Research in the Knowledge Sciences at the University of Calgary

by Ian Witten and Brian Gaines

Recherche en sciences de la connaissance à l'Université de Calgary

RÉSUMÉ: Le Knowledge Science Institute (KSI) de l'Université de Calgary a pour objectif d'effectuer des recherches en sciences de la connaissance, sur les technologies associées et sur leur impact économique et social. Ce survol se concentre sur la recherche en sciences de la connaissance effectuée par les professeurs d'informatique ou en collaboration étroite avec eux. Celle-ci se répartit parmi les domaines suivants: l'acquisition de connaissances pour les systèmes experts, l'apprentissage par ordinateur et la programmation par exemples, la théorie des systèmes et leur organisation (particulièrement la modélisation et la prédiction) et la programmation logique. La recherche dans les domaines associés à l'informatique, tels que l'interaction homme-ordinateur, les systèmes répartis, et les méthodes formelles pour la vérification de logiciel et de matériel, n'est pas couverte par cet article.

Since the late 1960's the Computer Science Department at the University of Calgary has enjoyed a major research presence in the field of human-computer interaction, and the last five to ten years have seen sharply increased emphasis on AI topics and methods. This trend was focused and strengthened by the establishment in 1985 of the Knowledge Science Institute (KSI), whose objectives are to study knowledge science, technology, economy and social impact, and to promote interdisciplinary research on knowledge-based systems. This has involved the provision of space, capital equipment, operating funds, the attachment of faculty, and the recruitment of research staff.

While the KSI operates as a network of multi-disciplinary research groups spanning several faculties and academic disciplines, this overview concentrates on knowledge science research undertaken by, or in close collaboration with, computer science faculty members. It is divided into the areas of knowledge acquisition for expert systems; machine learning and programming by example; system theory and organization (especially modeling and prediction); and logic programming. However, different projects mingle and overlap, and in practice no real distinction is made between those that involve artificial intelligence techniques and those that do not. Research in allied fields of computer science, such as human-computer interaction, distributed systems, and formal methods for software and hardware verification, is not covered here.

Research in Knowledge Engineering

A model of the development of the computing industry has been developed and used to account for past events and predict future trends, particularly fifth and sixth generation priorities. The industry came into being as a result of electronic device technology; thus enabling society to cope with an increasingly complex world. Underlying all developments in computing is a tiered succession of learning curves which make up the infrastructure of the computing industry. The model provides a framework for the industry based on this logical progression of developments. It links the framework empirically to key events in the development of computing. It links the framework theoretically to a model of economic, social, scientific and individual development as related learning processes with a simple phenomenological model. The fifth generation development program, with its emphasis on human-computer interaction and artificial intelligence, and the sixth generation program, with its emphasis on knowledge science, are natural developments in the foci of attention indicated by the model [9, 10].

Knowledge engineering represents a key technology within this model, and is a cornerstone of our activities.

Rapid Prototyping

Knowledge acquisition techniques are being developed for rapid prototyping of knowledge-based systems [14, 15, 16]. Problems of knowledge engineering have been recognized since the early days of expert systems and are impeding practical applications. It is difficult to elicit from experts the knowledge structures underlying their skills. This has heightened interest in the possibility of providing knowledge support systems to automate knowledge engineering as a process of direct interaction between domain experts and the computer [6, 8]. In 1980, we proposed that personal construct psychology could provide foundations for expert systems, particularly in systems that combined interactivity with database access and expert advice to provide decision support, and gave examples of algorithms and programs that extracted entailment rules from repertory grid data. In 1983, we reported further enhancements of these techniques and a preliminary experiment to validate them empirically as a knowledge engineering technique for priming expert systems. This work led to industrial studies of the methodology applied to the development of expert systems: Boeing Computer Services and Lockheed Software Technology Center have reported success in applications; and validation has been reported in a statistics domain.

Ian Witten and Brian Gaines are professors at the University of Calgary. Brian Gaines is founder and director of KSI.

