Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Winter 1996 No. 38 hiver 1996

An official publication of CSCSI, the Canadian Society for Computational Studies of Intelligence
Une publication officielle de la SCEIO, la Société canadienne pour l'étude de l'intelligence par ordinateur

Special Issue: Machine Learning

Edition speciale: L'apprentissage machine

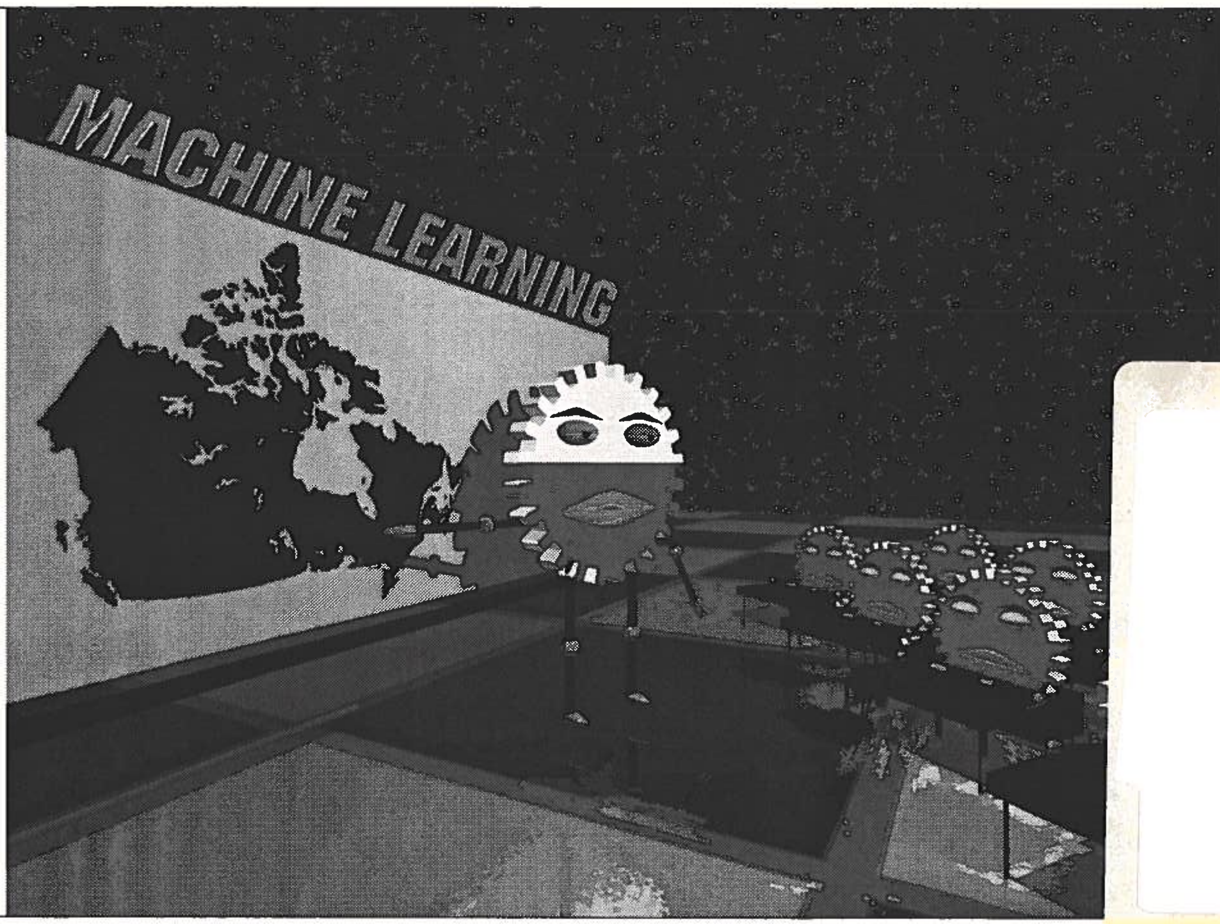

Cover graphic design by Paul Amirault

AI '96

Canadian Artificial Intelligence Conference

21- 24 May 1996

Program Chairperson
Gordon McCalla
Dept of Computer Science
University of Saskatchewan
Saskatoon, SK S7N 5A9
email: mccalla@cs.usask.ca
Phone: 306-966-4902
Fax: 306-966-4884

Program Committee
Marie-Michelle Boulet, Université Laval
Renato de Mori, McGill University
Mark Fox, University of Toronto
Lev Goldfarb, University of New Brunswick
Jim Greer, University of Saskatchewan
Howard Hamilton, University of Regina
Graeme Hirst, University of Toronto
Rob Holte, University of Ottawa
Dekang Lin, University of Manitoba
Jim Little, University of British Columbia
Joel Martin, National Research Council
Paul McFetridge, Simon Fraser University
Bob Mercer, University of Western Ontario
Evangelos Milios, York University
Peter Patel-Schneider, AT&T Bell Labs
Fei Song, University of Guelph
Andre Trudel, Acadia University
Qiang Yang, University of Waterloo

AI '96 is the eleventh biennial conference on Artificial Intelligence sponsored by the Canadian Society for Computational Studies of Intelligence. It will be held at the Holiday Inn on King, Toronto, Ontario during 21-24 May 1996 in conjunction with Vision Interface and Graphics Interface (AI/GI/VI/ '96).

AI/GI/VI '96 is a unique event with three Canadian research conferences that present the latest results in artificial intelligence, computer graphics and computer vision. Each conference offers three concurrent days of invited and submitted papers. For a single registration fee, conference participants can attend presentations in any of the three, promoting the exchange of knowledge among these important disciplines. Two days have been set aside for workshops and other events. A banquet and electronic theatre provide additional opportunities to meet speakers and other attendees for informal discussion in a social setting.

For further information about AI '96, please contact **Gordon McCalla**, the Program Chairperson.

http://ai.iit.nrc.ca/cscsi/conferences/ai96.html

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Winter 1996 | **No. 38** | **hiver 1996**

Canadian Artificial Intelligence is published three times a year by the Canadian Society for Computational Studies of Intelligence (CSCSI). *Intelligence Artificielle au Canada* est publiée trimestriellement par la Société canadienne pour l'étude de l'intelligence par ordinateur (SCEIO). Canadian Publications Mail Product Sales Agreement No. 507032.

ISSN 0823-9339

Contents | Contenu

Canadian Artificial Intelligence welcomes submissions on any matter related to artificial intelligence. Please send your contribution, electronic preferred, with an abstract, a photograph, and a short bio to:

Dr. Peter Turney
Co-Editor, *Canadian Artificial Intelligence*
Institute for Information Technology
National Research Council Canada
M-50 Montreal Road
Ottawa, Ontario, Canada
K1A 0R6 or — peter@ai.iit.nrc.ca
Telephone (613) 993-8564 Fax (613) 952-7151

Advertising rates are available upon request.

