Canadian Artificial Intelligence

Intelligence Artificielle au Canada

July juillet 1989 No. 20

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 ordi

PROFILE: A Conversation with Len Schubert
Connie Bryson
PROFILE: Une conversation avec Len Schubert

A Knowledge-Based Approach to Discrete Systems Modeling and Simulation
Lubomír Masár
Une approche de la modélisation et de la simulation des systèmes discrets basée sur les connaissances

Recent AI Activities at the University of Guelph
Mary McLeish, Amelia Fong Lochovsky, David K.Y. Chiu
Activités Récentes à l'Université de Guelph

Artificial Intelligence at Queen's
Janice Glasgow
Intelligence Artificielle à Queen's

International Computer Science Conference '88
Nick Cercone et al.
Conférence Internationale en Informatique '88

NEXPERT OBJECT

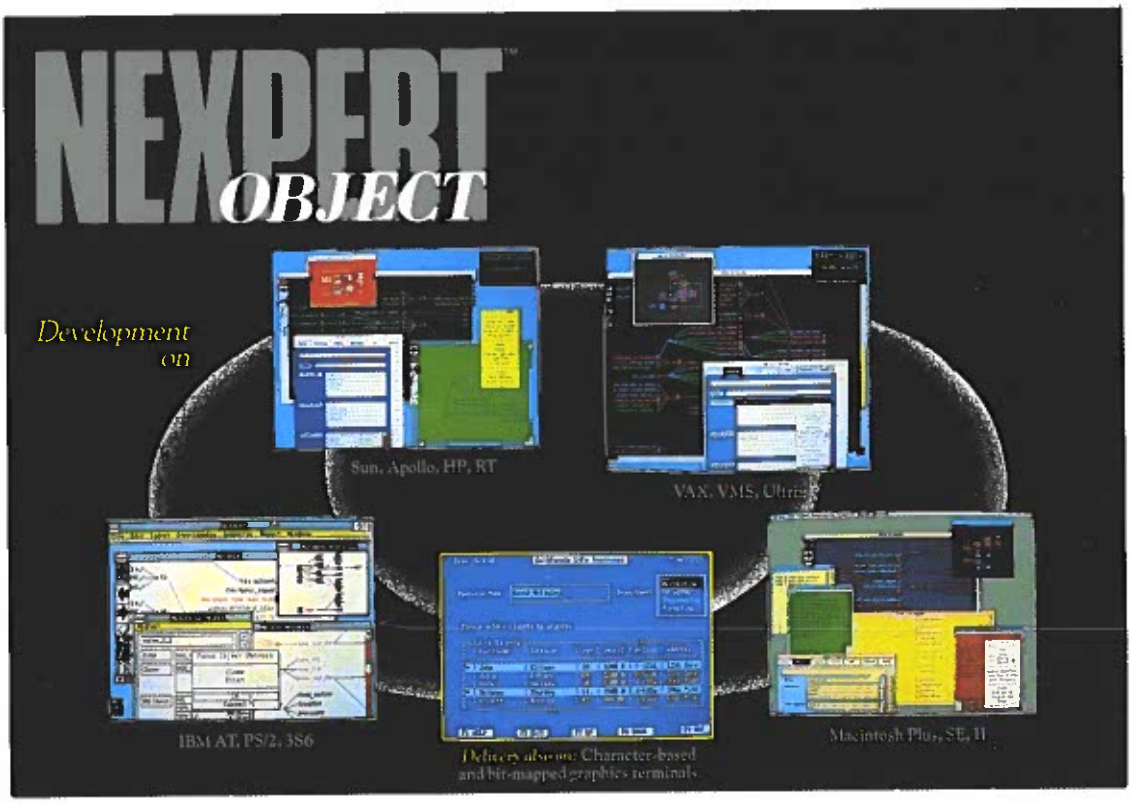

REASONS FOR NEXPERT'S SUCCESS

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. 20 July/Juillet 1989

Contents

Communications
Executive Notes 3
Notes from Members 4
Letters to the Editors 4

AI News 6

Feature Articles
PROFILE: A Conversation with Len Schubert 9
Connie Bryson
A Knowledge-Based Approach to Discrete Systems Modeling and Simulation 12
Lubomír Masár

Research Reports
Recent AI Activities at the University of Guelph 18
Mary McLeish, Amelia Fong Lochovsky, David K.Y. Chiu
Artificial Intelligence at Queen's 22
Janice Glasgow

Conference Reports
International Computer Science Conference '88 26
Nick Cercone et al.
Second International Workshop on Artificial Intelligence and Statistics 17
Barry de Ville

Publications
Book Reviews 32
Books Received 37
Computational Intelligence Abstracts 38
Technical Reports 39
World Watch 42

Conference Announcements 51

Communications
Notes administratives
Notes des membres
Lettres aux éditeurs

Nouvelles de l'IA

Gros Titres
PROFILE: Une conversation avec Len Schubert
Connie Bryson
Une approche de la modélisation et de la simulation des systèmes discrets basée sur les connaissances
Lubomír Masár

Rapports de Recherches
Activités Récentes à l'Université de Guelph
Mary McLeish, Amelia Fong Lochovsky, David K.Y. Chiu
Intelligence Artificielle à Queen's
Janice Glasgow

Rapports des Conférences
Conférence Internationale en Informatique '88
Nick Cercone et al.
Deuxième Congrès International en Intelligence Artificielle et Statistiques
Barry de Ville

Publications
Critiques de livres
Livres reçus
Résumés d'*Intelligence informatique*
Rapports techniques
Vue sur le monde

Annonces des Conférences

Canadian Society for Computational Studies of Intelligence

Founded 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 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 page dernière. 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.

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

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 submissions ne devraient 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.

Réclame:
Les prix pour les annonces et les trousses pour la presse sont disponibles sur demande. Écrivez à Marlene Jones 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
ou à / or to: 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
ou à / or to: 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

In Praise of Isolation

One of the established truisms of artificial intelligence is that research and/or development must be done in large groups with a critical mass of expertise. After all, isn't this the way that places like M.I.T., Stanford, and Carnegie-Mellon have achieved success over the thirty-plus years that the field has been in existence? Isn't it a fact that companies wanting to get into the field usually get a large number of scientists and engineers together to form the core of their AI R&D group? Isn't this what the CIAR has tried to do with its network of interlinked scholars, and what the PRECARN consortium has been set up to achieve for pre-competitive AI research?

While it is undeniable that many successes have been achieved by large groups working together in mutual synergy, there are also advantages to not being in a big group. Isolation has its benefits! Let me try to outline some of the reasons why.

One of the advantages of being in a group is that the increased knowledge shared by the group can help individuals in the group prune useless branches from their own research explorations. However, "group think" can sometimes prune promising branches which might lead to new discoveries. Many large groups are slaves to particular goals, often goals which have been set in order to maintain the large amount of funding needed to support the group. The individual's ability to follow up on interesting, but peripheral, ideas is restrained by the need to work towards the group's specific targets. The isolated researcher, on the other hand, is free to work on whatever he or she pleases. Freedom from group think and the lack of targetted goals thus places the isolated researcher in a position to take independent perspectives on the research discipline.

Of course, any researcher needs colleagues to bounce ideas off of, and the isolated researcher needs this too. Things like electronic mail, the telephone, cheap air travel, the postal service, workshops, conferences, journals, etc., make communication possible with colleagues elsewhere (well, maybe not the postal service). However, this is no substitute for daily discussions over coffee or at the pub. To achieve daily feedback, the isolated researcher may be forced to talk to people outside his or her own immediate research discipline, perish the thought! Paradoxically, the "drawback" of having no immediately relevant colleagues may in fact actually enhance the isolated researcher's ability to make a contribution, since such interdisciplinary perspectives are often the key to real innovation in a field, especially in AI.

While all of this may be true in many scientific fields, you may protest, AI has its own particular problems. In AI the body of knowledge is too large, the field moves too fast, and the need for expensive hardware and software is too great for a single researcher to achieve any success. I don't think these are relevant factors even in AI; in fact, I don't really think that AI is any different from other fields on these attributes.

The amount of knowledge needed to do AI is indeed vast. Even if the number of papers which need to be read to understand these issues is considerably less than the immense number of papers that are published each year (given how much AI research essentially re-invents the wheel), there is still a real problem for the isolated researcher who doesn't have easy access to others who can help in sorting out the relevant from the irrelevant. Moreover, the field is still quite interdisciplinary, which makes the task of understanding a broad range of issues even more difficult. This is perhaps the biggest problem for the isolated researcher. He or she can try to understand everything (and likely not have time to do anything), or accept being unable to keep up in every area of AI and attempt to actually achieve success in some specialization. Thus, the key for the isolated researcher is to pick a specific area of interest and stick to it.

The next problem faced by the isolated researcher is that the field is moving quickly. However, any perusal over time of conference proceedings, research journals, and books will certainly show that movement in AI is more of a cycle (hopefully an outward spiral) than it is a monotonic progress forward. Neural networks come and go and then come back; logic-based knowledge representation is "in", then "out", then "in" again; domain-specific

research is done, then general theories are de rigueur, then narrow, focussed investigations are again the way to go. The point is that the isolated researcher will likely be somewhat immune from the latest fad (not being as subject to group think) and will likely be able to keep plugging away at some particular approach through thick and thin, hence making real progress in his or her area of specialization.

Finally, the isolated researcher would seem to have a real problem in getting the kind of hardware and software environments which are needed to do AI. This is a problem which is rapidly vanishing. The cost of hardware and software useful for AI is diminishing rapidly, and a very nice programming environment is now affordable to even the most isolated researcher, as long as he or she has a modest access to resources. Even if the researcher cannot find resources (or does not wish to spend resources on hardware and software), the lack of an adequate computing environment has often stimulated the use of brain cells rather than computer memory cells and has led to interesting theoretical research. The imperative to "hack" often obstructs real progress in AI, which, as in any other discipline, must be concerned at least partially with the elaboration of good theoretical ideas.

There is, of course, nothing wrong with big groups doing AI. I just would like to challenge the idea that this is the only way to proceed. I think that at this time in Canadian AI it is wise to reflect on the need for diversity, the importance of not putting all of our resources into big, networked consortia of various sorts, leaving no room for the "little guy". Historically, the Canadian artificial intelligence R&D community was geographically fragmented into small, more or less independent research groups. Contrary to popular belief, this community of isolated researchers did fairly well for itself – we produced excellent research, graduated top-ranked students (the three most recent Computers and Thought award winners each have at least one degree from a Canadian university), formed the world's first national AI society (AISB and SIGART were international), hosted a conference some four years before AAAI even existed as an organization, and were instrumental in the development of a number of interesting AI products. When the CIAR came along to establish "big time" AI in this country, it found many highly qualified scholars of international reputation already plying their trade here. Canadian AI is still doing fairly well for itself, as the number of papers accepted at the upcoming IJCAI conference in Detroit written by Canadian authors attests (in fact Canada had the highest accept ratio of any country in the world). Let us happily welcome attempts to inject new funds into AI R&D, but let us not forget that often "in isolation there is strength".

Gordon McCalla
University of Saskatchewan

1989 General Meeting of CSCSI

The annual general meeting of CSCSI/SCEIO will be held in conjunction with IJCAI '89 (Detroit, Michigan). The meeting will be held in room M2-29 Cobo, 12:30 - 2:00 p.m., on Thursday, August 24, 1989. Please attend!

Notes from Members Letters to the Editors

As editors of the *Computational Intelligence* journal, we would like to respond to Ilie Popescu's letter in the April issue of *Canadian Artificial Intelligence* in which the suggestion is made that there be a refereed section in the magazine for applications-oriented research results. Dr. Popescu raises an important issue – the lack of Canadian outlets for practical AI research. However, he is incorrect when he says that *Computational Intelligence* is unsuitable since its "profile is strictly theoretical".

Computational Intelligence, in fact, has a broad mandate to publish research of both a theoretical and applied nature. To be sure, many of the papers which appear in the pages of *Computational Intelligence* are theoretical, but this is largely a reflection of the kind of submissions received by the journal. If we were to receive more papers with an applications focus, there would be more papers in the journal with a practical orientation.

It seems to us that, instead of setting up a refereed section of *Canadian Artificial Intelligence*, it would be better to encourage AI practitioners to submit papers to *Computational Intelligence* more frequently than they now do. The journal has no bias against practical research, as long as it is of international calibre and has something important and novel to say. Indeed, over the years we have published many papers of practical import and in the future would welcome an influx of good applications-oriented papers.

In sum, let's leave the pages of *Canadian Artificial Intelligence* free from the interference of referees and keep the magazine a turbulent, exciting forum for the interchange of ideas, half-baked or not. Refereed research has its place in the pages of a journal, not those of a magazine.

Gordon McCalla and Nick Cercone
Co-Editors of *Computational Intelligence*

(*Editor's Note*: To order *Computational Intelligence*, see the last page of this issue)

No Bugs in the Church-Turing Thesis

The usual attack on AI starts with the invocation of some mystical "intuition" which a system of computer+program "inherently" lacks and concludes that a computer program in principle can never do thus and so. Another, less common, attack starts from some real proficiency of function in the natural world outside of man which has not yet been equalled by a robot system. Then the conclusion is immediately announced that this or that part of nature outthinks in principle any computer program.

In their note in this quarterly, April 1989, entitled "Bugs in the Church-Turing Thesis", authors Thimbleby and Witten comment on the fine locomotor efficiency of a

millipede and deduce that nature computes far more efficiently than we ever can on our machines. The conclusion seemed to them to follow that the basic thesis of Turing and Church is in error. Their remarks belong to the second type of anti-AI approach. But mere complexity of structure or agility of function does not imply thought. I would like to emphasize and amplify this point.

Consider a snowflake. It has marvelous precision and intricacy but it is no testimony to the ability of water molecules to think up a design. Consider a liver cell, a huge storehouse of biological catalysts. When it grows to a certain size it divides. The two new cells each contain a copy of the same enormously complicated DNA strands. The replication of the DNA information depends on molecular properties and was not "thought through" by the liver cells. Even the wonderful liver cells of thinking Homo-sapiens are not considered by anyone to be capable of thought or computation. Furthermore, biological intricacies have arisen by evolution, not by conscious design.

The electrons of all of the hydrogen atoms of the universe seem to keep the same average distance from their protons. But the electron of a hydrogen atom does not carry out a quantum mechanical calculation in order to determine what its average distance from the proton should be. The motions of the electron relative to the proton are properties of the system. Unthinking nature does not "know".

The confusion of proficiency of locomotion in nature with computation leads naturally to some kind of insect worship, wherein a millipede or a bumblebee is endowed in some sense with enormous intelligence. Indeed the authors do not fear to make the essentially meaningless statement that "computers have an IQ commensurate with the common earthworm". It has been a while since I met an earthworm that could paraphrase natural language, play tournament-grade chess, deduce the structure of molecules from X-ray cystallographic data, etc.

Current machine+program systems may be inferior in some functional ability to a millipede rushing along or a liver cell, but they do something that the locomotion system of the millipede and the liver cell do not do. They have a representation of some part of the world and they manipulate this representation. The best systems can acquire data about the outside world, build an internal symbolic representation of part of the world, manipulate this representation, generate plans and even, based on the generated plan, provide output to devices which will produce some effect in the world. Such behaviour is not displayed by single molecules or by instinctual or reflex systems. It is variously called throught or computation.

Moreover, the authors seem to imply that copying the neuron mechanism of brain activity is the only way to produce intelligent behaviour. It has been said thousands of times that AI researchers are not trying to create life, only to build intelligent systems. The authors also confuse the proficiency of receptors and effectors hampered by their inability to perform quantum mechanical calculations in real time. A robot obeys the laws of classical Newtonian mechanics. Neither brains nor computers perform quantum mechanical calculations in order to direct motion. If we wanted to understand neuron function in detail at the molecular level we might ultimately need quantum mechanics, but that is the task of psychology, not of AI.

Malcohm Bersohn
University of Toronto

Is AI Science or Engineering?

At first glance this seems like a useless question. AI people have their agenda. They have to get on with their work. Labelling doesn't seem like a useful task.

Everybody knows the difference between science and engineering. They are both in principle studying something. Engineers study the properties of the artefacts of man, machines. Scientists study the rest of the universe. However the scientist can and usually does stop at acquiring understanding. The engineer must thereafter apply the understanding.

The dual role of the engineer tends to puzzle us. Many, if not most, of the people called engineers don't bother to study anything. They know which tables to look up information in and they are busy doing projects which are similar or identical to those which have been done before. The absence of novelty and study in the lives of most engineers causes confusion. We have to distinguish between fundamental engineering, which is concerned with acquiring basic understanding of machines in order to find and overcome limits to their capabilities, and practical engineering, which handles the details in any particular piece of machinery. Fundamental engineering is an enormous intellectual challenge. Engineering in everyday life often avoids the challenge. Engineering students frequently fit in with this idea of avoiding the intellectual challenge.

So we are confused, and in the Western world there is a strong cult of the prestige of science versus engineering. All kinds of engineering specialties are decorated as Materials Science, Soil Science, etc. The general contempt for engineering means that bright people usually avoid it, and the level of engineering expertise of the society as a whole more and more testifies to this. AI people and computer "scientists" generally are careful in this environment not to refer to themselves as engineers.

It is instructive to compare this situation with that in Japan. In Japan the opposite is true. Pure science is not prestigious. Sufficient understanding of nature is supposed to be obtained from a book or a lecture. There the great task of life is considered to be to apply the understanding, not increase it. Computer scientists are respectfully referred to as "system engineers". Bright young people in Japan want to become engineers. The level of Japanese engineering reflects this fact. AI in Japan is considered to be a sparkling facet of system engineering.

In any common sense of these words, AI is both science and engineering. AI people must be scientists; they must study the abstract properties of algorithms, the abstract properties of cognition, logic, language, etc. They must also be fundamental engineers, experts at understanding the properties of giant programs, because AI programs are eventually the largest of all. Finally they must have that attention to practical detail which characterizes good engineers, even those who may lack curiousity as to fundamentals. Without programmers' devotion to practical detail, AI programs won't actually do the real tasks of the world.

When AI people worldwide are proud of being both engineers and scientists, a certain bottleneck will have been overcome.

Malcohm Bersohn
University of Toronto

Prolog Standardization Working Group

Ottawa will be the place of the next meeting of the international working group for the standardization of the programming language Prolog, ISO/IEC JTC1 SC22 WG17. The meeting will be on October 11-13, 1989, in the Park Lane Hotel, Ottawa. The Canadian working group invites all Prolog experts to participate in the standardization effort. If you feel you can contribute, contact: Dr. Stan Szpakowicz, Dept. of Comp. Sci., U. of Ottawa, Ottawa, Ont, Canada K1N 6N5 (613) 564-2450

All those interested in attending the meeting (as observers or advisors to the Canadian delegation) may contact the local host: Andre Vellino, Computing Research Laboratory, Bell-Northern Research Ltd., P. O. Box 3511, Station "C", Ottawa, Ontario, Canada K1Y 4H7 (613) 763-7514 email: vellino@bnr.ca fax: (613) 763-4222

Canadian Space Agency

Minister of Regional Industrial Expansion, Harvie Andre, announced on March 1, 1989, the establishment of the Canadian Space Agency. "The Agency will have a staffing base of 300 person-years and will manage Canada's civilian space program, directing expenditures of almost $3 billion to the end of the century." Dr. Larkin Kerwin has been appointed president of the agency. Contact Kathy Kennedy, Industry, Science and Technology Canada, (613) 998-5263, for more information.

Economics of Space

The *Financial Post* reports (March 20, 1989) that Canada "...ranks high in commercial space exploitation." As an indicator of this commercialization, they report that earnings in the last 12 years by Canadian companies from space-related equipment have been more than (Canadian) government spending on space in the same period.

High Technology Week '89

Minister of State (Science and Technology) William C. Winegard announced on March 22, 1989, that October 15-22, 1989, would be the fourth annual High Technology Week in Canada. The Canadian High Technology Show, the largest general electronics trade show in Canada and the second largest in North America, will highlight High Technology Week.

Canada Scholarships Program

Minister of State (Science and Technology) William C. Winegard announced on April 7, 1989, details of the Cdn$80 million Canada Scholarships Program. The program, "... designed to recognize excellence and encourage more top students, especially young women, to study

Réunion du Groupe de Travail pour la Normalisation de Prolog

Ottawa sera le prochain lieu de rendez-vous pour le groupe de travail international pour la normalisation du langage de programmation Prolog, ISO/IEC JTC1 SC22 WG17. Le rendez-vous aura lieu du 11 au 13 octobre 1989, au Park Lane Hotel, à Ottawa. Le groupe de travail canadien invite tous les experts de Prolog à participer à cette tentative de normalisation. Si vous croyez pouvoir contribuer à cette tentative, veuillez contacter: Dr. Stan Szpakowicz, Department of Computer Science, U. of Ottawa, Ottawa, Ontario K1N 6N5 (613) 564-2450

Tous ceux qui veulent assister à la réunion (comme observateurs ou comme conseillers à la délégation canadienne) peuvent contacter la dèlègation locale: André Vellino, Computing Research Laboratory, Bell-Northern Research Ltd., P. O. Box 3511, Station "C", Ottawa, Ontario, Canada K1Y 4H7 (613) 763-7514 email: vellino@bnr.ca fax: (613) 763-4222

Agence Spatiale Canadienne

Le Ministre de l'Expansion Industrielle Régionale, M. Harvie André, a annoncé le premier mars 1989 la création de l'Agence Spatiale Canadienne. "L'Agence payera 300 personne-années, et dirigera la dépense de près de $3 milliards jusqu'à la fin du siècle." Le Dr. Larkin Kerwin a été nommé président de l'Agence. Pour plus d'information, veuillez contacter: Kathy Kennedy, Industry, Science and Technology Canada (613) 998-5263

L'Économie et l'Espace

Le *Financial Post* note (le 20 mars 1989) que le Canada"... se distingue hautement par son exploitation commerciale de l'espace". Comme indicateur de cette commercialisation, il note que les bénéfices retenus pendant les 12 dernières années par des sociétés canadiennes fabriquant de l'équipement utilisé dans l'espace, sont superieurs aux dépenses pour l'espace faites par le gouvernement canadien durant la même période.

Semaine de la Haute-technologie '89

Le Ministre de l'État (Science et Technologie), M. William C. Winegard, a annoncé le 22 mars 1989 que la semaine du 15 au 22 octobre 1989 sera la quatrième semaine annuelle de la haute technologie. L'Exposition Canadienne de la Haute Technologie, la plus grande exposition commerciale générale d'électronique au Canada et la deuxième en Amérique du Nord, sera le point culminant de la semaine.

Programme de Bourses du Canada

Le Ministre de l'État (Science et Technologie), M. William C. Winegard, a annoncé le 7 avril 1989 les détails du Programme Canadien de Bourses de Cdn$80 millions. Le programme, "...conçu pour reconnaître l'excellence et d'encourager les meilleurs étudiants, en particulier les jeunes femmes, à étudier la science et le génie," a attribué

science and engineering", awarded 2,500 scholarships worth Cdn$2,000 each in 1988. This year, 2,500 more scholarships will be awarded. The scholarships can be renewed for three years, for a maximum award of Cdn$8,000. A minimum of 50 percent of the scholarships are awarded to women. The awards are aimed at students in undergraudate programs.

Ursula Franklin Wins New Wiegand Award

Canadian scientist Dr. Ursula Franklin has received the first annual Wiegand Award for Canadian Excellence (see *Canadian AI*, January 1989, No. 18, for details of this award). She received the award at a ceremony at the U. of Waterloo on March 2, 1989. Ursula Franklin has made "an outstanding contribution to our understanding of the human dimensions of science and technology". Currently Director of the Museum Studies Program at the U. of Toronto, Franklin has a Ph.D. in experimental physics.

EVES Moves to ORA

The EVES research team (Ottawa) has moved from I.P. Sharp Associates to join a Canadian subsidiary of Odyssey Research Associates (ORA). This move is the result of the acquisition of I.P. Sharp by Reuters and Canadian government foreign ownership requirements. EVES group research includes automated deduction, language semantics, logic and computer security.

New Software Products

• Gold Hill Computers (Cambridge, MA) has started shipping GoldWorks II, a release of the GoldWorks expert system shell. GoldWorks II improves the user interface of GoldWorks and sells for US$7,500.
• Graphael (Waltham, MA) has announced the availability of G-LOGIS 3.0 on DEC VAX/VMS systems. G-LOGIS is an object-oriented programming language.
• IBM (White Plains, NY) has introduced a speech analysis product for the PS/2 intended for use by professionals dealing with speech and hearing impaired individuals. It analyzes attributes of speech, displaying feedback on those attributes.
• Human Intellect (San Mateo, CA) has released a version 2.0 of its $US69.95 Instant-Expert shell for the Macintosh.
• Natural Language Inc. (Berkeley, CA) has announced a version of its Natural Language software for the DECstation 3100, for US$17,500.
• Xerox (Mountain View, CA) has released Datacopy Acutext, text recognition software for the Macintosh. Aside from being able to capture text in a variety of fonts and sizes, coping with italics, subscripts, tabs and other features, it can also adapt to the document, learning its idiosyncrasies as it is scanned. The price is US$995.
• Dragon Systems (Newton, MA) has introduced a voice transcription system for the IBM AT called DragonDictate. Priced at US$9,000, the system includes software, a microphone, and an add-in board.
• Hecht-Nielsen Neurocomputers (San Diego, CA) has introduced the ExploreNet software package for Sun computers, which allows users to run its neural net paradigms without the use of its coprocessor. Price is US$3950.
• Hecht-Nielsen Neurocomputers (San Diego, CA) has introduced 12 new neural network paradigms to its Neurosoft software, which now has 17 paradigms.
• Parallel Logic Programming Ltd. (Twickenham, UK) has introduced PARLOG, a parallel logic programming language for the Macintosh and IBM PC. Price is US$150.
• Multilogic (Budapest, Hungary), a newly formed company, in conjuction with Brainware (Berlin, West Germany) and Densitron Computers (Biggin Hill, Kent, UK), has released CS (Concurrent Sequential) Prolog, a parallel Prolog running on transputers. Unlike other parallel Prologs (e.g., Parlog, Concurrent Prolog and Guarded Horn Clauses), CS Prolog offers full backtracking. CS Prolog is claimed to be particularly useful in the domain of simulation. The problem of updating a global time variable is bypassed by giving each independent CS Prolog process its own local time variable and updating them via the Time Warp algorithm. CS Prolog is based on the Edinburgh Syntax, with about 200 builtin predicates.

2500 bourses de Cdn$2000 chacune en 1988. Cette année, encore 2500 bourses seront attribuées. Les bourses sont renouvellables trois années de suite, pour une somme totale de Cdn$8000. Au moins la moitié des bourses seront attribuées à des femmes. Ces bourses sont destinées à des étudiants sous-gradués.

Ursula Franklin Gagne le Nouveau Prix Wiegand

La savante, Dr. Ursula Franklin, a reçu le premier Prix annuel Wiegand pour l'Excellence Canadienne (voir Canadian AI, janvier 1989, no. 18 pour les détails de ce prix). Ce prix lui a été remis à une cérémonie à l'Université de Waterloo le 2 mars 1989. Ursula Franklin a "contribué de manière exceptionnelle à notre comprehension des dimensions humaines de la science et de la technologie." Actuellement Directeur du Programme "Museum Studies" à l'Université de Toronto, Franklin a reçu son Ph.D. en physique expérimentale.

EVES Déménage chez ORA

Le groupe de recherche EVES (Ottawa) a déménagé de chez I.P.Sharp Associates pour se joindre à une filiale canadienne de Odyssey Research Associates (ORA). Ceci est le résultat de l'acquisition de I.P.Sharp par Reuters, et les limitations du gouvernement canadien sur l'investissement étranger. Le groupe EVES fait de la recherche en déduction automatique, en sémantique des langages et en sécurité logique et informatique.

... (continued from previous column)

Also available with CS Prolog is Alex, an expert systems shell. Costs range from UK£665 for a CS Prolog interpreter running on an IBM PC to UK£1440 for a system packaged with Alex, running on a transputer card in an IBM PC. Educational discounts are available.

• Symbolics (Burlington, MA) has introduced a neural network package called Plexi, a system which can be integrated with other Symbolics products such as the Joshua expert system tool and Statice database. Symbolics claims a processing speed of 65,000 connections per second. Plexi runs on all Symbolics workstations, including the newly introduced XL400 and MacIvory.

• Symbolics Inc. (Burlington, MA) has announced two new networking products. The first, Symbolics X-Windows, allows users on multiple X-Windows supported consoles to interact with a Symbolics system. The second, Symbolics-NFS, allows files to be transferred between a Symbolics system and other NFS systems.

• AICorp (Waltham, MA) has released its KBMS Developer's Workstation expert system for the IBM OS/2 and is planning a release for MS-DOS. Price is US$9500.

• Franz, Inc. (Berkely, CA) has introduced Allegro Composer, an interactive program development environment for Common Lisp. It has a Lisp object inspector, a debugger, and a profiler and allows for the simultaneous inspection of multiple processes.

• Inference Corp. (Los Angeles, CA) has announced three releases of its ART-IM product (Automated Reasoning Tool for Information Management): (1) a beta test version of ART-IM/MVS for IBM mainframes, (2) a version for IBM PCs and (3) ART-IM/VMS for the VAX.

• Neuron Data Inc. (Palo Alto, CA) has released Mac Nexpert*SQL, a link between Neuron Data's expert system shell, Nexpert Object, and Oracle Corp.'s relational database, Oracle, for use on the Mac II. This enables Nexpert Object to access data records in Oracle.

• Neuron Data Inc. has released three new products for its Nexpert Object expert system shell on the Macintosh. They are (1) AIVision, a tool for creating interactive graphic interfaces for expert systems developed with Nexpert Object, (2) Nexpert Macintosh Programmer's Workshop (NMW) Library, which allows developers to embed and customize all the functions of Nexpert Object and (3) Network Innovation Corp.'s CL/1 connectivity language, which allows applications such as Nexpert Object to transparently access data from remote systems.

• Brainware (Berlin, West Germany) has released version 2.0 of its expert system shell, Zeno Professional. Zeno programs are translated into C, making it portable and fast (1,000 rules per second, according to Brainware). IBM PC, VAX, and Unix versions are available.

New Hardware Products

• NeXT Inc. (Palo Alto, CA) has announced that the NeXT computer will be available for sale to corporate clients through Businessland Inc. (San Jose, CA). NeXT products were reported to be available in Canada starting in May 1989.

• Intel Corp. has announced the 486 microprocessor. The 25MHz chip is 386 compatible, and has a RISC core and built-in numeric co-processor.

• Intel Corp. has announced the i860 RISC chip. The chip contains cache memory, 3-D graphics, and floating point math functions.

• Apollo (Chelmsford, MA) has introduced the Series 10000 Visualization System, a RISC graphics supercomputer based on Apollo's PRISM RISC architecture. Apollo claims a performance of 100,000 3-D polygons per second, all for US$94,900.

• Togai Infralogic (Irvine, CA) is a new company whose goal is to explore the area of fuzzy logic. It plans on developing a microprocessor specifically to deal with fuzzy logic. Samples of the processor, called the FC110 Digital Fuzzy Processor, should be available in the summer for approximately US$90.

• Symbolics (Burlington, MA) has announced the XL400, a member of its 3600 family workstations. Symbolics claims the XL400, based on the Ivory processor, performs at approximately twice the speed of the 3650. The XL400 comforms to the 32-bit VME bus specification, allowing the use of standard peripherals.