The knowledge acquisition research focuses on the use of techniques based on personal construct psychology to automate knowledge engineering for expert systems. Knowledge support systems are being developed as interactive knowledge engineering tools, and the design criteria for such systems are being clearly defined. KITTEN is a knowledge support system developed for AI workstations which includes tools for interactive repertory grid elicitation and entailment analysis that have been widely used for rapid prototyping of industrial expert systems. It also includes tools for text analysis, behavioural analysis and schema analysis, that offer complementary and alternative approaches to knowledge acquisition. The KITTEN implementation integrates these tools around a common database with utilities designed to give multiple perspectives on the knowledge base [17, 18].

A wide range of tools developed elsewhere are also being investigated such as AutoIntelligence, ETS, KSSO, and NEXPERT; and validation studies are being undertaken of the performance of these tools in eliciting and structuring expert knowledge.

Hypermedia and Knowledge-based Sports Coaching

In the Knowledge Science and Neuro-Motor Laboratory in the Faculty of Physical Education, studies are being undertaken of knowledge-based sports coaching systems and the cognitive foundations of knowledge transfer [19]. The research focus is on eliciting, understanding, and applying the knowledge structures underlying the performance skills of athletes and the coaching skills of those involved in their selection and training. The abstract knowledge structures for sports coaching have been developed independently of specific sports. They have been shown to lead to significant analyses of a wide range of sports that guide the generation of coaching strategies and the individual acquisition of playing skills, and to be meaningful and comprehensible to both coaches and players. These knowledge structures are being applied through a variety of media, including books and computer-based hypermedia systems [20].

The new *Steps to Success* book series from Human Kinetics Publishers is based on the explicit use of the structures, and is available in handbooks for athletes and coaches. Initially ten sports are covered including archery, badminton, basketball, golf, gymnastics, tennis and swimming. In the research laboratory a range of hypermedia systems based on the knowledge structures are being developed, including an expert system for ice hockey coaching which was shown at the Winter Olympics in February 1988. This runs on a Macintosh using specialist knowledge acquisition, videodisc and CD-ROM control software. Film for the videodisc has been shot during a number of international games and associated training sessions.

Knowledge Processes in Society, Manufacturing and Human-Computer Interaction

A number of fundamental studies have been undertaken of the way in which knowledge develops in society through a variety of social and psychological processes [5, 7]. The results of these studies have direct implications for the design of expert systems to support the formation of expertise and to avoid problems of legal liability. The knowledge processes involved in manufacturing have been studied and modeled as part of the development of expert systems for computer-integrated manufacturing. The knowledge structures underlying human-computer interaction have also been studied in depth as part of the development of knowledge-based user interface management systems. We see major advances in a number of disciplines being possible through the development of overt knowledge structures available on shared knowledge bases to scientific and technological groups and communities. Our work in this area is now focusing on computer support of collaborative research, knowledge-based heterogeneous system integration, and the use of object-oriented knowledge bases to support creative thinking processes.

Knowledge Acquisition Workshops

The KSI has been heavily involved in organizing a series of workshops on knowledge acquisition for knowledge-based systems (KAW) attended by research workers worldwide. Three KAW's have been held in Banff, sponsored by the American Association for Artificial Intelligence and co-chaired by John Boose of Boeing Advanced Technology Center and Brian Gaines of the KSI. For the first KAW in November 1986, the intention was to hold a discussion-intensive meeting of some 35 highly involved researchers. In practice, over 120 papers were submitted and some 500 applications to attend were received from some 30 countries. This level of interest lead to a further AAAI KAW in Banff in October 1987 and another in November 1988, to an IEEE-sponsored KAW in England in September 1987, to an ACM-sponsored KAW in Germany in June 1988, and to a shorter KAW emphasizing integration issues at AAAI in St. Paul in August 1988. Further Knowledge Acquisition Workshops are already scheduled for Paris, Amsterdam, Tokyo and Banff in 1989 and 1990. These workshops have had a great impact on the literature on knowledge acquisition, resulting in some sixty journal papers and three books to date, and the discussion resulting from them has stimulated knowledge acquisition research worldwide.

