf. Sci., Ohio State Univ., Columbus, OH, USA). *Artif. Intell. Rev.* (UK), vol.3, no.1, p.3-22 (1989).

Research efforts to study computation and cognitive modeling on neurally-inspired mechanisms have come to be called connectionism. Rather than being brand new, it is actually the rebirth of a research program which thrived from the 40's through the 60's and then was severely retrenched in the 70's. Connectionism is often posed as a paradigmatic competitor to the symbolic processing tradition of artificial intelligence, and, indeed, the counterpoint in the timing of their intellectual and commercial fortunes may lead one to believe that research in cognition is merely a zero-sum game. This paper surveys the history of the field, often in relation to AI, discusses its current successes and failures, and makes some predictions for where it might lead in the future. (70 refs.)

771 Training of a neural network for pattern classification based on an entropy measure
C. Koutsougeras, C.A. Papachristou
(Dept. of Comp. Eng. & Sci., Case Western Reserve Univ., Cleveland, OH, USA).
IEEE International Conference on Neural Networks (IEEE Cat. No.88CH2632-8), San Diego, CA, USA, 24-27 July 1988 (New York, NY, USA: IEEE 1988), vol.1 p.247-54

A neural net model for pattern classification is introduced. Unlike models in which the network topology is specified before training, in this model the network expands during training. The proposed model introduces a novel type of unit (neuron) and a standard treelike feed-forward network topology. The simplicity of the interconnection pattern is a particular advantage over existing models. Internal representations are formed by separating hyperplanes. Selection of the hyperplanes and expansion of the network is based on an entropy measure which is appropriately defined. The weight vectors of all units with a certain layer are determined in a single presentation of the training set. (10 refs.)

777 Neural net pruning - why and how
J. Sietsma, R.J.F. Dow
(Mater. Res. Lab., DSTO, Melbourne, Vic., Australia).
IEEE International Conference on Neural Networks (IEEE Cat. No.88CH2632-8), San Diego, CA, USA, 24-27 July 1988 (New York, NY, USA: IEEE 1988), p.325-33 vol.1

A continuing question in neural net research is the size of network needed to solve a particular problem. If training is started with too small a network for the problem no learning can occur. The researcher must then go through a slow process of deciding that no learning is taking place, increasing the size of the network and training again. If a network that is larger than required is used, then processing is slowed, particularly on a conventional von Neumann computer. An approach to this problem is discussed that is based on learning with a net which is larger than the minimum size network required to solve the problem and then pruning the solution network. The result is a small, efficient network that performs as well or better than the original which does not give a complete answer to the question, since the size of the initial network is still largely based on guesswork but it gives a very useful partial answer and sheds some light on the workings of a neural network in the process. (3 refs.)

798 Distortion invariant character recognition by a multi-layer perceptron and back-propagation learning
A. Khotanzad, J.H. Lu
(Dept. of Electr. Eng., Southern Methodist Univ., Dallas, TX, USA).
IEEE International Conference on Neural Networks (IEEE Cat. No.88CH2632-8), San Diego, CA, USA, 24-27 July 1988 (New York, NY, USA: IEEE 1988), p.625-32 vol.1

A neural network based approach for distortion (translation, scale, and rotation) invariant character recognition is presented. To reduce the dimension of the required network, as well as to achieve invariancy, six distortion-invariant features are extracted from each image and are used as inputs to the neural net. These six continuous-valued features are derived from the geometrical moments of the image. A multilayer perceptron (MLP) with one hidden layer along with back-propagation training algorithm is used. The MLP is trained with twelve 64x64 differently oriented, scaled, and translated binary images of each of the twenty-six English characters. Its performance is tested using eight binary images from each character which were not used during training. Results of experimentation with different numbers of hidden layer nodes are presented. (8 refs.)

2.0 SYSTEMS AND TECHNIQUES

2837 Professionals and expert systems: a meeting of minds?
M. Newman
Dept. of Accounting & Bus. Finance, Manchester Univ., UK
Comput. Soc. (USA), vol. 18, no. 3, p. 14 - 27 (July 1988)

The effects computers have already had on general work practices are looked at. The following questions are addressed: do so-called intelligent medical expert systems actually display intelligent behaviour? Does the increasing use of similar expert systems pose a threat to the professions and their patients? In a deeper sense, does the threat of "intelligent" machines challenge the way people think about themselves? The author concludes that the AI community has attempted to raise the status of machinery and as a consequence has devalued humanity. (29 refs.)

2928 An expert system architecture and its application to the evaluation of scintigraphic image sequences
G. Sagere, H. Niemann
Erlangen-Nurnberg Univ., West Germany
Proceedings of the Symposium on the Engineering of Computer-Based Medical Systems (Cat. No. 88CH2606-2), Minneapolis, MN, USA, 8-10 June 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p. 82 - 8

A general framework for expert systems in the area of image understanding is described. It represents declarative knowledge in an associative network, allows the attachment of arbitrary procedural knowledge, and is used by a problem-independent control algorithm. The framework is used to implement a complete system for the diagnostic evaluation of scintigraphic images of the heart. The input to the system is an image sequence which is first preprocessed and segmented to obtain the left ventricle and its sectors. The contours of the ventricle and the sectors are the basis for knowledge-based analysis of the images. Analysis proceeds from the contour segments via objects and their motions to diagnostic descriptions of the image sequence. Experiments show that the system successfully completes analysis, and in particular does not make wrong diagnostic suggestions. (22 refs.)

1068 Expert system for selecting expert system shells
A. Martin, R.K.H. Law
(Sch. of Ind. & Bus. Studies, Warwick Univ., Coventry, UK). *Inf. Softw. Technol.* (UK), vol.30, no.10, p.579-86 (Dec. 1988).

A prototype expert system is described that advises end users in the selection of a suitable development environment for small- to medium-scale expert system applications. The system, running on a microcomputer, assesses the suitability of an application as a whole for expert systems techniques and recommends the five 'best' products from its knowledge base of 42 products. The paper takes a case study approach, using the system itself as a recursive discussion example, applying some of the concepts involved in building an expert system. It describes the operation of the system and suggests improvements; it includes a feature list for classifying expert system shells. (11 refs.)

1071 Why engineers should not use artificial intelligence
D.L. Parnas
(Queens Univ., Kingston, ON, Canada).
INFOR, (Canada), vol.26, no.4, p.234-46 (Nov. 1988).
(CIPS Edmonton '87 Conference, Edmonton, AB, Canada, 16-19 Nov. 1987).

It can be said that the most promising field within computer science is artificial intelligence, often simply known as AI. Some will interpret this as meaning that AI is a field that holds great promise. Others interpret this as meaning that AI is a field whose practitioners make great promises. Recently conferences, journals and newspaper articles have contained suggestions that AI offers special new techniques that can make drastic changes in the role of computer systems in the world. The paper presents a more skeptical view. It argues that (a) the terminology used in many AI discussions is poor, (b) that many techniques widely touted as revolutionary are ad hoc, 'cut and try', methods that will not lead to trustworthy products, (c) that many claims about AI and expert systems are exaggerated, and (d) that the fundamental research is more philosophical than practical. Most importantly, it concludes that many applications being tackled using ad hoc, heuristic methods can be solved using conventional systematic analysis and sound engineering practice.

1075 The need for user models in generating expert system explanations
R. Kass
(Center for Machine Intelligence, Pennsylvania Univ., Philadelphia, PA, USA), T. Finn.
Int. J. Expert Syst. Res. Appl. (USA), vol.1, no.4, p.345-75 (1988).

An explanation facility is an important component of an expert system, but current systems for the most part have neglected the importance of tailoring a system's explanations to the user. The article explores the role of user modeling in generating expert system explanations, making the claim that individualized user models are essential to produce good explanations when the system users vary in their knowledge of the domain, or in their goals, plans, and preferences. To make this argument, a characterization of explanation, and good explanation is made, leading to a presentation of how knowledge about the user affects the various aspects of a good explanation. Individualized user models are not only important, but also practical to obtain. A method for acquiring a model of the user's beliefs implicitly by eavesdropping on the interaction between user and system is presented along with examples of how this information can be used to tailor an explanation. (42 refs.)