• Siemens of West Germany, along with the Linguistic Research Department of the University of Texas, has developed a translation system for use in translating technical documents. The system, called METAL(Machine Translation and Analysis of Natural Languages), can adjust its translation to cope with different subjects. Claimed translation capabilities are 200 pages per 8-hour day.

Corporate Changes

• Graphael (Waltham, MA) has announced an alliance with Technology Applications Inc. to apply object-oriented database technology to engineering problems.

• Apple Computer (Cupertino, CA) has acquired Coral Software (Cambridge, MA), developers of Allegro Common Lisp and Pearl Lisp.

• Artificial Intelligence Ltd. (Hertfordshire, England) and Rank Xerox of Europe have combined forces to market knowledge development tools such as the LOOPS object-oriented environment and the en-vos Lisp Software Development Environment in the European market.

• Carnegie Group (Pittsburgh, PA) and Olivetti have formed a new company, Delos, which will provide expert system products and consulting services to European markets.

• Sun Microsystems (Mountain View, CA) and en-vos (Mountain View, CA) have signed an agreement which will bundle en-vos's Medley with Sun-4 workstations, for an initial price of US$31,990.

• Texas Instruments (Austin, TX) and Intellicorp are marketing Intellicorp's KEE expert system tool with TI's Macintosh II-based microExplorer.

• The U.S. Air Force Avionics Lab (Wright-Patterson, OH) has awarded contracts to Martin Marietta and MIT's Draper Labs for research on neural networks. The Marietta contract, worth US$820,000, will investigate the development of a robot arm modelled on the human arm. The MIT contract, for US$622,000, will investigate adaptive flight control, where the neural net learns aircraft control functions.

A Conversation with Len Schubert

by Connie Bryson

Une conversation avec Len Schubert

RÉSUMÉ: Il est difficile de persuader Len Schubert de parler de lui-même. Ce chercheur grandement reconnu en IA est trés modeste et n'apprécie pas les entrevues. Cependant, l'écrivain indépendante Connie Bryson a réussi à le joindre au téléphone dans son bureau à l'Université de Rochester. Schubert, agé de 48 ans, est en congé sans soldes de l'Université de l'Alberta.

Dans l'entrevue, Schubert a exprimé le fait qu'il n'était pas sûr d e son retour au Canada. Ceux qui se soucient de la fuite des cervaux pourraient lire ses commentaires sur l'attirance de faire de la recherche aux États-Unis. La fuite des cerveaux ne concerne pas Schubert; ses commentaires reflettent les idées d'un homme dédiéà sa recherche.

Getting Len Schubert to talk about himself is like pulling teeth. The highly regarded AI researcher is modest to a fault and a reluctant interview subject. However, free-lance writer Connie Bryson managed to catch up with him on the telephone from his office at the University of Rochester. Schubert, 48, is on unpaid leave from the University of Alberta.

In the interview, Schubert says he is unsure whether he will return to Canada. Those worried about the "brain drain" might make note of his comments on the attractions of doing research in the United States. But brain drain issues don't interest Schubert. His interest is in AI; his comments reflect the views of a man dedicated to his research.

Canadian AI: What motivates your research? Is it a fascination with subjects like temporal reasoning or are you looking farther ahead?

Schubert: I'm very much goal-oriented. I want to build a system that one can talk with, that understands in a human-like way, that can plan, reason, and answer questions. The effort to build a machine that would have a human-like intelligence, or better, is what I've wanted to be part of ever since I first heard about it 20 years ago. There are a lot of practical applications, and these do interest me, but what really excites me is the idea of a machine that has human-like versatility, generality, the ability to converse in plain language.

Canadian AI: Where does your work on temporal reasoning and taxonomic relationships fit in?

Schubert: These relationships play a very important role in human understanding, language understanding, intelligence.

To build this human-like machine, one needs a special computational mechanism to handle time relationships because they're very much a part of, for instance, a story line. In understanding a story, there are some things we infer automatically. My aim is to design some mechanism that would have this human-like efficiency in handling time relationships and the other kinds of relationships we handle very quickly, for example spatial or taxonomic relationships. We have no difficulty at all answering a question like "Mary is a girl. Is she a computer?" But if you analyze this from a logical point of view, there are a fair number of inference steps involved. However we know instantly that she is not a computer, apparently without any thought. That's what we're trying to duplicate in the machine – to be able to do this so quickly and in a way that essentially consumes no noticeable computational resources. It's difficult because we don't have complete access to what is really going on in the human mind.

All of my work on specialized methods for efficient inference is really part of a larger language understanding project, part of the effort to mechanize the kinds of inference, the kinds of knowledge representation that are needed to support language understanding.

What motivates my approach to AI research is a belief that language is a window to the mind. It is the best overt manifestation that we have at our disposal to determine what the language of the mind is, what knowledge representation is.

I'm taking it as a methodological assumption that language is in some sense very similar to the representations we are using when we think. And so, taking that as an assumption, it makes sense to tackle the problems of knowledge representation and learning from a linguistic point of view.

I'm interested in how we can represent a particular linguistic expresssion in a machine in a formal way, and in a way we can manipulate, to get the kind of inference processes that people are capable of. I think that is the key to the mind.

Canadian AI: Where would you say AI research stands now in terms of building this human-like machine? Are we at the beginning, the middle, the end?

Schubert: I think we're farther than the beginning but it's very hard to say.

I guess those of us working in the area have a worst-case conception of what might happen, a best-case conception, and an expected-case. The best case is what motivates us and keeps us going.

To keep working on this kind of difficult problem, day after day, you have to believe that success is just around the corner. But on the other hand, if you're at all realistic,

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

you also realize it may not be just around the corner. So you have to be of two minds.

It's likely to take more than eight years and less than a hundred to create this intelligent machine you asked me about. What keeps me optimistic is that it just might be eight years. But it's more likely somewhere in between those two.

Canadian AI: What are the major hurdles AI researchers have yet to overcome?

Schubert: Although research always reveals new challenges, at the same time there's been a pretty stable conception of what the major problems are in AI. One of those is extending knowledge representation to be more expressive than it is now, to be able to represent all the kinds of things people are capable of comprehending. We aren't there yet. We don't fully understand what is the right way to represent every kind of knowledge. We have yet to fully understand what the right language of the mind is and what set of specialized representations for space, time, taxonomies and so on may be needed.

There's also concept formation – how we arrive at concepts that are particularly appropriate for describing a thing, or for formulating and solving problems. This area was in limbo for about a decade. From the late '60s to the late '70s there was a feeling that we didn't know enough about knowledge representation to formulate learning or concept formation algorithms. That's changed now. There's been a great deal of work and considerable progress. But there's a long way to go before machines can form concepts as easily as people and divide up a problem domain the right way.

Canadian AI: What is your research style? Are you a loner or a team player?

Schubert: I guess I'm somewhere between those two extremes. It's important for me to have people around who are knowledgeable in those areas that impinge on my work, like linguistics, logic, philosophy, psychology.

One of the things that was very productive for me was the study group we ran at the University of Alberta – the Logical Grammar Study Group. It had people from linguistics, philosophy, psychology, and computer science. We worked through the recent literature on developments directly relevant to natural language understanding. I learned a great deal from my colleagues in different disciplines.

The difficulty with AI is that it is so interdisciplinary. It draws on so many different disciplines that no one person can really master all the relevant material. It's very important to have access to people who have mastery of these areas.

I've become a little of a jack of all trades and master of none. I'm not as knowledgeable about logic, formal semantics, or psycholinguistics as people who specialize in those areas. But one has to have at least a smattering of all these things.

I think it's important to maintain a broad perspective. Often the best ideas come into your research domain from outside. Some development in computer architecture or expert systems or wherever may generate ideas which can be explored in your domain. They may have a strong bearing on or even provide a key to solving problems which seem insurmountable.

Canadian AI: Has that ever happened in your work?

Schubert: Yes. For a long time after getting into AI I was not all that interested in working on language, even though I was very interested in ultimately building this generally intelligent conversing machine. I thought language was such a complex and in many ways arbitrary phenomenon that it would be best to postpone working on it until I had a much better understanding of knowledge representation in the abstract, inference, learning, things of that nature; I thought it would be best to work in those areas before tackling language, which seemed to me to be a mess.

That view, for me at least, was overturned by developments in linguistic theory which gradually infiltrated AI. Originally it was the work of the generative linguists, Chomsky in particular, which was influential to some extent in natural language processing. But their emphasis was mostly on syntax, and AI people were interested in understanding, not just in structural analysis of language. So that influence was somewhat limited.

Then came the work of Richard Montague and follow-on work by people like Gerald Gazdar. Their work was not really different in kind from what AI people had done all along, the difference was really only in degree and theoretical underpinning. The Montague- or Gazdar-based approach to semantic representation was a lot more elegant, simpler, direct and theoretically defensible as compared to the more ad hoc way AI people were doing this.

It's always possible to build a small system, if you troubleshoot it long enough and patch it long enough; eventually it will do something you can demonstrate and perhaps will have some practical value. But the problem is extensibility. If it's not built in a principled way, if it's not theoretically transparent, then you won't be able to enlarge it. You won't be able to deal with more complicated, more comprehensive domains. You won't get to a machine that has the same breadth of understanding, the same versatility as human beings.

The work of Montague and Gazdar showed me that language was not a large collection of arbitrary rules but a very subtle, elegant and beautiful structure. It changed my mind about the feasibility of doing research in computational linguistics in a theoretically well-founded way.

Canadian AI: You mentioned that you have a small robotics project at the University of Alberta. Is robotics a side interest for you or is it related to your natural language research?

Schubert: It's a little of both. I've always been interested in robots. Ultimately robotics is one of the ways one envisions AI being helpful to humanity. So that's always fascinated me – the idea of a machine that not only thinks in a human-like way but moves in a human-like fashion.

But apart from this interest in robots and their potential application, there is a connection between natural language processing and robotics. It's through the notion of intelligent planning.

Natural language is a planned activity. We say what we say to accomplish some purpose. If we're going to have human-like machines producing language, we'll need a planning system that governs the program running on that machine, one that will generate the actions that promote the goals of the speaker. In essence this is speech planning.

Now it seems to me that the kind of planning system required for driving natural language is of the same kind that's required to drive a robot. A robot, to live up to its name, needs to be goal-directed, needs to accomplish a particular task by devising its own method of accomplishing it. That's exactly what we do with language production – decide what we want to accomplish through what we say and then generate a plan that realizes those

conversational goals. I view robotics as a test bed for planning systems.

Canadian AI: When did you become interested in AI?

Schubert: I discovered AI while I was a doctoral student in aerospace studies at the University of Toronto. I was a heavy computer user at the time and started thinking that surely there must be more exciting things to do with the computer than crunching numbers and solving differential equations.

I was delighted to discover there was a new literature. AI caught my interest and I've been totally hooked on the subject ever since.

After graduating from U of T, I concentrated on making the switch to AI through postdoctoral studies. Johns Hopkins had a recently inaugurated computer science program, not really a department. James Slagle was there; giving lectures in AI; that was the big attraction for me. His AI texbook wasn't available in bookstores yet, so I got the information straight from the horse's mouth.

Canadian AI: Are there any differences in the nature of AI research in Canada and the U.S.?

Schubert: I don't really think so. I think research in every discipline is highly international because of the quality of communication worldwide. Ideas are exchanged pretty quickly, so we're all part of an international effort.

I don't see that much difference either in the nature of the research or the people themselves. The same ideas are being pursued in both countries; people are shuttling back and forth.

Canadian AI: Is there more pressure in the U.S. to work on things that have an applied orientation?

Schubert: From my restricted perspective, I don't think it's any different here than in Canada. There's been a general trend, and in some respects an unfortunate one, towards mission-oriented or applied research. But that's been happening in both Canada and the U.S. I'm not purely theoretically motivated by any means. I'm very interested in the real life import of what I do, what it ultimately is going to mean. But I don't want to justify everything I do in terms of two years hence. To have to predict "I will do this by such and such a date" or "This problem will be solved by then and it will have these particular benefits for the economy"— this isn't my orientation.

Canadian AI: What are the differences between doing AI research at a Canadian university and an American one?

Schubert: My limited time here at Rochester hardly makes me an expert on this subject, but there are some obvious differences. The biggest is the funding situation.

For Canadian academic research, virtually all funds come from NSERC and are quite limited, although they are fairly easily obtained. If you're a respectable researcher you'll get some money from NSERC. The amounts, while they do vary, are modest in comparison to funds available in the United States.

If I look over the lists of grants that people hold in this department, I would say their grants are several times larger — a factor of three would be about right — than what a comparable person in Canada would get.

Mind you, the overheads charged by the universities here tend to be quite high, about 60 percent. So the researcher is left with 40 percent. But even then I think you're several times ahead in the U.S.

For large, expensive projects, one tends to be better off in the U.S. The money comes in much larger chunks than it does in Canada and from a variety of places.

That's a general difference. A specific difference is in private and public funding of universities. The University of Rochester is private and so is able to be somewhat elitist, for instance in restricting itself to a graduate program in AI.

As a result teaching loads are lighter. Here at Rochester, the faculty teach two graduate courses. This compares with typically three, including at least one undergrad course, at the University of Alberta. That makes a big difference.

Canadian AI: Do these differences put Canadians at a disadvantage?

Schubert: Yes, although I guess CIAR has helped to some extent in providing funding for course relief. The disadvantage depends on what one's research is. If one is more or less purely theoretically oriented and wants to spend a minimum of time writing research proposals, then one may be better off in some places in Canada.

Proposals in Canada, although aimed at only modest funding, are also modest in size compared to American research proposals, which are like little books one has to write.

But for someone like me with a heavy computing demand, or if you want to build a robot, the funding situation in Canada is a definite disadvantage.

And it's not just equipment. What is perhaps even more important is highly expert technical help and secretarial help.

The most valuable resource to a researcher is time. So what makes a researcher productive is optimal use of his or her time, which means not being distracted with mundane paper shuffling or routine programming.

AI is a field where you can only demonstrate the validity of your ideas with very large programs that take a lot of person-hours to write. Expert technical help, someone with a master's degree in computer science, is needed to help with the development; someone who can take the abstract algorithms and translate them into working programs. This is absolutely crucial and is very difficult to get in Canada.

What happens in Canada is that a lot of tasks fall to graduate students. But there are severe limits to the kind of development you can entrust to graduate students. For one thing, master's students aren't around for that long. If you have an ongoing development, it takes a long time to train a master's student to understand it to the point where they can start to add to it or modify it. That in itself takes a year. By the time they get around to actually programming, it's time to graduate.

What tends to happen is that these students build some subset in isolation from the general project. The bits and pieces are often left lying around and aren't integrated.

On the other hand, doctoral students are around for a longer time, but they are supposed to be doing advanced research with a strong theoretical component. They too cannot be put to work on extensive, essentially routine, programming. So who will do this work?

You need full-time help to give continuity to a project, someone who has time to become familiar with a large, complex system. That's the main obstacle in Canada.

In the U.S., there are more department personnel simply because the departments have larger budgets. There are often highly quailified analysts and technicians funded by large continuing grants.

Canadian AI: You spent 16 years in a large computer science department at the University of Alberta and now you're in a much smaller department at the University of Rochester. Do you prefer the environment you're in now?

Schubert: I hadn't given much thought to whether I function best at a small university or a larger one. However it's now apparent to me that being in a small department, particularly a research-oriented one like Rochester, it's much easier to concentrate on the research work than it is in a large department with a large undergraduate teaching load.

I like the informality of a small department. Instead of having major meetings with minutes, you can have a get-together over pizza and make decisions that way.

Canadian AI: Are you an example of the brain drain? Do you think you'll stay in Rochester?

Schubert: I don't know. I have a kind of aversion to thinking about my own future. I much prefer thinking about AI; I always find things move along without my paying much attention.

As for the brain drain, I may just be one of those molecules diffusing back and forth across the border. This is happening all the time. I don't think I'm a special case.

A Knowledge-Based Approach to Discrete Systems Modeling and Simulation

Lubomír Masár

Une approche de la modélisation et de la simulation des systèmes discrets basée sur les connaissances

RÉSUMÉ: Le but de cet article est de suggérer une nouvelle manière d'aborder la modélisation et la simulation des systèmes discrets. Utilisant une similarité entre l'évolution d'un réseau de Petri et l'activité déductive d'un système (expert) à base de connaissances, nous considérons la simulation comme un traitement des connaissances (sur le plan des connaissances) et la modélisation comme un problème de représentation des connaissances (sur le plan symbolique). L'approche proposée est fondée sur le concept de Newell des connaissances, et elle offre un paradigme de simulation nouveau et original: la simulation est, en fait, un processus déductif non-monotonique d'un système (approprié) à base de connaissances dont la base de connaissances contient des connaissances "spécifiques au domaine d'application"; un tel processus est commandé par des méta-connaissances générales et spécifiques au domaine

1. Introduction

Determinatio est negatio
Spinoza

The evolution of Artificial Intelligence (AI) and its (more or less) impressive results have influenced many fields of "classical" computer science; many principles sprouted in AI have often stigmatized further development of those fields.

The area of discrete systems modeling and simulation is now characterized by a plethora of simulation languages, tailored to specific applications. Nevertheless a discrete dynamic system can be characterized by means of several fundamental concepts: states of the system change as a result of events; events may be structured into processes which are suitable abstractions of parallelism in dynamic systems; parallel processes are implemented through the mediation of a quasiparallel system running on a von Neumann computer – further conceptual primitives can be found more or less automatically in simulation practice.

In fact, simulation languages serve as tools for the representation of a system's model using a given formalism: thus, we can regard them as tools for knowledge representation. The purpose of simulation study *inter alia* consists of the acquisition of new information about the system under consideration. By the succession of events and/or simulation steps we are therefore collecting further information about the system. In substance, we may see the simulation as a process leading to the acquisition of new knowledge from that represented in the model of a simulated system. Thus, simulation itself may be thought of as an inference process over suitable knowledge – *mutatis mutandis* in the same way this process is conceived in AI. This idea is the central theme of our paper.

Mr. Masár graduated from the Technical University of Brno (Moravia) in 1983 with a M.Sc. in Computer Science for research in Petri net analysis and NP-completeness theory. He is currently writing his Ph.D. thesis on knowledge based techniques in discrete simulation and the relationship between modeling and knowledge representation. Mr. Masár's other research interests include logic programming and Petri net applications in non-monotonic reasoning.

2. Inference and simulation: how do they relate?

The idea loosely verbalized above – namely, the close relation between simulation and inference – forms a basis for the outline of the principal analogy which will steer us towards a knowledge-oriented approach to discrete sim-

ulation. The proposed approach is based on some ideas of Jávor and Ören (Jávor 1986, Ören 1986) and tries to apply Newell's concept of the knowledge level to the simulation methodology for discrete systems (cf. Newell 1982, Levesque 1984). It is presumed the reader is to some extent familiar with the Petri net formalism and we will not discuss Petri net theory (see Peterson 1981 for more details). Let us remember, however, the following points:

- A Petri net is a formal expression of the control structure of our parallel discrete system.
- A transition firing corresponds to the occurrence of an event in the system being considered.
- Events are timeless (i.e., their duration takes no time).

Furthermore, let us suppose a knowledge-based system (KBS) whose knowledge base (KB) contains knowledge encoded using the formalism of production systems: the knowledge base is built from chunks of knowledge in the form *if* <situation> *then* <action> (production rules) and from facts. Operations on such a KBS are performed in inference cycles (Laurent 1984) and their results cause the modification of the knowledge base. The table below roughly summarizes how (some) basic concepts in Petri net correspond to knowledge-based system formalisms.

Table 1.

Petri net vs. knowledge-based system

PN	KBS
system's states	contents of KB
state transition	inference in KB + its modification
conflict transition	conflict rules
transition choice	conflict resolution
token removing	knowledge retracting
token adding	knowledge asserting

The following definition proposes a formal expression of the analogy between PN and KBS.

Definition 1. Let $\pi = \langle P, T, R \rangle$ denote the ordinary Petri net with initial marking μ_0. Let $\mathbf{P} = \{Pr_i\}$, i=1,2 ... card (*T*), be the set of production rules of the knowledge-based system **KBS** and **D** be the base of facts in that system. Let us designate our knowledge-based system **KBS** as a pair **KBS** = (**P**,**D**). Production rules are defined as follows:

Pr_i = **if** C_i **then** A_i, where

$C_i = p_{k1} \& p_{k2} \& \ldots \& p_{kn}$,
n=card ($.t_i$), $p_{kj} \in .t_i$, j=1, 2, ... n, and

A_i = assert (p_{l1}) & assert (p_{l2}) & ... & assert (p_{lm}) &
retract (p_{k1}) & retract (p_{k2}) & ... & retract (p_{kn}),
m=card ($.t_i$), $p_{lr} \in t_i.$, r=1,2,...,m.

When the initial marking is $\mu_0 = (z_1, z_2, \ldots, z_s)$, s=card (*P*), then the factual base **D** will be the bag $\mathbf{D} = \{p_1, \ldots, p_1, p_2, \ldots, p_2, \ldots, p_s, \ldots, p_s\}$, where the number of occurences of p_i in **D** equals z_i.

We identify the knowledge-based system **KBS** as the ***production model*** (PM) of Petri net π.

The theorem below elucidates (perhaps obviously) the character of the relation between a Petri net and a production model.

Theorem 1. (*Functional equivalence of* **PN** *and* **KBS**). Let π and **KBS** be structures from Definition 1. Then equivalence below holds:

$$\forall t_i \in T : t_i \text{ can fire} \Leftrightarrow \mathbf{D} \vdash C_i$$

Proof.

I. $(t_i \Rightarrow \mathbf{D} \vdash C_i)$

If t_i is able to fire then

$p_j \in .t_i \; \mu(p_j) \geq 1$, $j=k_1, k_2 .., k_n$

Thus, a bag B corresponding to C_i is

$$B = \{p_{k1},\ldots,p_{k1}, p_{k2},\ldots,p_{k2}, \ \ldots \ , p_{kn},\ldots,p_{kn}\} \subseteq \mathbf{D}$$
$$\Downarrow \qquad \Downarrow \qquad\qquad \Downarrow$$
$$z_{k1}\text{times} \quad z_{k2}\text{times} \quad \ldots \quad z_{kn}\text{times}$$

ie. $\mathbf{D} \vdash C_i = p_{k1} \& p_{k2} \& \ldots \& p_{kn}$

q.e.d

II. $(\mathbf{D} \vdash C_i \Rightarrow t_{ii})$

If $\mathbf{D} \vdash C_i$ then at least $\{p_{k1}, p_{k2}, p_{k3}, \ldots, p_{kn}\} \subset \mathbf{D}$ and according to our definition must be
$z_{kl} > 0$, $l = 1,2,\ldots,n$,

ie. $\mu(p_j) \geq 1$ for all $j=k_1,k_2,\ldots,k_n$ and t_i can fire.

q.e.d.

Some remarks should be made at this point. It should be emphasized that a Petri net is not the model of a considered system in any way. It is only a model of the system's control structure (e.g., the control structure of a simulation program). The same holds for the production model: it serves more or less as a shell. (We are using production systems as a tool for endowing a knowledge-based system with Petri net primitives, even when this knowledge representation scheme does not present the only possibility.) Therefore, both Petri net and production models are highly abstract tools having sufficient power to describe a dynamic system at the selected level.

When considering the event realization as an inferential step, we recognize the apparent nonmonotonicity of such inference. Each event retracts a bit of (factual) knowledge from the KB and asserts some new ones.

Thus, knowledge in the KB perpetually changes in the course of inferences. The nonmonotonicity we are refer-

ring to is *inherent* in the Petri net as well as it is massive in comparison with other kinds of nonmonotonic reasoning (McDermott 1980, Etherington 1986).

Let us illustrate the use of the production model in a simple example of a system with parallel processes (according to Peterson 1981) whose control structure is depicted by means of the Petri net at Figure 1.

Production rules **P** are as follows:

if (Pr_1)
$p_1 \& p_4$
then
assert (p_2) & retract (p_1) & retract (p_4)

if (Pr_2)
$p_2 \& p_5$
then
assert (p_3) & retract (p_2) & retract (p_5)

if (Pr_3)
p_3
then
assert (p_1) & assert (p_4) & assert (p_5) & retract (p_3)

if (Pr_4)
$p_6 \& p_5$
then
assert (P_7) & retract (P_5) & retract (P_6)

if (Pr_5)
$p_4 \& p_7$
then
assert (p_8) & retract (p_4) & retract (p_7)

if (Pr_6)
p_8
then
assert (p_4) & assert (p_5) & assert (p_6) & retract (p_8)

Factual database corresponding μ_0 = (1,0,0,1,1,1,0,0) will be

$\{p_1, p_4, p_5, p_6\}$

Comparing the Petri net evolution with the application of production rules reveals that both events and inferences run indeterministically; i.e., in the production model the conflict resolution has to be solved. (Our production model is not commutative [Nilsson 1980]. Moreover, the authors think this issue should be studied separately in "production-model-like" production systems.) The following section deals with the problem of how to define and use metaknowledge mainly with respect to the inference control in production models.

3. The role of metaknowledge in production models

According to Definition 1, events in Petri nets are mapped onto inferences of the production model. Knowledge bases **P** and **D** are semantically non-oriented, i.e., although domain-specific knowledge represented in production rules *is* related, such relations are not explicitly expressed in the knowledge base (cf. Aikins 1983 and Nilsson 1980 for more detailed discussion of this issue).

When considering an event occurrence as a production rule application, we may see the knowledge about such rule as the representation of a strategy for using object-level knowledge (holding in **P** and **D**). We feel the need to specify the forms of metaknowledge about events/inferences in our production model (even though we do not think our classification scheme will be sufficient or the best one):

- *control knowledge* expresses the strategy for application of production rules (e.g., metaknowledge about the priority of an event within the framework of "its" process with respect to other processes)
- *structural knowledge* concerning the event's membership in "its" process (we mean explicit grouping of production rules which results in the notion of process)
- *time knowledge*: introducing non-primitive events into the production model, we get a novel approach to model time understanding – duration of an event is, in fact, additional knowledge about the event itself and, consequently, may be regarded as metaknowledge (cf. classical simulators "unconsciously" use just this metaknowledge in the strategy "from two or more possible events take the one which will occur first").

Let us dwell for a moment on the time knowledge. Its introduction gives some interesting consequences: in agreement with the general application of metaknowledge to inference control, the model time is really used to determine which event/inference to be next realized/applied in a consistent manner. Hence we may conceive this choice (i.e., model time value) as a matter of preference and, consequently, we need not distinguish the pri-

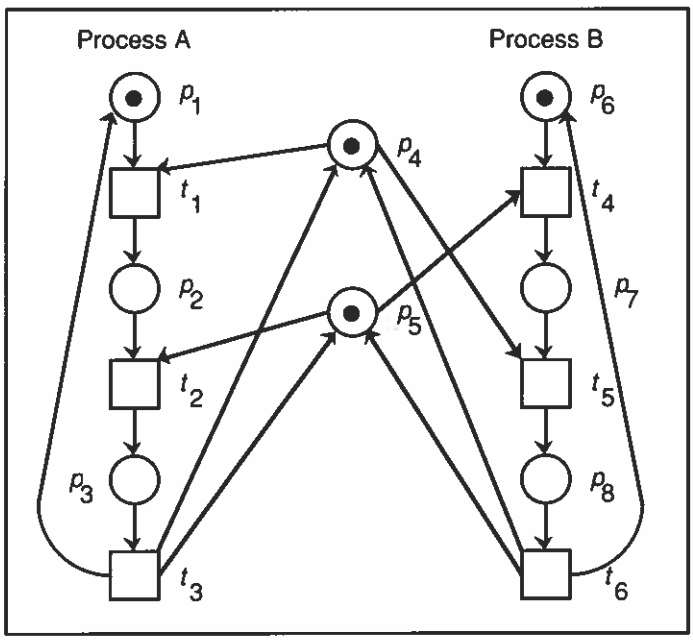

Figure 1

ority of an event from its occurrence moment. Both the priority and occurrence moment act equivalently as metaknowledge, both play functionally the same role (see section 4 of our paper).

In light of the above kinds of metaknowledge, we can mention once more the analogy following to some extent from Theorem 1 and from the proposed conception of discrete simulation. Interpretation of a Petri net is, in effect, the articulation of (various sorts of) metaknowledge in a suitable notation. Classically conceived simulation theory to some extent conceals this reality from us while the herein proposed approach distinctly demonstrates it. Petri nets *per se* only map conditions explicitly to events; production systems map situations explicitly to actions. Through an interpretation of a Petri net we supply additional knowledge about events as such, i.e., metaknowledge.

We will illustrate the use of the concept of metaknowledge in our simple example (Figure 1). In addition to points (1) through (3) we will assume that each event in Figure 1 is given some time $\tau_i \neq 0$, $\tau_i \neq \tau_j$ for $i \neq j$. The following metarule resolves the conflict set CS = $\{Pr_2, Pr_4\}$ (which now is not a true conflict, thanks to partial interpretation)

if (MR1)
CS ≠ 0 & (an ordering < exists in CS with respect to the occurrence moment)
then
(apply the $Pr \in$ CS which is maximal wrt. <)

Let us assign a higher priority to the process A than B. We can resolve conflict set CS =$\{Pr_1, Pr_4\}$ (which *does not* represent a true conflict in PN) simply by using the following metarule

if (MR2)
CS ≠ 0 & (an ordering << exists in CS with respect to the priority of realization)
then
(apply that $Pr \in$ CS which is maximal wrt. <<)

We can go further and append other metarules, e.g. ,

if (MR3)
CS ≠ 0 & (an ordering {| exists in CS with respect to the probability of realization)
then
(apply that $Pr \in$ CS which is maximal wrt. {|)

Lastly we should append a metarule (more exactly "metametarule"), claiming:

if (MMR)
(MR1 through MR3 are inapplicable)
then
(randomly choose a Pr from CS)

As for the precedence of the outlined metarules, there is the apparent order from MR1 to MMR; this order represents growing semantic strength of expressed metaknowledge with respect to the competence for inference control.

From our simple example we may now look at the problem of how to synchronize parallel processes (in a parallel system describable by a Petri net) as the problem of how to formulate and use various kinds of metaknowledge concerning the character of events within dynamic discrete systems using parallelism.

We feel that the metaknowledge representation embodied in MR1, MR2, MR3, and MMR is only an initial look at the complex task of how to control simulation by means of metaknowledge. Our aim has been to show how one might study simulation control (e.g., a calendar program and its activities) performed on a von-Neumann computer. The outlined method of control is aimed at highly modular and modifiable systems for the control of simulation experiments. Moreover there is the possibility of deriving metarules through inductive learning and building a metaknowledge base.

4. Simulation — an activity at the knowledge level

The above thoughts and the "functional" equivalence of PN and a production model show that the basic notion in our approach is that of knowledge. Let us briefly discuss some of the underlying principles most likely governing the knowledge-based techniques and their discrete simulation analogues: we abstract formal, syntactic aspects of the knowledge as such and try to look at the whole matter from the point of the knowledge level hypothesis presented by Newell (Newell 1982).

According to Newell's suggestion, we may consider a discrete system as an agent (cf. Newell 1982 and Levesque 1984). Goals of the agent are events in our model. An agent's knowledge is represented in its data structures and the agent's behaviour is ruled by the principle of rationality: actions are selected to attain the agent's goals. Hence, the knowledge and events are intimately linked by the principle: if the agent has knowledge that one of its actions will lead to one of his goals (events), then the agent will select that action (see the permissive nature of Petri net evolution).