Knowledge Engineering Training Program

An intensive one-year training program has been developed for knowledge engineers. Funded by a grant from Employment and Immigration Canada, the program was given to eleven professionals at graduate level from April 1987 through March 1988. The students, with degrees in a range of topics including computer science and psychology, were also employed as research assistants for a number of projects. Courses have been given in knowledge engineering, instructional design, computer hardware and telecommunications, software development and advanced software systems, videodisc development and human performance.

The training program has aroused much industrial interest and it is expected that the material will be widely disseminated and used for internal training programs and short courses. Arrangements are being made to repeat the program this year.

Research in Machine Learning and Programming by Example

A strong research interest in machine learning has emerged like a phoenix rising from the funeral pyre of adaptive control systems in engineering (some of us have

murky pasts). The phenomenon of adaptivity underlies much of our work. So does a central fascination with human interfaces and the communication of information between man and machine. Thus we stress teaching as much as learning; the provision of a congenial environment for casual users to transfer their knowledge and skill.

Robot Programming by Example

Studies are being undertaken to improve communication of human requirements to robot systems. The aim is to enable computer-controlled systems to be taught everyday tasks by people skilled mainly at the tasks. The people should not need, and will not have, the kind of expertise of today's programmers. Existing computer controlled systems — for example, robots and office systems — lack instructability and adaptability. People find it difficult to translate their task skill into explicit procedures for robot, office and other systems. The research aims to enable non-programmers to specify procedures, without having to do this explicit translation.

A high-precision robot arm has been acquired for studies of robot instructability. Initially the robot will be used for research into techniques for non-programmers to specify tasks to robots, as applied to hospital drug dispensing. It will be tested as part of the Calgary Foothills Hospital automatic, robot-based drug dispensing system. The system will select unit doses of drugs and put them into individual patient trays. In a small operation such as this, the pharmacist himself will have to be capable of setting up the dispensing tasks and modifying the robot's performance.

A number of existing programming-by-example systems have been studied and reconstructed in the laboratory. An MIT scheme for creating robot programs from example traces has been reimplemented. A number of shortcomings which were not apparent in the original research for learning object transforms from a user's physical positioning of the arm were addressed. A partially complete program leaves undefined those transforms not available from standard CAD data for a task, and the user can conveniently show the required movements at run time. In addition, the use of quaternions is being investigated as a compact and coordinate-independent representation for robot limb rotations, suitable for inducing functional relationships in tasks comprising straight line motions.

Programming by Example in the Human Interface

The user interface provides a multitude of opportunities for programming by example which are being investigated. Rather than concentrating solely on technical methods of inducing programs from examples — an endeavor which became rather sterile in the mid-1980's after an initial flush of enthusiasm — this research emphasizes the central role that interaction can play in clarifying users' intentions. We have studied the acquisition of procedural knowledge, or "know-how", from end users in the domain of interactive graphics. Experts express procedures constructively, using any of the tools available in the interactive drawing environment, and well-structured procedures are inferred using a variety of weak generalization heuristics. A second application domain being investigated is that of casual-user interaction with office computer systems using the metaphor of an automated office clerk, while a third is in the realm of text editing. Another example of this kind of work is the "reactive keyboard", invented some years ago at the U. of Calgary to help disabled people interact textually with computers.

A central theme underlying this work is the importance of the teacher's view of the learner and his capabilities. Teacher and learner share a common world and have the same capacity for action, an "equal-opportunity" situation. By eagerly carrying out predicted actions (which the teacher can undo if incorrect) the learner may encourage the teacher to satisfy appropriate felicity conditions and make explicit the conditions underlying his actions. In case of ambiguity, dialogue boxes serve to allow teachers to clarify their intentions.

Machine Learning

Many more technical topics in machine learning are being pursued. Our methodology is to begin by reconstructing influential learning systems developed elsewhere. We have reconstructed well-known version space systems, example-generating systems, statistical concept learners, explanation-based learning algorithms, planners, the exploratory learning system AM, and numerous schemes for function induction. In so doing we have learned important lessons and built a springboard for developing our own ideas [13, 25].