1087 A model-based framework for characterization of application domains for the expert system technology
M. Hadzikadic, D.Y.Y. Yun, P.-C. Ho
(Dept. of Comput. Sci. & Eng., Southern Methodist Univ., Dallas, TX, USA). Methodologies for Intelligent Systems. *Proceedings of the Second International Symposium,* Charlotte, NC, USA, 14-17 Oct. 1987 (New York, NY, USA: North-Holland 1987), p.64-72

Application domains potentially suitable for development using the expert system technology vary widely in terms of requirements, features, and utility. Heuristic classifications of application domains have been attempted with limited success. Yet, the need for better characterization of these domains increases with the popularity of the expert system technology. The need arises both from fitting the domains with appropriate representations and from matching applications with proper development tools. The paper endeavors to establish a set of primitives and compound processes that comprise a framework usable for describing a wide class of application domains. It differs from previous attempts of classification in that the primitives encompass more fundamental concepts, while the framework allows the capture of high-level abstractions and concrete representations of application domains. (9 refs.)

837 First generation expert systems: a review of knowledge acquisition methodologies
I.M. Neale
(Centre for Insurance & Investment Studies, City Univ., London, UK). *Knowl. Eng. Rev.* (UK), vol.3, no.2, p.105-45 (June 1988). [received: 27 Jan 1989]

Reviews a wide range of knowledge acquisition techniques in the context of attempts to achieve a systematic methodology. These have been poorly documented by expert system builders, who are often inclined to overvalue textbooks and the ways experts themselves claim they solve problems. No one method has a universal advantage; each has some value. Techniques should be selected to suit the domain, the task, the expert and the knowledge engineer. Knowledge acquisition involves creating a conceptual model of expert knowledge and reasoning, from analysis of data elicited by these techniques. A survey of the literature indicates increasing emphasis on tools for knowledge acquisition, used directly by experts. Several projects currently directed towards providing a proper epistemological foundation for knowledge acquisition are discussed and compared. None has yet produced a complete epistemologically sound methodology; however, recognition of the need to create a conceptual model at the knowledge level (rather than the symbol level) is an important advance. (245 refs.)

839 INNOVATOR: an expert system for new product launch decisions
S. Ram
(Dept. of Manage. Inf. Syst., Arizona Univ., Tucson, AZ, USA), S. Ram.
Appl. Artif. Intell. (USA), vol.2, no.2, p.129-48 (1988).

Expert systems have become an increasingly important area of research for marketing academicians and practitioners. This paper describes an expert system called INNOVATOR that screens new product ideas and provides an approve/reject/reevaluate decision. If desired, the system can explain the logical reasoning it used to arrive at the decision. INNOVATOR uses a rule base consisting of regular IF and fuzzy IF rules and a backward chaining strategy to make recommendations. The paper explains the knowledge (expertise) requirements of the system, as well as the knowledge elicitation process and the knowledge representation scheme. The system has been implemented on an IBM 4381 using a shell called Expert System Environment (ESE). INNOVATOR, in its current form, assists consultants in financial service organizations in deciding whether to offer specific product innovations to their customers. The system is operational and has been validated by comparing its decisions with those of experts in the field. (22 refs.)

533 System safety and legal issues in AI
W.A. Hyman
(Dept. of Ind. Eng., Texas A&M Univ., College Station, TX, USA), W.L. Johnston, S.Spar. *Comput. Ind. Eng.* (UK), vol.15, p.51-8 (1988). (10th Annual Conference for Computers and Industrial Engineering, Dallas, TX, USA, 23-25 March 1988).

Nowadays it is important for expert system developers to understand the legal context in which they are working and the legally viable allegations that may come to be made against them. Equally important is the inclusion of system safety techniques which are expressly intended to identify and address system hazards before they are the cause of injuries or property losses. Understanding and dealing with these issues now can prevent this area of potentially useful technological development from being buried under the burden of products which are poorly conceived, executed and marketed. (17 refs.)

3.0 APPLICATIONS

2926 Diagnostic algorithms & clinical diagnostic thinking
J. M. Tonkonogy, J. Armstrong
Dept. of Psychiatry, Univ. of Massachusetts Med. Center, Worcester, MA, USA
Proceedings of the Symposium on the Engineering of Computer-Based Medical Systems (Cat. No. 88CH2606-2), Minneapolis, MN, USA, 8-10 June 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p. 71-4

Diagnostic algorithms are evaluated as a model of medical reasoning in the differentiation of the two syndromes or diseases with overlapping clinical manifestations. Linear-discriminant analysis (LDA), nonlinear discriminant analysis (NLDA), and sequential-statistical analysis (SSA) are compared with the decision-making practices of physicians, as they differentiated 175 cases with anterior or posterior aphasia, 789 cases with cerebral infarction or intracerebral hemorrhage, and 200 cases of schizophrenia or schizophreniform organic psychosis. Comparisons showed that the clinical decision-making process used by physicians can be approximated by neural sets or parallel distributed processing. Implications for the implementation of computer diagnostic programs in clinical practice are discussed. (9 refs.)

2951 Schedule generation in a flexible manufacturing system: a knowledge-based approach
S. Shen
Jet Propulsion Lab., California Inst. of Technol., Pasadena, CA, USA), Yiu-Long Chang. *Decis. Support Syst.* (Netherlands), vol. 4, no. 2, p. 157 - 66 (June 1988)

A flexible manufacturing system employs computer controlled

machines which can be set up very quickly to perform diverse tasks. A job in an FMS, unlike a job in a traditional job shop, thus can go through alternative routes before its completion. Because of this increased flexibility, job scheduling in an FMS becomes dynamic and often demands real-time performance; despite this the problem is inherently NP-hard. The purpose of the paper is to develop a framework where real-time scheduling is achieved by utilizing past scheduling experience or knowledge, together with a reasonably good algorithm. The authors present an algorithm first and then discuss how to integrate scheduling knowledge and algorithms in a framework called frames. The frame approach solves a scheduling problem by using inheritance, demons and defaults, rather than executing algorithms, although algorithms are also maintained in frames and can be called whenever they are needed. The frame approach is considered as one of the most powerful knowledge representation constructs in artificial intelligence. (38 refs.)

471 Knowledge-based project control employing heuristic optimisation
T. J. Barber, J. T. Boardman
Inf. Technol. Res. Inst., Brighton Polytech., UK
IEEE Proc. A, Phys. Sci. Meas. Instrum. Manage. Educ. Rev. (UK), vol. 135, no. 8, p. 529-38 (Nov. 1988)

One particular problem of great commercial significance which professional project engineers confront is that of expeditious control of resources to achieve timely completion of complex plans. A wealth of expertise exists, little if any of which has been formalized, so as to permit the development of computer-based aids for intelligent project control. The authors describe a heuristic algorithm which has been developed in close consultation with expert project engineers. It has been integrated within an intelligent knowledge-based system (IKBS) and the resultant computer tool evaluated in a real-world project management situation. The algorithm incorporates six variables which have been carefully formulated to assess the merit of expediting selected activities in project plans within a hierarchical framework of control strategies. An empirical study of these heuristics, taken in conjunction with the IKBS, is presented, and some parallels are drawn with an expert project engineer's performance on a typical project plan. (10 refs.)

441 An expert system based algorithm for short term load forecast
S. Rahman, R. Bhatnagar
Dept. of Electr. Eng., Virginia Polytech. Inst. & State Univ., Blacksburg, VA, USA *IEEE Trans. Power Syst.* (USA), vol. 3, no.2, p. 392-9 (May 1988) [received: 17 Oct 1988]

Existing studies on 1-24 hr load forecasting algorithms are reviewed, and an expert-system-based algorithm is presented as an alternative. The logical and syntactical relationships between weather and load as well as the prevailing daily load shapes have been examined to develop the rules for this approach. Two separate, but similar, algorithms have been developed to provide 1-6 hr. and 24 hr. forecasts. These forecasts have been compared with observed hourly load data for a Virginia electric utility for all seasons of the year. The 1 hr. and 6 hr. forecast errors (absolute average) ranged from 0.869% to 1.218% and from 2.437% to 3.48% respectively. The 24 hr. forecast errors (absoluate average) ranged from 2.429% to 3.300%. (18 refs.)

452 Expert system for autonomous handling of elementary assembly operations
P. Levi, J. Majumdar, B. Wild
Forschungszentrum Informatik FZI, Karlsruhe, West Germany *Int. J. Prod. Res.* (UK), vol. 26, no. 10, p. 1671-80 (Oct. 1988)

This article addresses the new concept of a real-time expert system for elementary two-arm assembly operations under the inclusion of operation-based vision. Every operation is represented by a local blackboard which contains a recognition part, a world model and a measurement part. The total set of these behaviour-oriented modules is supervised and scheduled by an operation monitor, which is again a blackboard. The local-knowledge sources of each individual blackboard receive sensor signals as input and can directly produce a reaction without going to a higher level global inference engine. This local, short connection between sensor signal and reflex is decided on the basis of compiled rules which have the form of production rules. The main types of these rules are regulation tasks depending on the operation and analysis rules for situation assessments. Global plans are produced as assembly graphs and transmitted in parts to the operation monitor as task specifications. A 3-D CAD modeller is used to support the object representation and recognition. (15 refs.)