Arlene Merling
Managing Editor, *Canadian Artificial Intelligence*
Alberta Research Council
3rd Floor, 6815 - 8 Street N.E.
Calgary, Alberta, Canada
T2E 7H7 or — merling @arc.ab.ca
Telephone (403) 297-2608 Fax (403) 297-2505

Book reviews and candidate books to review should be sent to:

Graeme Hirst
Canadian Artificial Intelligence
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4 or — gh@ai.toronto.edu

Recycled / Recyclable

Stan Matwin

Special Issue: Machine Learning

Machine Learning (ML) is one of the most active areas of Artificial Intelligence. At the last International Joint Conference on Artificial Intelligence (IJCAI-95) in Montreal, ML was the largest represented subfield of Artificial Intelligence in terms of the number of papers in the technical programme. In the opinion of many, there is no intelligence without the ability to learn from a changing environment. In his 1991 paper [Schank 91], Roger Schank said while discussing different views on Artificial Intelligence: "The fourth view of Artificial Intelligence is one that I myself have espoused. It is a long-term view. It says that intelligence entails learning."

In Canada, historically there has been a fair amount of interest and good work in ML. The last two CSCSI conferences, in 1992 in Vancouver and in 1994 in Banff, both included ML workshops that attracted reasonably large numbers of participants. This interest inspired the editors of *Canadian Artificial Intelligence* to publish a special issue on ML. I agreed to act as the guest editor of this issue. In general, my idea was to produce a magazine issue which will give the readership a good idea of the scope and depth of machine learning work performed in Canada, or at least in those labs and groups that have answered the call for papers. The format was fairly open, and consequently it has resulted in a variety of styles of submissions. These can be roughly divided into two categories: longer position papers and shorter research summaries. I think both kinds of papers describe the Canadian scene. Research summaries also contain references to more technical articles that describe the work in more detail. As is to be expected in a collection of individual research topics, there is no single underlying emerging theme. It is interesting, however, that practically all papers mention specific applications as an important component of the work. This is a good sign of vigour and dynamism of our field in Canada. Applications often attract new, motivated students, bring in additional funds, and, last but not least, prove the usefulness of our work to the outside world.

I would like to thank all the contributors for their enthusiasm and help in preparing this special issue. We all hope that the readers will find it interesting, and that it may help promote new applications and attract new students to the field of ML.

[Schank 91] Schank, R., "Where's the AI," *AI Magazine*, vol. 12 No. 4, pp. 38-49.

Call for Nominations
CSCSI Distinguished Service Award

This award is presented biannually to an individual who has made outstanding contributions to the Canadian artificial intelligence community in one or more of the following areas:

- Community service
- Research
- Training of students
- Research/Industry interaction

The award, which will be presented at CSCSI '96, provides:

- honorary lifetime membership in CSCSI
- conference fees when the award is presented
- a token gift.

The first award was presented in 1992 to Professor John Mylopoulos of the University of Toronto. Alan Mackworth was the 1994 recipient.

Recommendations should include a brief (one page) summary of the nominee's qualifications for receiving the award. The final decision will be made by the CSCSI executive.

Recommendations for the award should be addressed to :

Stan Matwin
Computer Science Dept.,
University of Ottawa,
Ottawa, Ontario K1N 6N5
E-mail: stan@ csi.uottawa.ca

Deadline for receipt of nominations is **March 31, 1996.**

It is very much hoped that the winner will be present to receive the award in person, and we ask nominators to make discreet enquiries in advance to ensure that there is a reasonable likelihood of their nominee attending the conference.

Canadian Artificial Intelligence Magazine Goes Electronic!

Peter Turney and Suhayya Abu-Hakima

On August 23rd, 1995, at the CSCSI/SCEIO Annual Meeting in Montreal, the decision was made (by a large majority vote) to propel *Canadian Artificial Intelligence* into the future. Starting with this issue, *Canadian Artificial Intelligence* magazine is now available on the World Wide Web. The magazine will continue to be available in the printed form that you hold in your hands now, for those traditionalists among us.

All current members of CSCSI/SCEIO are entitled to World Wide Web access to *Canadian Artificial Intelligence*, at no extra charge. If you would like to try electronic access, please send an E-mail message to peter@ai.iit.nrc.ca,with the following information:

(1) your full name
(2) your E-mail address,
(3) a short user name (5-8 characters, for web access)
(4) a password (5-8 characters, for web access), and
(5) your CSCSI/SCEIO membership identification number, which can be found on the mailing label of the magazine.

When it is time for you to renew your membership, you will have the option of continuing to receive a printed copy of *Canadian Artificial Intelligence* by regular mail, at the current membership rates, or you may choose to use web access only, at reduced rates.

Our new rates are: Printed **Copy and Web Access (all current members):**

- regular member: $40
- CIPS member: $30
- student member: $15

Web Access Only (the new option, when renewing):

- regular member: $30
- CIPS member: $25
- student member: $15

The electronic form of *Canadian Artificial Intelligence* is a PostScript file. When you print the PostScript file on your local printer, the result should be essentially the same as the copy you are reading now.

As you can see, the printed form of *Canadian Artificial Intelligence* has also changed. *Canadian Artificial Intelligence* magazine is the major expense of CSCSI/SCEIO and the cost of the magazine currently exceeds the revenue from membership fees (see the Treasurer's Report in issue #36). Instead of increasing membership fees, we have decided to reduce our printing costs. We believe that you will find the quality acceptable. This change will enable us to balance our budget, which has been in deficit. Web access means reduced costs, reduced membership fees, faster distribution, and less harm to the environment. In these days of fiscal restraint, we believe you will all be happy to discover that the membership fees have not increased.

Call for Papers: "Software Agents"

Canadian Artificial Intelligence is seeking articles or research reports for an upcoming special issue on software agents (final papers due May 1, 1996).

Innes A. Ferguson
Institute for Information Technology
National Research Council Canada
M-50 Montreal Road.
Ottawa, Ontario K1A 0R6
Telephone (613) 993-8553
Email: innes@ai.iit.nrc.ca

Contributions must be in electronic format. Acceptable formats are Microsoft Word or Wordperfect, for Macintosh or PC.

Knowledge Mining in Databases: An Integration of Machine Learning Methodologies with Database Technologies*

Jiawei Han, Yongjian Fu, Krzysztof Koperski, Gabor Melli, Wei Wang, Osmar R. Zaîane

Résumé

Les chercheurs de notre groupe travaillent activement depuis des années à la découverte de connaissances dans les bases de données. Plusieurs résultats intéressants ont été publiés, et un prototype de système de découverte de connaissances, DBMiner (auparavant appelé DBLearn), a été développé et présenté dans plusieurs conférences. Notre recherche couvre un large spectre de la découverte de connaissances, dont 1) l'étude de la découverte de connaissances dans les bases de données relationnelles, orientées objet, déductives, spatiales, et actives, et dans les systèmes globaux d'information, et 2) le développement de méthodes diverses de découverte de connaissances, dont l'induction orientée attribut, le creusage progressif pour l'extraction de règles à niveaux multiples, l'extraction de connaissances guidée par méta-règle, etc. On a aussi étudié dans notre recherche diverses techniques pour la découverte de diverses sortes de connaissances, dont la généralisation, la caractérisation, la discrimination, l'association, la classification, le regroupement, etc., et l'application de la découverte de connaissances à la réponse intelligente aux requêtes, à la construction de bases de données à couches multiples, etc.