Methods for creating structural descriptions from examples involve searching a huge space of possible structural descriptions to check consistency with incoming examples. Efficient ways of representing such spaces in terms of their boundaries are being studied and extended to deal with infinite sets of descriptions (which often arise in practice). Ways of automatically constructing crucial examples which discriminate effectively between competing hypotheses are also under investigation. Shortcomings have been discovered in existing systems for example generation (in particular, the examples generated may be insufficiently general); these are being rectified in a new program which has been built.

Explanation-based learning (EBL) is a semantically-based, knowledge-intensive paradigm for machine learning which contrasts sharply with syntactic or "similarity-based" approaches mentioned above. The EBL technique has been redeveloped from the perspective of problem-solving, and studies have been made on how to embed a learning capability into standard inference engines. The work is being extended to the domain of single- and multi-agent planning. We reject the current trend to isolate learning from other activity and study it separately, preferring instead to integrate it into the very heart of the problem solving [12].

Function induction attempts to synthesize a functional relationship from input-output examples. The data may be structured, with numeric and non-numeric parts. Primitive functions are provided to the system, as are known constants which may be used in functions. We are investigating a number of different algorithms for function induction and comparing their performance both theoretically and practically.

Research in System Theory and Organization

Our fascination with adaptation, mentioned above, also manifests itself lower on the phylogenetic scale of "artificial intelligence" in research on text modeling, prediction and compression. The representation and

management of uncertainty is another common thread which runs through many projects. Also, a strong interest in autonomous and parallel systems colours much of our work.

Text Compression

Complementing the structural work on machine learning is a long-standing and intensive research effort in adaptive text analysis and compression, which grew out of earlier work on adaptive modeling. A text compression method developed at the U. of Calgary in 1982/83 still outperforms all known techniques, most of them by a considerable margin. It typically achieves 2.2 bits/character on mixed-case English, and in a recent Australian competition for a compression program both joint winners used variants of our method. A wide variety of compression techniques have been studied and previously unsuspected relationships between them have been found. The two dominant techniques for adaptive coding are statistical and macro methods, and these have been compared and related. Moreover, a very extensive collection of previously described compression methods has been implemented, tested, and assessed [24, 1].

Good compression requires good prediction of text, and studies are being made of the application of the predictive methods to various areas within both the human-computer interface and machine learning; this relates to the programming-by-example research mentioned above [23].

Modeling Uncertain Data

As our systems move more and more into the real world, uncertainty management becomes increasingly important. Rapid prototyping of expert systems involved inferring relationships from uncertain or "fuzzy" data, and the KITTEN system mentioned earlier contains a novel mechanism for fuzzy inference [11]. Within machine learning, efforts are being made to address the problems of creating learning systems that can recover from errors, for it is unlikely that they will find much practical application until they are robust enough to cope with the inevitable vagaries of real-life situations.

Theoretical work on uncertainty management includes a new, unifying framework that encompasses both probability and fuzzy logic, as well as two-valued classical logic, as special cases of a general "standard uncertainty logic" based on valuation of lattice. A separate project has considered the problem of acquiring uncertain rules from examples within a three-valued logic under the ambit of logic programming [2]. A firm theoretical foundation for three-valued logic programming has been established, including a Horn clause subset, procedural interpretation, extension of the standard Herbrand universe and least fixed-point semantics, and a proof of the correctness of SLD-resolution. Finally, the use of stochastic processes in neural nets has been investigated as a basis for representing uncertain inference.

Autonomous Systems

It is apparent that while many future automated systems will have to operate autonomously, little progress is being made towards this goal. We are looking at the phenomenon of autonomy from three directions[21]. The first regards autonomy as self-government or self-motivation. This is represented by some AI research on representing and using goals, together with physiological, psychological, and philosophical accounts of motivation and goal-seeking behaviour. The second concerns the biological independence of organisms which have the ability to maintain their own organization in a capricious environment. The third aims to construct instructable systems that can learn and carry out instructions autonomously but are not self-motivated.

Work in self-motivation is related to the machine learning research discussed above; in particular the reconstruction of AM. The most concrete examples of artificially autonomous systems of the second genre are various kinds of "deviant" software: worms and viruses; and we have performed experiments and undertaken extensive studies of these phenomena [22]. The third kind includes programming-by-example systems, where the user interface gives learning a human-like autonomy.