489 Software development issues for parallel processing
L. Russell, R. N. C. Lightfoot
Comp. Sci. Corp., Beltsville, MD, USA
Proceedings COMPSAC 88: The Twelfth International Computer Software and Applications Conference (Cat. No. 88CH2611-2), Chicago, IL, USA, 5-7 Oct. 1988 (Washington, DC, USA: *IEEE Comput. Soc. Press* 1988), p. 306-7

Parallel architectures for computers give rise to several issues with respect to software development. The issue to be resolved depends on the life-cycle phase in which the problem appears. During problem analysis and program design, the issue is best resolved by using an intelligent language environment, exemplified by the Rational R1000 ADA machine. During the post-implementation phase, an existing program can be restructured by using an expert tool that generates directed graph models of the program, analyzes the models and the source to identify potentially parallelizable code sequences, and through workstation window-oriented displays guides the user through the process of restructuring a program to run efficiently on a parallel architecture machine. An eight-step methodology for reworking existing Fortran programs for more efficient execution on parallel architecture machines is presented.

1123 Generating goal-oriented explanations
K.R. McKeown
(Dept. of Comp. Sci., Columbia Univ., New York, NY, USA). *Int. J. Expert Syst. Res. Appl.* (USA), vol.1, no.4, p.377-95 (1988).

Many researchers have shown the need for the production of expert system explanations that directly address the concerns of the user and yet few systems have taken this need into account. The author shows how to generate explanations that are tailored to the domain goals of the current user. The approach includes a method for deriving the user's goal from the discourse segment preceding the request for explanation, for representing different points of view in the underlying knowledge base, and for relating the derived goal and point of view to determine information that is relevant to include in the explanation.

Thus the domain goal influences the content of the generated explanation. This method for providing tailored explanations is embedded as part of a problem-solving system that can advise a university student about courses to take in an upcoming semester. The user converses with the system via a natural language interface. In addition to using natural language techniques to interpret user input, the system makes use of a natural language generator to produce its responses as opposed to the more usual use of templates. Examples are provided of different explanations that the system produces in response to the same request for an explanation when the user has different goals. (34 refs.)

1143 A survey of expert systems in manufacturing and process planning
T. Gupta, B.K. Ghosh
(Dept. of Ind. & Syst. Eng., Wisconsin Univ., Milwaukee, WI, USA). *Comput. Ind.* (Netherlands), vol.11, no.2, p.195-204 (Jan. 1989).

The issue of capturing the knowledge and skill of human experts in various problem domains and passing it on to future generations via the computer is an active area of research. Domain-specific expert systems are being developed to aid the human decision maker. The authors present various applications of expert systems to the problem domain of industrial and manufacturing engineering. These systems were developed with commercial as well as research interests. For each of the systems discussed, the authors focus on the problem definition, implementation scheme, and special features, if any. Subsequently, they focus on the continuing efforts for design and development of expert systems for process planning. Some of the popular approaches used are briefly discussed. (33 refs.)

1154 ILIAD as an expert consultant to teach differential diagnosis
H.R. Warner, P. Haug, O. Bouhaddou, M. Lincoln, H. Warner, Jr., D. Sorenson, J.W. Williamson, Chinli Fan
(Dept. of Med. Inf., Utah Univ., Salt Lake City, UT, USA). *Proceedings The Twelfth Annual Symposium on Computer Applications in Medical Care* (IEEE Cat. No.88CH2616-1), Washington, DC, USA, 6-9 Nov. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.371-6

A program (ILIAD) is described that uses knowledge frames representing diseases encountered in internal medicine to teach medical students about differential diagnosis. The ILIAD consultant utilizes a number of sophisticated inferencing mechanisms to mimic the strategy of a medical expert in working on a patient. Its knowledge is represented in both Bayesian and Boolean frames, which permit use of sensitivities and specificities to describe the relationship of a disease to its manifestations and provide a basis for explaining its conclusions. In addition to differential diagnosis, ILIAD provides advice regarding the most appropriate information to seek at each stage of the workup. ILIAD's knowledge base is also used to simulate patient cases and evaluate the problem solving performance of medical apprentices. It is presently being used by third-year medical clerks on the wards of three hospitals, and is being beta tested at eight additional sites. (7 refs.)

1155 DTREE: microcomputer-assisted teaching of psychiatric diagnosis using a decision tree model
M.B. First, J.B.W. Williams, R.L. Spitzer
(Dept. of Psychiatry, Columbia Univ., New York, NY, USA). *Proceedings The Twelfth Annual Symposium on Computer Applications in Medical Care* (IEEE Cat. No. 88CH2616-1), Washington, DC, USA, 6-9 Nov. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.377-81

The DTREE system was developed to provide clinicians with computer-assisted teaching of psychiatric diagnosis according to the American Psychiatric Association's Diagnostic and Statistical Manual of Mental Disorders, Third Edition, Revised (DSM-III-R). At the core of DTREE is an expert system that guides the user through the process of making a diagnosis. To best model the DSM-III-R system, a decision-tree design has been used in which questions are sequentially asked about a case, with the answer determining which question is asked next. Since the primary goal of an expert system used for teaching is to be able to provide explanations of what it is doing, annotated comments and additional explanatory text are available for each equation that is asked. DTREE is currently undergoing field testing in both clinical and educational settings. (21 refs.)

1164 OVERSEER: an expert system monitor for the psychiatric hospital
J.D. Bronzino, R.A. Morelli, J.W. Goethe
(Dept. of Eng. & Comput. Sci., Trinity Coll., Hartford, CT, USA). *Proceedings The Twelfth Annual Symposium on Computer Applications in Medical Care* (IEEE Cat. No. 88CH2616-1), Washington, DC, USA, 6-9 Nov. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.8-12

A knowledge-based system that monitors the treatment of psychiatric patients in real time, called OVERSEER, is described. Based on procedures and protocols developed in the psychiatric setting, OVERSEER monitors the clinical database and issues alerts when standard clinical practices are not followed or when laboratory results or other clinical indicators are abnormal. Written in Prolog, OVERSEER is designed to interface directly with the hospital's database, and, thereby utilizes all available pharmacy and laboratory data. Moreover, unlike the interactive expert systems developed for the psychiatric clinic, OVERSEER does not require extensive data entry by the clinician. Consequently, the chief benefit of OVERSEER's monitoring approach is the unobtrusive manner in which it evaluates treatment and patient responses, and provides information regarding patient management. (12 refs.)

1165 The Transfusion Advisor: a knowledge-based system for the blood bank
K.A. Spackman, M.J. Chabot, J.R. Beck
(Dartmouth-Hitchcock Med. Center, Hanover, NH., USA).
Proceedings The Twelfth Annual Symposium on Computer Applications in Medical Care (IEEE Cat. No.88CH2616-1), Washington, DC, USA, 6-9 Nov. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.18-21

The Transfusion Advisor, a knowledge-based system designed to facilitate the review of transfusion requests in the hospital blood bank, is described. An evaluation of a research prototype of the system, which draws conclusions about 12 hemostatic disorders and critiques the appropriateness of the use of frozen plasma, cryoprecipitate (cryo), and platelets, which are products most commonly used for the treatment of hemostatic disorders, has been performed. The evaluation shows that there is closer agreement of the expert system with the domain expert than with the clinicians ordering the transfusions, and that the use of the transfusion advisor to screen transfusion requests may result in a substantial decrease in the use of these blood products. (3 refs.)

1166 DISPO advisor-expert system for psychiatric disposition
P.Barta, W. Barta
(Johns Hopkins Hospital, Baltimore, MD, USA).
Proceedings The Twelfth Annual Symposium on Computer Applications in Medical Care (IEEE Cat. No.88CH2616-1), Washington, DC, USA, 6-9 Nov. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.22-5

An expert system designed to assist psychiatric residents is described. This microcomputer-based decision support system helps residents find the proper disposition for patients who come to the emergency room. The system uses an inexpensive, commercially available expert system shell, VP-EXPERT by Paperback Software, to match patients with inpatient and outpatient resources appropriate to their needs. The inference engine uses both forward and backward chaining, and interfaces with data stored in dBase III files. The system is currently in daily use by residents. (6 refs.)