Abstract

Active research has been conducted on knowledge discovery in databases by the researchers in our group for years, with many interesting results published and a prototyped knowledge discovery system, **DBMiner** (previously called **DBLearn**), developed and demonstrated in several conferences. Our research covers a wide spectrum of knowledge discovery, including (1) the study of knowledge discovery in relational, object-oriented, deductive, spatial, and active databases, and global information systems, and (2) the development of various kinds of knowledge discovery methods, including attribute-oriented induction, progressive deepening for mining multiple-level rules, meta-rule guided knowledge mining, etc. Techniques for the discovery of various kinds of knowledge, including generalization, characterization, discrimination, association, classification, clustering, etc. and the application of knowledge discovery for intelligent query answering, multiple-layered database construction, etc. have also been studied in our research.

Introduction

With the rapid growth of the number of databases and the tremendous amounts of data being collected and stored in databases, it is increasingly important to develop software tools to assist in the extraction of "information" or "knowledge" from data, understanding the implications of data in databases, and automatic construction of knowledge-bases from databases. The research into knowledge discovery in databases (or data mining) [1, 18] has attracted wide attention in both academia and industry.

The Knowledge Discovery research group in the **Database Systems Research Laboratory** of the School of Computing Science, Simon Fraser University, has been working in this promising research field for several years and has contributed to the field in the following aspects.

1. The development of a prototyped knowledge discovery system, **DBMiner** (previously named **DBLearn**) [8, 10], which integrates machine learning methodologies with database technologies and discovers different kinds of knowledge from large databases efficiently and effectively. The system may discover different kinds of knowledge, including characteristic, discriminant, association, and classification rules using a set of knowledge discovery methods, originated from our own research, including attribute-oriented induction [3], progressive deepening for mining multiple-level rules [6], meta-rule guided knowledge mining [2], etc.

2. The study of knowledge discovery in different kinds of databases [14,12, 17, 13, 20], including knowledge discovery in relational, object-oriented, deductive, spatial, and active databases, and global information systems, and the application of knowledge discovery for intelligent query answering [11], multiple-layered database construction [9], etc.

The remainder of the paper is organized as follows. A more detailed description of the **DBMiner** system is presented in Section 2. Knowledge mining in advanced database systems and knowledge discovery applications are outlined in Section 3. A summary and a discussion of our ongoing research is presented in Section 4.

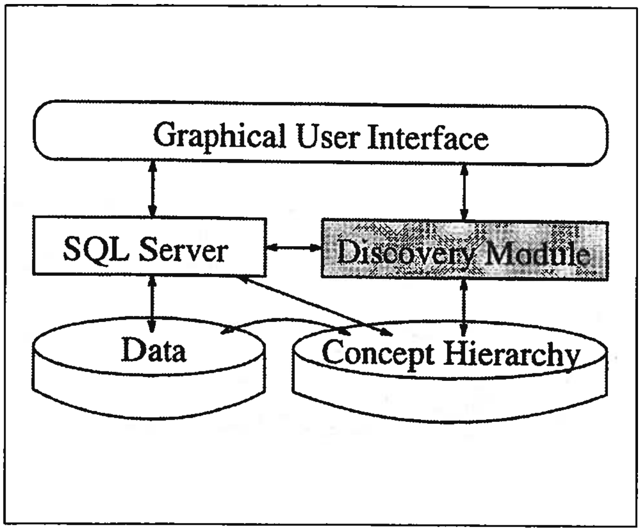

Figure 1. General architecture of ***DBMiner***

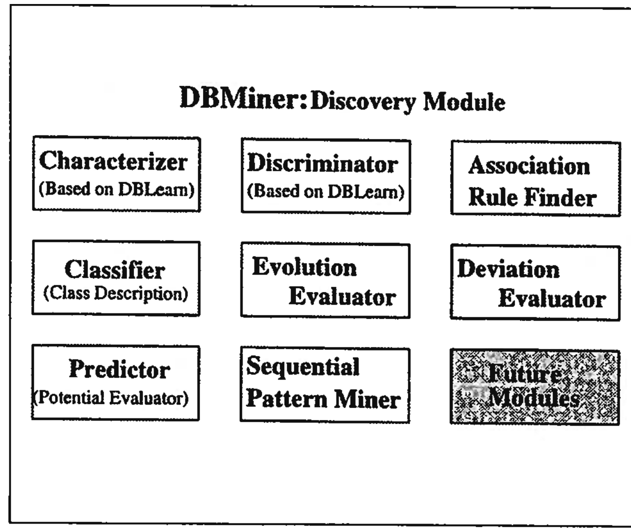

Figure 2. Function modules of ***DBMiner***

DBMiner: A database mining system

DBMiner, a comprehensive database mining system prototype, has been developed in Simon Fraser University. It started with an interesting induction method, attribute-oriented induction [3, 4], for learning characteristic rules and discriminant rules in relational databases. The method resulted in an early version of the system, **DBLearn** [4, 8]. Experiments with DBLearn were performed in NSERC (Natural Science and Engineering Research Council of Canada) research grant information database and in several large industrial databases with successful results and good performance. Further extensions and enhancements of the DBLearn system since 1993 have led to a new generation of the system: DBMiner [7]. DBMiner consists of several new functional modules besides the characterizer and discriminator in DBLearn. It performs dynamic adjustment of concept hierarchies and automatic generation of numeric hierarchies. It discovers different kinds of knowledge rules and generates different forms of outputs, including generalized relations, generalized feature tables, and multiple forms of generalized rules. Moreover, system performance has been improved, graphical user interfaces have been enhanced for interactive knowledge mining, and a client/ server architecture has been constructed for industrial applications.

The major features of the system include the integration of machine learning and database technologies, high speed and efficiency in analyzing large databases, interactive knowledge mining, and smooth integration with commercial relational database systems.

Figure 1 shows the general architecture of DBMiner which integrates a relational database system, such as a Sybase/ Oracle SQL server, with the discovery module. The core of the DBMiner system is the discovery module, which is further detailed in Figure 2. It consists of multiple functional modules, including characterized discriminator, association rule finder, classifier, evolution evaluator, etc.