Connectionist Architectures

The renewed interest in neural-like processing systems has stimulated a re-evaluation of older paradigms for system organization in terms of massively parallel architectures. For example, we have shown how neural-like systems can be related closely to other well-known formulations of computing such as production systems, finite-state automata, Markov decision processes, and Turing machines. We have studied how novel information representation schemes such as stochastic representations are eminently suited to massive parallelism; benefits of this approach are both theoretical and practical. We have made a detailed theoretical and empirical study of the Boltzmann machine from stochastic, Markovian and Bayesian viewpoints. We have designed a new VLSI architecture which serves as a building-block for connectionist systems and shown how this can be used to tackle such apparently sequential problems as the n-Queens puzzle and chess playing.

Research in Logic Programming

A major part of the research centers around logic programming, and theoretical and practical studies of programming in logic have been undertaken. (However, we have catholic tastes in language and our implementations are based on Prolog, Lisp — mostly Scheme — and Smalltalk in approximately equal proportions).

Logical Arithmetic

In the past, implementations of real arithmetic within logic programming have been non-logical. This is a fundamental deficiency which detracts seriously from the logical paradigm. The meaning of programs depends on the order in which statements are executed and on the details of the computer's handling of finite-precision arithmetic. We have discovered how a purely logical interpretation of arithmetic can be created by the use of interval arithmetic as the underlying representation [3]. This can be translated to an implementation within Prolog and invites a redefinition of the language's arithmetic capabilities to use intervals to represent numbers. As well as having a sound logical basis, the resulting system allows a very concise and powerful programming style and is potentially very efficient.

Parallel Prolog using Time Warp

Current research on parallelism in logic programming focuses on AND-parallelism, in which the goals of a conjunction are executed concurrently (the alternative is OR-parallelism which explores different ways of achiev-

ing a goal in parallel). The most difficult problem in implementing a distributed AND-parallel system is resolving conflicting variable bindings. Our scheme uses the ideas of "time warp", a mechanism first proposed to implement the virtual time paradigm in distributed simulation, to solve this problem. Simply put, bindings are given timestamps, and conflicts are resolved by comparing these timestamps.

AND-parallel implementations not using the virtual time approach either restrict the parallelism available, or adopt a "committed-choice" strategy and thereby sacrifice the full nondeterminism of standard Prolog. In contrast, our implementation preserves the semantics of sequential Prolog while allowing maximum parallelism. Two considerations motivate this choice: first, for many problems nondeterminism must be retained; and second, maintaining full compatibility with sequential Prolog allows application programs to be ported directly. Variants of the scheme which support non-terminating programs and provide smart backtracking in a sequential system are also envisioned. Another avenue of investigation is adapting Prolog for use in distributed simulation. This adaptation is also based on time warp, but requires parallelism and interprocess communication to be expressed explicitly.

Prolog Programming Environments

We also have a practical interest in constructing programming environments for Prolog. We have built interactive interfaces to the EMACS text editor, a system which displays debugging information graphically [4], a LINT-style utility that checks the internal consistency of Prolog programs, a library facility that allows easy construction of multi-module programs, documentation and typesetting aids in the style of Knuth's "literate programming", and are investigating the practical application of type inference schemes. We have also constructed a number of Prolog interpreters.

Facilities

The Computer Science Department in general, and the Knowledge Science Institute in particular, is well-equipped for research. The main departmental workhorses are several SUN-4 systems running UNIX (one dual-processor VAX-780 remains as a relic from the olden days). There are a number of high-performance workstations (Symbolics and Xerox Lisp machines for AI, IRISes for graphics) as well as many general-purpose SUN workstations. A network of Apollo workstations resides in the KSI, together with a robot arm and the usual assortment of Macintoshes and IBM PC's. A BBN Butterfly is used for work on parallel algorithms. We have access to sophisticated helmet-mounted eye-tracking equipment in the Faculty of Physical Education. The university possesses super-computing facilities (Cyber 205) and supports a large Honeywell Multics operation.