The behaviour of a model (all the time we are dealing with models running on a serial computer) is, therefore, determined by a global principle as opposed to the common view that behaviour of the model is determined bottom-up through local processing of a model's subsystems. Thus the event realizations (inferences) run at the knowledge level.

We can go a step further: a Petri net is the description of the model at the knowledge level. Consequently, the model building (*viz.*, modeling as such) is a symbol-level activity running at the symbol level: it consists of inventing appropriate knowledge representations and encoding knowledge into them.

The ideas presented here perhaps shed light on the issue of determinism. Indeterminism is sometimes the desirable feature of a model, and in this case its preservation is required whenever formalism is used for model building. Indeterminism is he corollary of incompleteness in our knowledge about the system being modeled and/or

COMDALE/X : The Intelligent Expert System Development Tool

Description

COMDALE/X is an expert system development tool with exceptional debugging and explanation capabilities. COMDALE/X is flexible and provides a highly effective, efficient, transparent and complete environment for the development and delivery of expert systems.

COMDALE/X is written in C and consists of the COMDALE/X Rule Compiler, the COMDALE/X Application Program and a number of utility programs. It provides you with powerful and efficient methods for organizing, representing, debugging and applying knowledge of any complexity, while maintaining its ease of use and user-friendliness. COMDALE/X runs on IBM PC compatible microcomputers under the DOS and QNX operating systems.

Features

- Fast development - there is no need to learn new and difficult programming languages or to deal with typical computer science-like programming found in other products
- The software is robust, designed specifically for industrial and commercial users
- Menu-driven interface
- Customization allows increased user friendliness
- Report generation
- Can call programs written in languages such as C, Fortran, Pascal and Lisp
- Ability to access data from external sources such as databases, spreadsheets, etc.
- Can be imbedded in your applications
- Extensive library of object code which can be used for integration by the system developer
- Object definition and manipulation with inheritance capabilities
- Define demons
- Knowledge base compiled for speed of execution
- Facilities for checking knowledge base syntax and organization
- Debugging facilities include a separate Debug Menu with features for rule tracing, setting breakpoints and saving consultation sessions, etc.
- Features to support knowledge base maintenance
- Extensive functionality allows for flexibility in inference control, knowledge representation and organization
- Forward chaining, backward chaining, depth-first, breadth-first and integrated combinations of these search strategies are supported
- Dynamic examination of inference procedure
- Query facilities allow you to request justification of recommendations and conclusions made
- You may ask questions about static and dynamic knowledge
- Explanations may be graphical images, text or a combination of both
- Ability to deal with uncertainty in knowledge and data
- Security features may be used to protect your compiled applications
- Complete documentation - Manual describes features, provides comprehensive examples and gives directions for effectively using COMDALE/X
- Knowledge base is 100% compatible with the COMDALE/C real-time expert system development tool

Simple Rule Syntax

Knowledge base files are easy to create and are self-documenting.

```
Rule 237
if ANY tank level is high
and pulp flowrate > = 600
or discharge pressure is_almost critical
then TEXT
You should check your discharge lines
for blockage. #
cf = 100
then ACTIVATE pumpdiag ; cf = 100
then APPLYRULE * * critical ;
cf = 100 #


Rule 71
if high_interest investment is required
and amount to_be_invested < 100000
and term chosen < = medium period
and deposit option is redeemable
then TELL term deposit
recommended ; cf = 90
then interest earned = amount
to_be_invested * term chosen * term
interest
cf = 100
then TELL interest earned # ;
cf = 100
then REPORT invest.doc ;
cf = 100 #
```

Comdale Technologies Inc.
833 The Queensway
Suite 202
Toronto, Ontario
Canada, M8Z 5Z1
Tel : 416-252-2424

Comdale®
An Expression of Knowledge

Areas of Application

- Equipment diagnosis
- Troubleshooting
- Maintenance
- Scheduling
- Design and configuration
- CAD/CAM
- Policy documentation
- Training
- Decision support
- Legal advisor
- Financial analysis
- Planning and forecasting
- Process control
- Others

the incompleteness of the knowledge level description itself (there need not exist any state of knowledge that would determine an action — owing to the system's inherent indeterminism). Sometimes the behaviour of discrete systems can be predicted by the knowledge level description; often it cannot. Nevertheless, at the symbol level "entire ranges of behaviour can be describable" (Newell 1982).

5. Conclusions

I did not try to coin "the better and more involving approach". I only aimed at dealing with the dynamic discrete system from the standpoint of artificial intelligence paradigms and sought to outline some useful principles sprouting at the boundary between AI and general system theory. If the reader finds some fragments of this thought interesting, maybe useful or even stating another paradigm, this paper has not been written in vain.

6. **Acknowledgement**

I am very grateful for substantial comments on an earlier draft provided by Milan Ceska. His opinions and rich experience set a mirror to my ideas and often kept me from fallacy.

7. References

[1] Aikins J., Prototypical Knowledge for Expert Systems, *Artificial Intelligence* 20 (1983), pp 163-210

[2] Etherington D., Formalizing Nonmonotonic Reasoning Systems *Artificial Intelligence* 331 (1987), pp 41-85

[3] Jávor A., Discrete Event Simulation In: *Proc. Europ. Congress on Simulation*, Prague 1987, pp. 71 - 75

[4] Laurent J-P., Control Structures in *Expert Systems Technology and Science of Informatics* 3, 3, 1984, pp. 147 - 162

[5] Levesque H., Foundations of Functional Approach to Knowledge Representation *Artificial Intelligence* 23 (1984) pp. 155 - 212

[6] McDermott D. & Doyle J., Non-Monotonic Logic I *Artificial Intelligence* 13 (1980) pp. 41 - 72

[7] Nilsson N., Principles of Artificial Intelligence, Tioga Publishing, Palo Alto, CA, 1980, 476 pp.

[8] Newell A., The Knowledge Level *Artificial Intelligence* 18 (1982), pp. 87 - 127

[9] Ören T., Simulation methodology In: *Proc. Europ. Congress on Simulation*, Prague 1987, pp. 145 - 149

[10] Peterson J., Petri Net Theory and Modeling of Systems, *Engelwood Cliffs, Prentics Hall,* 1981

Second International Workshop on Artificial Intelligence and Statistics

Barry de Ville

Deuxième Congrès International en Intelligence Artificielle et Statistiques

RÉSUMÉ: IA/Stat est toujours une nouvelle discipline. Elle est caractérisée par des représentations statistiques de l'incertitude dans les systémes d'IA, par des méthodes statistiques de formation de noyaux et de concepts et par des systèmes de consultation expert et d'autres méthodes d'IA en statistiques.

January 1989
Fort Lauderdale, Florida

The Second International Workshop on Artificial Intelligence and Statistics took place during the first week of January in Fort Lauderdale, Florida. Many prominent researchers from around the world who work on the interface of AI and Statistics (AI/Stat) attended, including Canadians, Europeans, Japanese and Australians. The workshop was organized by two AT&T Bell Laboratories researchers – William Gale and Daryl Pregibon. The last workshop resulted in the publication of Artificial Intelligence and Statistics, Gale (ed.), Addison-Wesley, 1986.

AI/Stat is still an emerging discipline. It is characterized by statistical representations of uncertainty in AI systems, by statistical methods of clustering and concept tion, and by expert consultation systems and other AI methods in statistics. Canadian papers were presented in the areas of knowledge acquisition and machine learning (M. McLeish, M. Cecile, University of Guelph; B. de Ville, E. Suen, KnowledgeWorks, Ottawa); knowledge representation (W. Oldford, Waterloo); and uncertainty representation (E. Neufeld, University of New Brunswick).

AI/Stat appears to be more heavily supported in Europe than in North America. The bi-yearly European Compstat conference is expected to have 30% of its papers in this area. There will be another European conference on "Expert Systems and Statistics", organized by David Hand and David Spiegelhalter, at the University of Edinburgh this July. For details, write Mr. P.R. Fisk, Department of Statistics, University of Edinburgh, The King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, Scotland.

Expect to see some of the papers emerging out of the AI/Stat workshop to appear in the new *Annals of Mathematics and Artificial Intelligence.*

Barry de Ville works for KnowledgeWorks Research Systems Ltd. in Ottawa

Recent AI Activities at the University of Guelph

Mary McLeish
Amelia Fong Lochovsky
David K.Y. Chiu

RÉSUMÉ: Ce rapport présente les travaux de trois groupes de recherche en IA dans le département de Computing and Information Science à l'Université de Guelph. David Chiu parle d'un projet en traitement de formes et d'intelligence-machine. Amelia Fong parle de plusieurs projets en reconaissance d'images. Mary McLeish parle de la gestion de l'incertitute et de systèmes diagnostiques médicaux, en plus d'autres projets. Le département vient de commencer un programme M.Sc. (août '86) dont un des mandats est de poursuivre des sujets multidisplinaires et appliqués.

This report outlines the activities of three groups working in AI related areas within the Department of Computing and Information Science at the University of Guelph. David Chiu discusses a project in pattern analysis and machine intelligence. Amelia Fong Lochovsky discusses several projects in Image Processing. Mary McLeish reports on her work on uncertainty management and medical diagnostic systems as well as several other projects. The department recently began an M.Sc. program (fall '86), with a particular mandate for applied and interdisciplinary topics.

Image Processing and Computer Vision Algorithms (Lochovsky)

Currently, there is much research attention in parallel architectures for image processing applications. In particular, pipeline architectures have become "extremely relevant for machine vision applications" such as in manufacturing inspection systems [11]. Traditional algorithms often do not take advantage of the inherent parallelism of the hardware. Recently, many important new algorithms have been developed for pipeline architectures [12]. I have published various fast algorithms on pipeline architectures and have proposed a method for analyzing both time and space complexities for algorithms on such architectures. [1,3,4,5,6]. Other research in this area includes work on component labelling algorithms [2].

Image Processing and Pattern Recognition Software Environment (Lochovsky)

Pipeline image processing systems are increasingly being used in industrial vision systems, in particular in manufacturing visual inspection. There has been little work done on the study of the appropriate software environment for rapid program development for these machines. There are general purpose pipeline image processing systems available commercially and they provide a readily available environment for developing new inspection applications before the design of special hardware is undertaken. However, most of these image processing systems do not have the software programming support necessary for a research environment. They provide precanned subroutine libraries which are inadequate when new algorithms have to be developed. Hence, researchers often have to develop the algorithms using low level machine specific code, which is hard to debug and also difficult to maintain.

My work in this area involves the design of a high level language environment providing software features which reflect the image processing applications as well as the hardware capabilities. Hence the user can program and interact with the powerful hardware in a much easier fashion. Abstraction of various hardware resources such as image storage capabilities and bus connections have been studied. A proprocessor was designed to support the high level programming environment and the abstract data structures. The preprocessor also analyzes the image processing tasks expressed as expressions and generates the hardware setup and connections necessary for multipass operations through the pipeline processor. A prototype preprocessor has been implementated [7,8,9].

Applications (Lochovsky)

Image processing and pattern recognition techniques have been applied to electron microscopic images. This is joint research work with Dr. G. Harauz of the Dept. of Molecular Biology & Genetics at the University of Guelph. Techniques and software (on the IRIS graphic workstation) have been developed for automatic detection of ribosomes. Ongoing research involves the detection of ribosomes with special characteristics and the recognition of properties for classification purposes. [10]

Pattern Analysis and Machine Intelligence (Chiu)

This project concentrates on developing techniques for extracting useful information from empirical data. The study is theoretical and methodological with real life applications and uses methods from mathematics, statistics, information theory, pattern recognition and artificial in-

Dr. Mary McLeish is an Associate Professor in the Departments of Computing and Information Science and Statistics. Dr. Amelia Fong Lochovsky is currently an Associate Professor in the Department of Computing and Information Science at the University of Guelph. Dr. David Chiu is an Assistant Professor in the Department of Computing and Information Sciences at the University of Guelph.

telligence. The work touches on issues in inductive learning, inference and knowledge acquisition. We have been examining two real life problems in the research.

The first problem involves analysis/synthesis of patterns from molecular sequences. This includes identifying signiture subsequences and phylogenetic structure for classification and inferencing. rRNA and tRNA sequences collected from a computer databank have been used. Currently, we are developing algorithms based on the random graph model which incorporates structural and probabilistic variation in the analysis. Other approaches using maximum entropy and event covering have also been evaluated.

The second problem is related to the first but incorporates time as a factor. Historical earthquake data in China has been used to study the phenomenon of earthquakes for long-term prediction. Because of the complexity of the phenomenon, a general theory of earthquakes which yields a practical model of explanation is still missing. We have adopted the pattern recognition approach to evaluate the phenomenon and model it as a system which evolves in time. We have performed several feasibility studies and have gathered some preliminary results.

Theoretical Issues in Uncertainty Management (McLeish)

This work investigates a number of ideas related to probabilistic logic as formulated by Nils Nilsson. The papers [16,18] study nonmonotonic reasoning in this context. This model allows for inconsistent vectors in Nilsson's model — thus capturing the flavour of default reasoning where it might be considered true that "all birds fly" and that "Tweety is a bird", but the logical or consistent conclusion that "Tweety flies" does not necessarily hold (e.g., Tweety is a penguin). Several ways of incorporating these vectors are considered and the implications for the max entropy and projection solutions are studied. Some comments are also made concerning computational problems.

A later paper [17] develops a relationship between Dempster-Shafer theory and probablistic logic. In this relationship, the subsets of the outcome set on which belief is assessed in Dempster-Shafer theory are related to sentences in the following manner. The outcome space becomes the set of possible worlds and those worlds in which the sentence is true form the subset of the outcome space representing that sentence. Using the matrix notation of Nilsson and the associated probability vector, the vector V representing the probabilities of the sentences being true become exactly the belief functions, Bel, of Dempster-Shafer theory. Instead of the usual Dempster-Shafer theory rules for belief function combination, the entailment schemes of Nilsson are employed. One important difference is that now the probability vector P is not over the set of possible worlds but is over its power set. It is shown that meaningful results can be obtained from the maximum entropy solutions which are defined everywhere for both monotonic and non-monotonic models. The results are compared with those which would have been obtained from Nilsson's original model and with ordinary Dempster-Shafer belief function combination rules. The results in the case of the appropriate non-monotonic reasoning model were found to be more realistic than those found in the previous papers.

Another more recent work concerns the answer to a question first posed more formally in the Stanford technical report by Nilsson on probabilistic logic. Two examples are given using the projection and maximum entropy methods of computing a specific (non-interval) entailment value. In one case the answers from the two methods are the same and in the other they differ. Nilsson makes a conjecture as to what might be, at least partially, the cause for the difference in this example and poses the question of what are the general conditions for equality. The paper [24] examines this issue through a study of generalized matrix inverses and gives the family of solutions to the underconstrained system. The paper provides a definite way of testing for equality of the two approaches through a row reduction of the sentence matrix. It also provides an alternate way to find the two solutions and some further examples.

Related to the general problem of non-monotonic reasoning is work currently submitted for IJCAI 89 [15]. The work by Pearl *et al.* on a calculus of extreme probabilities is considered in some different contexts. Probabilistic logic is investigated for its ability to deal with very small or large probability values and especially in the manner in which they arise in modelling default reasoning. It is again found to be impossible to obtain defined results in probabilistic logic if only consistent vectors are used. If appropriate beliefs are extreme, results can be obtained which give the desired results for non-monotonic reasoning. In the traditional "penguin" problem, the probability that Tweety flies is still of order epsilon when Tweety is a penguin but the probability that all birds fly is within epsilon of 1. The "ad hoc" Mycin formulas (later ones) for evidence combination are also investigated for extreme values and some interesting results found using odd's ratios.

Applied Project in Medical Diagnostic Systems (McLeish)

A project was initiated in 1987 by members of the Department of Computing and Information Science and the Ontario College of Veterinary Medicine to explore the possibilities for a medical expert system. The OCV hospital has been running a medical information system for several years which collects on-line medical data such as bacteriology, clinical pathology, parasitology and radiology often electronically generated by lab equipment. Other patient information is also stored such as age, sex, breed, presenting complaint, treatment procedures, diagnosis and outcome. The intention of the project was to make use of the data to aid the diagnostic process, rather than developing a system based largely on expert opinions. Such an approach has been suggested by Drs. Patil, Szolovits and Schwarz when they wrote (*New England Journal of Medicine*, 1987) that recent medical systems have been largely based on the assumption that to have expert capability, they must somehow mimic the behaviour of experts. "Earlier work using mathematical formalisms (decision analysis, pattern matching, etc.) were largely discarded ...". They suggest it is time to link the old with the new: "now that much of the AI community has turned to causal, pathophysiologic reasoning, it has become apparent that some of the earlier discarded strategies may have important value in enhancing the performance of new programs. The authors recognize the difficulty of this approach when they state that "an extensive research

effort is required before all these techniques can be incorporated into a single program".

We have begun by considering two prototype domains. The first concerns equine colic, a complaint of horses and often requiring emergency surgery. The decision of whether or not to operate is a difficult one, based on many parameters. A diagnostic chart is now used as a guideline in which occurrence symptoms in two columns are simply added to obtain some idea of the advisability of surgery. A system based on this was implemented in Q'NIAL but found to often give equivocal results and produced poor overall performance.

Data was available on 253 cases involving 50 parameters. Outcome information of the following type was available: whether surgery was performed, whether or not a surgical lesion was actually found and the final state of the animal (i.e., lived, died or euthanized). The data included a number of variable types (real, discrete, nominal) and missing values. Discriminant analysis and logistic regression both found some of the most significant variables to be ones which were very subjectively measured (pain, distension, etc.) and little information was discovered from the clinical pathology data. A PLS1 learning algorithm (due to L. Rendell) discovered some other important variables in the pathology data, which turned out to be most predictive for the outcome information: lived/died. Difficulties were found with Quinlan's algorithm, due to missing data and mixed data types. However our data has been sent to Australia to be tried on very recent versions of their methodology. Bayesian inductive inference (P. Cheeseman) produced interesting classes in the data, which were often useful in determining special situations in which surgery was indicated but perhaps the overall prognosis was bad anyway and euthanasia the more relevant solution. The results of these studies are given in [21,22].

Another approach, which has proven quite successful, has been fully implemented in two versions, one using Q'NIAL and one using C (in conjunction with Oracle). It involves a weight of evidence combination rule due to I. Good. The continuous variables have been essentially discretized using a fuzzy approach. The clinicians have provided membership functions and appropriate a-level cutoffs have been obtained from the data which maximizes a likelihood function. The evidence combination scheme provides a probability value indicating a level at which surgery is required.

Like many of the probability schemes used, strong independence assumptions are made. In an attempt to capture some of the variable interactions, methods which considered symptom groups, rather than individual symptoms, have been developed. Several search techniques are used to find a set of independent (pair wise) symptom groups of maximal weights of evidence covering an incoming case. This is a tricky procedure involving a balance between several factors, such as small sample sizes and error measurements. Three different heuristic methods have been compared and one selected as the best candidate so far as judged by performance illustrated on roc curve graphs. See references [23,26,13] for details.

The prototype involves about 75 symptoms (with the continuous variables discretized) making the computational complexity of the problem very high. It is implemented on a Sequent parallel machine with 8 CPUs, and parallel algorithms have been developed and implemented. The machine was purchased through an NSERC equipment grant award (principal investigator M. McLeish). Other special grants have been obtained from competitions for Excellence Funds through the University of Guelph and NSERC operating grants. A number of other grants have been and are being applied for including the Networks of Centers of Excellence with John Mylopolous and PRECARN.

A fuzzy approach has now been implemented based partly on the work of the Adlassnig group in Austria. A combination of C and Oracle provides a nice way to implement a fuzzy relational database. The type of sup-min combination rules needed to combine fuzzy relations can be accomplished within Oracle itself, although there are some situations in which invoking C improves efficiency. Further discussions of the details of redesigning the original hierarchical database into a relational model and the construction of the fuzzy relational model can be found in [20]. A later paper [25] describes a blackboard architecture approach to handling the various components of the system. Such a representation seems particularly suitable for a parallel implementation. Work on the colic domain has been conducted with graduate student Matt Cecile, Dr. P. Pascoe of OVC and Alex Lopez of CIS.

A new domain is now under investigation by a graduate student, Pauline Yao, in conjunction with Dr. T. Stirtzinger, Clinical Studies. This concerns liver disease in small animals. A Dempster-Shafer approach is being tried, but with the difference that the usual probability mass function is being estimated from a study of the data. This has advantages similar to the weight of evidence methodology in that expert opinions can be used to modify the resulting belief functions. (The earlier methodology is being compared with this technique.) The liver disease domain differs considerably from the other prototype in that there is a very large outcome class and symptom group set. The size of the available data sets is too small to rely entirely on traditional statistical and machine learning methods. A linear C code implementation is in place and a parallel algorithm is being designed for this usually exponential method. Several new graduate students have now joined the project.

Other Projects (McLeish)

An earlier write-up in *Canadian Artificial Intelligence* described the rather unusual application domain investigated by Karl Langton. This concerned the Marine Mammal Laboratory in Hawaii under the direction of L. Hermann investigating dolphin communication. Karl brought a number of AI tools to this lab, especially user interface facilities. A natural language system was developed to handle the dolphin's learned vocabulary and grammar. Some of the results of this work appear in [19,14].

William Hunt has developed an interesting knowledge representation and learning scheme for metaphorical reasoning. A hierarchical thesarus-based approach is used as the primary form of representation which includes three units called gestalts, schemata (modelled as frames) and concepts (see work by G. Lakoff). Heuristic search methods are used in the generation process. Using ART and Sun Common Lisp, a program has been developed to generate the metaphorical representations. A testing methodology uses human subjects in both the knowledge acquisition of concepts

and the evaluation of the computer-generated metaphorical representations and the results have been analyzed statistically. The methodology is able to learn from the source salient features of the target and is designed for everyday phrases and not just scientific analogies. For example, "books are treasure chests" (target: book, source: treasure chest) or "the wind is a cat".

Mary McLeish is currently on sabbatical leave, visiting Dr. D. Spieglhalter at the Medical Research Council in Cambridge, England. Dr. Spiegelhalter's work on the use of Bayesian networks and directed graphs for expert systems has been reported recently in a discussion paper in JRSS, Series B50, 1988. Actual working systems have been implemented in Lisp using object-oriented programming. A project is under way using Bayesian networks to model the effects of time in dealing with adverse drug reactions. Dr. Spiegelhalter will be giving a tutorial at IJCAI'89 on the use of Bayesian methods in expert systems.

References

[1] Fong Lochovsky, A., "Algorithms and Architectures for a Class of Non-linear Hybrid Filters using Pipeline Architectures", to appear in *Computer Vision, Graphics, and Image Processing*, 1989.

[2] Fong Lochovsky, A., "Algorithms for real-time com ponent labeling of images", in *Image and Vision Computing*, Vol. 6, No. 1, Feb., 88, pp. 21-28.

[3] Dinstein, I, and Fong Lochovsky, A., "Computing Lo cal Minima and Maxima of Digital Images in Pipeline Image Processing Systems Equipped with Hardware Comparators", *Proceedings of the IEEE*, Mar. 1988, pp. 286-287.

[4] Fong Lochovsky, A., "The Effect of Window Mask Shape on Algorithm Design in Pipeline Architectures", submitted to *VISION INTERFACE '89*, sponsored by the Canadian Image Processing & Pattern Recognition Society, June 1989.

[5] Fong Lochovsky, A., "Parallel algorithms and Implementation of a class of detail-preserving filters using pipeline architectures", to appear in *Computer Vision and Shape Recognition*, edited by A. Krzyak, T. Kasvand and C. Suen, published by World Scientific Publishing Co. Pte Ltd.

[6] Fong Lochovsky, A., "Analysis and Comparison of parallel algorithms for image enhancement in computer vision applications", in *IEEE Computer Society International Computer Science Conference '88 on Artificial Intelligence* Dec. 1988, pp. 473-480.

[7] Fong Lochovsky, A., "Implementation Issues of parallel algorithms using pipeline architecture in visual inspection", *Proc. of International Assoc. of Pattern Recognition Workshop on Computer Vision*, Tokyo, Oct. 1988, pp. 307-310.

[8] Fong (Lochovsky), A., "Increasing Software produc tivity of Pipeline Image processing Architectures in Manufacture Inspection Applications", *Proc. IEEE 1988 International Conference on Systems, Man & Cybernetics*, Beijing, China, Aug. 1988, pp. 1138-1142.

[9] Fong Lochovsky, Amelia C., "High level language abstraction in image processing machines", *Proc. of Second International Symposium on image processing, computer generated images, technology and applications*, ACM-SIGGRAPH & CESTA, Nice, France, April 1986, Vol. 1, pp. 417-424.

[10] Fong Lochovsky, A., and Harauz, G., "Automatic Selection of Ribosomes in Electron Microscopic Images", to be submitted to *Ultramicroscopy*.

[11] Sanz, J.L., Dinstein, I., and Petkovic, D., "Computing Multi-colored Polygonals Masks in Pipeline Architectures and its Applications to Automated Visual Inspection", CACM, April 1987, pp. 318-329.

[12] Sanz, J.L., and Dinstein, I., "Projection- based Geo metrical Feature Extraction for Computer Vision Algorithms: Algorithms in Pipeline Architectures", in *IEEE Transaction on Pattern Analysis and Machine Intelligence*, Jan. 1987, pp. 160-168.

[13] Gburzynski, P., and McLeish, M., "An Efficient Heuristic Algorithm N-space Partitioning", submitted to the *SIAM Journal of Computing*.

[14] Langton, K., McLeish, M., Herman, L., "Delphi: An Intelligent Interface for a Dolphin Communication Laboratory", submitted to the *Second International Conference on Industrial and Eng. ap. of AI and Expert Systems*.

[15] McLeish, M., "Extreme Probabilities and Non-Monotonic Reasoning in the Context of Probabilistic Logic", submitted to IJCAI '89.

[16] McLeish, M., "Probabilistic Logic: Some Comments and Possible Use for Nonmonotonic Reasoning", in "Uncertainty in Artificial Intelligence", *Machine and Pattern Recognition*, Vol. 5, 1988, pp. 55-63.

[17] McLeish, M., "Further Work on Nonmonotonicity in the Framework of Probabilistic Logic with an Extension to Dempster-Shafer Theory", *Machine Intelligence and Pattern Recognition*. (Edited by T. Levitt and L. Kanal.)

[18] McLeish, M., "Dealing with Uncertainty in Nonmonotonic Reasoning", *Proceedings of the Fourth Annual Conference on Intelligent Systems and Machines*, April 1986, pp. 1-5.

[19] Langton, K., Tarbox, B., Herman, L., McLeish, M., "User Modeling at a Dolphin Language Laboratory", *Proc. of the Fourth Conference on Artificial Intelligence Applications*, San Diego, 1988, pp. 400-404.

[20] McLeish, M., Cecile, M., and Lopez-Suarez, A., "Database Issues for a Veterinary Medical Expert System", *Proceedings of 4th International Conference on Scientific and Statistical Database Management, Rome*, June 1988, pp. 33-48.

[21] McLeish, M., "Exploring Knowledge Acquisition Tools for a Veterinary Medical Expert System", accepted for the *First International Conference on Industrial and Engineering Applications of AI and Expert Systems*, June 1988, pp. 778-788.

[22] McLeish, M., "Comparing Knowledge Acquisition and Classical Statistical Methods in the Development of a Veterinary Medical Expert System", *Interface 88 Proceedings* (Statistics and Computing), April 1988, accepted.

[23] Cecile, M., McLeish, M., Pascoe, P., Taylor, W., "Induction and Uncertainty Management Techniques Applied to Veterinary Medical Diagnosis", *Proceedings of the AAAI Uncertainty Management Workshop, Minnesota*, August 1988, pp. 38-49.

[24] McLeish, M., "A Note on Probabilistic Logic", *Proceedings of the American Association for Artificial Intelligence*, Minnesota, August 1988, pp. 215 - 219.

[25] Cecile, M., Lopez-Suarez, A., McLeish, M., "Using Large Databases of Statistical Information in the Design of a Veterinary Medical Diagnostic System", *Proceedings of the 1988 AAAI Workshop on Databases*

in Large AI Systems", Minnesota, August 1988, pp. 36-46.
[26] McLeish, M. (et al.), "Uncertainty Management and Knowledge Acquisition Techniques from a Medical Information System", accepted for the 2nd International Workshop on AI and Statistics (W. Gale, Bell Labs), Jan. 1989, submitted for *Annals of Math and AI.*

Artificial Intelligence at Queen's

edited by Janice Glasgow

Intelligence Artificielle à Queen's

L'IA est un sujet important à Queen's University. Pour l'enseignement, Queen's offre un dimplôme spécialisé en sciences cognitives. Ce programme est géré par les départements de psychologie et d'informatique. La recherche en IA à Queen's se fait en robotique et en perception, sur les systèmes basés sur les connaissances, et en mu-

AI is an active area of teaching and research at Queen's University. With respect to teaching, Queen's offers a specialized degree in Cognitive Science. This program is administered jointly by the Departments of Psychology and Computing and Information Science. It involves breadth of knowledge about psychology and computing with depth in the areas of AI and cognitive psychology. There is also a variety of graduate AI related courses offered at the university.

Research in AI at Queen's is concentrated in the areas of robotics, perception, knowledge-based systems and music. Following is a description of some of the current AI projects at Queen's.

Robotics and Perception

Researchers:
Department of Computing and Information Science:
Roger Browse
Randy Ellis
Dorothea Blostein
Department of Psychology:
Susan Lederman
Terry Caelli
Department of Electrical Engineering:
Mohammed Bayoumi

Most existing robotic systems are only capable of operating within a completely understood layout of the workspace. Providing that each situation is a replication of a prototype environment, the robot may perform its actions effectively. Robots of the future will be expected to operate in imprecise and changing environments, and will therefore require sophisticated perceptual capabilities. Artificial intelligence projects within the Queen's Robotics and Perception Laboratories are directed at issues in computational perception and support for intelligent robots.

Terry Caelli and Roger Browse are investigating the problem of image segmentation using techniques such as controlled relaxation and spatio-chromatic feature filters derived from supervised learning conditions and separated feature maps as proposed in feature integration theory. Generalized texture discrimination shares many aspects of the problems of segmentation, and so this topic is also addressed, particularly with respect to human performance in conditions of controlled texture differences. In addition, Dorothea Blostein is investigating the use of texture information in the derivation of surface orientation. Browse and Caelli research the questions of object recognition and part identification utilizing both canonical view methods and object model approaches.

In the experimental psychology division of the Robotics and Perception Laboratories at Queen's, Susan Lederman and her collaborator Roberta Klatzky (University of California at Santa Barbara) have demonstrated that human recognition of objects by touch is fast, accurate, and flexible. Their work over the past few years constitutes one of the first attempts known to provide a systematic theoretical framework and comprehensive experimental approach aimed at modelling the intelligent exploration and recognition of objects by human touch (with and without vision). They have used a variety of experimental methodologies successfully in their work, all of them involving the analysis of human hand movements and/or other performance measures during human tactile object exploration and identification. Such techniques have permitted them to address the following kinds of questions: Are there clear associations between the way a human manually explores an object and the type of object knowledge required (shape, hardness, texture, etc.)? Since this has been shown to be the case, what is the nature of these links? Why are particular hand movements used to extract particular kinds of information about three-dimensional objects? What factors affect the order in which intelligent exploration and information processing about objects occurs, and how? What is the nature of the representations of objects in memory, and the processes (data-driven and model-driven) by which these representations are created when touch is used (with and without vision).