The functionalities of the first three modules are described as follows:

- The **characterizer** [4] discovers a set of **characteristic rules** from the relevant set of data in a database. A characteristic rule summarizes the general characteristics of a set of user-specified data.
- A **discriminator** [4] discovers a set of **discriminant rules** from the relevant set(s) of data in a database. A discriminant rule distinguishes the general features of one set of data, called the *target class,* from some other set(s) of data, called the *contrasting class(es).*
- An **association rule finder [6]** discovers a set of association rules (in the form of $A_1 \wedge \cdots \wedge A_i \rightarrow B_1 \wedge \cdots \wedge B_j$) at multiple concept levels from the relevant set(s) of data in a database. For example, it may find from a large set of transaction data an association rule, such as if a customer buys *(one brand of)* milk, s/he usually buys *(another brand of)* bread.

In the development of other functional modules, **attribute-oriented induction** [4, 7] also plays an essential role. It integrates a machine learning paradigm *learning-from-examples* [15] with set-oriented database operations and substantially reduces the computational complexity of database learning processes. Moreover, the generalized relation can be further analyzed by integration with other machine learning methods [7], including ID-3 [19], Cluster-2 [16], etc.

The system also performs *automatic generation of conceptual hierarchies* for numerical attributes and *dynamic conceptual hierarchy adjustment* [5] for all the attributes based on the statistical distribution of the set of relevant data, which produces desirable generalized results.

DBMiner offers both graphical and SQL-like interfaces [7]. For example, to characterize *Computer Science* grants in the *NSERC94* database in relevance to discipline and

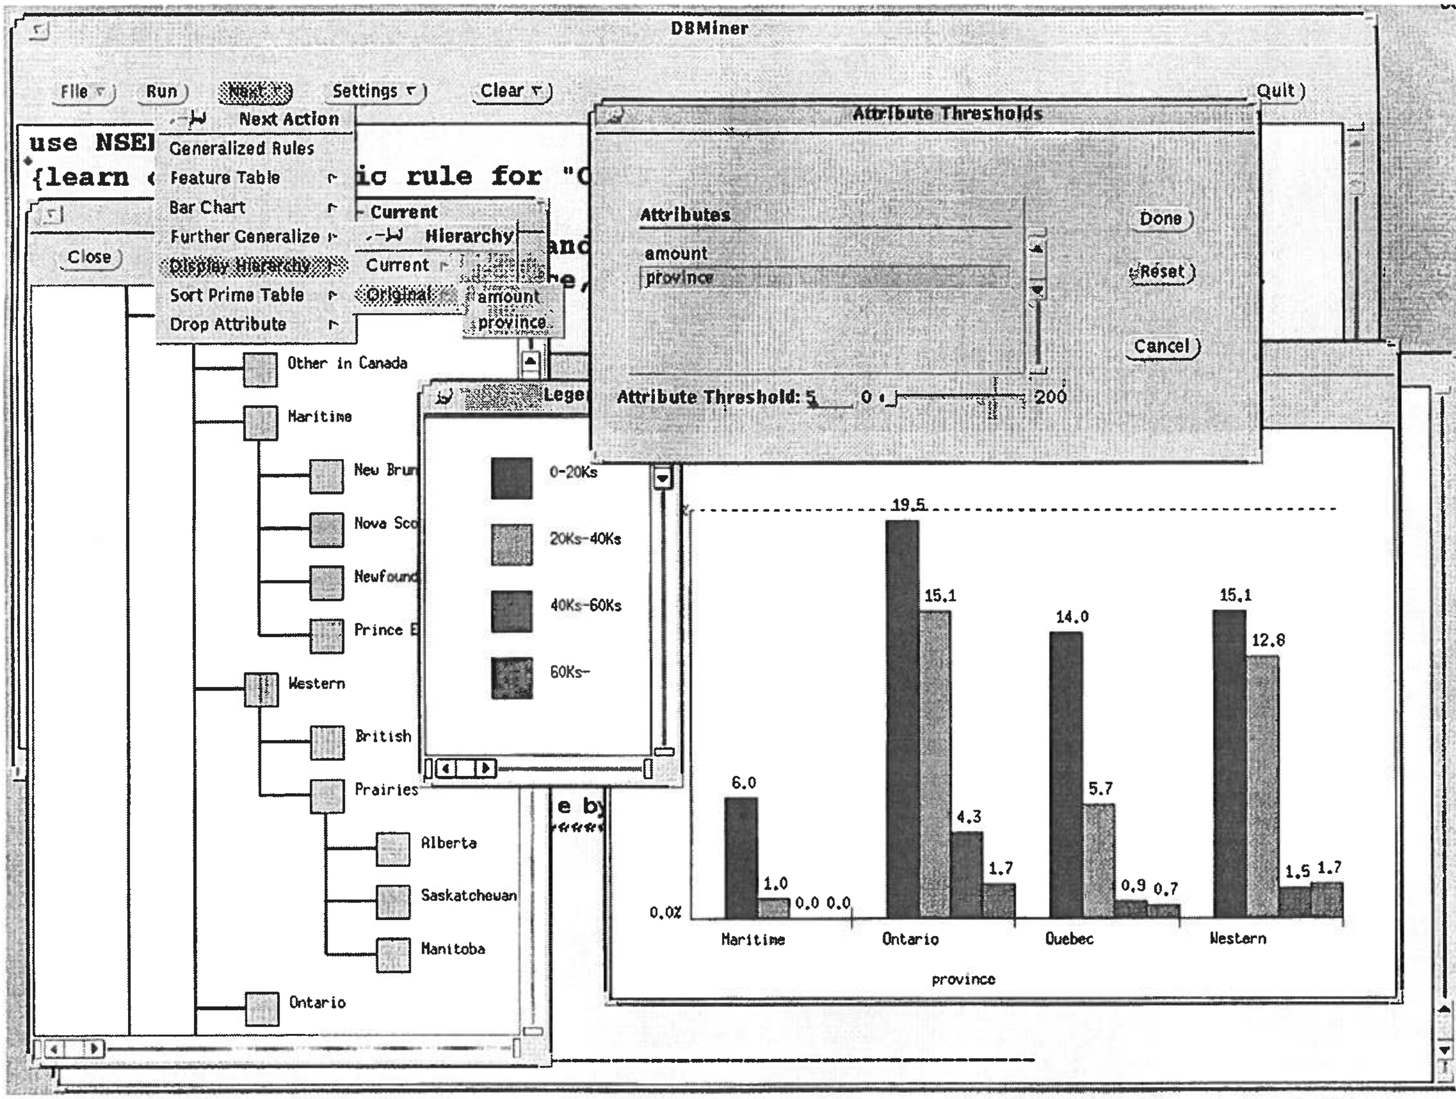

Figure 3: A snapshot of the execution of ***DBMiner.***

amount categories and the distribution of count% and amount%, the data mining query is as follows:

use NSERC94
characterize "CS_Discipline_Grants"
from award A, grant_type G
where A.grant_code = G.grant_code **and** A.disc code = "Computer"
in relevance to disc_code, amount, **percentage** (count), **percentage** (amount)

To process this query, the system first obtains the relevant set of data by processing a relational database query, then generalizes the data using an **attribute-oriented induction** approach [3, 4], and presents different forms of outputs. The output outlines the number or amount distribution of computer science (research) grants according to discipline categories (such as *theory, AI, database,* and so on). The output forms include generalized relations, generalized feature tables, bar charts, and generalized rules.