1198 Towards knowledge-based identification of mineral mixtures from reflectance spectra
K. Ali
(Basser Dept. of Comput. Sci., Sydney Univ., NSW, Australia), C. Horsfall, R. Lister.
Knowl.-Based Syst. (UK), vol.2, no.1, p.5-13 (March 1989).

GOLD/M is an expert system that identifies some mineral mixtures in rocks by searching for characteristic features in the infrared reflectance spectra of the rocks. Some of these minerals indicate subsurface gold deposits. The system is implemented in a shell based on subjective Bayesian reasoning, as used in the Prospector expert system. An earlier system, called GOLD, could identify only the one mineral in a pure sample. GOLD/M identifies up to two minerals in a rock consisting of a mixture of minerals. This paper discusses the extensions made to cope with mineral mixtures, and discusses problems that emerged with the Prospector formalism. (14 refs.)

1239 A high-speed hardware unit for a subset of logic resolution
D. Wong
(Dept. of Electr. Eng., Stanford Univ., CA, USA).
Proceedings of the 21st Annual Workshop on Microprogramming and Microarchitecture - MICRO '21 (IEEE Cat. No. 88TH0236-0), San Diego, CA, USA, 30 Nov. - 2 Dec. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), p.73-8

The authors present a hardware design for doing a subset of resolution, one of the key steps in logical inference. The hardware can perform resolution directly without any macroprogramming or microprogramming. Given two input clauses in the appropriate format, the hardware produces the resolvent clause via hardwired circuitry in a manner similar to the way an ALU performs addition. Unlike other designs that use AND and OR parallelism, this approach uses the parallelism in the unification and substitution of arguments. Prolog examples are considered, although the hardware is useful for other logic applications as well. (6 refs.)

890 Development of a process planning system using knowledge engineering and geometric processing
N. Uemura
(Production Syst. Dev. Lab., NEC Corp., Tokyo, Japan), S. Yokoi, Y. Hisatomi, K. Inagaki.
NEC Res. Dev. (Japan), no. 91, p.111-15 (Oct. 1988).

Knowledge engineering and geometric processing are key technologies for the computerization of process planning. A prototype of an automated process planning system for machining has been developed, and the system achieves the integration of CAD and CAM. The expertise for process planning is incorporated into the system as rules. The rules can reference geometric data directly. Sixteen samples of mechanical components were actually manufactured using the data from the process planning system with CAD and CAM, and the process plans by the system were found to be as good as those made by experts. (4 refs.)

920 MICROBE - a prototype expert system for treatment planning in bacterial infections
P. Nykanen, P. Kaatiala
(Med. Eng. Lab., Tech. Res. Centre of Finland, Tampere, Finland), J. Rantanen, J. Lumino, P. Gronroos.
Scandinavian Conference on Artificial Intelligence.
Proceedings of the SCAI '88, Tromso, Norway, 9-11 March 1988 (Amsterdam, Netherlands: IOS 1988), p.253-62

The prototype MICROBE for treatment planning in bacterial infections is described. The tasks of the systems is to help inexperienced doctors find suitable treatment for infections and to identify casual organisms. The development environment is Symbolics Lisp machine and Knowledge Engineering Environment by IntelliCorp. The system is object-oriented and it uses frames as the knowledge representation method. The system utilizes data from bacteriological databases in a VAX/VMS-environment. (9 refs.)

938 ALFA: automated load forecasting assistant
K. Jabbour, J.F.V. Riveros, D. Landsbergen, W. Meyer
(Syracuse Univ., NY, USA).
IEEE Trans. Power Syst. (USA), vol.3, no.3, p.908-14 (Aug. 1988). [received: 10 Jan 1989]

ALFA, an expert system for forecasting short-term electricity demand is presented. ALFA is in operation at the Energy Management System center at Niagara Mohawk Power Corporation in upstate New York, generating, in real time, hourly load forecasts for up to 48 hours in advance. ALFA uses an extensive 10-year historical database of hourly observations of 12 weather variables and a rule base that takes into account daily, weekly, and seasonal variations of load, as well as holidays, special events, and load growth. A satellite interface for the real-time acquisition of weather data, and the machine-operator interface are also discussed. (7 refs.)

939 Rule-based relay modelling for transient-stability studies
G.T. Vuong, G. Paris
(Inst. de Recherche d'Hydro-Quebec, Varennes, Que., Canada).
IEEE Trans. Power Syst. (USA), vol.3, no.3, p.1306-9 (Aug. 1988). [received: 10 Jan 1989]

A novel approach to relay modelling for transient-stability studies is presented. It is based on a simulation language which offers many characteristics of AI (artificial intelligence) technology. It is close to natural, declarative, rule-based, and inclusive of concepts used in protective relaying. The proposed option gives users great autonomy, which would lead to better productivity since relay models and/or production schemes can be developed within a few hours instead of the many weeks required under normal circumstances. Moreover, systems involving many coordinated relays can be easily described. (6 refs.)

943 Expert systems and computer aided design: a productive merger
S.M. Wagner
(US Air Force Wright Aeronaut. Labs., Wright-Patterson AFB, Dayton, HO, USA), W.H. Shaw, Jr.
1988 IEEE Engineering Management Conference - Engineering Leadership in the 90's (from AI to ZJ) (IEEE Cat. No.88CH2651-8), Dayton, OH, USA, 24-26 Oct. 1988 (New York, NY, USA: IEEE 1988), p.78-84

An application of expert system concepts is presented for computer-aided design (CAD) environments to show how artificial intelligence can act as a design assistant. This concept is demonstrated by a recent, successful software development effort to implement a digital circuit CAD tool enhanced by such an expert system for educational use. The discussion centers on the use of graphics, simulation, and expert systems to construct an effective CAD workstation running on personal computer systems. The authors note the use of graphics technology to provide schematic capture which, when coupled with a database system, provides the necessary environment for a rule-based expert system. The CAD tool demonstrates the use of design heuristics, metaknowledge, and integrated circuit data to construct a forward chained reasoning mechanism for detecting design errors. The concept is extendable to other graphics oriented CAD tools. Operating experience with the tool is described and summary performance measures, which serve as possible predictors of other candidate systems performance, are presented. Issues influencing the merger of knowledge-based models and CAD are also briefly discussed. (25 refs.)

961 Expert system applications at the United States National Aeronautics and Space Administration
J. Liebowitz
(Dept. of Manage. Sci., George Washington Univ., Washington, DC, USA).
Int. J. Comput. Appl. Technol. (Switzerland), vol.2, no.1, p.26-9 (1989).

One of the main developers, funders, and users of expert systems technology in the United States is the National Aeronautics and Space Administration (NASA). With the movement from the space age to the space shuttle age to the space station age, artificial intelligence technology, especially expert systems, is expected to play a significant role in space-related activities in the years to come. The paper surveys some of these NASA-supported expert systems applications. (13 refs.)

994 Symbolic Neuroengineering for natural language processing: a multilevel research approach
M. Dyer.
Report CSD-880022, Univ. California, Comp. Sci. Dept., Los Angeles, CA, USA (Sept. 1988), 22 pp.

Natural language processing (NLP) research has been built on the assumption that natural language tasks, such as comprehension, generation, argumentation, acquisition, and question answering, are fundamentally symbolic in nature. Recently, an alternative, sub-symbolic paradigm has arisen, inspired by neural mechanisms and based on parallel processing over distributed representations. The assumptions of these two paradigms are compared and contrasted, resulting in the observation that each paradigm possesses strengths exactly where the other is weak, and vice versa. This observation serves as a strong motivation for synthesis. A multilevel research approach is proposed, involving the construction of hybrid models, to achieve the long-term goal of mapping high-level cognitive function into neural mechanisms and brain architecture. Four levels of modeling are discussed: knowledge engineering level, localist connectionist level, distributed processing level, and artificial neural systems dynamics level. The two major goals of research at each level are to explore its scope and limits and to find mappings to the levels above and below it. The capabilities of several NLP models, at each level, are described, along with major research questions remaining to be resolved and major techniques currently being used in an attempt to complete the mappings. Techniques include: forming hybrid systems with spreading activation, thresholds and markers to propagate bindings; using extended back-error propagation in reactive training environments to eliminate microfeature representations; transforming weight matrices into patterns of activation to create virtual semantic networks; using conjunctive codings to implement role bindings; and employing firing patterns and time varying action potentials to represent and associate verbal with visual sequences.