Another project in the lab addresses the nature of preattentive and attentive processing of stimuli presented to the skin. To continue and expand the initial work done by Lederman, a custom-designed apparatus is currently being constructed, in conjunction with Tom Moore of the Mechanical Engineering department. This equipment, known as Q'HAND for Queen's Haptic Attention Device, permits stimuli (e.g., surfaces varying in texture, material, compliance, form, etc.) to be delivered to any of several fingers on each hand under computer control. This has rarely been attempted in studying tactile information processing to date and as such offers an opportunity for exploring these issues in considerably greater depth.

Janice Glasgow is an Associate Professor in the Department of Computing and Information Science at Queen's. Her research interests include programming language semantics and AI.

There are exciting and potentially significant spinoffs of the work on human object perception and recognition by Lederman and Klatzky (described above) in industry. They are attempting to apply knowledge of the structure and function of biological perceptual systems (touch, with and without vision) to the design of flexible perceptual systems for sensate dextrous robot hands, capable of intelligent exploration, perception/recognition, grasping, and manipulation of objects. The first successful robotic application of Lederman and Klatzky's work with humans has been implemented by Sharon Stansfield under the supervision of their collaborator Ruzena Bajcsy, who heads the perceptual robotics GRASP lab at the University of Pennsylvania. Stansfield developed a tactile perceptual system capable of recognizing a considerable number of common objects, using a stereo vision system to locate the object and a Puma 560 arm equipped with a Lord LTS-200 array force sensor on the wrist to actively explore the objects.

An important aspect of this novel collaborative team is the strong emphasis on active sensing and perceptual exploration by sensate robotic systems, which can be used effectively in a variety of situations. For example, the current practice by NASA is to replace an entire instrumentation panel of the space shuttle when a single component becomes defective. This is not a cost-efficient procedure, however. An alternate solution is to have a sensate robot actively explore the panel to locate and subsequently replace the faulty part. Carrying into space only small components, as opposed to relatively massive sections such as the entire panel, would conserve valuable space aboard the shuttle. Data-driven robotic exploration can be used to determine the material composition of soil on unfamiliar planets.

Randy Ellis's work in robotics principally involves how to design and use tactile sensors to enhance the perceptual and manipulation capabilities of advanced robotic systems. His current work concentrates on three problems: the propagation and management of uncertainty in perception; the theoretical specification of tactile sensors; and issues in force and tactile feedback for tele-operation systems.

The uncertainty project is driven by the observation that sensor readings, regardless of their source, inherently contain errors. In a model-based recognition system, these errors propagate as uncertainties through the signal-processing, initial matching, data interpretation, and possibly into a stage of planning a new sequence of actions to aid in the recognition of the object. Using a computational geometry approach, he has found tight analytic bounds on the uncertainty in a tactile recognition paradigm and has used these results to produce a robotic system which reasons about how to plan the "best next path" for subsequent contacts with an object in a robot workspace.

One frustrating problem in tactile perception is that sensor design and sensor use have been proceeding along divergent paths. Ellis is investigating the possibility of specifying parametric relations between robot limb dynamics, tactile sensor characteristics, and task specifications, with the goal of being able to determine what sensor properties are needed to accomplish a given task on a given robot. Ideally, such specifications could then be used to guide the design of more advanced tactile sensors than are currently available.

One ambitious project of the lab is tele-operation. It is a collaborative effort with Browse, Bayoumi, Ellis and Lederman. The goal of the project is to provide sophisticated, multisensory feedback to a human operator who is guiding a remote robot through its motions. Ellis is investigating the design and control of a force-feedback position controller, which will permit the operator to "feel" the forces acting on the end effector of the robot as the robot imitates the operator's motions. He and Lederman are also investigating issues in the human's interaction with the system, to find convenient and powerful presentations of multisensory information and to characterize the operator's performance in common manipulation tasks.

Knowledge-Based Systems:

Researchers:
Computing and Information Science:
Janice Glasgow
Mike Jenkins
Mechanical Engineering:
Brian Surgenor

The Nial Project at Queen's University is involved in the design, implementation and application of the Nested Interactive Array Language, Nial. One of the principal application areas of Nial is artificial intelligence, where the powerful operations for dynamic manipulation of both symbolic and numeric data permit straightforward implementation of AI techniques.

Research in the area of knowledge-based systems falls into two categories: tools and applications.

AI Tools:

The Nial AI Toolkit is one product of the on-going research of the Nial Project. The philosophy of this work is to provide building blocks from which tailored knowledge-based systems can be constructed. The goal is to provide direct implementations of basic techniques that can be understood well enough to be modified by the programmer constructing a particular application.

One tool that has been developed is a logic programming environment. The goal of this research was to integrate the functional style of Nial with the declarative capabilities of logic. A variety of theorem provers have been implemented, including ones with semantic unification and heuristic search strategies. These provers vary in both flexibility and efficiency.

Semantic unification replaces the need for built-in predicates in the logic programming system. Since any Nial expression can be a term in the logic, both arrays and function applications on arrays can be represented in the language. The underlying array theory semantics of Nial provides a basis for treating arrays as logical terms in the language. The approach being used is based on an axiomatization of array theory in which arrays are constructed using two fundamental functions of the theory: reshape and hitch.

A major scientific result of research in this project was

the development of an environment in which both databases and logic programs can be represented and manipulated. One reason for this success is based on the mathematical formalism of array theory, on which Nial is based. A roster database is an array structure that can be manipulated using functions of array theory. Similarly, logic clauses can be represented as embedded arrays and mathematical functions (unify and resolve) can be expressed in array theory to define control strategies for the logic. Thus a single semantic entity can be used for both forms of representation. The result of this research is the ability to explore the advantages and disadvantages of separating factual information, such as that stored in a conventional database, from knowledge or general rules that relate these facts.

Other tools that have been developed are a frame language and forward and backward rule interpreters.

Applications:

The programming language Nial has been used at several universities and industries to develop prototype knowledge-based systems. These include systems for aircraft maintenance (Brunel Univ., England), well logging (Alberta Research Council), fuzzy sewage and system control (IBM and Technical University of Denmark), automatic software testing (IBM Toronto), and insurance policy screening (two insurance companies). As well, Queen's University is playing an active part in the Computer Aided Process Engineering Lab, a subgroup of the Manufacturing Research Corporation Centre of Excellence for Ontario. In this project we are collaborating with chemical and mechanical engineers from the Universities of Waterloo, Toronto and Queen's in developing prototype knowledge-based systems for process engineering.

The Nial Project at Queen's is currently developing two prototype expert systems under contract from industry:

a) Automated Data Management System for Mobile Servicing Station

This system is being developed in collaboration with Spectrum Engineering of Peterborough and is funded by the Strategic Technology for Automation and Robotics (STEAR) program. Major tasks of this system include fault diagnosis, sensor data handling, sensor trend analysis and reliability maintenance and reporting. Because of the need to manage both data and knowledge in this system, Nial has proven to be a useful implementation language.

b) Depreciation Knowledge-Based System

The capital recovery group in Bell Canada have described their requirement for a computer system to assist them in doing life analysis of the company's property. This project involves Bell and the Nial Project at Queen's in a research contract to explore the development of such a system using knowledge-based systems technology and the programming language Nial. This system will interact with an existing process developed by Bell Canada that automates much of the routine computations required to do life analysis. This new system is expected to make suggestions on what curves and lives best represent the expected mortality based on inputs by the user of past history and factors that will affect future lifetime trends.

c) Control and Operation of a Thermal System

It has been suggested that AI can provide the means to significantly improve the performance of process operations by providing the operator with timely expert advice on optimal control and procedures for a given mode of operation. Brian Surgenor is currently investigating the performance of a prototype knowledge-based system applied to a physical analog of a superheater steam temperature system.

Artificial Intelligence and Music:

Researchers:
Department of Computing and Information Science:
Dorothea Blostein
Eli Blevis
Peter Farrett

The work of Peter Farrett involves symbolic logic as applied to knowledge-base modeling. In particularly, he is considering monotonic and nonmonotonic models. The domain of discourse is the "composer's assistant", a tool for music composition which models the behaviour of the end user. This is an interdisciplinary field and merges advanced mathematics, musical knowledge, and computer science. It is hoped that this research will extend the capability of knowledge-base system technology in current computer music applications.

Eli Blevis is involved in the construction of a programming language that allows composers to express what they do in a manner that is meaningful to them. This path has led him to consider the construction of the language from a first semantic principle: what is a music logic and how does it differ from logics that are intended to model notions of truth or reality? The importance of this work for computer science is hopefully a demonstration of the influence of modern thought about logics to application-specific programming language design and semantics.

Successful computer vision depends on combining low-level primitive-extraction processes with high-level knowledge-intensive interpretation processes. In order to study the modeling of world knowledge, and the use of these models in image interpretation, it is helpful to focus on restricted image domains that have well defined semantics.

Dorothea Blostein is investigating methods for extracting the semantics of two-dimensional data representations such as graphs, charts, sheet music or wiring diagrams. Such diagrams are governed by strict rules and conventions. Since these conventions are well defined in an informal way, research can focus on how to formalize these notational conventions in a representation suitable for image interpretation. For concreteness, she is focusing on a particular type of diagram (sheet music), but with the goal of developing general representation and interpretation techniques which can be adapted to a variety of notations. This problem has important practical

applications. For example, a utility company which wishes to computerize its record keeping may have huge quantities of diagrams (showing locations of pipes and cables, manholes, etc.) which need to be entered into a computer system. This is a tedious and time-consuming task for humans; it is clearly preferable for the computer to be able to interpret images of the data directly.

Current Funding of Described Research

NSERC (Operating, Infrastructure and Equipment); NRC (STEAR Program); Manufacturing Research Corp. of Ontario (MRCO); Information Technology Research Centre (ITRC); U.S. Office of Naval Research (ONR); Bell Canada.

References

[1] R. Bajcsy, S. Lederman, R. Klatzky, "Machine systems for exploration and manipulation: A conceptual framework and method of evaluation". Submitted.

[2] E.B. Blevis, M.A. Jenkins, "Computing Similarity in Music Presentation" at the *First International Workshop on Artificial Intelligence and Music, Gessellschaft fur Mathematik and Datenverar- beitung*, St. Augustin, West Germany 1988.

[3] E.B. Blevis, M.A. Jenkins, J.I. Glasgow, "Motivations, Sources, and Initial Design Ideas for CALM: A Composition Analysis/generation Language for Music". *Proceedings of the First Workshop on A.I. and Music at AAAI-88* St. Paul, Minnesota 1988.

[4] E. Blevis, "A Basis for Effective Logic Programming in Nial" Master's Thesis Queen's University at Kingston 1986 *Thesis* .

[5] E.B. Blevis, M.A. Jenkins, E.P. Robinson, "On Seeger's Music Logic" INTERFACE: *Journal of New Music Research*. (To appear, Summer 1989.)

[6] D. Blostein, N. Ahuja, "A Multi-Scale Region Detector" Computer Vision, Graphics, and Image Processing Accepted June 24, 1989.

[7] D. Blostein, N. Ahuja, "Shape from Texture: Integrating Texture — element Extraction and Surface Estimation", Submittted to *Pattern Analysis and Machine Intelligence.*

[8] D. Blostein, N. Ahuja, "Representation and Three-Dimensional Interpretation of Image Texture: An Integrated Approach". *Proceedings First International Conference on Computer Vision* London, England June 1987.

[9] D. Blostein, "Recovering the Orientation of Textured Surfaces in Natural Scenes" Co-ordinated Science Laboratory, University of *Illinois Technical Report UILU-ENG-87-2219*, 1987.

[10] R. Chau, J. Glasgow, M. Jenkins, "A Framework for Knowledge Based Systems" in Nial *Proceedings of the 6th Annual IEEE Phoenix Conference on Computers and Communication* , Feb. 1987.

[11] R. Chau, "A Model for Knowledge Information Processing Systems" in *Nial Master's Thesis* Queen's University at Kingston, Jan. 1987.

[12] R. Chau, J. Glasgow, M. Jenkins, "Fuzzy Information Management Using the Roster Model". *Proceedings of 21st Hawaii International Conference on System Sciences,* January 1988.

[13] R.E. Ellis, "Uncertainty Estimates for Polyhedral Object Recognition" *Proceedings of the IEEE Conference on Robotics* 1989 (in press).

[14] R.E. Ellis, "A Tactile Sensing Strategy for Model-based Object Recognition" Computer and Information Science, University of Massachusetts Ph.D. Thesis, published as *COINS Technical Report* 87-96, 1987.

[15] R.E.Ellis E.M. Riseman, A.R. Hanson, "Tactile Recog - nition by Probing: Indentifying a Polygon on a Plan" *Proceedings of the AAAI Conference* pp. 632-637, 1986.

[16] R.E. Ellis, "An Approach to the Integration of Vision and Touch for Robot Control" *Proceedings of the IFAC Symposium on Robot Control* pp. 343-348, 1985.

[17] J. Glasgow, M. Jenkins, C. McCrosky, "User Defined Parallel Control Strategies" *Proceedings of the Second IEEE Symposium on Logic Programming* , 22-28 Boston 1985.

[18] J. Glasgow, M. Jenkins, C. McCrosky, H. Meijer, "Expressing Parallel Algorithms in Nial" *Parallel Computing* , Feb. 1988.

[19] J. Glasgow, M. Jenkins, "Array Theory, Logic and the Nial Language" *Proceedings of the International Conference on Computer Languages* Miami, Florida, October 1988.

[20] L. Hache, "NFL: A Frame Representation Language in Nial Master's thesis" Queen's University at Kingston 1986 *Nial Frames Thesis.*

[21] M. Jenkins, J. Glasgow, "A Logical Basis for Nested Array Data Structures" Queen's University at Kingston *Technical Report # 87-205* (Revised Feb. 1988) December 1987.

[22] M. Jenkins, J. Glasgow, C. McCrosky, "Programming Styles in Nial IEEE Software Engineering" 46-55 1986. Also available as *Proceedings of the 19th Hawaii International Conference on Systems Sciences.*

[23] M. Jenkins, W. Jenkins, Q'Nial AI Toolkit Nial Systems Ltd., Kingston, Ontario 1987.

[24] M. Jenkins, J. Glasgow, E. Blevis, R. Chau, E. Hache, D. Lawson, "The Nial AI Toolkit" *Proceedings of Avignon '88 Eighth International Workshop on Expert Systems and their Applications*, June 1988.

[25] R. Klatzky, S. Lederman, C. Reed, "There's more to touch than meets the eye: relative salience of object dimensions for touch with and without vision" *Journal of Experimental Psychology*: General 1987 116(4) 356-369.

[26] S. Lederman, R. Klatzky, "Hand movements: A Window into haptic object recognition"*Cognitive Psychology* 1987 19(3) 342-368.

[27] S. Lederman, R. Browse, R. Klatzky, "Haptic process ing of spatially distributed information" *Perception & Psychophysics* 1988 44(3) 222-232.

[28] S. Lederman, R. Browse, "The physiology and psychophysics of touch" (pp. 71-91). In P. Dario (Ed.), Sensors and sensory systems for advanced robots. NATO ASI Series F: *Computer and Systems Science* Springer-Verlag: Berlin 1988.

[29] S. Lederman, R. Klatzky, "Haptic exploration and object representation". In Goodale, M. (Ed.) *Vision and action*: the control of grasping. New Jersey: Ablex, in press.

International Computer Science Conference '88

Conférence Internationale en Informatique '88

RÉSUMÉ: Le but de cette conférence était de rassembler des universitaires et des industriels de l'Est et de l'Ouest pour discuter de problèmes en informatique. Cette conférence rappelait les conférences canadiennes CSCSI/SCEIO: c'était amical et beaucoup des résultats présentés dans les sessions techniques étaient le travail d'étudiants gradués ou de nouveaux Ph.D.

19 - 21 December 1988, Hong Kong

The International Computer Science Conference '88, Artificial Intelligence: Theory and Applications, was held 19-21 (22-23 for tutorials) December 1988 at the Excelsior Hotel in Causeway Bay in Hong Kong. This first international conference held in Hong Kong devoted to computer science was sponsored by Gilman Office Machines, IBM World Trade Corporation, and NCR (Hong Kong) Limited. The purpose of the conference was to bring together people from academia and industry of the East and West who are interested in problems related to computer science. This conference was reminiscent of the Canadian CSCSI/SCEIO conferences — it was casual and friendly and many of the results presented in the technical sessions were primarily the work of advanced-standing graduate students and/or newly graduated Ph.D.s. In all, some 400 papers were submitted and 97 refereed papers were selected for presentation in parallel sessions and inclusion in the proceedings. The keynote address, delivered by Dr. K. Furukawa, of Japan's Fifth Generation Computer Project (ICOT Research Centre) was entitled "Towards a Unified Computation Scheme for Knowledge Information Processing". Invited speakers included Dr. Dan Bobrow, from Xerox PARC, USA, who spoke on "Moving Beyond Object-Oriented Programming", Professor Han Adler, from UC Berkeley, USA, who spoke on "Combinational and Geometrical Approaches to Linear Programming", and Professor Frank Harary, from New Mexico State University, USA, who spoke on "Presentations of a Hypercube and their Roles".

ICSC'88 covered a three-day technical program, followed by three one-day tutorials (in two days) on Machine Learning, Machine Translation and Natural Language Processing. The technical program covered the following areas: AI Architecture, Automatic Reasoning, Distributed and Parallel Computing, Logic Programming, Natural Language, Machine Learning and Neural Networks, Programming Languages and Complexity, Pattern Recognition and Vision, Expert and Knowledge-based Systems, AI Implementation Issues, Software Engineering and Implementation Issues, and User Interfaces and the Programming Environment.

A small but excellent selection of invited speakers contributed positively to the overall quality of the conference. We present details of three of these speakers.

Keynote Address

Dr. Koichi Furukawa gave the keynote address at the conference. Dr. Furukawa is the director of the First Laboratory of ICOT—the well-known Japanese Fifth Generation Research Project. First Laboratory is responsible for pursuing fundamental research into Logic Programming Systems (as opposed to pursuing the applications of Logic Programming to Natural Language processing, databases, and so forth, which are the concerns of some of the other laboratories at ICOT). Within the Logic Programming community Furukawa's research group is well respected and its work is widely known.

Dr. Furukawa reported on the achievements of his laboratory during the first seven years of the ten-year ICOT project. One of these has been the research that his group has done in the areas of meta-programming, partial evaluation, and program transformation. Another is the development of CIL—Complex Indeterminate Language—a way of programming by specifying complex constraints. But the main achievement, according to Furukawa, of these past seven years has been the invention and development of a concurrent Logic Programming language known as FGHC (Flat Guarded Horn Clauses).

Meta-programming—which is what metacircularity allows you to do—is a prerequisite for the program transformation technique known as partial evaluation. Furukawa's group uses partial evaluation meta-interpreters to "compile" very high-level non-procedural knowledge representations into lower lever, more verbose, but faster Prolog code. The process is called partial evaluation because, in broad terms, the meta-interpreter "runs" the high-level code as far as it can, until it runs into something that is input dependent. During the talk, Furukawa showed how traditional expert system rules could be partially evaluated to produce fast Prolog code. Though techniques like this have many obvious merits—they are easy to understand and to modify—they have not been very prominent until the advent of logic programming.

Nick Cercone is Director of the Centre for Systems Science at Simon Fraser University and is currently co-editor of Computational Intelligence *and a past president of CSCSI/SCEIO.*

This is, Furukawa claimed, largely because logic programming makes building meta-interpreters very easy.

Constraint logic programming is a very active area in Logic Programming today. Furukawa's group has developed two constraint logic programming languages, one called CIL and the other called CAL. CIL, or Complex Intermediate Language, was born at ICOT from research in discourse understanding systems that are based on situation semantics. CIL turns out to be a general knowledge representation language that is particularly suited to language processing.

The idea of concurrent Prolog is first due to K.L. Clark and S. Gregory in the UK around 1981, and was later championed by E. Shapiro. ICOT invited many of these key people to visit the Tokyo laboratory—in fact Shapiro, Clark, Gregory, and K. Kahn all visited, most at around the same time in 1983—and from this interaction GHC and eventually Flat GHC were born at ICOT. These languages forego backtracking—they are sometimes called "committed choice languages." What they give a user in the place of backtracking is a clear and simple notation for specifying concurrent programs—and the promise of speed. FGHC is not intended as general purpose Logic Programming language. What it is, instead, is a vehicle for quickly implementing operating systems, utilities, and general purpose Logic Programming languages. ICOT regards it as the "assembler level" language from which the rest of the personal logic programming machine software will be built.

The ICOT group has already built a 64 processor prototype version of a parallel inference machine, and an operating system is being written for this machine in a language that is an extension of FGHC.

Our personal impression of the ICOT research group is that they are a very talented bunch. They have moved from being students of logic programming and its uses in artificial intelligence to researchers working at the leading edge of Logic programming research in a very short space of time.

Invited Speakers

The invited speech for December 20th was given by Dr. Daniel Bobrow, of Xerox, PARC. Bobrow was concerned with describing the implementation of Object-Oriented Programming in Lisp, methods for extending a Lisp implementation of Object-Oriented Programming and methods of maintaining and extending Object-Oriented programs without changing existing code.

A Lisp implementation of Object-Oriented Programming retains the Lisp property of meta-circularity. Programs are data, and data can be programs. This means that Object-Oriented Programming is not simply an add-on to Lisp, but that Lisp itself becomes Object-Oriented. Lisp primitives are implemented as objects, so that a particular string, for example, is an instance of the class String and inherits its structure and behaviour from this class.

A second consequence of metacircularity is that classes themselves may be implemented as objects. A class is then an instance of a superclass. The generic structure and behaviour of all classes is described by the superclass. This means that high-level changes and extensions to Object-Oriented programs are very easy to make.

This latter point highlights an attraction of Object-Oriented Programming. Ideas which have been implemented as objects are easily extended. In addition to high-level extensions, the property of objects that they are built up out of smaller pieces makes it possible to make other programmatic changes without changing existing code. If it is determined that a function is too general, that it is doing too much work and should instead be implemented as several specialized versions, its component methods can be split and assigned to these specialized versions by new superordinate dispatch code. The dispatch code is the only physical change made to the program. Conversely, it may be desirable to combine the behaviour of several different functions into a single function. An Object-Oriented approach permits one to combine the methods of other objects into a new object.

The extensions which have been developed for a Lisp implementation of Object-Oriented Programming are focused on developing intelligent dispatch code. One problem in introducing new techniques into a programming language is that the user must know whether a particular function is implemented traditionally or by the new technique. In Lisp, the user must know whether a function is implemented procedurally, in which case it is funcalled, or as a method of an object, in which case it is called using send. This is true only if Object-Oriented Programming is not integrated into Lisp. Integration into Lisp is done by generating dispatch code which looks like a function call. Lisp, rather than the user, determines how a particular function is implemented.

Other extensions to dispatch code have developed the ability to select the appropriate method. Traditionally, methods are selected by the object which receives a message. Corresponding to a particular piece of behaviour is one and only one method. The extensions described by Bobrow allow objects to have multimethods. The simplest extension generates dispatch code, which selects the appropriate method on the basis of both the object and accompanying arguments. An object may have associated with it several methods; the appropriate method is chosen on the basis of the number of arguments in the call to the object. This is then extended by selection methods on the basis of the identity of the arguments. Arguments are evaluated and the appropriate method is selected associating methods with possible values of the arguments.

Han Adler (UC Berkeley) presented an invited talk entitled "Algebraic, Combinatorial, and Geometric Approaches to Linear Programming". (There are no notes for this talk in the proceedings.) He surveyed the state of the art of solving large linear programs from both a practical and theoretical level. This field still exhibits a fascinating dichotomy between theory and practice. No one knows why some algorithms are so successful in practice, nor why the theoretically good algorithms perform so badly in practice.

On the theoretical side, Adler reviewed the relaxation method (Agmon, Motzkin, Schoenberg 1954), the simplex method (Fourier, Kantorovich, Koopman, and Dantzig), Katchiyan's ellipsoid method, and Karmarkar's interior point method (1984). Katchiyan's contribution was significant because it not only showed that linear programs could be solved in polynomial time, but also—and this was a revelation to the mainstream research community in this area according to Adler—because Katchiyan showed how to correctly pose the complexity problem for linear programming by using the length of the input. Karmarkar, while at AT&T's Bell Labs, introduced the interior point method that looks like it may lead to an algorithm that is not only good theoretically, but is also good in practice.

On the practical side it turns out that no one knows the

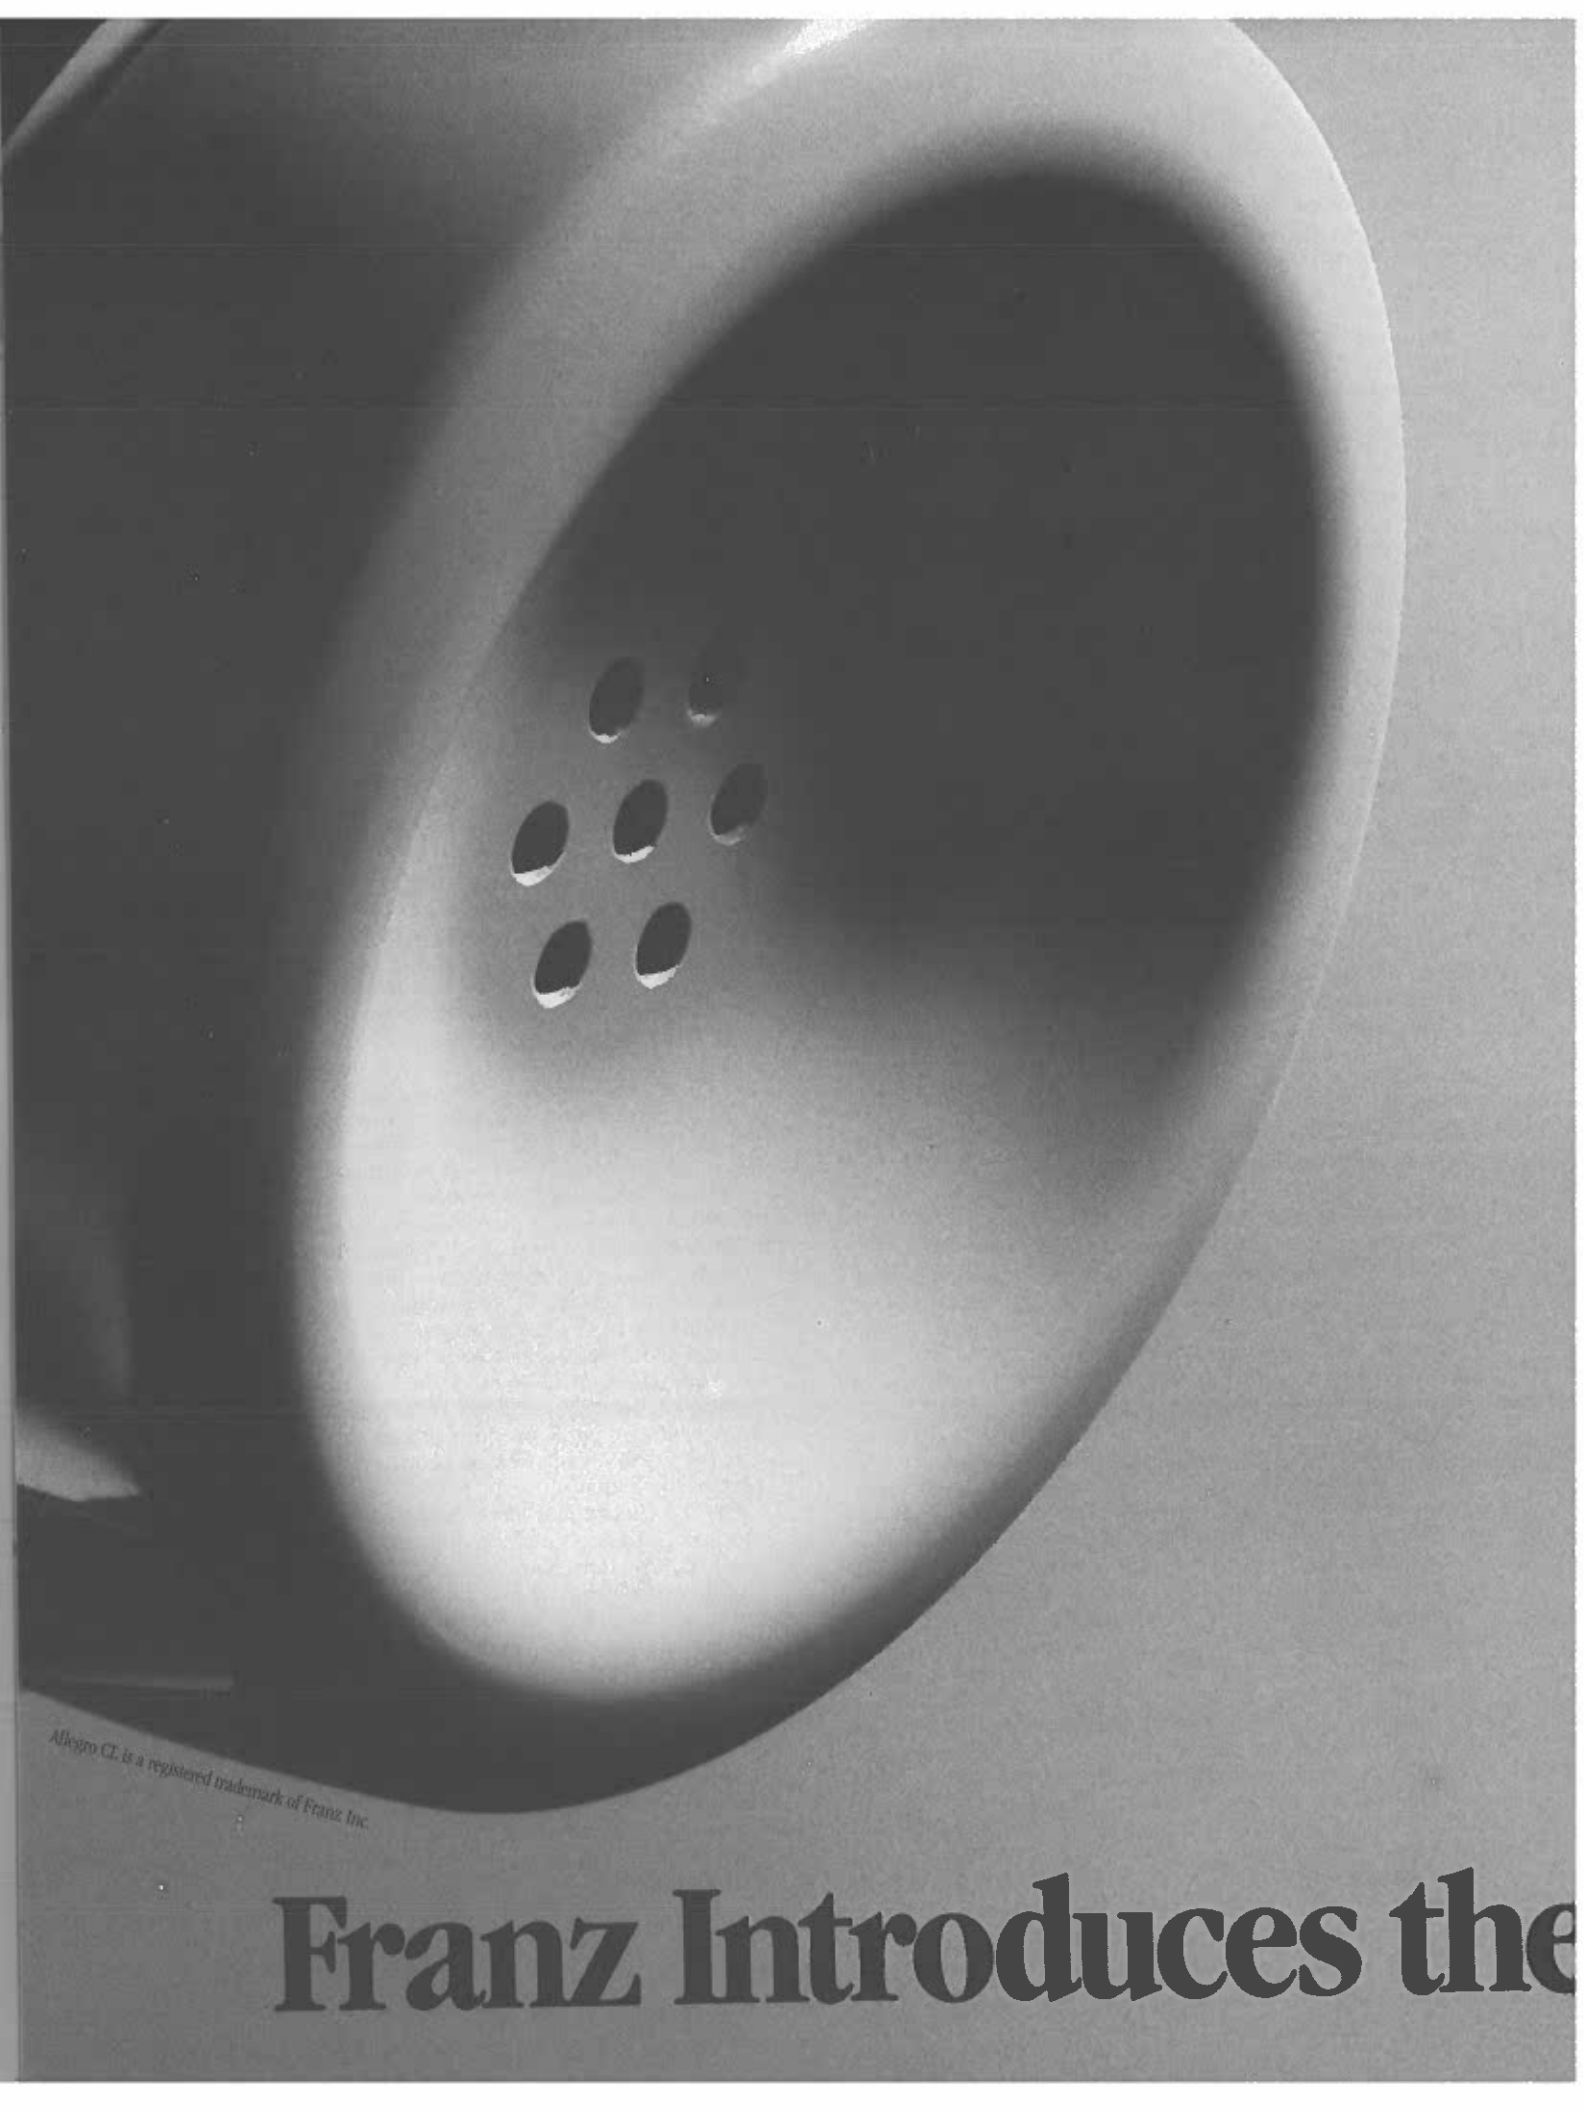
Allegro CL is a registered trademark of Franz Inc.
Franz Introduces the

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.

properties of the relaxation method. Currently derivatives of the simplex method are still the most popular methods for solving large linear programs, but why they are so successful remains largely a mystery. Katchiyan's method, on the other hand, turns out to be very unattractive as a practical method. But Karmarkar's interior point method is promising in this regard according to Adler.