A snapshot of the execution of DBMiner is presented in Figure 3. The DBMiner system is accessible with the Internet address: http ://www.dbg.sfu.ca/dbl/dbminer or http ://fas.sfu .ca/cs/research/groups/DB/dbminer .

Knowledge discovery in advanced database systems and knowledge discovery applications

Besides knowledge discovery in large relational databases, investigations have also been performed on efficient and effective methods for knowledge discovery in object-oriented databases [12], spatial databases [13, 14, 17], active databases [12], deductive databases [2], transaction databases [6], and global information systems [20]. Three of them are outlined below to convey the ideas.

For knowledge discovery in object-oriented databases [12], techniques have been studied on generalization of complex objects, attributes and methods, class and aggregation hierarchies, spatial and multimedia data, etc. After generalization of complex objects and structures into relational-like data, most techniques developed for data mining in relational systems can be applied for knowledge mining in object-oriented databases.

For knowledge discovery in spatial databases [13, 14, 17], a set of techniques have been developed in our research, including nonspatial-data-dominant generalization [14], spatial-data-dominant generalization [14], spatial data clustering [17], and spatial feature association [13].

For resource and knowledge discovery in global information systems (Internet) [20], a multiple-layered database structure has been proposed, which first transforms the low-level, highly unstructured, widely distributed, heterogeneous global information-base into relatively structured higher layer databases by various kinds of generalization techniques and then constructs a multiple-layered information-base to facilitate resource and knowledge discovery. Preliminary experiments on a small subset of a networked information-base have demonstrated some limited success of the approach [20].

Knowledge discovery has strong application potential. Besides the use of discovered knowledge for decision making, process control and knowledge-base construction, we have investigated its application in intelligent query answering [11] and multiple-layered database construction [9]. Database queries can be answered intelligently by analyzing the intent of query and providing generalized, neighbourhood or associated information using stored or discovered knowledge. Knowledge discovery substantially broadens the spectrum of intelligent query answering and provides discovered knowledge and knowledge discovery tools, which include generalization, data summarization, concept clustering, rule discovery, query rewriting deduction, lazy evaluation, construction and application of multiple-layered databases, etc.

A summary of our ongoing work

Our research progress on knowledge discovery in databases has greatly motivated our ongoing work in this direction. These ongoing studies are illustrated as follows.

• Further enhancement of the power and efficiency of the knowledge discovery mechanisms [7], including the improvement of rule quality and system performance for the existing functional modules, the development of techniques for mining new kinds of rules, etc.

• Further development of high-level, user-friendly interfaces for interactive knowledge mining. This includes multiple platforms of knowledge mining interfaces, including X-window-oriented, PC-window-oriented, and Netscape (WWW) oriented interfaces, and visual presentation of the discovered knowledge.

• Extension of data mining techniques to advanced and/or special purpose database systems, including further studies on spatial data mining, knowledge mining in heterogeneous databases, and knowledge discovery in global information systems. This may lead to the generation of a set of new, special-purpose data mining systems, such as **GeoMiner, WebMiner,** etc.

References

[1] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R.Uthurusamy. *Advances in Knowledge Discovery and Data Mining.* AAAI/MIT Press, 1995.

[2] Y. Fu and J. Han. Meta-rule-guided mining of association rules in relational databases. In *Proc. 1st Int'l Workshop on Integration of Knowledge Discovery with Deductive and Object-Oriented Databases (KDOOD'95),* Singapore, Dec. 1995.

[3] J. Han, Y. Cai, and N. Cercone. Knowledge discovery in databases: An attribute-oriented approach. In *Proc. 18th Int. Conf. Very Large Data Bases,* pages 547-559, Vancouver, Canada, August 1992.

[4] J. Han, Y. Cai, and N. Cercone. Data-driven discovery of quantitative rules in relational databases. *IEEE Trans. Knowledge and Data Engineering,* 5:29-40, 1993.

[5] J. Han and Y. Fu. Dynamic generation and refinement of concept hierarchies for knowledge discovery in databases. In *Proc. AAAI'94 Workshop on Knowledge Discovery in Databases (KDD'94),* pages 157-168, Seattle, WA, July 1994.

[6] J. Han and Y. Fu. Discovery of multiple-level association rules from large databases. In *Proc. 1995 Int. Conf. Very Large Data Bases,* Zurich, Switzerland, Sept. 1995.

[7] J. Han and Y. Fu. Exploration of the power of attribute-oriented induction in data mining. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining.* AAAI/MIT Press, 1995.

[81 J. Han, Y.Fu, Y. Huang, Y. Cai, and N. Cercone. DBLearn: A system prototype for knowledge discovery in relational databases. In *Proc. 1994 ACM-SIGMOD Conf. Management of Data,* page 516, Minneapolis, MN, May 1994.

[9] J. Han, Y. Fu, and R. Ng. Cooperative query answering using multiple-layered databases. In *Proc. 2nd Int. Conf. Cooperative Information Systems,* pages 47-58, Toronto, Canada, May 1994.

[10] J. Han, Y. Fu, and S. Tang. Advances of the DBLearn system for knowledge discovery in large databases. In *Proc. 1995 Int. Joint Conf. on Artificial Intelligence,* Montreal, Canada, Aug. 1995.

[11] J. Han, Y. Huang, N. Cercone, and Y. Fu. Intelligent query answering by knowledge discovery techniques. In *IEEE Trans. Knowledge and Data Engineering (accepted),* 1995.

[12] J. Han, S. Nishio, and H. Kawano. Knowledge discovery in object-oriented and active databases. In

F. Fuchi and T. Yokoi, editors, *Knowledge Building and Knowledge Sharing*, pages 221-230. Ohmsha, Ltd. and IOS Press, 1994.

[13] K. Koperski and J. Han. Discovery of spatial association rules in geographic information databases. In *Proc. 4th Int'l Symp. on Large Spatial Databases (SSD '95)*, Portland, Maine, Aug. 1995.

[14] W. Lu, J. Han, and B. C. Ooi. Knowledge discovery in large spatial databases. In *Far East Workshop on Geographic Information Systems*, pages 275-289, Singapore, June 1993.

[15] R. S. Michalski. A theory and methodology of inductive learning. In Michalski et al., editor, *Machine Learning: An Artificial Intelligence Approach, Vol. 1*, pages 83-134. Morgan Kaufmann, 1983.

[16] R. S. Michalski and R. Stepp. Automated construction of classifications: Conceptual clustering versus numerical taxonomy. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 5: 396-410, 1983.

[17] R. Ng and J. Han. Efficient and effective clustering method for spatial data mining. In *Proc. 1994 Int. Conf. Very Large Data Bases*, pages 144-155, Santiago, Chile, September 1994.