988 On the simultaneous interpretation of real world image sequences and their natural language description: the system SOCCER
E. Andre, G. Herzog, T. Rist
(Univ. des Saarlandes, Saarbrucken, West Germany).
ECAI 88. *Proceedings of the 8th European Conference on Artificial Intelligence*, Munich, West Germany, 1-5 Aug. 1988 (London, UK: Pitman 1988), p.449-54

The aim of previous attempts at connecting vision systems and natural language systems has been to provide a retrospective description of the analyzed image sequence. The step from such an a posteriori approach towards simultaneous natural language description reveals a problem which has not yet been dealt with in generation systems. Automatic generation of simultaneous descriptions calls for the application of an incremental event recognition strategy and for the adequate coordination of event recognition and language production. In order to enable free interaction between these processes, it is useful to implement them in parallel. The system SOCCER is presented, which is based upon such a concept. Short sections of soccer games have been chosen as the domain of discourse. In analogy to radio reports, the system generates a description of the game which it is watching and which the listener cannot see. (24 refs.)

592 An expert system for workshop simulation
M.S. Eid
(Sch. of Eng., Moncton Univ., NB, Canada), C. Poirier.
Comp. Ind. Eng. (UK), vol.15, p.91-7 (1988).
(10th Annual Conference for Computers and Industrial Engineering, Dallas, TX, USA, 23-25 March 1988).

Presents an expert system that can generate SLAM commands for a job-shop planning application where a user with no SLAM or simulation knowledge can interact directly with a microcomputer to run the SLAM application. The expert system is in the form of a user-friendly conversational program which automatically translates the problem parameters into SLAM commands, executes the SLAM program, and provides an unsophisticated simulation user with the analysis results. The paper gives an example in the form of an auto-body repair shop. (5 refs.)

591 Knowledge base design for law enforcement
A.B. Badiru
(Sch. of Ind. Eng., Oklahoma Univ., Norman, OK, USA), J.M. Mathis, B.T. Holloway. *Comp. Ind. Eng.* (UK), vol.15, p.78-84 (1988). (10th Annual Conference for Computers & Industrial Engineering, Dallas, TX, USA, 23-25 Mar. 1988).

Describes a project involving the development of an expert system for suspect identification in armed robbery incidents. A prototype of the system, named AREST (armed robbery eidetic suspect typing), has been completed and is described. The investigative process can be greatly simplified, expedited, and enhanced by the use of expert systems. An expert system can serve as an aid that handles the preliminary crime analysis and, thus, frees the human officer to concentrate on higher-level analysis worthy of his expertise. (5 refs.)

598 Legal expert systems and discretionary decision making
C.D. Hafner, D.H. Berman
(Center for Law & Comput. Sci., Northeastern Univ., Boston, MA, USA). Electro/88 Conference Record, Boston, MA, USA, 10-12 May 1988 (Los Angeles, CA, USA: Electron. Conventions Manage. 1988), p.36/1/1-6

Public perception of unfairness and capriciousness in the legal system has led to increasing interest in 'determinate' rules, which ensure greater consistency but which severely limit the discretion of individual decision makers. The authors suggest that artificial intelligence offers a new approach to developing and promulgating legal norms and guidelines, that could increase the consistency of legal decisions by providing relevant and persuasive information to decision makers, rather than by limiting their discretion. Using the example of a judge sentencing a criminal, they show how medical expert systems might be adapted for use as legal 'consultants', providing judges with the ability to find out, for situations of interest to them, what range of actions is considered 'normal' or 'reasonable'. Although the danger of over-reliance on computers must be recognized and addressed, this danger must be weighed against the potential for improving fairness in areas where discretion has been shown to foster gross inconsistency and discriminatory practices. (17 refs.)

612 Expert systems for insurance and investment
A.R. Butler
(Colonial Mutual Group, London, UK). Third International Expert Systems Conference, London, UK, 2-4 June 1987 (Oxford, UK: Learned Inf. 1987), p.173-82

Under the flag of the UK Government's Alvey programme, eighteen insurance companies, one insurance broker, two firms of consulting actuaries, two firms of general consultants, three insurance professional organizations, four academic institutions, and one contractor have formed a collaborative venture to develop two demonstrable prototype expert systems for insurance work. The first is a system for assessing fire risks on clothing trade premises. The second is a system for assisting with buy/no buy decisions in the equity investment area. The paper gives the essential background to the work, explains the technical approach of the equity selection system, discusses what has been learned from the club's operations so far, and indicates the direction of future work. (1 ref.)

628 School bus routing and scheduling: an expert system approach
Der-San Chen, H.A. Kallsen
(Alabama Univ., Tuscaloosa, AL, USA), R.C. Snider.
Comp. Ind. Eng. (UK), vol.15, p.179-83 (1988).
(10th Annual Conference for Computers and Industrial Engineering, Dallas, TX, USA, 23-25 March 1988).

Proposes an expert system approach to routing and scheduling school buses for a rural school system. The expert system is programmed in TURBO Prolog for use on an IBM/XT and is applied to rural county school systems in Alabama. The routing problem is concerned with the determination of a stop-to-stop route to be traversed to each school by each bus whereas the scheduling problem, the determination of times at all bus stops for each bus. A bus may be used for multiple runs. Each route is

designed in such a way that the bus capacity, student riding time, school time window and road condition constraints are satisfied while attempting to minimize the number of buses required in operation, minimize the fleet travel time and balance the bus loads. The knowledge base contains factual knowledge such as road map, school locations, bus capacities stop locations, number of students at each stop, and drivers' homes. It also contains procedural knowledge such as heuristics for finding a route and scheduling multiple runs for a bus subject to various constraints. The inference engine or control program chooses the appropriate heuristics used in constructing efficient routes and schedules with respect to various objectives or goals. A user interface includes the graphic display of road maps and determined routes. (8 refs.)

631 Training NASA satellite operators: an expert system consultant approach
J. Liebowitz
(Dept. of Manage. Sci., George Washington Univ., DC, USA), P. Lightfoot.
Educ. Technol. (USA), vol.27, no.11, p.41-7 (Nov. 1987).
[received: 07 Nov 1988]

Computer-based training is an area of focus that the National Aeronautics and Space Administration (NASA) sees as deserving attention. Specifically, the ability to have intelligent computer-assisted instruction (ICAI) for helping to train individuals is a worthwhile consideration for NASA to explore. The authors describe an expert system prototype built for training spacecraft operators on satellite power subsystems contingency operations. It acts as a 'consultant' to help the spacecraft operator know what to do in case something goes wrong with a satellite's power subsystem. This prototype uses production rules because: they can be used to represent strategies; they can provide a clear, simple, transparent structure; and they are amenable to automatic manipulation. Future directions for further development and research are also addressed. (11 refs.)

637 A microcomputer-based expert system prototype for helping the layperson find specific information on breast cancer questions
J. Liebowitz
(Dept. of Manage. Sci., George Washington Univ., DC, USA).
Comput. Educ. (UK), vol.12, no.4, p.465-70 (1988).

Medical expert systems have been around since the early 1970s. MYCIN, PUFF, CADUCEOUS and CASNET are examples of medical expert systems that have been built; however, their use, like most medical expert systems, is for the physician. There is a growing need for educating the public in certain areas of medicine. Laypersons want to be better informed. This paper focuses on a microcomputer-based expert system which is being developed to help the layperson be more informed on general and specific questions dealing with breast cancer. (6 refs.)

641 Soil classification: an ESIE approach
S. Malasri
(Dept. of Civil & Archit. Eng., Miami Univ., Coral Gables, FL, USA).
CoED (USA), vol.8, no.4, p.43-7 (Oct.-Dec. 1988).

An expert system for soil classification is developed using the ESIE expert system development tool. The system identifies the type of soil in accordance with the AASHTO recommendation from laboratory results of a soil sample. ESIE is an expert system inference engine distributed under the shareware marketing concept. It has several limitations but can be used for an inexpensive rapid prototype development for several types of problems. The paper describes ESIE features and its application in solving the soil classification problem. (6 refs.)

651 Automated reasoning: a new concept in power system security analysis
S.S. Shah, S.M. Shahidehpour
(Dept. of Electr. & Comput. Eng., Illinois Inst. of Technol., Chicago, IL, USA).
Proceedings of the International Workshop on Artificial Intelligence for Industrial Applications: IEEE AI '88
(Cat No.88CH2529-6), Hitachi City, Japan, 25-27 May 1988 (New York, NY, USA: IEEE 1988), p.58-63

The authors consider a specific application of expert systems for power systems security analysis. Once a severe disruption has occurred in the power network, a series of tasks have to be accomplished before the operator can make any decisions. It is suggested that a computer-driven decision making mechanism will help the operator to perform his duties with minimum flaws and maximum efficiency. Several aspects of the application of automated reasoning in power system security analysis have been addressed. The adjustments of real and reactive power injections, phase-shifters, and tap changing transformers have been considered as proper corrective actions in establishing the heuristic rules. The implementation of the idea has been tested on a six-bus system, and the results have been compared with those of analytical approaches. The proposed methodology has also been applied to the IEEE 30-bus system, and results are presented. (17 refs.)