In a recent visit to Berkeley, Karmarkar and Adler experimented with this method. After modifying the code so that all of the "safety checks" were dropped, they discovered that linear programs were solved in typically 30 to 40 iterations. But in theory the lack of these "safety checks" means that the worst-case running time of this algorithm may no longer be polynomial. The fact that in practice the number of iterations is roughly the same—about 30 or 40—regardless of the size of the problem, is also unexplained. Thus there is still no complete answer to the question of whether a linear programming algorithm can be both good in practice and good theoretically.

Submitted Papers

Submitted papers were presented in 24 sessions, many of them concurrently. Obviously we did not attend all presentations. Canadian participation was particularly significant in this aspect of the conference. Six of the 24 sessions were chaired by Canadian researchers (3 others were chaired by expatriates). Thirteen of the 97 submitted papers accepted were from Canadian authors (including 4 from Simon Fraser authors). We present details of several selected sessions.

Session 2a: Automatic Reasoning

This session contained several interesting papers. C.H. Tzeng's paper "A Study on the Dempster-Shafer Theory Through Shafer's Canonical Examples" is interesting because it represents another step in a line of research that shows how the Dempster-Shafer theory can be understood within the traditional Bayesian framework. Dempster and Shafer developed their theory as a generalization of the Bayesian framework, a generalization that could gracefully represent the lack of statistical knowledge. H. Kyburg, Jr.'s work in 1987 is a predecessor of Tzeng's in that it also attempted to show how to reinterpret the Dempster-Shafer theory in a traditional Bayesian framework.

Tzeng takes Shafer's message models—which Shafer proposed as a way of giving the Dempster-Shafer theory a clearer mathematical foundation—and shows how to reinterpret them within the Bayesian framework. In Tzeng's approach both the plain text messages and the coded messages are governed by prior probability distributions. Given a randomly coded message, the basic probabilities are taken to be the usual conditional probabilities. Under certain restrictions on the prior distributions, the basic probabilities reduce to Shafer's basic probabilities.

Tzeng continues by defining the notion of independent message models and shows how to combine the evidence from two independent message models using Bayes' rule. If the basic probabilities happen to reduce to Shafer's basic probabilities, then the general combination rule reduces to Dempster's well-known combination rule. Thus Shafer's models are specializations of Bayesian models—Bayesian models with certain implicit assumptions.

N.Y. Foo & A.S. Rao's Belief Revision in a Microworld takes the Alchouron, Gardenfors, and Makinson metatheory for belief revision and applies it to the blocks world. Actions are defined in terms of the basic Gardenfors operators of revision, expansion and contraction of theories. States of the blocks world are captured as canonical theories. Laws of the world are coded as the most entrenched formulas (using Gardenfors's entrenchment relation). The authors show that this approach handles qualification and that laws with pre-conditions are equivalent to laws with post-conditions in the blocks world.

Session 3a: Distributed and Parallel Computing

This session was primarily concerned with concurrency control. Simon S. Lam's paper entitled "Protocol Conversions via Projections" was a revised and shortened version of an earlier paper; it presented a formal model for accurately translating messages from one communication protocol into another. An alternative algorithm to two-phase locking and timestamp ordering was presented in the paper entitled "The Multiphase Concurrency Control Algorithm for Distributed Databases". The authors, Chihping Wang and Victor Li, claim that their algorithm, which makes guesses at a viable order for execution of pending transactions before engaging in negotiation between competing sites, outperforms the alternative algorithms when transmission delay and transaction size are small. In "Detecting Tasking Communication Deadlocks in Concurrent Ada Programs", authored by J. Cheng and K. Ushijima, a method was presented for detecting and diagnosing the causes of a variety of previously undetectable deadlocks in concurrent Ada programs at run-time, just before they occur. Theo Ungerer presented "Language Abstractions for Concurrency Control", which he co-authored with Eberhard Zehender. Three common sequential control constructs from high-level programming languages (e.g., the CASE construct) are generalized into parallel control constructs which may be used to control concurrency at all levels from the program level to the instruction level.

Session 5b: Natural Language

The session was concerned with natural language generation. Cecile Paris's paper titled "Responding to a user's domain knowledge: Incorporating the user's level of expertise into a generation system" contained her arguments against a simple bipolar view of the user as either novice or expert in favour of one which recognizes varying levels of expertise. She reported a natural language generation system which varied individual sentences in an explanation according to the user's level of knowledge about the topic of each sentence. Experts are given information about structure; novices are given information about function. Polguère presented the Iordanskaja and Polguère paper "Semantic processing for text generation", which described a representation and architecture for text generation. The novelty lay in a single conceptual representation which was transformed into many different semantic representations each of which represented either small pragmatic differences or even different representation in different languages. The paper by Gailly and Ribbens, "Hermes: A written French generator focussing on quantifier scope expression", was concerned with generating correct and intelligible quantified expressions. That by Sprenger and Gerlach, "Expectations and Propositional Attitudes - Pragmatic Issues in WISBER", described how to deal with ellipsis and indirect queries by

first interpreting the user's intentions and goals.

Session 5c: Natural Language

This session focused mainly on abductive methods for interpreting natural language. The odd man out was Daniel Radzinski, who presented his paper "Mandarin Chinese and Context-Freeness", in which he argues that it isn't. In "A Prolog-like Inference System for Computing Minimum-Cost Abductive Explanations in Natural Language Interpretation", authored and presented by Mark Stickel, interpretation of a NL utterance is held to be the least-cost conjunction of provable facts and assumptions needed to make the utterance true. A system for computing assumptions and costs was given. In a similar vein, Gary Hall presented "A Solution to the MAP problem in Natural Language Interface Construction", wherein a method was given for computing least-cost interpretations for unknown and ambiguous NL expressions contained in database queries, based on the data dependencies embedded in the structure of the database. Another abductive approach was presented in "Maintaining Consistency and Plausibility in Integrated Natural Language Understanding", written by Xuemin Liu, Toyoaki Nishida, and Shuji Doshita. Based on the ATMS system of deKleer, this approach uses a measure of plausibility to choose among competing consistent sets of assumptions about the meaning of unknown NL expressions.

Session 9a: Expert and Knowledge-based System

G. Burle presented a paper, "Using Qualitative Reasoning for Experiment Planning in Performability Evaluation of Computer Systems", which demonstrated how the process-oriented approach to qualitative reasoning could be applied to support experiment planning. Wray Buntine began his discussion by considering a (Bayesian) probabilistic analysis of default reasoning before running out of time. His paper was entitled "Default and Likelihood Reasoning and Qualitative Probabilistic Reasoning". Pierre Berlandier presented a self-contained talk entitled "Reflexive Constraints for Dynamic Knowledge Bases". NRCC researcher Sherman Lang presented "Graph Synthesis for Knowledge Acquisition in Mobile Robots"— a description of the synthesis of a geometric world model for use in navigation by a mobile robot. Finally, Lori Lamel formalized the information used in spectrogram reading as "A Knowledge Based System for Stop Consonant Identification" — an expert system application. Polite questions were asked of all speakers.

Session 9c: Expert and Knowledge-based System

H.Q. Tan presented a paper, "Automatic Equation Derivation and Code-Generation in Engineering Analysis", which described a package which takes equations generated by MACSYMA and VAXIMA and transforms them into FORTRAN. P. McFetridge described how a natural language interface can automatically and dynamically generate knowledge structures required to answer queries to databases by examining input queries for presupposed information about the database in "Knowledge Acquisition in SystemX: A natural language interface to relation databases". Leung and Nijssen presented "Designing relational databases through examples", described a system which accepted English language descriptions and, after querying the user about the elements of the English sentences, generated a relational database.

Other Sessions

Other sessions left undescribed include sessions on AI architecture, automatic reasoning, distributed and parallel computing, logic programming, natural language, sessions machine learning and neural networks, programming language and complexity, pattern recognition and vision, expert and knowledge-based systems, AI implementation issues, software engineering and implementation issues, and user interface and programming environments.

The Banquet et al.

Simply put, the banquet dinner, held in the Marina Room of the Excelsior Hotel on 20 December 1988, was exquisite. A traditional 12-course Chinese banquet which included such delicacies as "double-boiled pigeon", "garoupa", "shark's fin soup", "intoxicated prawns", etc., was much more than anyone has a right to expect. Most frequently heard comment — "I'm sure glad I don't have to do the dishes". A rather lengthy program preceded the dinner, which at the time appeared to be cruel and unusual punishment since the menu was made available to participants beforehand. In all, the lunches and banquet, which were part of the conference fare, were truly excellent.

Summary

Approximately 200 people attended this first-ever computer science conference held in Hong Kong. Although the organizers expressed some disappointment at what they believed to be a small turnout, this conference compares very favourably with other "first time" conferences held in other popular sites, e.g., the first international IEEE-sponsored Data Engineering Conference held in Los Angeles in 1985. Even after extensive promotion, an audience of approximately 300 attended. Five years later, Data Engineering is a highly visible and well-attended conference.

International Computer Science Conference '88 was an enjoyable experience overall. The technical program was good, albeit it somewhat varied given the conference theme. The Excelsior Hotel in Causeway Bay proved to be an ideal host site in Hong Kong. Facilities and amenities were first rate, service was excellent and the conference discount for room rates was generous and affordable. The program co-chairpersons, Francis Chin and Jean-Louis Lassez, and the general chairman, K.W. Ng, deserve credit for the program and smooth operation of the conference.

Book Reviews

edited by Graeme Hirst

Foundations of Deductive Databases and Logic Programming

Jack Minker
[University of Maryland]

San Mateo, CA: Morgan Kaufmann, 1988, 746 pp
Paperback, ISBN 0-934613-40-0, US$36.95

Reviewed by
Ranabir Gupta
Simon Fraser University

The Book in Essence

Foundations of Deductive Databases and Logic Programming is a selection of papers taken from among those presented at a workshop of the same name, convened by Jack Minker in August 1986. Attendance being by invitation, the most promising ongoing research of prominent investigators was represented. After being reviewed and fortified with background material, this work was intended to serve as the foundation for further research afield. Thus, the description on the cover informs the reader that the theory underlying such diverse fields as stratified databases, parallel query processing, update semantics and circumscription is addressed. The volume is aimed at graduate students and researchers in deductive databases, logic programming, "artificial intelligence", "expert systems", and theorem proving.

In keeping with these objectives, much substance of value has been added to the collection of papers. The outstanding features include the attention to background, open research problems, and references. Papers have been fleshed out with survey material. In addition, topics treated at depth merit dedicated survey papers. An extended introduction provides extensive references to alternate approaches taken in the various fields and also points out good survey papers in those areas.

The introduction, as well as some papers, pointedly discusses directions for future research and lists many open problems. Most of the research problems suggested, however, are in the nature of extensions (such as investigating special cases) to larger issues either addressed, or referred to, in the papers. Little space is devoted to outlining general long-term directions. A graduate student searching for a master's thesis topic would find these pastures green.

The 23-page subject index is comprehensive. Though there is no separate bibliography in the manner of the *Readings in* series (Morgan Kaufmann), the references, some internal to the book, are adequate.

Highlights

The book is a collection of 18 papers. The average length of 39 pages indicates the depth of the treatment. The papers are organized under three sections:

1. Negation and stratified databases, with six papers;
2. Fundamental issues in deductive databases and implementation, with seven; and
3. Unification and logic programs, with five.

Part 1: Negation and stratified databases

The section on negation and stratified databases is well-focused. A review of the model theory of semantics of classical schemes for expressing negation in databases appears in the paper by Shepherdson. This comprehensive survey emphasizes the closed-world assumption semantics of negation and those of negation-as-failure as a means of practical realization. The search for expressive logics, despite the apparent impossibility of finding generic models for languages more powerful than Horn-clause logic, is well motivated. Thereafter, well-known approaches taken are introduced and related semantically to one another. For example, the program classes for which negation-as-failure is complete with respect to Clark's completed database are outlined. The reader is referred to other papers for many proofs.

The theory of stratified databases is developed independently in the two papers by Apt et al. and Van Gelder. Together, they summarize the earliest work on stratification. The first paper gives an excellent survey and defines the models of stratified programs in terms of the fixed points of nonmonotonic operators. The second develops syntactic methods to prevent infinite recursion and deadlocks. "Tight tree" semantics is developed to support the idea of restricting proofs involving negation as failure. "Tight tree" semantics is related to iterated fixed-points semantics.

Two subsequent papers by Lifschitz and Przymusinski respectively apply techniques debated in knowledge representation circles to the stratified database domain. Essentially, pointwise circumscription and prioritized circumscription are applied in the presence of negated database information. The latter paper defines the intuitively pleasing class of "perfect models" of locally stratified databases. Problems with circumscription as a general tool include the cumbersome model-definition methods and bad complexity characteristics. This area is replete with open problems, which I found inadequate mention of.

The last paper (Topor and Sonenberg) is rather atypical of the section in that it defines a subclass, namely "allowed stratified databases", which is tractable and stable under updates. Some decision problems related to this subclass are left open.

Part 2: Fundamental issues in deductive databases and implementation

The issues tackled by the papers in this section fall into several categories. The presentation, with the exception of the topic of integrity constraints, is therefore not so self-contained as in the first section.

"Foundations of semantic query optimization for deductive databases" (Chakravarthy et al.) partially introduces the groundwork in the area. It concentrates on the non-query-dependent case where part of an integrity constraint subsumes the body of a non-recursive intensional database rule. During "semantic optimization", the integrity constraints are used to extend the rules. The extended

rules can be used to process queries faster by checking untenable conditions beforehand.

The idea of manipulating intensions links this to the next paper (Imielinski). The latter fits into recent work done by Imielinski and others on overcoming the problem of unsafe queries through a complexity-tailored hierarchy of answers. Some levels of the hierarchy comprise finite representations of infinite sets as predicates/rules. Little previous work is directly relevant to this fresh and intriguing line of thought.

Sadri and Kowalski, in "A theorem-proving approach to database integrity", propose a straightforward extension to SLDNF resolution to check consistency with integrity constraints (denials). The viewpoint is conventional Prolog programming search-space reduction.

Another valuable paper is the timely "Performance evaluation of data-intensive logic programs" (Bancilhon and Ramakrishnan). For four different types of queries, using analytical methods on the ten best-known recursive query processing algorithms, factors which influence performance are identified. It is shown that the characteristics of the database and query can cause orders-of-magnitude variations in performance of different algorithms. This comprehensive survey clearly explicates the terms of reference of the tests.

The section is completed by papers on view update semantics in dynamic logic based on perfect models(Manchanda and Warren), superjoin algorithms in anticipation of the move of deductive databases out of main memory to secondary storage (Thom, Ramamohanarao, and Naish) and an approach to compiling semantics of negation based on the generalized closed-world assumption (Henschen and Park). The latter uses a logic based on the truth value "indefinite" in addition to true and false. The last three papers stand isolated in their subfields. Though peppered with references, they are unsupported by other papers in this collection.

Part 3: Unification and logic programs

"Unification revisited" (Lassez, Maher, and Marriot) sets the theme of the last section. It proposes a framework for integrating many forms of unification, viewing unification as a general process of determining "compatibility" of objects. Mathematical properties of such a process are explored, and a good set of references is provided. As the introduction notes, however, other than the problems in nonstandard unification, the area is quite closed.

In "Logic programming and parallel complexity", on the other hand, Kanellakis lists many open problems related to parallel unification. A problem's optimizability and parallelizability are shown to be related to its membership in complexity class NC. Stage functions are introduced in this good starting paper in the field. "Equivalences of logic program" (Maher) deals with the topic in a theoretical light. Strengths of equivalence notions are compared in cryptic mathematical terms. Sagiv, in "Optimizing datalog programs", deals with program transformations yielding "uniformly equivalent" programs through a decidable, normal-case polynomial process. This limited equivalence is practically attractive because of the equivalence results that hold for some local rule modifications. In some cases, the "optimization" actually yields a worse program. Sagiv also exposes open problems with regard to termination for more involved notions of equivalence.

In the last paper of the book, van Emden and Szeredi discuss the conversion of AND to OR control using fold/unfold transformations, a topic with implications in parallel processing, though not significantly in general logic programming. While this section's unifying thread is its interest in equivalences of terms or programs, the theme takes different forms in each paper.

Looking beyond

Evidence that most papers in the book are individually influential lies in the numerous references to these papers in recent conference proceedings. A layperson's measure may be the figures gleaned from a general-purpose citation index. One such index abstracts 27 publications in computing science, most only peripherally concerned with deductive databases or logic programming. This source indicates 22 of the 31 authors shared 100 references to their work in 1987. Of these, 56 were directly related to the work included in *Foundations of Deductive Databases and Logic Programming*, a fairly impressive citation record. References were also made to other work presented at the workshop but not included in this book.

Understandably, and commendably, the book does not presume to please everyone. The viewpoint is decidedly theoretical. The goals rarely include discussing details of applications or implementation. Neither do data-structuring issues or any of the "heuristic" methods familiar to "expert systems" designers find room. Builders of "expert systems", as commercially understood, will not find more here than discussion on theoretical bounds on expressibility and the complexity in systems that employ, for example, negation as failure.

Thus, the *Foundations* are not laid for laypersons. With rare exceptions, familiarity with logic and model-theoretic terminology is assumed. Moreover, the book imposes the outlook of logic programming on deductive database issues. Seldom does a flavour akin to "incrementally extending set-based relational database capabilities" seep through.

It is worth mentioning that recent interest in assimilating functional programming, higher-order and nonstandard logics into the logic programming fold is not reflected herein. Neither do topics peripheral to deductive databases, such as semantic data modeling and function symbols' incorporation, find mention here.

A book of 750-odd pages inevitably excludes some work even in its field of specialization; and deductive databases is an expanding field. Should one want to look beyond, the following comments may be helpful. Work in negation and related issues from a database perspective is being carried out by the LDL group; though mentioned in passing, papers by Naqvi, Zaniolo et al. are conspicuous by their absence. Frequently quoted authors of some significant work on integrity constraints and their use in query optimization include Soisalon-Soininen, Nicholas, and King. Fundamental work in update logics, including that by Abiteboul, Cosmadakis, and Bancilhon, miss mention in the limited treatment of update semantics in this volume.

Some references to lambda-terms and higher-order unification and its use in program transformation may be found in *Logic Programming: Proceedings of the Fifth International Conference and Symposium. Readings in Knowledge Representation* (Morgan Kaufmann) might be a useful guide into literature on circumscription and nonmonotonic operators. For database-oriented views on this area, the reader is referred to publications such as *Readings in Database Systems and Artificial Intelligence* (Morgan Kaufmann) and proceedings of ACM-PoDS, ACM-SIGMOD, and conferences on expert database sys-

tems. Nevertheless, within its restricted domain, this book more than lives up to its name.

Given that most individual papers in the book have much to recommend them, what about the book as a whole? Would I, the initiate, rely on this book to guide me from the basics up to the frontiers? Yes, if my interests lay in one of the subfields below. I would need to chase many pointers but could be assured of not overlooking something significant. Well, would I, the researcher, think that this book contains nearly all of the state-of-the-art in research in 1986? Yes, if my field were stratified databases. A qualified yes if my field were handling negation, or use of integrity constraints in optimal query evaluation, perhaps in parallel logic programming. Otherwise, I would expect to find some good papers, not a compendium. Does this book contain a good amount of truly seminal work? Not apparently, but good candidates abound. But would I say it contains a good mix of foundational material and surveys? Emphatically so! On the whole, this is a book worth keeping within easy reach, given compatible interests.

Ranabir Gupta is a Ph.D. candidate in the School of Computing Science at Simon Fraser University. He is pursuing research in deductive databases and semantic data modeling under the supervision of Dr. Jiawei Han.

The Artificial Intelligence Debate: False Starts, Real Foundations

Stephen R. Graubard (editor)

Cambridge, MA: The MIT Press, 1988, 311 pp
Paperback, ISBN 0-262-54074-2, US$9.95

Reviewed by
Peter Turney
University of Toronto

This book is a collection of 14 essays that were originally published in the Winter 1988 issue of *Daedalus*. The intended audience appears to be people outside AI research who wish to know what is happening in the field. An acquaintance with AI research, perhaps through other popular books, is helpful but not required. Most of the authors are well known in AI. They approach AI from a variety of perspectives, including computer science, neuroscience, philosophy, sociology, media technology, engineering, physics, mathematics, and biology.

Taken individually, the papers are generally engaging and well written. On the other hand, there is little unity to this collection. Several papers do share a concern with the conflict between connectionism and symbolicism, but there are exceptions. The book's title, *The Artificial Intelligence Debate*, is misleading, in that there is no single topic of debate, and indeed many papers are not part of any debate at all. I would prefer a more accurate title, for example, *Symbolic and Connectionist Approaches to Artificial Intelligence*.

Introductions, Histories, Speculations, and Opinions

The first paper is by Seymour Papert. Readers of the second edition of *Perceptrons* (Minsky and Papert 1988) or Minsky's *Society of Mind* (1985) will not be surprised by what Papert says here: There is no "universal mechanism" for intelligence. A general algorithm for learning or problem solving that can be applied to all or most domains has long been a holy grail for AI researchers. Papert thinks that their faith is misplaced.

In the next paper, Hubert Dreyfus and Stuart Dreyfus argue against the possibility of intelligent machines. They exploit the conflict between connectionism and symbolicism. The symbolic approach is attacked with the same arguments that connectionists employ, and vice versa.

A large part of this paper is a history of the roots of connectionism and symbolicism. The difference between these two approaches to AI can be expressed with contrasting pairs of words. Table 1 lists some of the pairs that are mentioned in this essay. When I see these words together in a table, rather than distributed throughout an essay, I am struck by how arbitrary it is to group some of them together. For example, there are parallel algorithms for logic and problem solving; there are symbolic algorithms for learning, pattern recognition, and statistics. Much of the work in machine learning constitutes a proof that this table is seriously defective. The net result of this paper is to make me aware of the dissonance of the distinction between connectionism and symbolicism. This result is contrary to the authors' intentions.

Symbolicism	Connectionism	Page
mind	brain	15, 16
atomism	holism	16, 17, 22
reductionism	anti-reductionism	16, 17, 22
discrete	continuous	18, 23
thinking	pattern recognition	24
logicalstatistical		16, 20
serial	parallel	21
designed	evolved	19
innate	tabula rasa	18, 19
philosophical	neurological	16, 18
problem-solving	learning	16
digital	analog	18, 23
static	dynamic	19, 34

Table 1: Symbolicism versus Connectionism

In the next essay, Robert Sokolowski asks an intriguing question, "will writing turn out to be a four-thousand-year prelude to artificial intelligence?" (p. 50). The idea is that AI is a natural extension of the written word. This suggests the possibility that computers may be able to do verbal thinking (translation, for example) but not nonverbal thinking (passion, for example).

It is common to argue that a machine cannot be intelligent, because there is a certain ability that humans have but machines could never have. Arguments of this type differ only in what they choose for that certain ability. This paper is interesting in that verbal ability has been ceded to computers, which seems to me to be a major concession.

Pamela McCorduck presents an entertaining travelogue, taking us on a tour through the history of AI. The tour makes a stop at Homer's *Iliad* and passes by Descartes, Hume, Locke, Newton, and Julien Offray de la Mettrie. We visit Dartmouth College in 1956, where an interesting conference takes place. We meet Patrick Winston, Allen Newell, Herbert Simon, John McCarthy, Marvin Minsky,

Roger Schank, and Edward Feigenbaum.

In the introduction to "Neural Nets and Artificial Intelligence", Jack Cowan and David Sharp state that they will concentrate on the problem of "the representation of external events inside neural nets" (p. 85). I expected a critical comparison of different representations, because the title of this book led me to expect a debate. In fact, the article is a short history of attempts to make neural nets learn and remember..As such, the article succeeds admirably. The historical treatment is fair and balanced. The bibliography is reasonably large (85 entries) and well selected.

The constructive part of "Real Brains and Artificial Intelligence", by George Reeke and Gerald Edelman, is the overview of Edelman's (1986) theory of Neural Darwinism. Edelman believes that groups of neurons evolve in the brain, much as organisms evolve in the world. The unfortunate part of this paper is the rather shrill attack on conventional approaches to AI, both symbolicist and connectionist. The authors' strongest criticism is directed at the "brittleness" of symbolic AI algorithms; their inability to handle domains just slightly different from their intended domains (p. 149).

I find it strange that the authors do not mention John Holland's (1986) genetic algorithms. Holland is also concerned with the application of the theory of evolution to the problem of brittleness in AI. This paper makes it appear that the authors are isolated geniuses, fighting against the ignorant dogma of orthodox AI. In fact, there have been several conferences devoted to genetic algorithms.

Anya Hurlbert and Tomaso Poggio discuss machine vision. They claim that traditional AI has denigrated vision, denying that it involves intelligence. Yet vision has been more difficult to program "than the most sophisticated mathematical reasoning" (p. 214). The authors believe, contrary to traditional AI, that machine vision is the epitome of intelligence. The first sentence of the paper expresses this theme very well: "Vision is more than a sense; it is an intelligence" (p. 213). The authors view machine vision as a science that builds on the best aspects of symbolic and connectionist approaches to AI, yet stands apart from both of these approaches (p. 217). They argue persuasively that both symbolic and connectionist approaches can benefit greatly from the lessons which machine vision has to teach.

Sherry Turkle's paper argues that the "scientific culture" of psychoanalysis has much in common with the "scientific culture" of AI, and that the two fields are converging, ultimately to become allies. Her paper is grounded firmly in the sociological approach to the analysis of science. Turkle does not mention the controversy surrounding this approach. In the philosophy of science, there are two strongly polarized camps. One camp believes that science is essentially a social enterprise, governed by fashion, cultural climate, and political pressures. This camp has its roots in Thomas Kuhn's *The Structure of Scientific Revolutions* (1962). The second camp believes that progress in science is a matter of competition between theories, where, ideally, the winner maximizes rationally desirable features, such as predictive ability. For an example of this view, interested readers might turn to Larry Laudan's *Progress and itsProblems* (1977).

Hilary Putnam points out that AI research has not *yet* achieved anything that sheds much light on the nature of intelligence. He argues that the brain may be too complex for AI research to succeed. He concludes that it is still too early to say anything significant about the possibility of AI. It is hard to disagree with Putnam, but his conclusion seems excessively cautious. Sometimes it is better to risk being wrong than it is to withhold judgement.

Daniel Dennett replies to Putnam's paper. In essence, Dennett says, "we shall see" (p. 291). But to wait and see is just what Putnam proposes to do.

John McCarthy discusses "Mathematical Logic in Artificial Intelligence". Once again, the title of the book misled me, and it took me a few pages to realize my error. This essay is an *introduction* to (not a *defence* of) the application of formal logic to AI. It discusses different ways in which logic can be applied to AI and the problems that lead to current research in nonmonotonic logic. I am sympathetic to the view presented here, that AI researchers and philosophers have much to learn from each other.

My strongest criticism of this collection is directed at the misleading title it has been given. Also, very little effort was spent on the transformation of the special issue of *Daedalus* into this book: there is no preface, no index, no glossary, no conclusion, and no commentary to tie the articles together.

Apart from these criticisms, this is an interesting collection. It is suitable for people who want to know more about AI but want to see neither Lisp code nor expressions in first-order logic.

References

Edelman, G.E., *Neural Darwinism: The Theory of Neuronal Group Selection*, Basic Books, New York, New York, 1987.