Software resources include Prolog, Scheme, and Smalltalk implementations on most machines (including, for many researchers, Macintoshes at home). The NEXPERT expert system shell is used for expert system prototyping and development. Important components of our support environment include Gnu EMACS, TeX, and electronic mail. Virtually all computers are fully networked with high-speed (10Mbits/s) links. In the summer of 1988 most large departmental research projects, including the AI research, are moving to newly-available space, which is being specially modified for our needs, and will be extensively wired with network and other communication facilities.

Researchers

Computer Science faculty engaged in this research include John Cleary, Brian Gaines (founder and director of the KSI), Paul Kwok, Bruce MacDonald, Mildred Shaw, and Ian Witten. Faculty from other departments who are associated with the KSI are Bruce Clark, Institute for Computer-Assisted Learning; Larry Katz, Faculty of Physical Education; Doug Norrie, Department of Mechanical Engineering; and Joan Vickers, Faculty of Physical Education. Notable visitors who have spent extended periods in the KSI recently include John Andreae, University of Canterbury, New Zealand; Tim Bell, University of Canterbury, New Zealand; Harold Thimbleby, University of York, UK; and Brian Woodward, Knowledge Science Institute.

References

[1] Bell, T.C., Cleary, J.G. and Witten, I.H. *Text Compression,* Prentice Hall, Englewood Cliffs, NJ, in press.

[2] Cleary, J.G. "Acquisition of uncertain rules in a probabilistic logic" *International Journal of Man-Machine Studies, 27(2),* 1987, pp 145-154.

[3] Cleary, J.G. "Logical Arithmetic" *Future Computing Systems, 2(2),* 1987, pp 125-149.

[4] Dewar, A. and Cleary, J.G. "Graphical display of complex information within a Prolog debugger" *International Journal of Man-Machine Studies, 25(5),* 1986, pp 503-521.

[5] Gaines, B.R. "Foundations of knowledge engineering" *Research and Development in Expert Systems III,* Bramer, M.A. (editor), University Press, Cambridge, 1986, pp 13-24.

[6] Gaines, B.R. "An overview of knowledge acquisition and transfer" *International Journal of Man-Machine Studies, 26(4),* 1987, pp 453-472.

[7] Gaines, B.R. "A conceptual framework for person-computer interaction in distributed systems" *IEEE Transactions on Systems, Man & Cybernetics,* in press.

[8] Gaines, B.R. "Knowledge acquisition systems for rapid prototyping of expert systems" *INFOR*, in press.

[9] Gaines, B.R. and Shaw, M.L.G. "From timesharing to the sixth generation: the development of human-computer interaction part I" *International Journal of Man-Machine Studies, 1(1),* 1986, pp 1-27.

[10] Gaines, B.R. and Shaw, M.L.G. "Foundations of dialog engineering: the development of human-computer interaction part II" *International Journal of Man-Machine Studies, 1(2),* 1986, pp 101-123.

[11] Gaines, B.R. and Shaw, M.L.G. "Induction of inference rules for expert systems" *Fuzzy Sets and Systems, 18(3),* 1986, pp 315-328.

[12] Krawchuk, B.J. and Witten, I.H. "Explanation-based learning: its role in problem-solving" *Journal of Experimental and Theoretical Artificial Intelligence*, in press.

[13] MacDonald, B.A. and Witten, I.H. "A framework for knowledge acquisition through techniques of concept learning" *IEEE Transactions on Systems, Man, and Cybernetics* (special issue on knowledge engineering), in press.

[14] Shaw, M.L.G. "Interactive elicitation and exchange of knowledge in group problem solving" *International*