670 Use of a rule-based system for process control
J.A. Bernard
(MIT Nucl. Reactor Lab., Cambridge, MA, USA). *IEEE Control Syst. Mag.* (USA), vol.8, no.5, p.3-13 (Oct. 1988).

A rule-based, digital, closed-loop controller that incorporates fuzzy logic has been designed and implemented for the control of power on the 5-MW Massachusetts Institute of Technology (MIT) Research Reactor under both steady-state and transient conditions. A comparison is made of the rule-based and analytic approaches. Differences in the division of labour between plant engineers and control specialists, the type of knowledge required and its acquisition, the use of performance criteria, and controller testing are discussed. The design, implementation, and calibration of rule-based controllers are reviewed, with specific examples taken from the completed work on the MIT Research Reactor. The possible role of rule-based technology in process control is evaluated. It is proposed that since rule-based systems are generally more robust than their analytic counterparts, they should be used both as backups to analytic controllers and as a means of improving the man-machine interface by providing human operators with the rationale for automatic control actions. (18 refs.)

CALL FOR PAPERS

The Third International Conference on Industrial & Engineering Applications of Artificial Intelligence & Expert Systems IEA/AIE-90

July 15-18, 1990 Mills House Hotel, Charleston, South Carolina

Sponsored by:
Association for Computing Machinery/SIGART
The University of South Carolina
The University of Tennessee Space Institute (UTSI)

In Cooperation with:
American Association for Artificial Intelligence
Canadian Society for Computational Studies of Intelligence
IEEE - Computer Society
International Association of Knowledge Engineers
International Center for the Application of Information Technology
International Neural Network Society

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

Pattern Recognition
Vision
Sensor fusion
Computer Aided Manufacturing
Computer Aided Design
Robotics
Planning/Scheduling
Diagnostics Systems
Intelligent Interfaces
Intelligent Databases
Autonomous Systems
Knowledge Representation
Knowledge Acquisition
Machine Learning
Natural Language Processing
Neural Networks
Intelligent Tutoring
Reasoning under Uncertainty
Distributed and Parallel Architectures
Qualitative Models
Blackboard Systems
Industrial Expert Systems

Please submit by Dec 1, 1989 four copies of an extended abstract (4 - 6 double spaced pages) to the Program Chairman. Authors will be notified by Feb 1, 1990 and final copies of papers will be due April 1, 1990.

Dr. Manton M. Matthews
Program Chairman, IEA/AIE-90
Department of Computer Science
University of South Carolina
Columbia, SC 29208
phone: (803)-777-3285
CSNET: matthews cs.scarolina.edu

Dr. Moonis Ali
General Chairman, IEA/AIE-90
MS 15
The University of Tennessee
Space Institute
Tullahoma, Tennessee 37388
(615) 455-0631, ext 236

The proceedings will be published by ACM and will be available at the conference. A few copies of the proceedings of earlier conferences are available by contacting Nancy Wise at (615)-455-0631 ext 236.

Upcoming Conferences

In Canada

CANCELLED
CIPS Conference on Telepresence and Remote Interaction
30 October - 1 November 1989, Edmonton, Alberta

Address your concerns to Patrick Slinn, CIPS Edmonton Section, Box 1881, Edmonton, AB T5J 2P3

Vision Interface '90
14 - 18 May 1990, Halifax, Nova Scotia

Being held jointly with Graphics Interface '90. Topics include: Range image analysis and processing; Motion representation and analysis; Scene analysis; Text understanding and signature verification; Segmentation, feature extraction and pattern analysis; Neural networks and special purpose architectures; Speech analysis and understanding. Submission deadlines: 4 copies of full paper due Oct 31. Tutorial proposals due Nov 15.
Contact: H. D. Cheng, School of Comp. Sci., Technical U. of Nova Scotia, PO Box 1000, Halifax, Nova Scotia, B3J 2X4.

In the United States

3rd Conference on
Industrial and Engineering Applications
of AI and Expert Systems
15 - 18 July 1990, Charleston, South Carolina

See announcement on page 56 for details

Conference on Logic Programming
16 - 19 October 1989, Cleveland, Ohio

Topics include: Applications of LP; LP and databases; Theory of logic and functional programming; Parallel execution of logic programs; Implementation of LP systems; Inference machines.
Contact: Ewing Lusk, Mathematics and Comp. Sci. Division, Argonne National Lab, Argonne, IL 60439.

Symposium on AI and Mathematics
3 - 5 January 1990, Fort Lauderdale, Florida

This is the first of a biennial series featuring applications of mathematics in AI as well as AI techniques and results in mathematics. Submission dealine: Oct 20.
Contact: Dr. Fred Hoffman, Organizing Chair, Int'l Symp. on AI and Mathematics, Dept. of Mathematics, Florida Atlantic U., Boca Raton, Florida 33431, (407) 367-3345.
email: hoffmanf@servax.BITNET.

3rd Conference on Theoretical Aspects of Reasoning about Knowledge
4 - 7 March 1990, Pacific Grove, California

Topics include: Semantic models for knowledge and belief; Decision procedures and proof systems; Resource-bounded reasoning; Probabilistic knowledge; Minimal knowledge proof systems; Analyzing distributed systems via knowledge; Knowledge acquisition and learning; Knowledge and commonsense reasoning; Knowledge, planning, and action; Knowledge in economic models.
Contact: Moshe Vardi, IBM Research, Almaden Research Center, K53-802, 650 Harry Rd., San Jose, CA 95120-6099.
(408) 927-1784. email: vardi@almvma.BITNET, vardi@ibm.com.

6th IEEE conference on AI Applications
5 - 9 March 1990, Santa Barbara, California

Papers describing significant unpublished results are solicited along three tracks: Engineering/Manufacturing; Business/Law/Decision Support; Enabling Technology (knowledge acquisition, representation, reasoning, searching, learning, software life cycle, consistency maintenance, verification/validation, project management, user interface, integration, problem solving architectures). There will be a best student paper award of $1,500.
Contact: Se June Hong (Room 31-206), IBM T.J. Watson Research Center, PO Box 218, Yorktown Heights, NY 10598.
(914) 945-2265. Fax: (914) 945-2141. Telex: 910-240-0632.
email: CSNET: HONG@IBM.COM.

AAAI 1990 Spring Symposium Series
27 - 29 March 1990, Palo Alto, California

There will be approximately 10 symposia, each limited to about 40 participants, selected on the basis of statements of interest submitted to the program chairs.
Contact: Hector Levesque, Dept. of Comp. Sci., U. of Toronto, Toronto, Ontario M5S 1A4. email: CSNnet: hector@ai.toronto.edu.

2nd Workshop on User Modeling
30 March - 1 April 1990, Honolulu, Hawaii

Topics include: Acquisition of user and student models; Plan recognition; Representation of user models; User stereotypes; Dialog planning and response tailoring; Levels of user expertise; Student modeling and tutoring strategies; Shell systems for user modeling; Conceptual models, mental models. Submission materials: 4 copies of extended abstract or full paper (min 6 pages double-spaced). Submission deadline: Nov 30, 1989.
Contact: Wolfgang Wahlster, Dept. of Comp. Sci.,
U. of Saarbruecken, D-6600 Saarbruecken 11, West Germany.
Phone +49 681 302 2363.
email: wahlster%fb10vax.uni-saarland.dbp.de.

Applications of AI VIII
16 - 20 April 1990, Orlando, Florida

Topics include: Expert/knowledge-based systems; Knowledge Acquisition; Natural Lanuages; Robotics (perception, planning, control,vision); Autonomous Systems and Vehicles; Planning and scheduling; Manufacturing; Diagnosis; Learning; Aeorspace.
Contact: SPIE Technical Program Committee/Orlando'90,
PO Box 10, Bellingham, WA 98227-0010. (206) 676-3290.
Telex: 46-7053.

2nd Conference on Innovative Applications of AI

1 - 3 May 1990, Washington, DC

Topics include: Natural language, Knowledge acquisition; Machine learning; Computational vision; Speech; Robotics. Submission material: 5 copies of max 8 single spaced pages. Submission deadline: Dec 12, 1989.
Contact: IAAI-90, AAAI, 445 Burgess Dr., Menlo Park, CA 94025. (415) 328-3123.