Holland, J.H., "Escaping Brittleness: The Possibilities of General-Purpose Learning Algorithms Applied to Parallel Rule-Based Systems", in *Machine Learning: An Artificial Intelligence Approach*, Volume II, edited by Michalski, R.S., Carbonell, J.G., and Mitchell, T.M., Morgan Kaufmann Publishers, Los Altos, California, 1986.

Kuhn, T.S., *The Structure of Scientific Revolutions*, University of Chicago Press, Chicago, Illinois, 1962.

Laudan, L., *Progress and its Problems: Towards a Theory of Scientific Growth*, University of California Press, Berkeley, California, 1977.

Minsky, M.L., *The Society of Mind*, Simon and Schuster, New York, New York, 1985.

Minsky, M.L., and Papert, S.A., *Perceptrons: An Introduction to Computational Geometry*, expanded edition, MIT Press, Cambridge, Massachusetts, 1988.

Peter Turney is a postdoctoral fellow in the Philosophy Department of the University of Toronto. He is currently working on the problem of the representation of scientific theories. He has a Ph.D. in Philosophy from the University of Toronto.

Practical Planning: Extending the Classical AI Planning Paradigm

David E. Wilkins
[SRI International]

San Mateo: Morgan Kaufmann Publishers, 1988, xiii+205 pp
(The Morgan Kaufmann series in representation and reasoning)
Hardbound, ISBN 0-934613-94-X, US$34.95

Reviewed by
Peter Davies
Expert Solutions

This book covers the AI planning system SIPE (System for Interactive Planning and Execution Monitoring), which has been the vehicle for the author's research at SRI over the last several years. SIPE, which was developed on Symbolics hardware in Zetalisp, extends the classical approach to planning in several ways: it can reason about resources, can post and use constraints, and can reason about different world states, using a deductive causal theory. The system has been used to generate correct plans in many domains, including the standard blocks world and extensions, cooking, aircraft operations, travel planning, construction of objects in a machine shop, and an indoor mobile-robot domain (SRI's robot Flakey).

Of great interest is the SIPE project's focus on efficiency, so that, unlike most earlier systems, it can be employed in real-world applications. Previous work in planning has been largely theoretical and systems have not been efficient enough to be practical. It is not clear whether any of the plans mentioned above were part of "real" applications. However, SIPE has been used (subsequent to the publication of the book) to build a prototype system for the scheduling of packaging lines at one of the world's largest breweries.

Chapter 1 sets the stage by discussing the "classical" approach to AI planning, which assumes a state-based representation, and contrasts this with other approaches (procedural and event-based representations). Wilkins restricts his attention to problems that address the frame problem, which refers to reasoning about how the world changes as scheduled events occur. The chapter also discusses the essence of planning, capabilities of planning systems, and the important systems developed to date. Finally, it includes a description of the focus of the SIPE system, which is the balance between epistemological and heuristic adequacy — i.e., can the problem be represented adequately and still be practically usable?

Two important central features of SIPE, discussed in chapter 2, are its ability to plan at different levels of abstraction and nonlinear plans. Also detailed here are assumptions and limitations of the system. In chapter 3 we see how knowledge is represented in SIPE, as a combination of frames and first-order logic. An operator description language is used to provide an easy understandable way to specify operators. Also discussed in this chapter are plots, which are step-by-step instructions for performing actions, and how plans are specified.

Chapters 4 through 11 detail other aspects of the SIPE system. Each chapter uses examples, mostly from the mobile-robot (Flakey) project. These examples serve well to illustrate the ideas being discussed. Topics covered are: hierarchical planning, constraints, the truth criterion, deductive causal theories, plan critics, resources, search, and adding reactive capabilities.

The last two chapters look at heuristic adequacy and compare SIPE with the other major planning systems developed to date. I found this to be the least satisfying part of the book. Wilkins's treatment of SIPE's efficiency and performance comparisons with earlier systems is vague and inconclusive. Of course, as Wilkins himself states, this is a difficult task. Data do not exist on planning systems attacking problems of similar complexity. However, since this efficiency is the main focus of the book, I would have expected a more thorough attempt at addressing this issue.

Despite the complaints mentioned above, I found the book to live up to its claim to be an "in-depth examination of this classical planning paradigm". It is well written, well organized, and very interesting. Wilkins provides a comprehensive description of all facets of the SIPE system. The book is a thorough treatment of this important and difficult area of AI research and will be particularly useful to those interested in applying the technology. The interested reader should also consult the special issue on planning of *Computational Intelligence* (Volume 4, Number 4, November 1988), of which Wilkins is the guest editor.

A claim is made that this book is accessible to those with a "basic background in AI". Perhaps the adjective "formal" should be added, since a cursory exposure to the technology of expert systems will not suffice. Some academic training in computer science or related disciplines is required. In particular an understanding of formal grammars and/or the rudiments of mathematical logic is desirable. However, it is hard to imagine that someone without this knowledge will be interested in the volume in the first place.

Peter Davies is a consultant, developer, and researcher in expert systems and other AI applications. He is the principal of Expert Solutions, a Toronto-based AI company.

Briefly Noted

Exploring Artificial Intelligence:
Survey Talks from the National Conferences on Artificial Intelligence
Howard E. Shrobe
[Symbolics Inc]
Morgan Kaufmann Publishers, 1988, xii+693 pp
(Distributed in Canada by John Wiley and Sons Canada Ltd)
Paperback, ISBN 0-934613-67-2
Hardbound, ISBN 0-934613-69-9, Cdn$63.95

Sixteen of the survey talks given at the 1986 and 1987 AAAI conferences that were not included in the proceedings of the conferences are collected in this book. The areas of AI covered are: intelligent tutors; learning; natural language; planning and search; reasoning; and AI architecture and systems. The papers are:

"Intelligent Tutoring Systems", by Beverly Woolf
"An Introduction to Explanation-Based Learning", by Gerald DeJong
"Knowledge-Based Natural Language Understanding", by Wendy Lehnert
"Natural-Language Interfaces", by C. Raymond Perrault and Barbara Grosz
"Reasoning about Plans and Actions", by Michael Georgeff
"Search: A Survey of Recent Results", by Richard Korf
"Qualitative Physics: Past, Present, and Future", by Kenneth Forbus
"Model-Based Reasoning: Troubleshooting", by Ran-

dall Davis and Walter Hamscher
"Artificial Intelligence Techniques for Diagnostic Reasoning in Medicine", by Ramesh Patil
"Evidential Reasoning Under Uncertainty", by Judea Pearl
"Temporal Reasoning in Artificial Intelligence", by Yoav Shoham and Nita Goyal
"Nonmonotonic Reasoning", by Raymond Reiter
"A Survey of Automated Deduction", by Woody Bledsoe and Richard Hedges
"Symbolic Computing Architectures", by Howard Shrobe
"The Common Lisp Object System: An Example of Integrating Programming Paradigms", by Daniel Bobrow
"Artificial Intelligence and Software Engineering", by David Barstow

Playing Fast and Loose with Time and Space
P.S. Mueller
Meadowbrook Press, 1989, 105 pp
Paperbound, ISBN 0-88166-153-8, US$4.95
(Distributed by Simon and Schuster as ISBN 0-671-67740-3;
or can be ordered by mail from the publisher, 18381 Minnetonka Blvd, Deephaven, MN 55391, U.S.A.)

P.S. Mueller's AI-oriented cartoons have been seen frequently in the pages of *Canadian AI*. In this collection, the high-tech themes take second place to outrageous puns ("croissant wrench", "You may already be a whiner") and other general absurdities — as illustrated. Mueller's humour is an acquired taste; if you don't hate this stuff you'll probably love it.

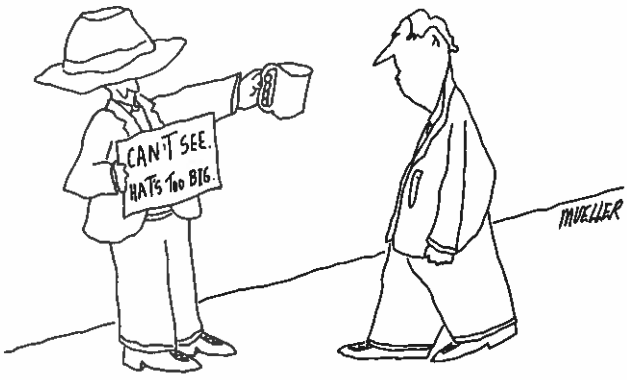

Readings from *AI Magazine*
Robert Engelmore
Menlo Park, CA: American Association for Artificial Intelligence, 1988, x+665 pp
Paperbound, ISBN 0-929280-01-6, US$74.95

This book contains all the feature articles published in The *AI Magazine* in its first five years, Spring 1980 to Winter 1985. The articles are grouped thematically and indexed. Not included are advertisements, research-in-progress items, and the like.

There is much of value in here, and as many of the early issues of the magazine are hard to come by, this collection will be valuable. Four articles cover the early history of the major AI labs; quite a number report on major AI projects of the time. The AAAI Presidents' addresses, usually major overviews of state of the field, are all included.

Unfortunately, there is also much that need not have been republished or which is simply out of date and has been superseded by the authors' later work. For example, pre-publication excerpts from *The Handbook of Artificial Intelligence* may have been of interest to the magazine's readers in 1980, when they were still current and the handbook was not yet available, but their inclusion in this collection is pointless. One would have appreciated an attempt by the editor to have separated the hard to get and enduring from the widely distributed or ephemeral.

—G.H.

Proceedings of the 1988 Connectionist Models Summer School
David Touretzky, Geoffrey Hinton, and Terrence Sejnowski (editors)
[Carnegie Mellon University, University of Toronto, and University of California, San Diego, resp.]
Morgan Kaufmann Publishers, 1989, xi+527 pp
(Distributed in Canada by John Wiley and Sons Canada Ltd)
Paperback, ISBN 1-55860-033-7

Fifty-six papers written by students and faculty are collected in this record of the second connectionist summer school at Carnegie Mellon University. The papers have all been revised under faculty guidance, and some describe work done during the ten-day-long school, which was attended by 50 students chosen from over 200 applicants.

Books Received

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.

Proceedings, First International Conference on Principles of Knowledge
Representation and Reasoning (KR '89), Toronto, May 1989

Ronald J. Brachman, Hector J. Levesque, and Raymond Reiter (editors)
[AT&T Bell Labs and University of Toronto]
San Mateo, CA: Morgan Kaufmann Publishers, 1989, ix+520 pp
Paperbound, ISBN 1-55860-032-9
Distributed in Canada by John Wiley and Sons Canada Ltd

Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence
Allan Collins and Edward E. Smith (editors)
[BBN Laboratories and University of Michigan]
San Mateo, CA: Morgan Kaufmann Publishers, 1988, ix+661 pp
Paperbound, ISBN 0-55860-013-2
Distributed in Canada by John Wiley and Sons Canada Ltd

Naive Semantics for Natural Language Understanding
Kathleen Dahlgren
[IBM Los Angeles Scientific Center]
Boston: Kluwer Academic Publishers, 1988, x+258 pp
(The Kluwer international series in engineering and computer science)
Hardbound, ISBN 0-89838-287-4, US$52.50

Relational Models of the Lexicon: Representing Knowledge in Semantic Networks
Martha Walton Evens (editor)
[Illinois Institute of Technology]
Cambridge, England: Cambridge University Press, 1988, ix+300 pp
(Studies in natural language processing)
Hardbound, ISBN 0-521-36300-4

English-Japanese, Japanese-English Dictionary of Computer and Data-Processing Terms
Gene Ferber
Cambridge, MA: The MIT Press, 1989, 470 pp
Hardbound, ISBN 0-262-06114-7, US$75.00

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

Attribute-Value Logic and the Theory of Grammar
Mark Johnson
[Brown University]
Stanford, CA: Center for the Study of Language and Information, 1988, xi+162 pp
(CSLI lecture notes 16)
Hardbound, ISBN 0-937073-37-7, US$37.50; Paperbound, ISBN 0-937073-36-9, US$15.95
(Distributed by the University of Chicago Press)

Sparse Distributed Memory
Pentti Kanerva
[NASA Ames Research Center]
Cambridge, MA: The MIT Press, 1989, xxii+155 pp
Hardbound, ISBN 0-262-11132-2, US$24.95

Speaking: From Intention to Articulation
Willem J.M. Levelt
[Max-Planck-Institut für Psycholinguistik]
Cambridge, MA: The MIT Press, 1989, xvii+566 pp
(ACL-MIT Press series in natural-language processing)
Hardbound, ISBN 0-262-12137-9, US$39.95

Machine Translation: How Far Can It Go?
Makoto Nagao
[Kyoto University]
Oxford: Oxford University Press, 1989, xii+150 pp
Hardbound, ISBN 0-19-853739-5

Computational Intelligence

Abstracts for 5(1) April 1989

Constructive Belief and Rational Representtion
Jon Doyle
Massachusetts Institute of Technology

It is commonplace in artificial intelligence to divide an agent's explicit beliefs into two parts: the beliefs explicitly represented or manifest in memory, and the implicitly represented or constructive beliefs that are repeatedly reconstructed when needed rather than memorized. Many theories of knowledge view the relation between manifest and constructive beliefs as a logical relation, with the manifest beliefs representing the constructive beliefs through a logic of belief. This view, however, limits the ability of a theory to treat incomplete or inconsistent sets of beliefs in useful ways. We argue that a more illuminating view is that belief is the result of rational representation. In this theory, the agent obtains its constructive beliefs by using its manifest beliefs and preferences to rationally (in the sense of decision theory) choose the most useful conclusions indicated by the manifest beliefs.

Compiling General Linear Recursions by Variable Connection Graph Analysis
Jiawei Han
Simon Fraser University

Compilation is a powerful preprocessing technique in the processing of recursions in knowledge-based systems. This paper develops a method of compiling and optimizing complex function-free linear recursions using a varible connection graph. It shows that a function-free recursion consisting of a linear recursive rule and one or more non-recursive rules can be compiled to (1) a bounded recursion, in which recursion can be eliminated from the program, or (2) an n-chain recursion, whose compiled formula consists of one chain, when n = 1, or n synchronized compiled chains, when n > 1. The study is based on a classification of linear recursions and a study of the com-

pilation results of each class. Using the variable connection graph, linear recursions are classified into six classes: acyclic paths, unit cycles, uniform cycles, non-uniform cycles, connected components and their disjoint mixtures. Recursions in each class share some common properties in compilation. Our study presents an organized picture for the compilation of general function-free linear recursions. After compilation, the processing of complex linear recursions becomes essentially the processing of primitive n-chain recursions or bounded recursions to which efficient processing methods are available.

Representing Defaults with Epistemic Concepts
Kurt Konolige
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, California
94025

Karen Myers
Department of Computer Science
Stanford University
Stanford, California
94305

Reasoning about defaults – implications that typically hold, but which may have exceptions – is an important part of commonsense reasoning. We present some parts of a theory of defaults, concentrating on distinctions between various subtle ways in which defaults can be defeated and on inferences which seem plausible but which are not correct in all cases. In order to represent this theory in a formal system, it is natural to use the epistemic concept of self-belief. We show how to express the theory by a local translation into autoepistemic logic, which contains the requisite epistemic operators.

Learning and Classification of Monotonic Ordinal Concepts
Arie Ben-David, Leon Sterling and Yoh-Han Pao
Department of Computer Engineering and Science and Center for Automation and Intelligent Systems Research
Case Western Reserve University
Cleveland, Ohio

Ordinal reasoning plays a major role in human cognition. This paper identifies an important class of classification problems of patterns taken from ordinal domains and presents efficient, incremental algorithms for learning the classification rules from examples. We show that by adopting a monotonicity assumption of the output with respect to the input, inconsistencies among examples can be easily detected and the number of possible classification rules substantially reduced. By adopting a conservative classification criterion, the required number of rules further decreases. The monotonicity and conservatism of the classification also enable the resolution of conflicts among inconsistent examples and the graceful handling of "don't knows" and "don't cares" during the learning and classification phases. Two typical examples in which the suggested classification model works well are given. The first example is taken from the financial domain and the second from machining.

Technical Reports

University of Toronto

Telos: A Knowledge Representation Language for Requirements Modelling
M. Koubarakis, J. Mylopoulos, M. Stanley, M. Jarke
KRR-TR-89-1

Telos is a knowledge representation language intended for software development applications. Its main features include (a) integrated facilities for representing and reasoning about temporal knowledge, (b) a powerful mechanism for extending the language through user-defined meta-attributes and (c) embedded inference facilities for query evaluation and integrity checking such as those used in deductive databases. The language was designed to offer a uniform representational framework and a functional view of knowledge bases in the spirit of KRYPTON. We present a brief overview of the main features of Telos and sketch its formalization. Also, we describe a prototype implementation of the language and discuss some applications in building software development environments.

Logic and the Complexity of Reasoning
Hector J. Levesque
KRR-TR-89-2

What does formal logic have to do with the computational task of ordinary common sense reasoning? Very little, it would seem, given the apparent difficulty we all have in actually doing logic, on the one hand, and the wholesale deviations from pure logical inference in our everyday thinking, on the other. We argue that this impression is mistaken and may result from attempting to study cognition in computational terms without the constraints of computational complexity. Ignoring these constraints can lead, among other things, to an overly simplistic and rigid view of logic that sacrifices computational realism to mathematical elegance and generality of expression. We illustrate this idea by considering a number of forms of reasoning that deviate from classical logical inference, but are nonetheless both logically coherent and computationally tractable.

All I Know: A Study in Autoepistemic Logic
Hector J. Levesque
KRR-TR-89-3

Current approaches to formalizing non-monotonic reasoning using logics of belief require new metalogical properties over sets of sentences to be defined. This research attempts to show how some of these patterns of reasoning can be captured using only the classical notions of logic (satisfiability, validity, implication). This is done by extending a logic of belief so that it is possible to say that *only* a certain proposition (or finite set of them) is believed. This research also extends previous approaches to handle quantifiers and equality, provides a semantic account of certain types of non-monotonicity and, through a simple proof theory, allows formal derivations to be generated.

Ordering information:

Requests for copies of the fore-named publications should be addressed to: KRR Technical Reports Secretary,10 King's College Road, Room 3202, University of Toronto, Toronto, Ontario, M5S 1A4

Research in Natural Language Understanding at the University of Toronto
Graeme Hirst
February 1989

An overview of current and recent research in NLU, including a bibliography and abstracts of theses.

A Processing Model of Warlpiri Syntax and Implications for Linguistic Theory
Barbara Brunson
Technical report CSRI-208, August 1988

Much of the development of the current Government-Binding (GB) theory of syntax has progressed independently of concerns raised in theories of language processing. Similarly, models of syntactic processing are often proposed that lack any underpinning in syntactic theory. The work described in this report focuses on the language Warlpiri, an Australian aboriginal language with properties that are difficult to reconcile with most theories of Universal Grammar – properties such as free word order and discontinuity. This language is studied from the two-fold perspective of establishing a linguistically and computationally sound processing model. This forces the linguistic model to be sufficiently precise to satisfy the demands of implementation as well as forcing the implementation to proceed in a linguistically principled way.

This report presents a portion of Warlpiri grammar in a revised GB-based account, addressing the issues of parsability, as well as more theoretical syntactic issues, that together force a reassessment and parametrization of certain linguistic principles. In particular, a revised version of theta theory and the notion of thematic identification are readily interpreted into processing strategies that extend naturally to deal with adjuncts and non-subcategorized arguments in a wide range of languages. The complementary nature of the syntax and morpho-syntax in the satisfaction of syntactic principles as well as in the construction of syntactic representations is addressed, as is the crucial relevance of prosodic information for preserving determinism in the parsing algorithm.

The Meaning Triangle as a Tool for the Acquisition of Abstract, Conceptual Knowledge
Stephen Regoczei and Graeme Hirst
Technical report CSRI-211, May 1988

The meaning triangle is presented as a useful diagramming tool for organizing knowledge in the informant-analyst interaction-based, natural language-mediated knowledge acquisition process.

In concepts-oriented knowledge acquisition, the knowledge explication phase dominates. During the conceptual analysis process, it is helpful to separate verbal, conceptual, and referent entities. Diagramming these entities on an agent-centered meaning triangle clarifies for both informant and analyst the ontological structure that underlies the discourse and the creation of discourse domains.

The Realization of Natural Language with Pragmatic Effects
Murray Watt
Technical report CSRI-215, August 1988

Realization is the process that takes a conceptual representation of a text to a linguistic representation. Realization is not a simple mapping from a text plan to natural language, because it adds pragmatic content to the text. This thesis includes a principled account of what aspects of meaning should be derived from the text plan and what should be derived from control in realization. From this account a model of realization is developed that distinguishes between texts with the intent of generating an optimal realization. The model is implemented in a system called Eloquence that organizes the choices available to the realization process and measures the appropriateness of texts for a given set of pragmatic goals. Eloquence separates text planning from realization. However, processing constraints apply when most appropriate and do not necessarily reflect module boundaries.

Generating Lexical Options by Matching in a Knowledge Base
Mara Miezitis
Technical report CSRI-217, October 1988

The goal of this thesis is to explore the development of LOG, a Lexical Option Generator, to provide *different* lexical options that may be used to express the *same* information. In addition to syntactic information, LOG provides pragmatic and stylistic information with these options. The lexical units supplied by LOG may be words or phrases – including, possibly *idioms*. LOG requires the *matching* of input information to a knowledge base of *situations* specifying the state of affairs that license the use of a particular lexical unit in output text. A representative survey of matching techniques is given, highlighting those aspects of the matchers desirable for LOG. The matching technique employed by LOG, using a *magnetization* process for directing information to situations likely to match, is presented in detail.

Intelligent Diagnosis of Ungrammaticality In Computer-Assisted Language Instruction
Mark Catt
Technical report CSRI-218, October 1988

We describe an approach to grammatical error diagnosis in computer-assisted language instruction (CALI). Our prototype system, Scripsi, employs a model of the linguistic competence of the second language learner in diagnosing ungrammaticality in learners' writing. Scripsi not only detects errors, but hypothesizes their cause and provides corrective information to the student.

Scripsi's grammatical model reflects the results of research in second language acquisition, which has identified *language transfer* and *rule overgeneralistion* as the chief sources of error in learner language. Thus, in characterizing the learner's "transitional competence", we model not only the grammar of the learner's native language, but also the strategies that give rise to overgeneralization. Although the approach is language-independent, our implementation targets French-

speaking and Chinese-speaking learners of English.

The computational realization of the model assumes that linguistic behaviour is *rule-governed*. We have adopted a rule-oriented grammatical formalism in which the processes of transfer and overgeneralization are readily interpreted. Linguistic rules are expressed in a feature-based grammatical framework closely related to the Standard Theory of transformational grammar. We have extended the shift-reduce parsing algorithm in order to accommodate context-sensitive and transformational aspects of the formalism.

We argue that the development of expertise in intelligent grammatical diagnosis is a prerequisite for the next generation of CALI tools – genuinely communicative systems capable of interacting linguistically with the student.

Lexical Cohesion, the Thesaurus, and the Structure of Text
Jane Morris
Technical report CSRI-219 October 1988

In text, lexical cohesion is the result of chains of related words that contribute to the continuity of lexical meaning. These lexical chains are a direct result of units of text being about the same thing. Finding text structure involves finding units of text that are about the same thing. Hence, computing the chains is useful since they will have a correspondence to the structure of the text. Determining the structure of text is an essential step in determining the deep meaning of the text. In this thesis, a thesaurus is used as the major knowledge base for computing lexical chains. Correspondences between lexical chains and structural elements are shown to exist. Since the lexical chains are computable, and exist in non midomain-specific text, they provide a valuable indicator of text structure.

On 'Extracting Knowledge from Text': Modelling the Architecture of Language Users
Stephen Regoczei and Graeme Hirst
Technical report CSRI-225, January 1989

We propose a unified view of natural language understanding and knowledge acquisition. Knowledge is not "extracted" from a text, but rather is *added* to the text by a "cogniting" agent. The text, and whatever is contained in it, serves only as a triggering mechanism. This process of addition is *concept cluster attachment*. This can be generalized to artificial notations such as mathematical formulas and diagrams (even Rorschach tests!) and general signing such as facial expressions and gestures. We develop a minimal three-level architecture for a cogniting agent, consisting of verbal, conceptual, and subconceptual levels.

We further propose that natural language understanding and knowledge acquisition from text require expertise. We discuss how this expertise may be acquired and incorporated into an expert system and incrementally build up the architecture of the theoretical version of such an expert system, which we call LUKES. We discuss its implementation as LOGOS using the sortal analysis tool SORTAL.

Ontological Assumptions in Knowledge Representation
Graeme Hirst
From the Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning, Toronto, May 1989

If knowledge representation formalisms are to be suitable for semantic interpretation of natural language, they must be more adept with representations of existence and non-existence than they presently are. I review the philosophical background and exhibit some ontological problems for KR. I then look at the shortcomings of current approaches, including several intentional formalisms and the work of Hobbs. The Meinongian theory of Parsons is considered. Lastly, I present a naive ontology for knowledge representation, identifying about nine distinct kinds of existence.

Ordering information:

Requests for free copies of any of the forenamed publications should be addressed to: Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Ontario M5S 1A4

L'échéance pour le numéro d'octobre est le 15 août

Deadline for the October issue is 15 August

Please complete the readers survey in April '89 issue

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

2558 Putting into practice Moreau's extended generalized modus ponens
A.M. Petieau, A. Moreau, D. Willaeys
Lab. d'Automat. Ind. et Valenciennes et du Hainaut-Cambresis, France
Uncertainty and Intelligent Systems. *Proceedings of the 2nd International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems: IPMU '88*, Urbino, Italy, 4-7 July 1988 (Berlin, W.Germany: Springer-Verlag 1988), p. 373-84

Whereas uncertainty is often treated by expert systems, imprecision is not in current use, though it has already been studied and tools exist to treat it. The authors approach the extended generalized modus ponens, which treats simultaneously imprecision and uncertainty of facts and rules, from a theoretical point of view. They look at its implementation and study the results. (10 refs.)

2779 Higher order probabilities and intervals.
H. E. Kyburg, Jr.
Dept. of Philos., Rochester Univ., NY, USA
Int. J. Approx. Reason. (USA), vol. 2, no. 3, p. 195 - 209 (July 1988)

Many researchers have felt uncomfortable with the precision of degrees of belief that seems to be demanded by the subjective Bayesian treatment of uncertainty. Various responses have been suggested. The most common one has been to incorporate higher order probabilities in systems that reason in beliefs. These probabilities concern statements of first-order probability. Thus a first-order probability (e.g., the probability of heads on the next toss of this coin is 1/2) is the subject of a second-order probability; for example, the probability is .9 that the probability of heads on the next toss of this coin is 1/2. This approach is explored and is found to be epistemologically wanting, although there are important intuitions about beliefs that are captured by it. Furthermore, this approach may, in some circumstances, be computationally attractive. The author also briefly explores a number of other approaches, including taking probabilities to be intervals and construing probability values as fuzzy sets. (22 refs.)

2781 Choosing uncertainty representations in artificial intelligence.
T. S. Levitt
Adv. Decision Syst., Mountain View, CA, USA
Int. J. Approx. Reason. (USA), vol. 2, no. 3, p. 217 - 32 (July 1988)

Research in automation in machine vision, robotic planning, medical diagnosis, and many other fields gives rise to sources of uncertainty in inference and reasoning that are beyond conventional notions of measurement error. For a given domain of interest, an uncertainty representation is defined such that statements in the representation language model states of the domain and can provide problem solutions. Five criteria are proposed for selecting uncertainty representations for artificial intelligence: mathematical soundness, efficacy of domain models, appropriate representation languages, efficient computation, and choice of control mechanism. These issues are explored in an application of determining a building's location from multiple sensor returns. Finally, choice of uncertainty representations is seen to depend on choices for uncertain inference and control. (10 refs.)

2795 Decision trees as probabilistic classifiers
J. R. Quinlan
Sch. of Comput. Sci., New South Wales Inst. of Technol., Sydney, NSW, Australia
Proceedings of the Fourth International Workshop on Machine Learning, Irvine, CA, USA, 22- 25 June 1987 (Los Attos, CA, USA: Morgan Kaufmann 1987), p. 31 - 7

Decision trees are a widely known formalism for expressing classification knowledge and yet their straightforward use can be criticized on several grounds. Results are categorical and so do not convey potential uncertainties in classification. Small changes in the attribute values of a case being classified may result in sudden and inappropriate changes to the assigned class. Missing or imprecise information may apparently prevent a case being classified at all. The paper outlines extensions to the way a case is classified by a decision tree that address these shortcomings. (6 refs.)

2824 The craft of engineering knowledge: some practical insights
K. D. Bimson, L. B. Burris
Lockheed Software Technol. Center, Austin, TX, USA
Proceedings - 2nd International Symposium on Artificial Intelligence and Expert Systems, Berlin, West Germany, 20 - 24 June 1988 (Berlin, West Germany: AMK Berlin 1988), p. 201 - 14 vol. 1

Knowledge engineering is a highly creative business. As such, it tends to be more heuristic than algorithmic, more technique than technology. People learn the techniques of knowledge engineering by practice. A case study of this process is presented, tracing the knowledge engineering work performed in conjunction with the building of a knowledge-based software project management system. Knowledge engineering is defined to be a cyclical, phased approach to capturing the knowledge of a domain expert in an automated system. The phases of the knowledge engineering life cycle include the ways in which knowledge is gathered, analyzed, represented, and implemented in a prototype. The transformation of project management knowledge is described as it evolves through this life cycle, from knowledge expressed by management experts to implemented knowledge base structures. Emphasis is placed on the first phase, knowledge gathering and analysis. (9 refs.)

8 How humans process uncertain knowledge: an introduction for knowledge engineers
R. F. Hink, D. L. Woods
AI Mag. (USA), vol. 8, no. 3, p. 41-53 (Fall 1987). [received: 31 Aug 1988]

The question of how humans process uncertain information is important to the development of knowledge-based systems in terms of both knowledge acquisition and knowledge representation. This article reviews three bodies of psychological research that address this question: human perception, human probabilistic and statistical judgement, and human choice behaviour. The general conclusion is that human behaviour under uncertainty is often suboptimal and sometimes even fallacious. Suggestions for knowledge engineers in detecting and obviating such errors are discussed. The requirements for a system designed to reduce the effects of human factors in the processing of uncertain knowledge are introduced. (71 refs.)

2825 The problem of extracting the knowledge of experts
from the perspective of experimental psychology
R. R. Hoffman
Proceedings - 2nd International Symposium on Artificial Intelligence and Expert Systems, Berlin, West Germany, 20 - 24 June 1988 (Berlin, West Germany: AMK Berlin 1988), p. 215 - 29 vol. 1

The first step in the development of an expert system is the extraction and characterization of the knowledge and skills of an expert. This step is widely regarded as the major bottleneck in the system development process. To assist knowledge engineers and others who might be interested in the development of an expert system, the author presents a working classification of methods for extracting an expert's knowledge, some ideas about the types of data that the methods yield, and a set of criteria by which the methods can be compared relative to the needs of the system developer. The discussion highlights certain issues, including the contrast between the empirical approach taken by experimental psychologists and the formalism-oriented approach that is generally taken by cognitive scientists. (71 refs.)