[18] G. Piatetsky-Shapiro and W. J. Frawley. *Knowledge Discovery in Databases.* AAAI/MIT Press, 1991.

[19] J. R. Quinlan. *C4.5: Programs for Machine Learning.* Morgan Kaufmann, 1992.

[20] O. R. Zaiane and J. Han. Resource and knowledge discovery in global information systems: A preliminary design and experiment. In *Proc. 1st Int'l. Conf. on Knowledge Discovery and Data Mining (KDD '95)*, Montreal, Canada, Aug. 1995.

**Research is partially supported by the Natural Sciences and Engineering Research Council of Canada under the grant OGP0037230, by the Networks of Centres of Excellence Program (with the participation of PRECARN association) under the grants IRIS:HMI-5 and IRIS:IC-2, and by a research grant from the Hughes Research Laboratories.*

***Jiawei Han** is a Professor of Computing Science, Simon Fraser University. His major research interests include database and knowledge-base systems, knowledge discovery in databases, deductive and object-oriented databases, spatial and multi-media databases, logic programming, and artificial intelligence. He is the leader of the Canadian IRIS project IRIS:HMI-5 ("Data Mining and Knowledge Discovery in Large Databases"), and has served or is currently serving in the program committees of over 20 international conferences, including ICDE '95 (also, program committee vice-chairman), DOOD '95, ACM-SIGMOD '96, and VLDB '96.*

***Yongjian Fu.** PhD. student, Computing Science, Simon Fraser University. He received MSc. and BSc. degrees at Zhejiang University, China, and worked as a lecturer there before joining SFU. He is the major developer of the DBMiner system and works on research into knowledge discovery from databases and applications of knowledge discovery systems.*

***Krzysztof Koperski.** Ph.D. student, Computing Science, Simon Fraser University. His major research focus is on spatial data mining. He is also interested in spatial reasoning and spatial object-oriented databases.*

***Gabor Melli.** M.Sc. student, Computing Science, Simon Fraser University. He is an expert on systems and has contributed to the system support of the relational database system engine. His current research focuses on automatic prediction of attribute-value behaviour.*

***Wei Wang.** M.Sc. student, Computing Science, Simon Fraser University. He has been implementing the GUI interface for the DBMiner system on PCs and is performing research on knowledge discovery in heterogeneous databases.*

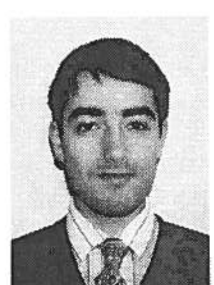

***Osmar R. Zaiane.** PhD. student, Computing Science, Simon Fraser University. He holds an M.Sc. degree in Computer Science from Laval University (Quebec, Canada) where he worked on mobile databases stored on smart-cards. He also holds an M.Sc. degree in Electronics from Paris XI University (Paris, France) where he worked on Natural Language Interfaces for textual databases. His current research focuses on knowledge discovery in global network information systems.*

Classification conceptuelle incrémentale et applications

Guy Mineau, Rokia Missaoui, Hafedh Mili, Robert Godin

Abstract

Over the past few years, we have explored a family of combinatorial structures as a basis for conceptual clustering. Those structures find a unifying theoretical foundation in the notion of Galois lattice (concept lattice). This paper is a brief overview of our work in this area concerning in particular the definition of several variations, the incremental generation of these structures and applications.

Résumé

Nous avons exploré au cours des dernières années une famille de structures combinatoires comme base de méthodes de classification conceptuelle. Ces structures trouvent un fondement théorique unifiant dans la notion de treillis de Galois (treillis de concepts). Cet article fait état des travaux que nous avons effectués dans cette direction concernant la définition de diverses variantes de la notion de treillis de Galois, la génération incrémentale de ces structures et diverses applications.

1. Introduction

Le treillis de Galois d'une relation binaire (ou treillis de concepts) est à la base d'une famille de méthodes de classification conceptuelle. Introduite par Barbut et Monjardet, cette notion a été popularisée par Wille qui a utilisé le treillis de Galois comme base de l'analyse formelle de concepts (Wille, 1982). Wille et son équipe ont initialement trouvé des applications en intelligence artificielle pour la représentation et l'acquisition de connaissances (Ganter, Stahl & Wille, 1986; Wille, 1989). Wille propose de considérer chaque élément du treillis comme un concept formel et le graphe (diagramme de Hasse) comme une relation de généralisation/spécialisation entre les concepts. Le treillis est donc perçu comme une hiérarchie de concepts. Chaque concept est une paire composée d'une extension représentant un sous-ensemble des instances de l'application et d'une intention représentant les propriétés communes aux instances. Diverses variantes de cette structure de base ont été proposées en incorporant des représentations plus riches du domaine tel que l'utilisation d'attributs (Ganter et al., 1986; Godin, Saunders & Gecsei, 1986), de graphes conceptuels (Godin, Mineau & Missaoui, 1995c; Godin, Mineau, Missaoui & Mili, 1995d; Mineau & Godin, 1994) et de connaissances taxinomiques (Godin & Mili, 1993; Mineau & Godin, 1994) . La recherche de représentations plus compactes de la hiérarchie a conduit à introduire le concept de treillis d'héritage qui élimine certains types de redondance inhérente à la structure de treillis de Galois (Godin & Mili, 1993; Godin, Mineau, Missaoui, St-Germain & Faraj, 1995b; Godin et al., 1995d). Ceci peut être combiné à la notion de treillis élagué (Godin & Mili, 1993; Godin et al., 1995b; Godin et al., 1995d) qui ne conserve que le sous-ensemble des concepts dont les propriétés ne sont pas toutes obtenues par héritage. Ces variations peuvent être combinées de façon orthogonale produisant une large famille de structures dont l'intérêt a été démontré dans une variété importante de domaines d'application.

Du point de vue du domaine de l'apprentissage machine, la construction automatique d'une de ces structures peut être vue comme une méthode de classification (regroupement) conceptuelle (Michalski, Stepp & Diday, 1981). Comme des algorithmes incrémentaux efficaces ont été développés, la construction incrémentale peut être considérée comme une méthode incrémentale de formation de concepts (Fisher, Pazzani & Langley, 1991). Ces méthodes font partie de l'apprentissage non supervisé parce que les concepts ne sont pas prédéterminés et les instances utilisées pour l'apprentissage ne sont pas pré-classifiées (Carbonell, 1990). Contrairement aux hiérarchies produites d'autres méthodes telles que UNIMEM (Lebowitz, 1987), COBWEB (Fisher, 1987) et CLASSIT (Gennari, Langley & Fisher, 1990) , le treillis de Galois et ses variantes ont une structure formellement définie qui ne dépend pas de divers paramètres, de l'ordonnancement des instances ou de particularités algorithmiques. Une autre caractéristique distinctive est que la hiérarchie n'est pas nécessairement réduite à une structure d'arbre. La hiérarchie obtenue met en évidence de façon exhaustive les regroupements potentiellement intéressants par rapport aux observations. En contraste avec les méthodes habituelles de formation de concepts qui cherchent à produire un nombre limité de concepts en optimisant le contraste et la cohérence entre les classes, le nombre de concepts dans un treillis de Galois est souvent plus grand et la qualité des concepts plus variable. Carpineto et Romano (1993) ont démontré l'utilité du treillis pour la découverte et la prédiction de classes à partir de plusieurs corpus de données du domaine de l'apprentissage machine.