Journal of Personal Construct Psychology, 1987.
[15] Shaw, M.L.G. "An Interactive knowledge-based system for group problem solving" *IEEE Transactions on Systems, Man, and Cybernetics*, 1987.
[16] Shaw, M.L.G. and Gaines, B.R. "Interactive elicitation of knowledge from experts" *Future Computing Systems, 1(2),* 1986, pp 151-190.
[17] Shaw, M.L.G. and Gaines, B.R. "Advances in interactive knowledge engineering" *Research and Development in Expert Systems III,* Bramer, M.A.(editor), University Press, Cambridge, 1986, pp. 111-122.
[18] Shaw, M.L.G. and Gaines, B.R. "KITTEN: Knowledge initiation and transfer tools for experts and novices" *International Journal of Man-Machine Studies, 27 (3),* 1987, pp 251-280.
[19] Vickers, J.N. *Instructional Design in Sport and Physical Activity: A Knowledge Structures Approach,* Human Kinetics, Champaign, Illinois, 1987.
[20] Vickers, J.N. and Gaines, B.R. "A Comparison of Books and Hypermedia for Knowledge-based Sports Coaching" *Microcomputers for Information Management, 5(1),* 1988, pp 29-44.
[21] Witten, I.H. "In search of 'autonomy'." *Future Computing Systems, 2(1),* 1987, pp 35-53.
[22] Witten, I.H. "Computer (in)security-infiltrating open systems" *Abacus, 4(4),* 1987, pp 6 - 25.
[23] Witten, I.H. and Cleary, J.G. "Foretelling the future by adaptive modeling" *Abacus, 3(3),* 1986, pp 16-36.
[24] Witten, I.H., Neal, R., and Cleary, J.G. "Arithmetic coding for data compression" *Communications of the Association for Computing Machinery, 30(6),* 1987, pp 520-540.
[25] Witten, I.H. and MacDonald, B.A. "Using concept learning for knowledge acquisition" *International Journal of Man-Machine Studies*, in press.

AI Research and Development at CompEngServ, of the CEMTECH Group Ltd.

by Archie Bowen

Recherche et développement en IA à CompEngServ, du Groupe CEMTECH Ltée.

RÉSUMÉ: CompEngServ Ltée., une compagnie entièrement contrôlée par des intérêts canadiens qui fait maintenant partie du Groupe CEMTECH Ltée., dispose d'une expertise dans la conception et l'implantation de logiciel temps réel et de logiciel d'intelligence artificielle. Les projects complétés ou en cours incluent des systèmes pour la gestion de la recofiguration des communications, la résolution de conflits en aviation, la résolution de l'interférence électromagnétique, la mesure de la qualité de l'air intérieur, le chargement en billets automatisé pour les caissiers automatiques, la simulation de jeux de guerre et le diagnostic de systèmes de communication.

CompEngServ Ltd., now part of the CEMTECH Group Ltd., is wholly Canadian owned and has special expertise in the design and implementation of real-time and artificial intelligence software. CompEngServ was originally established in 1975 by Dr. B.A. Bowen in Ottawa to service industry and government.

Our professional employees are in the fields of engineering, computing science and psychology. All expert systems and AI projects are located in our Ottawa office.

Completed Expert System Projects

Our expert systems group has completed several projects since 1983; these are discussed below.

Planning Systems — *Communications Reconfiguration Management*

The Reconfiguration Management Expert System was designed and implemented for a U.S. Air Force early warning radar surveillance system. This knowledge-based system uses intelligent databases and rules derived from the experience of senior managers. It supports the anticipated extensive modifications to the external flight data communications requirements of this new radar technology during the design, implementation and operational phases.

Dr. Bowen is the Chief Operating Officer of the CEMTECH Group Ltd., Ottawa, Ontario.

The system represents a unique integration of expert systems and databases — the state of the art in the use of a frame-based intelligent database and a knowledge base.

This system was implemented in Goldworks, a Lisp-based tool containing an extensive database, help menus and about 250 rules. The number of rules is augmented by the intelligence of the object-oriented database, and is not a true indication of the inferencing complexity.

Another consideration when building the expert system was the capacity to create scenarios for training new operators; situations could be proposed and analyzed as a method of providing a tutorial. Each prepared situation is analyzed, explained and the rules used to create the correct configuration given to the trainee.

Transportation-Related Systems — *Aircraft Conflict Resolution*

This work involved a detailed analysis of the way in which air traffic controllers operate. The results of this study were then analyzed further to determine which functions were best executed by classical software and which by knowledge bases. The findings were then organized into an overall algorithm which was implemented and tested against cases proposed by controllers.

The system was implemented in TIMM, an inductive reasoning tool. This provided high performance (ap-