3rd Workshop on Nonmonotonic Reasoning

31 May - 3 June 1990, South Lake Tahoe, California

Topics include: General theories of defeasible inference; Comparison of formal systems; Applications to planning, commonsense reasoning; Knowledge update and truth maintenance; Relation to probability models; Theories of inheritance with exceptions; Argument-based systems; Proof theory, complexity, and automation. Attendance by invitation only, limited to 40 and some students. Submission material: 5 copies of 4000 word abstract. Submission deadline: Dec 17, 1989.
Contact: Kurt Konolige, SRI International EJ272,
333 Ravenswood Ave., Menlo Park, CA 94025. (415) 859-2788.
email: konolige@ai.sri.com.

AIENG 90
5th Conference of Applications of AI in Engineering

17 - 20 July 1990, Boston, Massachusetts

Topics include: Design; Diagnosis; Planning; Process Control; Robotics; Tutoring; Sensing and interpretation; Constraint reasoning; Learning; Qualitative models; Tools; User interfaces; Problem solving; Representation. Submission material: draft paper. Submission deadline: Oct 13, 1989.
Contact: Sandra Elliott, Conference Secretary, Computational Mechanics Institute, 25 Bridge St., Billerica, MA 01821,
(508) 667-5841. Fax: (508) 667-7582.

AAAI-90
8th Conference on AI

29 July - 3 August 1990, Boston, Massachusetts

Topics include: AI and education; Automated reasoning; Cognitive modeling; Commonsense reasoning; Impacts of AI technology; Knowledge acquisition and expert system design methodologies; Knowledge representation; Machine architectures and computer languages for AI; Machine learning; Natural languages; Perception and signal understanding; User interfaces; Philosophical foundations; Robotics. Submission material: 6 hard copies, max 11 pages plus bib and abstract. Submission deadline: Feb 20, 1990.
Contact: AAAI-90, AAAI, 445 Burgess Dr., Menlo Park, CA, 94025-4596.

Workshop on Term Subsumption Languages in Knowledge Representation

18 - 20 October 1989, New Hampshire

Focus in knowledge representation formalisms and systems based on term subsumption languages in the tradition of KL-ONE. All aspects of such formalisms and systems may be considered, including formal specification and properties of term subsumption languages, design and implementation of knowledge representation systems incorporating term subsumption languages, and use of these systems. This will be a working meeting, restricted to 35 participants.
Contact: Peter Patel-Schneider, AT&T Bell Labs, 600 Mountain Ave., Room 3C-410A, Murray Hill, New Jersey 07974.

Outside North America

2nd Symposium on AI

23 - 27 October 1989, Monterrey, Mexico

Topics include: Knowledge acquisition and representation; Machine learning; Evaluating knowledge engineering tools; Verification and validation; Constraint directed reasoning; Uncertainty management; Natural language; Truth maintenance systems; Managing ES projects.
Contact: David Garza, Centro de Investigacion en Informatica, ITESM., Suc. de Correos "J", C.P. 64849 Monterrey, N.L. Mexico.
Phone: (83) 58 2000 ext. 5133.
email: BITNET: ISAI@TECMTYVM.

Neuro Nimes 89
2nd Workshop on Neural Networks and Applications

13 - 16 November 1989, Nimes, France

Topics include: Algorithms on associative memory, self-learning, perception, control; Specialized languages for describing and simulating neural networks; Simulation techniques and neural computers; Tools for design and evaluation; electronic, optical, and molecular implementations.
Contact: Jean-Claude Rault, EC2, 269-287, rue de la Garenne 92000 Nanterre, France. phone (33.1) 47 80 70 00.
Telex: 612 469 Fax: (33.1) 47 80 66 29.

TENCON 1989 - AI and Neural Networks

22 - 24 November 1989, Bombay, India

TENCON is the premier IEEE International Conference sponsored by Region 10 (Australia, China, Hong Kong, Indian subcontinent, Japan, Korea, New Zealand, Singapore, etc.). Topics include: Expert systems technology and its application to information processing; Knowledge representation; Learning; Languages; Logic; Search techniques; Neural networks; Robotics.
Contact: V. Seshadri, AT&T Bell Laboratories, MT3G122,
200 Laurel Ave., Middletown, NJ 07748, USA. (201) 957-6516.
Fax: (201) 957-7545. email: ...!att!mtfmi!sesh.

AI in Industry and Government

23 - 25 November 1989, Hyderabad, India

Topics include: Recent trends and impact of AI technology; Application of AI techniques in areas such as banking systems, design and manufacturing, health management, law, marketing, planning and control, project management, sensing and interpretation, training.
Contact: The Director, National Centre for Expert Systems, Institute of Public Enterprise, Osmania U. Campus, Hyderabad 500 007, India. Phone: 868145. Telex: 0425-7064 IPE IN.

Knowledge-Based Computer Systems

11 - 13 December 1989, Bombay, India

Topics include: Learning; Logic programming; AI and engineering; Natural language understanding; AI systems and software; Pattern recognition; Intelligent tutoring systems; Reasoning; Knowledge representation; Speech; Vision.

Contact: Dr. S. Ramani, Program Committee, National Centre for Software Technology, Gulmohar Cross Road No. 9, Bombay 400 049, India. Phone: (91-22) 620-1606.
email: uunet!shakti!ikbcs. Telex: 11-78260 NCST IN.

6th Israeli Conference on AI and Computer Vision
26 - 27 December 1989, Tel-Aviv, Israel
Topics include: AI and education; AI languages, logic programming; Automated reasoning; Cognitive modeling; Expert systems; Inductive inference, learning and knowledge acquisition; Knowledge theory, logics of knowledge; Natural language processing; Planning and search; Image processing and pattern recognition; Image analysis and computer vision; Visual perception; Robotics.
Contact: *Vision:* Dr. Y. Yeshurun, 6th IAICV, Dept. of Comp. Sci., Tel-Aviv U., 69978 Tel-Aviv, Israel.
AI: Dr. J. Rosenschein, Dept. of Comp. Sci., The Hebrew University, 91904 Jerusalem, Israel.

Workshop on Meta-programming in Logic
4 - 6 April 1990, Leuven, Belgium
Topics include: Foundations of meta-programming (MP); Design and implementation of language facilities for MP; Knowledge representation with MP; Applications of MP. Submission material: 4 copies. Submission deadline: Nov. 30, 1989.
Contact: Maurice Bruynooghe, META-90, Dept. of Comp. Sci., Katholieke Universiteit Leuven, Celestijnenlaan 200A, B-3030 Heverlee, Belgium. email: maurice@cs.kuleuven.ac.be (program). dannyd@cs.kuleuven.ac.be (organization). Fax: +32-16-205308. Telex: 23674.

7th Conference on Logic Programming
18 - 22 June 1990, Jerusalem, Israel
Topics include: Applications; Deductive databases; Implementation; Languages; Machines; Parallel execution; Programming methodology and tools; Relations with AI; Theory and foundations. Submission materials: 5 copies, in English, must contain a 200-250 word abstract, max 4000 words. Submission deadline: Dec 1.
Contact: David Warren, Dept. of Comp. Sci., U. of Bristol, Bristol BS8 1TR, UK. email: warren@compsci.bristol.ac.uk.

10th Conference on Automated Deduction
23 - 27 July 1990, West Germany
Topics include: Theorem proving; Decision procedures; Logic programming; Unification; Program verification and synthesis; Inference systems; Term rewriting; Deductive databases; Applications. Submission materials: 6 copies of: max 5000 word research papers; 600 word summaries of working systems; 1500 word problem sets. Submission deadline: Nov. 27, 1989.
Contact: Mark E. Stickel, AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025.
email: Stickel@AI.SRI.COM.

13th Conference on Computational Linguistics
20 - 25 August 1990, Helsinki, Finland
16 - 18 Tutorials August.
Submit a topical paper on some critical issue in computational linguistics, or a project note with software demo. Submission material: 12,000 characters for topical paper, 6,000 for a project note. Submit via email, or 5 paper copies.
Submission deadline: Dec. 1, 1989.
Contact: COLING 90 Program Committee, Hans Karlgren, KVAL, Skeppsbron 26, S-111 30 Stockholm, Sweden.
Phone: (468) 789-6683. Fax: (468) 796 9639.
Telex: 15440 KVAL S.
email: COLING@COM.QZ.SE or COLING@ZCOM.BITNET.