2. Complexité et algorithmes incrémentaux

Dans le pire cas, le treillis de Galois (G) peut croître exponentiellement avec le nombre d'instances et de propriétés. Cependant, lorsqu'il y a une borne supérieure, disons k, sur le nombre de propriétés reliées à une instance, ce qui est le cas dans les applications habituellement considérées, la taille du treillis est bornée linéairement par rapport au nombre d'instances, n (Godin et al., 1986) : $|G| \leq 2^k n$. De plus, malgré le facteur exponentiel en k, des expériences avec de nombreuses applications et des résultats théoriques dans un cadre de distribution aléatoire uniforme montrent que le rapport $|G|/n$ est assez stable et de beaucoup inférieur à 2^k. En pratique, dans des applications de recherche documentaire, nous avons toujours observé que $|G| \leq k' n$, où k' est la valeur moyenne du nombre de propriétés par instance. De plus, des estimations théoriques toujours dans un cadre de distribution uniforme suggèrent une croissance linéaire par rapport à k'. Des analyses détaillées de la complexité du treillis sont présentées dans (Godin, 1989; Godin, Missaoui & April, 1993).

Pour plusieurs applications la complexité du treillis de Galois est raisonnable. Cependant pour d'autres applications, il est avantageux de considérer une structure moins riche qui peut être obtenue à un coût moindre. Dans cette perspective, la variation appelée treillis élagué comporte une complexité de beaucoup inférieure pour certaines configurations de données tout en conservant une représentation complète des données de base. Le nombre d'éléments dans un treillis élagué (GE) est borné par le nombre total de propriétés (Mineau & Godin, 1994). Avec une borne supérieure k sur le nombre de propriétés par instance, le nombre total de propriétés est bornée par kn et en conséquence : $|GE| \leq k$ n. Dans plusieurs applications, la croissance asymptotique de $|GE|$ devient logarithmique en n conformément à la loi de Bradford-Zipf (Brookes, 1977).

De nombreux algorithmes ont été proposés pour générer le treillis de Galois à partir d'une relation binaire (Bordat, 1986; Carpineto & Romano, 1993; Chein, 1969; Fay, 1975; Ganter, 1984; Godin, Missaoui & Alaoui, 1995; Malgrange, 1962; Norris, 1978). Parmi ces algorithmes, ceux proposés dans (Carpineto & Romano, 1993; Godin et al., 1995) permettent la mise à jour incrémentale et du treillis et du graphe de Hasse en incorporant les nouvelles instances une par une, ce qui est important pour plusieurs applications. L'algorithme de Bordat (Bordat, 1986) construit aussi le graphe de Hasse mais n'est pas incrémental. L'algorithme de Norris est incrémental mais ne construit pas le graphe de Hasse (Guénoche, 1990). Nous avons expérimenté plusieurs implémentations d'algorithmes incrémentaux, et des données empiriques provenant de plusieurs applications ont démontré que l'adjonction d'une nouvelle instance peut être faite en temps linéaire par rapport au nombre d'instances lorsque l'hypothèse d'une borne supérieure sur le nombre de propriétés reliées à une instance est respectée (Godin et al., 1995b; Godin, Missaoui & Alaoui, 1991; Godin et al., 1995; Godin et al., 1993). Avec cette hypothèse, l'analyse dans le pire des cas montre aussi une complexité linéaire. Une présentation détaillée des algorithmes incrémentaux ainsi qu'une comparaison avec la performance d'algorithmes non incrémentaux sont données dans (Godin et al., 1995). Un résultat surprenant est qu'une variante simple d'un algorithme incrémental s'est avérée meilleure que les algorithmes non incrémentaux testés dans la plupart des cas lorsque l'on cumule le temps pour mettre à jour incrémentalement le treillis (Godin et al., 1995).

Le treillis élagué peut être généré facilement à partir du treillis de Galois. Cependant, il est possible de générer le treillis élagué plus efficacement directement. Les premiers algorithmes proposés ont été des algorithmes non incrémentaux (Mineau, Gecsei & Godin, 1990). Récemment, nous avons proposé un algorithme incrémental pour générer le treillis élagué directement (Godin et al., 1995c). Diverses applications ont montré que la génération incrémentale du treillis élagué est toujours plus efficace que pour le treillis complet (Godin, Mili, Arfi, Mineau & Missaoui, 1995a; Godin et al., 1995b; Godin et al., 1995c). La différence peut être très variable et dépend essentiellement de la différence entre le nombre d'éléments des deux structures. Dans certains cas cette, différence est limitée (Godin et al., 1995a); dans d'autres cas, elle est très importante (Godin et al., 1995c). Un autre algorithme pour la mise à jour incrémentale du treillis élagué a été proposé dans (Dicky, Dony, Huchard & Libourel, 1994).

3. Applications

Les dernières années ont vu apparaître un volume croissant de littérature sur des applications novatrices des méthodes de classification conceptuelle présentées précédemment (Carpineto & Romano, 1993; Ganter & Wille, 1989; Godin & Mili, 1993; Godin & Missaoui, 1994; Godin et al., 1993; Guénoche & Van Mechelen, 1993; Krone & Snelting, 1994; Mephu Nguifo, 1993; Mineau, Godin & Missaoui, 1993; Oosthuizen, Bekker & Avenant, 1992). Dans cette section, nous survolons brièvement les applications et projets de recherche à l'intérieur desquels nous avons expérimenté ces approches.

Une première application que nous avons considérée est la recherche documentaire (Godin, Gecsei & Pichet, 1989; Godin et al., 1986). Dans ce contexte, le treillis est généré à partir de la relation d'indexage entre documents et termes d'index. Le graphe de Hasse est utilisé comme structure de base pour supporter une interface de navigation permettant à l'utilisateur de graduellement élargir ou spécialiser sa requête en terme des sous-ensembles de documents et de termes présents dans le treillis. Des expériences contrôlées ont mis en évidence le potentiel de cette méthode par rapport à la recherche booléenne et la navigation dans une classification hiérarchique pour le problème de recherche d'un sujet particulier dans le contexte d'un système public (Godin et al., 1993) et d'un système personnel. (Godin, Davidson, Missaoui & Mili, 1993).