2827 Limits of Logic
T. Vamos
Proceedings - 2nd International Symposium on Artificial Intelligence and Expert Systems, Berlin, West Germany, 20 - 24 June 1988 (Berlin, West Germany: AMK Berlin 1988), p. 25 - 62 vol. 2

The responsibility of a clear view about the possibilities of logical computer reasoning formulated by humans has grown to an immense intellectual task. The author examines developments in the logic of Aristotle, and of the "Age of Reason", and argues that the retrospective-historical view is much more than a demonstration of erudition; it has very practical consequences. The first is an obligation to be more modest, careful with fancy renamings of old ideas, a compulsion for deeper thinking about the mechanisms of thinking. The second, which follows from the first, is a caution to investigate the limits of general methods and theorems, to elaborate more real-life problems and analyze their results and applicability against the background of fashionable overstatements and to derive stimuli for investigation from these experiences. The ideas arising in an analysis of modern logic are often the same as those found in Aristotle, only in a clearer form. (58 refs.)

2829 Metalogic machines: a comparison of concurrent logic languages
G. A. Ringwood
Dept. of Comput., Imperial Coll., London, UK
Proceedings - 2nd International Symposium on Artificial Intelligence and Expert Systems, Berlin, West Germany, 20 - 24 June 1988 (Berlin, West Germany: AMK Berlin 1988), p. 113 - 46 vol. 2

The unsuitability of von Neumann architectures for AI applications has prompted the design of specialist parallel AI machines. Both top-down and bottom-up design methodologies have drawbacks. The middle-out strategy, working both up and down from a machine language as a means of representing and processing of knowledge, provides a symbiosis of software and hardware. The longest established and most well-founded method for the representation and manipulation of kowledge is logic. A notable result of the last decade's work on mechanical theorem proving was that a subset of predicate logic, namely Horn clauses, can form the foundation of a programming language. However, the execution model of PROLOG, the first language based on Horn clauses, was designed for efficient evaluation on von Neumann machines. An alternate execution model has given rise to a new class of Horn clause languages, concurrent logic languages, for which a new generation of parallel machines is being built. After describing the deviation of this generation from PROLOG, a comparison of its notable members is given. (30 refs.)

2832 Sixth generation computing by the year 2001: characteristics, opportunities, and challenge
T. Gomi
Appl. AI Syst. Inc. Ottawa, Ont., Canada
Proceedings - 2nd International Symposium on Artificial Intelligence and Expert Systems, Berlin, West Germany, 20 - 24 June 1988 (Berlin, West Germany: AMK Berlin 1988), p. 211 - 312 vol. 2

The paper discusses the emerging computing technology now loosely called sixth generation (6G) computing. The discussion includes tentative definitions of the concept, opportunities promised, and the technical and social challenges foreseen. Topics investigated include the Japanese human frontier science program and the advanced telecommunications research project as its first implementation, and the European Commission's BRAIN project. Topics also include the emerging neural networks or neural computer concept actively pursued in the United States, Europe, and Japan as a subfield of artificial intelligence. In addition, several of the advanced concepts in the existing national and international fifth generation computer R&D programs border those in sixth generation computers. Some examples of such advanced concepts found in the strategic computing initiative, ESPRIT, and Alvey are also surveyed. Challenges put forth by sixth generation computing are expected to impact almost all aspects of human life at a greater rate than that of fifth generation computing. Predictions are made on some of the more important implications. (52 refs.)

9 Thinking backward for knowledge acquisition
R. D. Shachter, D. E. Heckerman
AI Mag. (USA), vol. 8, no. 3, p. 55-61 (Fall 1987). [received: 31 Aug 1988]

This article explains the direction in which knowledge bases are constructed for diagnosis and decision making. When building an expert system, it is traditional to elicit knowledge from an expert in the direction in which the knowledge is to be applied, namely, from observable evidence toward unobservable hypotheses. However, experts usually find it simpler to reason in the opposite direction — from hypotheses to observable evidence — because this direction reflects casual relationships. Therefore, it is argued that a knowledge base should be constructed following the expert's natural reasoning direction, and then the direction facilitates knowledge acquisition in deterministic domains and is essential when a problem involves uncertainty. This concept is illustrated with influence diagrams, a methodology for graphically representing a joint probability distribution. Influence diagrams provide a practical means by which an expert can characterize the qualitative and quantitative relationships among evidence and hypotheses in the appropriate direction. (16 refs.)

25 Sensor data fusion
L. F. Paul
Tech. Univ., Lyngby, Denmark
J. Intell. Robot. Syst., Theory Appl. (Netherlands), vol. 1, no. 2, p. 103-16 (1988)

This article reviews some knowledge representation approaches devoted to the sensor fusion problem, as encountered whenever images, signals, or text must be combined to provide the input to a controller or to an inference procedure. The basic steps involved in the derivation of the knowledge representation scheme are: locate a representation based on exogeneous context information; compare two representations to find out if they refer to the same object/entity; merge sensor-based features from the various representations of the same object into a new set of features or attributes; and aggregate the representations into a joint fusion representation, usually more abstract than each of the sensor-related representations. The importance of sensor fusion stems first from the fact that it is generally correct to assume that improvements in control law simplicity and robustness, as well as better classification results, can be achieved by combining diverse information sources. The second element is that spatially distributed sensing, or otherwise diverse sensing, does indeed require fusion as well. (25 refs.)

11 A tentative comparison of numerical approximate reasoning methodologies
D. Dubois, H. Prade
Univ. Paul Sabatier, Toulouse, France
Int. J. Man-Mach. Stud. (UK), vol. 27, no. 5-6, p. 717-28 (Nov - Dec. 1987) [received: 07 Sep 1988] (Cognitive Engineering in Dynamic Worlds, Ispra, Italy, 11-14 Nov. 1986)

A critical discussion of approximate reasoning methods in artificial intelligence is proposed. The focus is on numerical approaches based on certainty factor, probability, possibility or evidence theory. The discussion is organized around three topics, namely, knowledge representation, inductive versus deductive reasoning, and control strategies. The authors attempt to outline a tentative classification of emerging trends in uncertain reasoning and to point out problems which are not solved yet or are sometimes overlooked by proponents of a single approach. (48 refs.)

20 Monitoring performance of devices using a coupled probability-possibility method
M. Ragheb, L. Tsoukalas
Illinois Univ., Urbana, IL, USA
Int. J. Expert Syst. Res. Appl. (USA), vol. 1, no. 2, p. 111-30 (1988)

A methodology is proposed for inferencing and decision making under uncertainty in rule-based systems where the knowledge about a given system is both probabilistic and possibilistic. In this approach, uncertainty is considered as consisting of two parts: randomness, meaning the uncertainty of occurrence of an object; and fuzziness, expressing the imprecision of the meaning of the object. The analysis uses the concepts of the probability of a fuzzy event and of information granularity. Application to the monitoring of the performance of engineering devices is demonstrated, where the propagation of the coupled probabilistic and possibilistic uncertainty is carried out over a model-based system using the rule-based paradigm. The approach provides a quantification for combining the concepts of reliability from probability theory, and of performance level from the theory of fuzzy subsets. Applicability of the methodology is demonstrated for a manufacturing system and a nuclear reactor's engineering safety feature. (13 refs.)

284 A new normative theory of probablistic logic
R. A. Aleliunas
Simon Fraser Univ., Burnaby, BC, Canada
Proceedings of the Seventh Biennial Conference of the Canadian Society for Computational Studies of Intelligence, Edmonton, Alta., Canada, 6-10 June 1988 (Toronto, Ont., Canada: Univ. Alberta Press 1988), p. 67 - 74

By probabilistic logic it means a normative theory of belief that explains how a body of evidence affects one's degree of belief in a possible hypothesis. A new axiomatization of such a theory is presented which avoids a finite additivity axiom, yet which retains many useful inference rules. Many of the examples of this theory do not use numerical probabilities. The author indicates the general reasons why recent theories of decision making and belief (nonmonotonic logic, the Dempster-Shafer theory, etc.) do not represent anything really new. (23 refs.)

316 Hybrid approaches [pattern recognition]
H. Bunke
Inst. fur Inf. und Angewandte Math., Bern Univ., Switzerland
Syntactic and Structural Pattern Recognition. Proceedings of the NATO Advanced Research Workshop, Sitges, Spain, 23-25 Oct. 1986 (Berlin, West Germany, Springer-Verlag 1988), p. 335-61

Most of the pattern recognition techniques developed fall into one of three major categories: statistical (or decision theoretic), structural, and artificial intelligence based approaches. Each of the different methods has its strength and its limitations. For overcoming these limitations, the three methods are sometimes mixed, resulting in a hybrid approach. The author discusses various pattern recognition methods with a particular emphasis on the question of how different methods are related with each other and how they can be combined into a hybrid approach. Decision theoretic methods, syntactic methods, structural prototypes, and relaxation are considered. Three important approaches to knowledge representation are discussed, namely formal logic, production systems, and semantic nets. Further topics on search, control, and system organisation are addressed. (88 refs.)

2.0 SYSTEMS AND TECHNIQUES

2056 Trends in Information Technologies and their Implications for Mental Health Research
L. Travis
Wisconsin U., Madison, Wisconsin, USA
Comput. Hum. Serv. (USA), vol. 2, nos. 3-4, p. 37-59, Fall-Winter 1987. [received: 13 June 1988] (Working Conference on Mental Health Computer Applications Research, Madison, WI, USA, 10-11 June 1985)

Some of the main developments in computer science and computer technology that will influence — or should influence — mental-health research are identified. These information technologies are: computer-generated graphics and image files for computer test administration; sophisticated PC systems for mental health-related instruction; computer communication networks; AI-developed natural language analysis techniques for narrative clinical progress rates; and expert system development for various mental health applications. (25 refs.)

2855 Experiments in knowledge elicitation [Alvey project]
M. Burton, N. Shadbolt
Dept. of Psychol., Nottingham Univ., UK
AISB Q. (UK), no. 65, p. 11 - 12 (Summer 1988)

Knowledge elicitation for IKBS is a considerable problem. Difficulties include a range of cognitive features of human memory which contrive to frustrate the knowledge engineer. In order to aid access to various forms of expertise, a number of knowledge elicitation techniques are available. These were evaluated by performing laboratory-based experiments and real-world case studies. The particular techniques which have been studied are: structured interviews; protocol analysis; card sorts; and laddered grids. The first two techniques are well known and probably the most popular KE techniques. The second two methods are contrived techniques, adapted from psychology, which aim to provide a "conceptual map" of

the expert's knowledge without explicitly asking for anything of the sort. Results of the study and further project work is discussed.

2857 Design methods in three Alvey expert systems clubs
H. Roberts
Manchester Bus. Sch., UK
AISB Q. (UK), no. 65, p. 18 (Summer 1988)

A study of three Alvey Expert Systems Awareness Clubs confirms that the clubs have been successful in deploying IKBS techniques in their fields of common interest. However, there are important differences between design methods for systems built as awareness exercises and those used for operational applications. Moreover, these differences appear to be largely irreconcilable within the context of precompetitive research. As a result, the methods of development used in the Alvey clubs are limited and require extension to make them more suitable for use within industry as a whole. The clubs studied were Aries (insurance), Dapes (data processing) and Planit (Planning). Each club funded a software house to build a prototype expert system in their field of interest. The systems built by Aries were an underwriting system and an equity selection system. Dapes developed two systems to help with the control of computer networks. Planit constructed an interactive planners assistant. (1 ref.)

43 Knowledge acquisition approaches in expert systems development
J. Liebowitz
George Washington Univ., Washington, DC, USA
Interface, Comput. Educ. Q. (USA), vol. 10, no. 2, p. 13-17 (Summer 1988)

Knowledge elicitation/acquisition is the biggest bottleneck in expert systems development and there are several reasons for this. In order to overcome these problems, various knowledge acquisition techniques are being used and developed. These methods include interviewing and automated knowledge acquisition techniques. The author surveys both types and focuses on one particular approach — the rating/repertory grid. The Expertise Transfer System (ETS) automated this method.

55 Surveying projects on intelligent dialogue
J. A. Mason
Dept. of Comput. Sci. & Math., Atkinson Coll., York Univ., North York, Ont., Canada), J. L. Edwards.
Int. J. Man-Mach. Stud. (UK), vol. 28, no. 2-3, p. 259-307 (Feb-March 1988). [received: 21 Oct 1988]

Four projects developing intelligent dialogue systems are surveyed, including work at Bolt Beranek and Newman, at Carnegie Mellon University (the XCALIBUR project), at the University of Hamburg (the HAM-ANS project), and at SRI International (the KLAUS project). The projects are compared using an evaluation method involving eight aspects of intelligent dialogue: control, models, connectivity, modality, form, knowledge representation and inferencing, knowledge acquisition, and external information sources and targets. The evaluation method used is proposed as a standard method for comparing and rating intelligent dialogue systems. (30 refs.)

74 Artificial intelligence techniques in man-machine communication
J. -P. Haton, M. -C. Haton
CRIN/INRIA, Vandoeuvre-les-Nancy, France). Comput. Stand. Interfaces (Netherlands), vol. 8, no. 1, p. 37-40 (1988)
Symposium on Standards and Interfaces in Personal Computing, Brussels, Belgium, May 1987)

The development of new communication media between man and machine (e.g., speech or vision) constitutes one of the major challenges of artificial intelligence. The paper addresses the present issues of the field, especially in relation to the techniques of knowledge-based systems: expert systems in interpretation problems, object-oriented representations of knowledge, and multi-knowledge-based architectures. It is illustrated by examples from the fields of speech, image and signal understanding. (6 refs.)

78 Personal Consultant Easy and Personal Consultant Plus [artificial intelligence packages]
R. Sun
Dept. of Comput. Sci. & Inf. Sci., Iona Coll., New Rochelle, NY, USA
Intell. Instrum. Comput. (USA), vol. 6, no. 5, p. 133-8 (May 1988).

Reviews the AI packages Personal Consultant Easy and Personal Consultant Plus from Texas Instrument. These are expert system shells introduced by TI in August 1984 but are only now gaining popularity. PC Easy is a low-cost tool for fast prototyping. PC Plus, on the other hand, is larger and more powerful, aiming at real applications of reasonable size, to be used on the IBM PC AT and above (or compatible). The knowledge base created with PC Easy is upwardly compatible with PC Plus. This allows users to upgrade their software as their needs grow. Both systems are for development purposes. The end user can interact with the finished system through a run time diskette, which is optional from TI. (3 refs.)

100 Acquiring design and analysis knowledge for knowledge-based systems
J. H. Boose, J. M. Bradshaw
Boeing Adv. Technol. Center, Boeing Comput. Services, Seattle, WA, USA
Intelligent CAD Systems I: Theoretical and Methodological Aspects, Noordwijkerhout, Netherlands, 21-24 April 1987 (Berlin, West Germany: Springer-Verlag 1987), p. 128-45

A workbench for acquiring design and analysis knowledge to build knowledge-based systems is presented. Eliciting and modeling such knowledge from a human designer is a major problem when building knowledge-based systems for CAD problems. Aquinas, an extended version of the Expertise Transfer System (ETS), combines ideas from psychology and knowledge-based systems research. Aquinas interviews design and analysis experts and helps them model, analyze, test, and refine their knowledge. Expertise from multiple designers or other knowledge sources can be represented and used separately or combined. User consultations are directed by propagating information through hierarchies. Aquinas delivers knowledge by creating knowledge bases for several different expert system shells. Help is given to the expert by a dialog manager that embodies knowledge acquisition heuristics. Aquinas contains many techniques and tools for expertise transfer; the techniques combine to make it a powerful testbed for rapidly prototyping portions of many kinds of complex knowledge-based systems. (24 refs.)

323 The development of an intelligent interface to a computational fluid dynamics flow-solver code
A. D. Williams
NASA, Lewis Res. Center, Cleveland, OH, USA
Comput. Struct. (UK), vol. 30, no. 1-2, p. 431-8 (1988). (Symposium on Advances and Trends in Computational Structural Mechanics and Fluid Dynamics, Washington, DC, USA, 17-19 Oct. 1988)

Researchers at the NASA Lewis Research Center are currently developing an "intelligent" interface to aid in the development and use of large computational fluid dynamics flow-solver codes for studying the internal fluid behaviour of aerospace propulsion systems. The paper discusses the requirements, design, and implementation of an intelligent interface to Proteus, a general purpose, three-dimensional, Navier-Stokes flow solver. The interface is called PROTAIS to denote its introduction of artificial intelligence (AI) concepts to the Proteus code. (22 refs.)

329 The knowledge acquisition bottleneck: time for reassessment?
J. Cullen, A. Bryman
Dept. of Social Sci., Loughborough Univ. of Technol., UK
Expert Syst. (UK), vol. 5, no. 3, p. 216-25 (Aug. 1988)

Knowledge acquisition has long been considered to be the major constraint in the development of expert systems. Conventional wisdom also maintains that the major problem encountered in knowledge acquisition is in identifying the varying structures and characteristics of domain knowledge and matching these to suitable acquisition techniques. With the aid of the first substantial systematic analysis of a sample of expert systems applications developed in the real world, the authors describe what is actually going on in terms of knowledge acquisition. In the light of the evidence, it is argued that a reappraisal of the conventional approach to knowledge acquisition is necessary.

331 The validation of expert systems — contrasts with traditional methods
P. N. Finlay, G. J. Forsey, J. M. Wilson
Dept. of Manage. Studies, Loughborough Univ. of Technol., UK
J. Oper. Res. Soc. (UK), vol. 39, no. 10, p. 933-8 (Oct. 1988)

Recent research has demonstrated that little validation is carried out during and after the building of operational research models for decision support. In a subject area that has been of great interest for more than a decade, this reflects the fact that validation is not yet firmly embedded into a methodology of mathematical modelling that underpins conventional decision support systems. With the interest now being shown in expert systems, it is important to devise an appropriate methodology of validation for such systems. This paper contributes to this development by describing the types of validation that should be performed with decision support systems and compares and contrasts the features found in the validation of mathematical models with those required of validation of expert systems. (19 refs.)

342 FLOSEL: expert selection of flowmeters
J. E. Lycett, T. Porter, D. Maudsley
Div. of Instrum. & Control., Sch. of Inf. Eng., Teesside Polytech., UK
Eng. Appl. Artif. Intell. (UK), vol. 1, no. 1, p. 37-40 (March 1988) [received: 18 Oct 1988]

The development of an expert system as a decision support tool for the instrument design engineer engaged in the specification of process instrumentation is considered. The paper describes how a system was built for decision support in the selection of flowmeters. The model, of particular use in front-end specification, uses fundamental knowledge given by the Lomas method for flowmeter selection and further detailed knowledge from many sources including standards, company procedures and domain expert opinion. Two methods of implementation are examined and the paper discusses how this front-end module may be incorporated in a larger system. (9 refs.)

350 A hybrid approach to finding language errors and program equivalence in an automated advisor
Xueming Huang, G. I. McCalla
Dept. of Comput. Sci., Saskatchewan Univ., Saskatoon, Sask., Canada
Proceedings of the Seventh Biennial Conference of the Canadian Society for Computational Studies of Intelligence, Edmonton, Alta., Canada, 6-10 June 1988 (Toronto, Ont., Canada: Univ. Alberta Press 1988), p. 161-8

This article presents a model of a language expert which recognizes language bugs and determines program equivalence in the context of the SCENT programming advisor being developed at the University of Saskatchewan. A hybrid approach, which combines the advantages of knowledge-based and theorem proving techniques, has been used in the language expert. Knowledge-based debugging is based on a debugging graph which is intended to recognize common program constructs and common bugs which frequently appear in students' programs. Empirical studies of the performance of the language expert show that it has achieved high performance in detecting and identifying language errors. (17 refs.)

351 Knowledge acquisition techniques for knowledge-based systems
M. L. Shaw
Dept. of Comput. Sci., Calgary Univ., Alta., Canada
Proceedings of the Seventh Biennial Conference of the Canadian Society for Computational Studies of Intelligence, Edmonton, Alta., Canada, 6-10 June 1988 (Toronto, Ont., Canada: Univ. Alberta Press 1988), p 169 - 76

This paper introduces the concept of knowledge support systems as the integration of interactive knowledge acquisition systems and expert systems shells. A prototype knowledge support system that has been implemented is described with examples of some of the knowledge acquisition and application tools provided. It includes knowledge acquisition and transfer through interactive repertory grid elicitation; exchange techniques for understanding and agreement between experts; and different forms of analysis that have been widely used for comparing data between experts and for rapid prototyping of expert systems. An evaluation model for knowledge support systems is presented, and the results of initial validation studies are reported, showing the extent to which experts agree with each other and with themselves at a later date. (30 refs.)

370 Expert system in computer manufacture
B. Pogson
Digital Equipment Corp., London, UK
Factory 2000: Integrating Information and Material Flow (Publ. No. 80), Cambridge, UK, 31 Aug. - 2 Sept. 1988 (London, UK: IERE 1988), p. 195-202

Digital Equipment Corporation manufactures a wide range of VAX computers, each configurable from a wide range of components. From the initial processing of an order, through the configuration of the machine to the sourcing of its components, a variety of expert systems, linked in a "knowledge network" cooperate with inventory and product description databases to control the manufacturing process. The configuration system, XCON, is the world's first and largest commercially successful expert system, and its history is summarized. A methodology for the management of expert system projects has grown out of the experiences with over fifty expert system development projects within Digital. This ten-stage methodology uses prototyping to incrementally build expert systems under continuous review.

3.0 APPLICATIONS

2467 Review of Expert Systems Technology
E. Turban
Dept. of Syst. Sci., U. of S. California,
Los Angeles, California, USA
IEEE Trans. Eng. Manage. (USA),
vol. 35, no. 2, p. 71-81, May 1988

An overview of expert systems technology is provided from a development point of view. The intent is to provide assistance in the practical aspects of constructing expert systems. Emphasis is placed on the benefits and limitations of expert systems, the various phases of the development processes, and the software tools used to expedite the development. Special

attention is given to the interface with engineering management. (20 refs.)

2698 Expert Systems in Japan
K. Ishii, S. Hayami
UPU Inc., Tokyo, Japan
IEEE Expert (USA), vol.3, no.2, p. 69-74, Summer 1988

A background and overview of AI research in Japan is provided. Industrial AI research in the steel, construction, and energy-related industries is examined, revealing a trend toward addressing design and productivity, as opposed to the diagnosis and cost-saving approach dominating U.S. expert systems.

2897 Developing an expert system for financial analysis.
E. Blocher ,
North Carolina Univ., Raleigh, NC, USA
K. A. Scalf.
Int. J. Policy Inf. (Taiwan), vol. 12, no. 1, p. 101 - 15 (15 June 1988)

This article describes the development of an expert financial analysis system, ANSWERS, which is designed to facilitate the review and analysis of financial statement data by accountants, auditors, managers, analysts, business owners, and financial advisors. The system reviews financial data and identifies unusual trends or relationships using financial statement ratios, reasonableness tests, and statistical projections. Any signficant trends or relationships are interpreted for the user, and expert advice is provided concerning the causes and/or effects associated with it. (7 refs.)

2948 PACIES: a part code identification expert system
Yu-Tong Chen
GM Tech. Center, Warrent, MI, USA
R. E. Young
IIE Trans. (USA), vol. 20, no. 2, p. 132 - 6 (June 1988)

This paper discusses an expert system to generate part codes for the selection of small part feeding and orienting devices for use in automatic assembly. The PACIES system embodies the specific knowledge and experience of human experts to determine a part code consistent with the UMASS coding system to facilitate the handling and orienting of small parts for use in automatic assembly. The system is implemented on an IBM/PC class microcomputer running MS-DOS using the micro-PROLOG programming language and consists of a rule set containing 163 rules. (9 refs.)

2949 Dynamic scheduling in cellular manufacturing systems: a framework for networked decision making
M. J. Shaw
Illinois Univ., Urbana-Champaign, IL, USA
J. Manuf. Syst. (USA), vol. 7, no. 2, p. 83-94 (1988)

This paper describes a dynamic scheduling method for cellular manufacturing systems. Essentially an artificial intelligence-based approach for the networking environment, the method is adaptive to the changes in the manufacturing environment and can take into account such information as loading factor, unexpected breakdowns, and new job arrivals. The major features of the method are a distributed task assignment mechanism executed through the communication network for intercell scheduling and a knowledge-based system for cell level scheduling. A main advantage of such a method is that the scheduling on both levels can be dynamically executed in real time. (25 refs.)

2950 Multi-pass expert control system — a control/ scheduling structure for flexible manufacturingcells
S. Y. D. Wu, R. A. Wysk
Lehigh Univ., Bethlehem, PA, USA
J. Manuf. Syst. (USA), vol. 7, no. 2, p. 107 - 20 (1988)

A control structure is outlined which takes advantage of both expert system technology and discrete event simulation. The simulation is used as a prediction mechanism to evaluate several possible control alternatives provided by the expert system. A performance measure is obtained from the simulation for each of the suggested alternatives. A control effector is then employed to affect the physical control of the cell based on the measure. The integration of the expert control system, the simulation, and the control effectors form a system called a multi-pass expert control system (MPECS). MPECS is designed for the control and scheduling of flexible manufacturing cells. Experiments to evaluate the performance of MPECS and have yielded advantages of between 2.3% and 29.3%, compared to single-pass, single-rule priority dispatching and multipass, multi-rule dispatching procedures. (20 refs.)

117 Machinery selection using expert systems and linear programming
D. E. Kline, D. A. Bender, B. A. McCarl, C. E. Van Donge
Texas A & M Univ., College Station, TX, USA
Comput. Electron. Agric. (Netherlands), vol. 3, no. 1, p. 45-61 (Sept. 1988)

Information on a Farm-level INtelligent Decision Support system, FINDS, is presented. This system is designed to aid in farm machinery selection through its use and interpretation of a farm management linear program (LP) which has been used in extension workshops to assist farm producers with farm-level decisions. Considerable expertise is required to interpret this farm management LP and it has not been used to its potential because there is a shortage of experts to work with individual producers. FINDS incorporates the expertise needed to use the LP for selecting farm machinery. The goal of FINDS is to help the user find machinery sets that can increase farm profit. The knowledge base of FINDS was structured using IF-THEN rules and LISP functions. With the rules and functions, FINDS automatically runs the LP model, extracts and interprets the relevant output, and recommends farm machinery changes that can increase farm profitability. (14 refs.)

2953 Sources and standards for computerized materials, property data and intelligent knowledge systems
J. G. Kaufman
Nat. Mater. Property Data Network, Columbus, OH, USA
Eng. Comput. (USA), vol. 4, no. 1-2, p. 75-85 (1988)

The explosion of interest in computerized material property databases and intelligent knowledge systems for flexible manufacturing processes has resulted in two significant developments: the incorporation of the National Materials Property Data Network, (MPD Network), and the formation of ASTM Committee E-49 on Computerization of Material Property Data. The MPD Network has the mission of providing engineers and scientists with easy on-line access to worldwide source of reliable, well-documented material property data. The status of its pilot network and plans for the future are outlined. ASTM Committee E-49 has the mission of developing and providing guidelines and standards to aid the builders, maintainers, and suppliers of databases in meeting the needs of industry and in establishing compatible and consistent sources capable of sharing data. The author also reports on the early progress of the Technical Committee of the MPD Network in establishing guidelines for the evaluation of data, recognizing that there are a number of aspects of data source building and management requiring audit as well as the data itself. A brief overview of other major material property data- base building activities in the United States is also provided. (27 refs.)

2989 NIDX— a real-time intrusion detection expert system
D. S. Bauer, M. E. Koblentz
Bell Commun. Res., Morristown, NJ, USA
Proceedings of the Summer 1988 USENIX Conference, San Francisco, CA, USA, 21-24 June 1988 (Berkeley, CA, USA: USENIX Assoc. 1988), p. 261-73

The design of the knowledge-based prototype Network Intrusion Detection Expert System (NIDX) for the Unix System V environment is described. NIDX combines knowledge describing the target system, history profiles of users' past activities, and intrusion detection heuristics forming a knowledge-based system capable of detecting specific violations that occur on the target system. Intrusions are detected by classifying user activity from a real-time audit trail of Unix system calls, then determining, using system-specific knowledge and heuristics about typical intrusions and attack techniques, whether or not the activity is an intrusion. The authors describe the NIDX knowledge base, Unix system audit trail mechanism and history profiles and demonstrate the knowledge-based intrusion detection process. (11 refs.)

143 Knowledge-based production scheduler reacts in real time
R. Reynolds
Autom. Div., Westinghouse Electr. Corp., Pittsburgh, PA, USA
I&CS (USA), vol.61, no. 7, p. 45-7 (July 1988)

Generating daily production schedules can be a very challenging proposition involving diverse and changing requirements. Factors such as resource availability, product delivery dates, production levels, and project priority all affect the generation of a daily schedule. This article concerns a project at Westinghouse Electric Corp. to build a computer-generated scheduling program, using a highly sophisticated expert system development tool, the Automated Reasoning Tool from Inference Corp. (Los Angeles, CA). The knowledge-based production scheduler, known as the master scheduling unit (MSU), is being developed for a Westinghouse customer who is building an automated aircraft parts manufacturing facility at an existing plant. The goal is to construct and monitor the execution of daily production schedules that reliably fulfil the requirements of the customer's manufacturing facility and revise and maintain the production schedules in response to changing shop conditions. It will also minimize the number of late orders while maximizing machine use.

155 Doc in a disk? [medical diagnostic software]
D. Essin
Univ. of Southern California Med. Center in Los Angeles, CA, USA
Database Searcher (USA), vol. 4, no. 7, p. 21-2 (July - Aug. 1988)

This paper reviews two medical diagnostic packages, Family Care and House Call. Both of these programs provide the layman with a computerized compendium or database of medical information that they can use at home to obtain information about health and medical care. Family Care concentrates on childhood health problems, while House Call covers a wide range of general medical topics, including some pediatrics. The former program breaks its knowledge base into four general categories: accidents, general, skin problems, and newborns. Family Care works on IBM PCs or Apple Macintoshes and costs $180. House Call, on the other hand, is more like a computerized version of a home health encyclopedia. The program's database consists of more than 400 information frames. Each one presents the symptoms, related problems, diagnosis (what tests, etc., normally establish the diagnosis), and usual treatment for a single medical problem. House Call costs around $80 and runs on either Apple IIs or IBM PCs.

189 An expert system for verification and tuning of simulation models
R. DeMori
Sch. of Comp. Sci., McGill Univ., Montreal, Que., Canada
R. Prager
Artificial Intelligence in Engineering: Tools and Techniques, Cambridge, MA, USA, Aug. 1987 (Southampton, UK: Comput. Mech. Publications 1987), p. 171-75

This article describes an on-going program of research in application of artificial intelligence in the domain of aerodynamic simulation. An engineer plays two roles in the preparation of a model for a real-time simulation. First, the model is built using the best available design data and coded to execute in a digital computer. Then the engineer must verify the fidelity of the model by comparing the performance of the simulator with data recorded during tests of the modeled aircraft. The results of the verification may demonstrate a need to tune the model to better simulate the aircraft in some areas. The goal of this research is to produce a knowledge-based system which acts as an assistant to an engineer during verification and tuning of real-time simulation models. Analysis of the problem has made it clear that current AI techniques are inadequate to meet this goal. Original contributions of this research are pointed out. The authors then describe the domain and the problem in more detail and discuss the relationships with other recent research. Next the design of the system is presented, showing the major reasoning components and their functionality. Then some of the methods for performing the required reasoning are described; the techniques employed are drawn from automatic speech recognition, qualitative physics and signal interpretation. To illustrate the performance of the KBS, an example of testing and reasoning is described. The current implementation and results are discussed. (15 refs.)