1st Japanese Knowledge Acquisition for Knowledge-based Systems Workshop
25 - 31 October 1990, Kyoto and Tokyo, Japan
Topics include: Computerized and manual transfer and modeling of expertise; Apprenticeship, explanation-based, and other learning systems; Methods to capture design knowledge and requirements; Issues in cognition and expertise that affect the knowledge acquisition (KA) process; Validation of KA techniques and the role of KA techniques in validating knowledge-based systems; Extracting and modeling knowledge from multiple sources, including text; Integration of KA techniques with hypermedia, DBMS, simulators, spreadsheets. Submission material: 5 copies of draft paper (max 20 pages). Submission deadline: February 28, 1990.
Contact: Brian Gaines, Dept. of Comp. Sci., U. of Calgary, Calgary, AB. email: gaines@calgary.cdn.

IJCAI-91
24 - 30 August 1991, Sydney, Australia
Contact: Prof. Barbara Grosz, Conference Chair, IJCAI-91, Aiken Computation Lab 20, Harvard U., 33 Oxford Street, Cambridge MA 02138, USA. (617) 495-3673.
email: grosz@harvard.harvard.edu.

IJCAI-93
29 August - 3 September 1993, Chambergy, France
Contact: Jean-Pierre Laurent, Local Arrangements Chair, IJCAI-93, Université de Savoie, BP 1104, F73001 Chambergy, France. Phone: (+33-79) 961-062. email: jplaure@imag.fr.

CSCSI/SCEIO Application and Order Forms

Use the forms below to subscribe to the journal ***Computational Intelligence***, to order publications and to join the Canadian Society for Computational Studies of Intelligence (CSCSI/SCEIO), with which you will receive ***Canadian Artificial Intelligence***. Complete the form of you choice and send it to CIPS (which administers membership for the CSCSI/SCEIO) at the address below, with the appropriate fee. (NOTE: Those residing outside of Canada who wish to order the 1988 CSCSI/SCEIO Conference Proceedings must mail the order form to Morgan Kaufmann Publishers.)

CIPS, 243 College Street (5th floor), Toronto, Ontario, CANADA M5T 2Y1

CSCSI/SCEIO Membership
Canadian Artificial Intelligence back issues

❑ I wish to join CSCSI/SCEIO and receive Canadian Artificial Intelligence ($35.00 Cdn.).

❑ I am a student (subtract $10.00 Cdn.)
and/or
❑ I am a member of CIPS (subtract $10.00 Cdn.)

❑ Please send me the following back issues of Canadian Artifical Intelligence ($10.00 Cdn. each including postage, $15.00 Cdn. for issue #3).*

Name _____

Mailing Address _____

_____

* If an issue you request is out of print, a photocopy will be provided. Issue #3 (March 1985) includes the supplement Towards a Canadian Fifth Generation Research Plan.

Computational Intelligence
Subscriptions

❑ Please enter my subscription to Computational Intelligence at the CSCSI/SCEIO non-institutional member discount rate ($20.00 Cdn.). ‡

Name _____

Mailing Address _____

_____

‡ Computational Intelligence subscriptions are filled by its publisher, the National Research Council of Canada. CIPS only certifies your eligibility for the discount and fowards your order.

CSCSI/SCEIO Conference Proceedings

Please send me the following CSCSI/SCEIO conference proceedings (indicate number of copies desired):

❑ # of copies ___ Edmonton, 1988 (Within Canada: members – $35.00 Cdn.; non-members – $40.00 Cdn. Add $5.00 Cdn. for postage. Mail to CIPS.) (Outside Canada: $35.00 US. Postage within U.S.: $2.00 US for the first copy, $.75 US for each additional copy. Postage outside U.S.: $3.00 US for the first copy, $2.00 US for each additional copy. Within California: add 7% sales tax. Mail to: Morgan Kaufmann Publishers, Order Fulfillment Center, P.O. Box 50490, Palo Alto, CA 94303, USA)

❑ # of copies ___ Montréal, 1986 ($30.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

❑ # of copies ___ Saskatoon, 1982 ($25.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

❑ # of copies ___ Victoria, 1980 ($25.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

Name _____

Mailing Address _____

_____

National Research Council Canada
Conseil national de recherches Canada

Computational Intelligence Intelligence informatique

Computational Intelligence, the official journal of the Canadian Society for Computational Studies of Intelligence, is a quarterly journal first published in 1985 by the National Research Council of Canada. It contains high-quality theoretical and experimental research papers in computational (artificial) intelligence, by encouraging contributions from the following fields: knowledge representation; natural language understanding; computational vision; applications of artificial intelligence; logic programming; theorem proving; learning; cognitive science; problem solving and planning; languages and tools for artificial intelligence; speech understanding; game playing; philosophical implications; and foundations of artificial intelligence. Three special issues were published in 1988: "Taking Issue: an inquiry into Computer Understanding"(February), "AI in France"(May), and "Planning"(November). The Journal is international in content and distribution and is quickly becoming one of the leading AI journals in the world

La revue Intelligence informatique, organe officiel de la Société canadienne pour l'étude de l'intelligence par ordinateur, est publiée quatre fois par an par le Conseil national de recherches du Canada depuis 1985. Elle renferme des articles de qualité dans le domaine de l'intelligence artificielle, tant au niveau théorique qu'expérimental, et encourage la publication de communications dans les domaines suivants: la représentation des connaissances, la compréhension des langages naturels, la vision computationnelle, les applications de l'intelligence artificielle, la programmation logique, la démonstration de théorèmes, l'apprentissage, la science cognitive, la résolution et la planification de problèmes, les langages et les outils de l'intelligence artificielle, la compréhension de la parole, les jeux, la portée philosophique et les fondements de l'intelligence artificielle. Trois numéros spéciaux ont été publiés en 1988: Forum : "An Inquiry into Computer Understanding" (février), "AI in France" (mai) et "Planning" (novembre). La revue est internationale tant par l'origine des auteurs qui y contribuent que par sa distribution et est rapidement en voie de devenir une publication de pointe dans le domaine de l'IA.

Computational Intelligence is published quarterly. Subscription rates for 1990 are $48.00 for Individuals and $120.00 for Institutions.

La revue Intelligence informatique est publiée 4 fois par an. En 1990, l'abonnement individuel coûte 48$ et l'abonnement collectif 120$.

For further information contact

Subscription Information
Research Journals, M-55
National Research Council of Canada
Ottawa, Ontario, Canada
K1A OR6

Pour tout information complémentaire, communiquez avec

Abonnement
Revues scientifiques, M-55
Conseil national de recherches Canada
Ottawa (Ontario)
K1A 0R6

Canada

AI System Development Tools

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

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

Nexpert
Full scale expert system development tool for PCs with rich graphics support. Object hierarchy with multiple inheritance. Rules support pattern matching, integrated forward and backward chaining, automatic goal generation, non-monotonic reasoning. Open architecture allows integration with external programs. For use on IBM AT, 386, PS/2 and Macintosh II, $6,100. For use on DEC VAX 2000, II, III under VMS/UIS, and UNIX workstations, $9,760.

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

Savvy PC
Sophisticated applications generator which has a powerful pattern recognition facility. Three major components: the Savvy Database Manager to store, manipulate, and retrieve information in plain-language style; the Savvy Language & Tutorial to allow customization of database applications; and the Savvy Retriever to translate English queries using pattern recognition. Runs on IBM PC or compatibles. $415.

LPA PROLOG Professional
First PROLOG compiler to allow development of large applications on IBM PC AT. Provides a general purpose high-level programming language with a superb range of functions and data handling capabilities. PC Dual Compiler System $1,215. flex (Forward chaining Logical EXpert) system toolkit (ms-dos) $485. MacProlog Color Edition $725.

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

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

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

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

Delphi Common Lisp (DCL)
Powerful new standard object-oriented programming facility. Compact, portable implementation of the complete Common Lisp standard. "Multithread" facility to create and control concurrent threads of program execution. Binary license for Sun 2, 3/50 and 3/100, $975. Sun 3/60, 3/200, and 386i, $1,945. Sun 4, $2,930. Apollo DN, $1,950.

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

ENQUIRY AND ORDER: Call (613) 592-3030 or send Purchase Order to Applied AI Systems, Inc.
Shipping and Handling charges and provincial sales tax (Ontario residents only) apply.

Applied AI Systems, Inc. is an authorized dealer of the above AI software/hardware products.
Prices are subject to change without notice.

Applied AI Systems, Inc.
Suite 602, Gateway Business Park
300 March Road
KANATA, Ontario, Canada K2K 2E2
Tel. (613) 592-3030
Fax (613) 592-2333