211 Experience in the development of an expert system for fault diagnosis in a commercial scale chemical process *P. S. Dhurjati, D. E. Lamb, D. Chester*
Foundations of Computer Aided Process Operations *Proceedings of the First International Conference*, Park City, UT, USA, 5-10 July 1987 (Amsterdam, Netherlands: Elsevier 1987), p. 589 - 625

The experience gained in the development of FALCON, an expert system for fault diagnosis in a commercial scale chemical process, is discussed. The first task was to define a target process for the development of the expert system. Next, a minimum fault set was defined and the adequacy of the existing instrumentation for diagnosing the faults in the minimum fault set was assessed. A dynamic mathematical simulation of the process was developed and verified with plant data. The faults in the minimum fault set were induced in the simulation and the results were initially used to extract knowledge about the reasoning processes and heuristics used by expert engineers. The knowledge base is a hybrid consisting of both quantitative and qualitative diagnosis methods. Fundamental engineering principles based on simple quantitative models are used first and any gaps in such quantitative knowledge are then filled with heuristics. A real-time inference engineer was written in Common Lisp. The FALCON system has a user interface which is the primary means of interaction between the operator and the expert system. The FALCON system has been extensively and successfully tested both on simulated plant faults as well as on actual plant data. (13 refs.)

224 Experience in using MUMPS for AI applications
R. D. Semmel
Appl. Phys. Lab., Johns Hopkins Univ., Laurel, MD, USA
MUG Q. (USA), vol. 18, no. 1, p. 1-5 (1988). (MUMPS Users' Group Meeting 1988 - MUMPS: A Lasting Competitive Edge in Computing, New Orleans, LA, USA, 13-17 June 1988)

Designing interfaces to large databases is a complex activity that must take into account semantic information and domain - specific knowledge. Attempting to encode such information and knowledge in programs using classical approaches has been fraught with difficulty and, as a result, frequently is avoided. Artificial Intelligence (AI) research has produced some general techniques that, when supplemented with domain knowledge, can be utilized to great advantage to solve problems that previously defied codification efforts. A MUMPS based relational database interface known as the in-

telligent navigational assistant (INA) has been built that exploits some of these techniques. The paper discusses the general problem scenario, provides some basic database and AI background, and briefly describes how the query information module of the INA incorporates AI techniques to generate SQL queries. (6 refs.)

397 Expert systems in business: a British perspective
M. A. Bramer
Sch. of Comput. & Inf. Technol., Thames Polytech., London, UK
Expert Syst. (UK), vol. 5, no. 2, p. 104-17 (May 1988)
[received: 26 Oct 1988]

In the early 1970s, research into artificial intelligence (AI) in Britain was almost wiped out following the publication of an unfavourable report to government. In the 1980s, AI emerged from obscurity with expert systems as its leading edge and has now become the focus of considerable activity, with an ever-growing volume of practical applications developed or under development in British industry and business. This paper reviews the growth of business interest in expert systems from a British perspective and surveys the principal work carried out to date. Technical issues highlighted by this work are considered and prospects for future development are discussed. (78 refs.)

229 Toward intelligent dialogue with ISIS
J. L. Edwards
Artificial Intelligence Manage. & Dev. Corp., Toronto, Ont., Canada
J. A. Mason
Int. J. Man-Mach. Stud. (UK), vol. 28, no. 2-3, p. 309-42 (Feb - March 1988) [received: 21 October 1988]

A proposal for the design of intelligent dialogue systems is developed within the context of the work on SDMS, a spatial database management system, part of the ISIS project at the Canadian Defence and Civil Institute of Environmental Medicine. The approach uses the Models aspect of Edwards and Mason's (1988) methodology as an organizing principle for design, including the two central notions of explicit-model design and a self-referencing model configuration. Further, a design space is presented that combines the aspects and dimensions of the methodology in order to provide clarity of focus for the developer, who must consider a multitude of options in designing and implementing intelligent systems. Aspects of the methodology particularly important to the design and development of intelligent dialogue systems are discussed, namely, connectivity, control, form and modality. The design methodology is elaborated through discussions and recommendations for how it might be applied in building future versions of SDBMS. (20 refs.)

230 Program translation via abstraction and reimplementation
R. C. Waters
Artificial Intelligence Lab., MIT. Cambridge, MA, USA
IEEE Trans. Softw. Eng. (USA), vol. 14, no. 8, p. 1207-28 (Aug. 1988)

An abstraction-and-reimplementation paradigm is presented in which the source program is first analyzed in order to obtain a programming-language-independent abstract understanding of the computation performed by the program as a whole. The program is then reimplemented in the target language based on this understanding. The key to this approach is the abstract understanding obtained. It allows the translator to benefit from an appreciation of the global features of the source program without being distracted by what are considered irrelevant details. Knowledge-based translation via abstraction and reimplementation is described as one of the goals of the Programmer's Apprentice project. A translator which translates Cobol programs into Hibol (a very-high-level business data processing language) has been constructed. A computer which generates extremely efficient PDP-11 object code for Pascal programs has been designed. (38 refs.)

242 Semi-automatic parallelization of Fortran programes
H. P. Zima, H.-J Bast, M. Gerndt, P. J. Hoppen
Inst. fur Inf. III, Bonn Univ., West Germany
CONPAR 86. Conference on Algorithms and Hardware for Parallel Processing. Proceedings, Aachen, West Germany, 17-19 Sept. 1986 (Berlin, West Germany: Springer-Verlag 1986), p. 287-94

The authors describe the design of an interactive, knowledge-based system for the semi-automatic transformation of Fortran 77 programs into parallel programs for a new supercomputer. The system is characterized by a powerful analysis component, a catalog of MIMD and SIMD transformations, and a flexible dialog facility. It contains specific knowledge about the parallelization of an important class of numerical algorithms. (20 refs.)

401 Information and decision support systems for risk analysis
K. Fedra
Adv. Comp. Applications Project, Int. Inst. for Appl. Syst. Anal., Laxenburg, Austria
Safety of Computer Control Systems 1988 (SAFECOMP '88): Safety Related Computers in an Expanding Market. Proceedings of the IFAC Symposium, Fulda, West Germany, 9-11 Nov. 1988 (Oxford, UK: Pergamon 1988), p. 53-9

Technological risk is a fact of life in a technological society. Risk analysis and risk management, designed to control and minimize these risks, involve numerous information processing aspects of considerable complexity and, as a consequence, induct the use of computers. Computer systems designed to provide information and decision support, not only during operation, but primarily in the design and planning stage of hazardous installations and the handling of hazardous materials, that extensively use databases, models, and increasingly, expert systems technology, face problems of validity and reliability and the proper use and interpretation of results. Some of these questions are discussed in the context of a number of information and decision support systems for industrial risk and hazardous substances management, developed at the International Institute for Applied Systems Analysis (IIASA) for various governmental and industrial clients. (18 refs.)

426 A decision support system for village health workers in developing countries
G. Porenta, B. Pfahringer, M. Hoberstorfer, R. Trappl.
Appl. Artif. Intell. (USA), vol. 2, no. 1, p. 47-63 (1988)

Ambulatory health care in developing countries is often provided by medical or paramedical personnel commonly called village health workers. Medical problems within their realm of competence include diarrhea, infestation with worms, diseases affecting the eyes and skin, and several kinds of infectious diseases. Also, assessment of malnutrition is among the tasks village health workers have to perform. A knowledge-based system is designed and implemented to provide decision support for those tasks. Expanded decision networks are used to represent the diagnostic processes for diarrhea and infestation with worms. The diagnostic procedures for diseases affecting the skin and the eyes are covered with a rule-based approach using two certainty factors. Treatment schemes and procedural knowledge about drug prescription are represented in a frame structure. The hybrid knowledge engineering tool VIE-KET is used to construct a modular consultation system that provides entry points for diagnosis, therapy, and drug prescription. Different aspects of the design and implementation of the system are discussed, and two sample consultation sessions are presented. (12 refs.)

Appel aux communications

Conférence Canadienne d'Intelligence Artificielle

CSCSI/SCEIO–90

Université d'Ottawa, Ottawa, Ontario, Canada
23-25 mai 1990

CSCSI/SCEIO-90 est la huitième conférence biennale d'intelligence artificielle commanditée par la Société Canadienne pour l'Étude de l'Intelligence par Ordinateur. On sollicite des communications décrivant des résultats originaux de recherches théoriques ou appliquées dans le domaine de l'intelligence artificielle. Les sujets suivants sont particuliérement recherchés:

Applications de l'IA	Architectures et Languages	Acquisition des connaissances
Apprentissage	Démonstration de théorèmes	Compréhension du langage naturel
Modélisation cognitive	Programmation automatique	Perception (vision, toucher, parole)
Planification	Raisonnement par défaut	Raisonnement à base de règles
Portée sociale de l'IA	Raisonnement temporel	Raisonnement de sens commun
Robotique	Recherche automatique	Représentation des connaissances

Veuillez faire parvenir 4 exemplaires des communications au président du comité de programme avant le 9 janvier 1990. La décision d'accepter ou de rejeter la soumission sera communiquée par courier au premier auteur (ou à l'auteur désigné) avant le 19 février 1990. Des copies prêtes à la reproduction des papiers acceptés devront être reçues au plus tard le 20 mars 1990.

Tous les exemplaires de la communication doivent être clairement lisibles et imprimés sur du papier format A4 ou 8-1/2x11 po. avec des marges d'au moins 1-1/2 po. et des fontes 12 points. (Les styles "article" de LaTeX ou Scribe avec des fontes 12 points sont conformes à ces spécifications et devraient être utilisés si possible.)

Les communications ne doivent pas excéder 5000 mots, incluant sommaire, bibliographie et l'espace pour les figures, tables et diagrammes. Les soumissions doivent être accompagnées d'une page portant le titre, les noms et adresses de tous les auteurs ainsi qu'un court sommaire. Les communications doivent contenir une explication concise de la contribution originale faite à la recherche en IA. Les papiers se verront allouer huit (8) pages dans les actes de la conférence.

Si la communication a été soumise à d'autres conférences, textuellement ou dans ses points essentiels, ce fait doit être clairement indiqué sur la page de titre.

Toutes les soumissions seront revues par le comité de programme. Les critères de choix sont la valeur globale, l'importance de la recherche, ainsi que la qualité de la présentation. Si une communication est acceptée, l'auteur, ou l'un des co-auteurs, devra être présent à la conférence pour en donner présentation.

Le journal *Artificial Intelligence* offre un prix pour la meilleure communication de la conférence qui sera sélectionné par le comité de programme. Cette communication sera ensuite publiée par le journal *Artificial Intelligence*.

Veuillez faire parvenir
les communications à:
Peter F. Patel-Schneider
CSCSI-90 Program Chair
AT&T Bell Labs
600 Mountain Avenue
Murray Hill, New Jersey 07974
U. S. A.
(201) 582-3399
pfps@research.att.com

Pour de plus amples informations,
veuillez contacter:
Ken Charbonneau
National Research Council
Conference Services Office
Building M-19
Montreal Road
Ottawa, Ontario K1A 0R6
CANADA
(613) 993-9009 (Téléphone)
(613) 957-9828 (Fax)

Président:
Dick Peacocke
CSCSI-90 General Chair
Bell Northern Research
Box 3511, Station C
Ottawa, Ontario K1Y 4H7
CANADA
(613) 765-2629
richard@bnr.ca

Upcoming Conferences

In Canada

CSCSI 90
23 - 25 May 1990, Ottawa, Ontario
See announcement on pages 50 and 52 for details

Introduction to AI
31 July - 4 August, Toronto, Ontario

This intensive 5-day course is sponsored by the Information Technology Research Centre (ITRC), a non-profit corporation designated as a Centre of Excellence. The course, taught by Dr. E. Milios and Dr. R. Greiner, of the University of Toronto, is adapted from a 4th-year undergraduate course. The focus is on the practical aspects involved in developing AI-based systems and includes theoretical foundations, case studies, and hands-on experience in a Lisp environment. Emphasis will be on expert systems and vision systems.
Contact: Rosanna Reid, (415) 978-8558.

In the United States

Summer Computer Simulation Conference
24 - 27 July 1989, Austin, Texas

Topics include: Knowledge-based simulation theory; Intelligent simulation systems; Knowledge-based simulation tools; Knowledge-based systems (KBS) using simulation; Knowledge representation for simulation; Intelligent simulation control architectures; Applications of simulation techniques to KBS; Interactions between conventional simulations and KBS.
Contact: Society for Computer Simulation, ATTN: Group XIII, PO Box 17900, 4838 Ronson Court, Suite 'L', San Diego, CA 92117-7900.

CIE: Computers in Engineering Conference
30 July - 2 August 1989, Anaheim, California

Technical papers are invited in all areas relevant to the utilization of computers in the engineering profession, from research and development to applications, education, business and management issues and challenges.
Contact: (for AI, ES, KBS, Design Theory) Dr. G. Gabriele, Rensselaer Polytechnic Institute, Dept. of Mechanical Engineering, 110 Eight St., Troy, NY 12180-3590. (518) 276-2601.
(for Computers/Robotics in Education, Teaching CAD, Computer Aided Learning Systems) Dr. K. Tamma, Mechanical Engineering Dept., U. of Minnesota, 111 Church St., Minneapolis, MN 55455. (612) 625-1821.

3rd Workshop on Qualitative Physics
9 - 11 August 1989, Palo Alto, California

The workshop will emphasize both theoretical and practical aspects of qualitative and naive physics. Topics include: New qualitative reasoning methods; Complexity analyses of fundamental algorithms; Implementation techniques and performance studies; Applications of qualitative physics; Temporal reasoning and representation; Automated modeling.
Contact: Dr. Patrick Hayes, Xerox PARC, 3333 Coyote Hill Rd., Palo Alto, CA 94304. (415) 494-4749
Email: hayes.pa@xerox.com.

4th Conference on Computers and Philosophy
10 - 12 August 1989, Pittsburgh, Pennsylvania

The conference will focus on intellectual, ethical and practical problems in all aspects of the relationships between philosophy and computers. Papers on AI are especially invited. There will be a showcase of instructional software for use in philosophy classes.
Contact: Leslie Burkholder, Center for the Design of Educational Computing, Carnegie Mellon U., Pittsburth, PA 15213. Email: lb0q@andrew.bitnet.

IASTED Conference on Expert Systems
16 - 18 August 1989, Honolulu, Hawaii

Topics include: Knowledge acquisition and representation; ES design and tools; Hard/software; Logic programming; Planning; Learning and adaptation; Modeling and simulation; Self-organization; Pattern recognition; Speech; Vision, Robotics; Applications.
Contact: ESNN-IASTED, Secretariat, PO Box 25, Stn. G., Calgary, AB, Canada T3A 2G1. Phone (403) 270-3616. Fax: (403) 270-8855. Tlx: 038-26670

IJCAI Workshops
20 - 25 August 1989, Detroit, Michigan

AI and Music: 20 Aug
Contact: Otto Laske, 926 Greendale Ave., Needham, MA 02192. (617) 449-0781. Email: laske@bu-cs.bu.edu.
Parallel Algorithms for Machine Intelligence: 20 Aug. Contact: Prof. Vipin Kumar, Computer Science Dept., Univ. of Texas at Austin, Austin, Texas, 78712. (512) 471-9571 Email: kumar@cs.utexas.edu.
Constraint Processing: 20 Aug
Contact: Rina Dechter, Dept. of Comp. Sci., Technion — Israel Institute of Technology, Haifa 32000, Israel.
Model-Based Reasoning: 20 Aug
Contact: Ethan Scarl, Boeing Computer Services, PO Box 24346, M/S 7L-64, Seattle, Washington, 98124-0346. (206) 865-3255
Knowledge Discovery in Databases: 20 Aug
Contact: Gregory Piatetsky-Shapiro, GTE Laboratories, 40 Sylvan Road, Waltham MA 02254. (617) 466-4236. Email: gps0%gte.com@relay.cs.net.
Conceptual Graphs: 20 - 21 Aug
Contact: Janice Nagle, 1641 E. Old Shakopee Rd., Bloomington, MN 55425
Symbolic Problem Solving in Noisy, Novel, and Uncertain Task Environments: 20 - 21 Aug
Contact: Mike Coombs. (505) 646-5757. Email: mcoombs@nmsu.edu.
Lexical Acquisition: 21 Aug
Contact: Dr. Uri Zernik, General Electric — Research and Development Center, PO Box 8, Schenectady, NY 12301 (518) 387-5370. Email: zernik@crd.ge.com.

Call for Papers

Canadian Artificial Intelligence Conference
CSCSI/SCEIO–90

University of Ottawa, Ottawa, Ontario, CANADA
23–25 May 1990

CSCSI/SCEIO-90 is the eighth biennial conference on Artificial Intelligence sponsored by the Canadian Society for the Computational Studies of Intelligence. Contributions are requested describing original research results, either theoretical or applied, in all areas of Artificial Intelligence research. The following areas are especially of interest:

Applications of AI	Automatic Programming	Common Sense Reasoning
Cognitive Modelling	Default Reasoning	Knowledge Acquisition
Learning	Knowledge Representation	Machine Architectures and Languages
Planning	Rule-based Reasoning	Natural Language Understanding
Robotics	Social Implications of AI	Perception (Vision, Touch, Speech)
Search	Temporal Reasoning	Theorem Proving

Submit four (4) complete copies of papers to the program chair. Papers must be received by 9 January 1990. Notification of acceptance or rejection will be mailed to the first author (or designated author) by 19 February 1990. Camera ready copy of accepted papers will be due 20 March 1990.

All copies of the paper must be clearly legible. Print papers on 8-1/2" by 11" or A4 paper with at least 1-1/2" margins, using 12-pt type. (The LaTeX or Scribe "article" styles, with "12pt" type, conform to these requirements and should be used if possible.) Papers must not exceed 5000 words in length, including abstract and bibliography and allowances for figures, tables, and diagrams. Papers must have a separate title page, containing the title of the paper, the names and addresses of all authors, and a short abstract. Papers must contain a concise statement of the original contribution made to Artificial Intelligence research. Papers will be allocated eight (8) pages in the conference proceedings. Papers which are also being submitted to other conferences, whether verbatim or in essence, must have this fact clearly indicated on the title page.

All submissions will be reviewed by the program committee. Acceptance will be based on the overall merit and significance of the reported research, as well as the quality of the presentation. As a condition of acceptance, the author, or one of the co-authors, will be required to present the paper at the conference.

The journal *Artificial Intelligence* has offered a best paper prize for the conference. This paper will be offered publication in the journal. Selection of the best paper will be done by the program committee.

Send papers to:
Peter F. Patel-Schneider
CSCSI-90 Program Chair
AT&T Bell Labs
600 Mountain Avenue
Murray Hill, New Jersey 07974
U. S. A.
(201) 582-3399
pfps@research.att.com

Send general inquiries to:
Ken Charbonneau
National Research Council
Conference Services Office
Building M-19
Montreal Road
Ottawa, Ontario K1A 0R6
CANADA
(613) 993-9009 (Voice)
(613) 957-9828 (Fax)

General Chair:
Dick Peacocke
CSCSI-90 General Chair
Bell Northern Research
Box 3511, Station C
Ottawa, Ontario K1Y 4H7
CANADA
(613) 765-2629
richard@bnr.ca

Automating Software Design: 21 Aug
Contact: Michael Lowry, Kestrel Institute, 3260 Hillview Ave., Palo Alto, CA 94304. (415) 493-6871. Email: lowry@coyote.stanford.edu.
AI and Simulation: 21 Aug
Contact: Dr. Jerzy Rozenblit, Dept. of Electrical and Computer Engineering, U. of Arizona, Tucson, AZ 85721. (602) 621-6177. Email: rozenblit@arizevax.bitnet.
Knowledge Acquisition: 22 Aug
Contact: R. Glenn Wright, Prospective Computer Analysts, Inc., 1215 Jefferson Davis Hwy. Suite 309, Arlington, VA 22202. (703) 892-2511
Object-Oriented Programming in AI: 22 Aug
Contact: Scott Woyak, EDS Research and Development, 3551 Hamlin Rd, Fourth Floor, Auburn Hills, MI 48057. (313) 370-1669. Email: sww@edsdrd.eds.com
Knowledge, perception, and planning: 22 Aug
Contact: Ernest Davis, Courant Institute of Mathematical Sciences, New York U., 251 Mercer Street, New York, NY 10012. (212) 998-3123. Email: davise@acf3.nyu.edu.
Intelligent Interfaces: 22 Aug
Contact: Yigal Arens, USC/ISI, 4676 Admiralty Way, Marina del Rey, CA 90292.
Blackboard Systems: 23 Aug
Contact: Daniel Corkill, Dept. of Comp. and Info., Sci., U. of Massachusetts, Amherst, Mass. 01003, Email CorkCCS.UMass:Edu (413) 545-0156
AI Aids for the Hearing Impaired: 23 Aug
Contact: Prof. Oscar Garcia, The George Washington U., Dept. of Electrical Engineering and Comp. Sci., Washington, DC 20052. Email: garcia@a.isi.edu.
Transferring ES Technology in Business: 24 Aug
Contact: Prof. Patrick Lyons, Dept. of Management, St. John's University, Jamaica, NY 11439
Connectionist AI: 24 & 25 Aug
Contact: Lokendra Shastri, Computer and Information Science, U. of Pennsylvania, Philadelphia, PA 19104
Integrated Human-Machine Intelligence in Aerospace: 21 August
Contact: Valerie Shalin, Honeywell Systems & Research Center, 3660 Technology Dr., MN65-2500, Minneapolis, MN 55413. (612) 782-7672. shalin@src.honeywell.com.
AI in Manufacturing: Aug 20 Diagnostic Systems for Manufacturing; Aug 21 Concurrent Engineering Design; Aug 22 Manufacturing Planning; Aug 23 Manufacturing Scheduling; Aug 24 Integrated Architectures for Manufacturing . Contact: Claudia Mazzetti, AAAI Office, 445 Burgess Dr., Suite 100, Menlo Par, CA 94025 (415 328-3123

IJCAI 11th Joint Conference on AI

20 - 26 August 1989, Detroit, Michigan

Topics include: AI tools and technologies; Machine architectures, languages, shells; Search methods; Knowledge acquisition, learning, analogy; Real-time performance; Parallel and distributed processing; Cognitive modeling; Planning, scheduling, and reasoning about actions; Natural language, speech understanding and generation; Perception, vision, robotics; ITS; Design, manufacturing, control. Contact: IJCAI 89, c/o AAAI, 445 Burgess Drive, Menlo Park, CA 94025-3496.

Annual General Meeting of the CSCSI/SCEIO
Place: IJCAI '89 Detroit, Michigan
Time: Thursday, August 24, 1989, 12:30 — 2:00 p.m.
Room: M2-29 Cobo
Everybody is welcome to attend!

AAAI 9th Distributed AI Workshop

12 - 14 September 1989, Orcas Island, Washington

Topics include: Intelligent agents; Coordination; Negotiation; Distributed reasoning; Distributed interfaces; Technology platform.
Contact: Miroslav Benda, Knowledge Systems Lab, Boeing, M/S 7L-64, PO Box 24346, Seattle, WA 98124. (206) 865-3244.

4th IEEE Symposium on Intelligent Control

24 - 26 September 1989, Albany, New York

Topics include: Machine learning; Adaptive and self-organizing control; Hierarchical/fuzzy/linguistic control; Knowledge-based control systems; Qualitative reasoning in intelligent control; Models of approximate reasoning; Evidential reasoning and control; Cooperative and antagonistic multiple-agent systems; Neural networks as a control tool; Expert databases for intelligent control; Robot vision; Space and underwater exploration.
Contact: Prof. K. P. Valavanis, Robotics Lab, Northeastern U., Dept. of ECE, Boston, MA 02115. (617) 437-2164.

Conference on Logic Programming

16 - 19 October 1989, Cleveland, Ohio

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

Workshop on Term Subsumption Languages in Knowledge Representation

18 - 20 October 1989, New Hampshire

Focus is 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.
Submission material: statement of interest (max 4 pages), outlining past contributions and current interests in the area. Most important, indicate your potential contributions to the workshop. Include an individual or group bibliography and a recent, relevant paper.
Limited financial help may be available.
Contact: Peter Patel-Schneider, AT&T Bell Labs, 600 Mountain Ave., Room 3C-410A, Murray Hill, New Jersey 07974

5th Conference on AI & Ada

16 - 17 November,1989, Fairfax, Virginia

Focus: scientific foundations and software technologies for future AI embedded systems that can be characterized as: real-time, intelligent, implemented in Ada, missions critical, operational for long periods of time, and distributed. Topics: AI applications in Ada; AI environments in Ada; Ada environment tools to support AI; Analysis of language features for AI; AI languages implemented in Ada; Managing AI/Ada interface.

mented in Ada; Managing AI/Ada interface.
Contact: AIDA '89, Dept. of Comp. Sci, George Mason U., 4400 University Dr., Fairfax, VA 22030-4444. (703) 323 - 2713

Neural Information Processing Systems — Natural and Synthetic

27 - 30 November 1989, Denver, Colorado

Topics include: Neurobiological models of development, cellular information processing, synaptic function, learning, and memory; Design and evaluation of net architectures to perform cognitive or behavioural functions and to implement conventional algorithms; Training paradigms for static and dynamic networks; Analysis of capability, generalization, complexity, and scaling; Applications to signal processing, vision, speech, motor control, robotics, knowledge representation, cognitive modeling and adaptive systems; Implementation and Simulation: VLSI or optical implementations of hardware neural nets.
Contact: Kathy Hibbard, NIPS89 Local Committee, Engineering Center, Campus Box 425, Boulder, CO 80309-0425.

IASTED Conference on Expert Systems

12 - 14 December 1989, Los Angeles, California

IASTED is the International Association of Science and Technology for Development. Topics include: Knowledge acquisition and representation; ES design procedures and tools; Logic programming; Reasoning; Advisory systems; User interface; Testing and maintenance; Neural networks. Special one-day conferences on ES in banking; ES in medicine, bio-engineering and health care.
Submission deadline: September 15, 1989.
Contact: Expert Systems LA, IASTED Secretariat, PO Box 25, Stn. G, Calgary, AB, Canada T3A 2G1. (403) 270-3616. Fax: (403) 270-8855. Telex: 03-826670.

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.

Outside North America

Modeling Autonomous Agents in a Multi-Agent World

16 - 18 August 1989, London, England

The purpose of the workshop is to stimulate exchange and discussions on the AI models of an autonomous agent having to interact with other autonomous entities. Distributed AI occurs as a special case of the problem, but the main concern deals with complementary aspects of DAI by allowing the agents to have unrelated goals.
Contact: Prof. Jean-Pierre Meuller, Université de Neuchatel, Institut de Mathèmatiques et Informatique, Chantemerle, 20, CH-2000 Neuchatel, Switzerland

IEEE ICIP '89 Image Processing Conference

5 - 8 September 1989, Singapore

Topics include: AI vision techniques; VLSI implementation; Image restoration/enhancement; Machine vision; Video communications; Office image processing; Image pattern recognition; Biomedical imaging; Remote sensing; System architecture.
Contact: Technical Program Chairperson, ICIP '89, c/o Meeting Planners, 100 Beach Road, #33-01, Shaw Towers, Singapore 0718, Republic of Singapore.
Email: (Dr. Cho-Huak TEH) eletehch@nusvm.bitnet, or chteh@nuseev.bitnet.

AI and Cognitive Science

14 - 15 September 1989, Dublin, Ireland

Topics include: AI tools and technologies; Natural language, machine translation, speech understanding and generation; Perception, vision, robotics; Intelligent tutoring systems; Design, manufacturing and control; Cognitive modeling; Commonsense reasoning; Knowledge representation; Planning, scheduling, reasoning about actions; Knowledge acquisition, learning, analogy; Automated deduction; Philosophical foundations, social implications.
Contact: Alan Smeaton, Chairperson, AI/CS — '89, School of Computer Applications, N.I.H.E., Dublin 9, Ireland.

5th Alvey Vision Conference

25 - 28 September 1989, Reading, U.K.

Topics include: Image processing and feature analysis; Practical applications of machine vision; Object recognition and identification; Scene analysis and image interpretation; Reconstruction of 3-d shape; Computational issues in visual perception; Sensors and robotic vision; Architectures for vision systems.
Contact: G. D. Sullivan, Dept. of Comp. Sci., U. of Reading, PO Box 220, Whiteknights, Reading, U. K. RG2 6AX
Phone: (0734) 318603. Email: AVC89@uk.ac.reading.

2nd Symposium on AI

23 - 27 October 1989, Monterrey, New 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; Applications and future trends.
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.

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: Advances in expert systems; Learning; Logic programming; AI and engineering; Natural language understanding; AI systems and software; Pattern recognition; Intelligent tutoring systems; Reasoning; Knowledge representation; Speech; Vision.
Submission deadline: August 15, 1989.
Submission material: 2 copies of paper, 2000-5000 words, with abstract of 200-300 words.
Contact: Dr. S. Ramani, Chairperson, 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; Applications; 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.

13th Conference on Computational Linguistics
20 - 25 August 1990, Helsinki, Finland
Tutorials August 16 - 18.
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. Tlx: 15440 KVAL S. Email: COLING@COM.QZ.SE or COLING@ZCOM.BITNET.

IJCAI-91
24 - 30 August, 1991, Sydney Australia
Contact: Prof. Barbara Grosz, Conference Chairperson, 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, Chambery, France
Contact: Prof. Jean-Pierre Laurent, Local Arrangements Chairperson, IJCAI-93, Université de Savoie, BP 1104, F-73001 Chambergy, France. Phone: (+33-79) 961-062.
Email: jplaure@imag.fr.

Veuillez répondre au questionnaire dans la revue avril '89

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 your 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* (Cdn$35).
( ) I am a student (subtract Cdn$10). **AND/OR** ( ) I am a member of CIPS (subtract Cdn$10).

( ) Please send me the following back issues of *Canadian Artificial Intelligence* (Cdn$10 each including postage, except Cdn$15 for#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 (Cdn$20).‡

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 forwards your order.

CSCSI/SCEIO Conference Proceedings

( ) Please send me the following CSCSI/SCEIO conference proceedings (indicate number of copies desired):
( ) Edmonton, 1988 (Within Canada: Members: Cdn$35; Non-members: Cdn$40. Add Cdn$5 for postage. Mail to **CIPS**.) (Outside Canada: US$35. Postage within U.S.: US$2 for the first copy, US$.75 for each additional copy. Postage outside U.S.: US$3 for the first copy, US$2 for each additional copy. Within California: add 7% sales tax. Mail to: **Morgan Kaufmann Publishers**, Order Fulfillment Center, PO Box 50490, Palo Alto, CA 94303, USA.)
( ) Montréal, 1986 (Cdn$30. Postage within Canada: Cdn$5. Outside Canada: Cdn$7. Mail to **CIPS**.)
( ) Saskatoon, 1982 ( ) Victoria, 1980 (Cdn$25. Postage within Canada: Cdn$5. Outside Canada: Cdn$7. Mail to **CIPS**.)

Name .....

Mailing Address .....

National Research Council Canada
Conseil national de recherches Canada

Computational Intelligence Intelligence informatique

Editors Directeurs scientifiques
N. Cercone / G. McCalla

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 1989 are $46.00 for Individuals and $114.00 for Institutions

La revue Intelligence informatique est publiée 4 fois par an. En 1989, l'abonnement individuel coûte 46$ et l'abonnement collectif 114$.

For further information contact

Subscription information
National Research Council of Canada
Ottawa, Ontario
K1A 0R6

Pour tout information complémentaire, communiquez avec

Abonnement
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