Dans une optique similaire, la notion d'espace de connaissance a été proposée comme support au repérage d'images représentées par des graphes conceptuels (Mineau & Godin, 1992). La structure d'espace de connaissance peut être définie à partir de la notion de treillis d'héritage élagué généré à partir des graphes conceptuels.

Un domaine d'application dans lequel nous avons oeuvré récemment est le génie logiciel. Dans un vaste projet de développement d'outils pour un atelier de génie du logiciel, le projet mobilisateur Le Macroscope Informatique, fruit d'une collaboration entre l'industrie, des universités et le Centre de Recherche en Informatique de Montréal (CRIM), nous avons utilisé le treillis de Galois et une variante, le treillis élagué, pour la réutilisation (Godin et al., 1995b). La classification conceptuelle sert comme outil de repérage et d'abstraction d'artefacts plus génériques avec une meilleure réutilisabilité. Dans la même veine, une approche plus ambitieuse à la réutilisation consiste à identifier des familles larges de composantes reliées entre elles qui peuvent être extraites et réutilisées en bloc. Dans (Mineau et al., 1993), la notion d'espace de connaissance est appliquée pour induire une hiérarchie de modèles génériques à partir de modèles existants. Le regroupement conceptuel devient un outil d'assistance à l'analyse du domaine, discipline qui est de plus en plus considérée comme primordiale au processus de réutilisation (Prieto-Diaz & Arango, 1991).

Dans le contexte du développement orienté objet, le treillis de Galois et le treillis élagué peuvent servir comme base pour des outils de génération et de réorganisation des hiérarchies de classes à partir de la spécification de leurs propriétés (Dicky et al., 1994; Godin & Mili, 1993). Les algorithmes que nous avons développés permettent de générer la hiérarchie de façon efficace et incrémentale si nécessaire. Dans le cadre du vaste projet de recherche et de développement portant sur le développement de logiciel orienté objet dans le contexte d'applications de gestion de réseau, le projet IGLOO (InGénierie du Logiciel Orienté Objet), résultat d'une collaboration industrie-université-CRIM, nous avons développé un outil qui permet d'appliquer les diverses méthodes de classification au problème de conception des hiérarchies de classes en fonction des propriétés spécifiées (Godin, Mili, Arfi, Mineau & Missaoui, 1994). Diverses métriques peuvent être calculées pour évaluer les hiérarchies obtenues. L'outil peut aussi extraire l'interface d'un ensemble de classes de l'environnement Smalltalk-80™ et produire la hiérarchie d'interface à partir des algorithmes de génération du treillis et du treillis élagué (Arfi, Godin, Mili, Mineau & Missaoui, 1994).

Un autre domaine d'application prometteur est la découverte de connaissances à partir d'une base de données. À cet effet, le treillis de Galois est une représentation efficace pour l'extraction de règles d'implication entre les propriétés d'une relation (Guigues & Duquenne, 1986; Wille, 1992). Des algorithmes ont été développés pour générer les règles à partir du treillis. Pour éviter la reconstruction des règles à chaque mise à jour de la base de données, un algorithme permettant de modifier le treillis et les règles correspondantes de façon incrémentale a été proposé (Godin & Missaoui, 1994). L'analyse des algorithmes permet d'entrevoir l'application à des volumes de données suffisamment larges pour être intéressants en pratique.

Remerciements

Ces travaux ont été supportés en partie par des subventions provenant du Conseil National de la Recherche du Canada, du Centre de Recherche en Informatique de Montréal dans le cadre du projet mobilisateur Le Macroscope Informatique géré par le Groupe DMR Inc., et du Ministère de l'Industrie, du Commerce, de la Science et de la Technologie du Québec, pour le projet IGLOO organisé par le Centre de Recherche en Informatique de Montréal.

Références

Arfi, A., Godin, R., Mili, H., Mineau, G. & Missaoui, R. (1994). A Tool for the Automatic Generation of the Interface Hierarchy for the Smalltalk-80 Class Library. In *Proceedings of the Acfas Colloquium on Object Orientation in Databases and Software Engineering - Colloque Orientation Objet en Bases de Données et Génie du Logiciel*, Montreal: pp. 61-78.

Barbut, M. & Monjardet, B. (1970). *Ordre et Classification. Algèbre et Combinatoire, Tome II.* Hachette.

Bordat, J. P. (1986). Calcul Pratique du Treillis de Galois d'une Correspondance. *Mathématiques et Sciences Humaines*, **96**, 31-47.

Brookes, B. C. (1977). Theory of the Bradford Law. *Journal of Documentation*, **33**(3), 180-209.

Carbonell, J. G. (1990). Introduction: Paradigms for Machine Learning. In J. G. Carbonell (Eds.), *Machine Learning: Paradigms and Methods,* (pp. 1-9). Cambridge, Mass.: The MIT Press.

Carpineto, C. & Romano, G. (1993). GALOIS: An Order-Theoretic Approach to Conceptual Clustering. In *Proceedings of the Machine Learning Conference*, pp. 33-40.

Chein, M. (1969). Algorithme de Recherche des Sous-Matrices Premières d'une Matrice. *Bull. Math. Soc. Sci. Math. R.S. Roumanie*, **13**, 21-25.

Dicky, H., Dony, C., Huchard, M. & Libourel, T. (1994). Un algorithme d'insertion avec restructuration dans les hiérarchies de classes. In *Proceedings of the Langages et Modèles à Objets*, Grenoble:

Fay, G. (1975). An Algorithm for Finite Galois Connexions. *Journal of Computational Linguistic and Languages*, **10**, 99-123.

Fisher, D. (1987). Knowledge Acquisition via Incremental Conceptual Clustering. *Machine Learning*, **2**, 139-172.

Fisher, D. H., Pazzani, M. J. & Langley, P. (Ed.). (1991). *Concept Formation: Knowledge and Experience in Unsupervised Learning*. San Mateo, CA: Morgan Kaufmann.

Ganter, B. (1984). *Two Basic Algorithms in Concept Analysis.* Preprint 831, Technische Hochschule Darmstadt.

Ganter, B., Stahl, J. & Wille, R. (1986). Conceptual Measurement and Many-Valued Contexts. In W. Gaul & M. Schader (Eds.), *Classification as a Tool of Research,* (pp. 169-176). Amsterdam: North-Holland.

Ganter, B. & Wille, R. (1989). Conceptual Scaling. In F. Roberts (Eds.), *Applications of Combinatorics and Graph Theory to the Biological and Social Sciences,* (pp. 139-167). New York: Springer-Verlag.

Gennari, J. H., Langley, P. & Fisher, D. (1990). Models of Incremental Concept Formation. In J. Carbonell (Eds.), *Machine Learning: Paradigms and Methods,* (pp. 11-62). Amsterdam, The Netherlands: MIT Press.

Godin, R. (1989). Complexité de Structures de Treillis. *Annales des Sciences Mathématiques du Québec*, **13**(1), 19-38.

Godin, R., Davidson, M., Missaoui, R. & Mili, H. (1993). Experimental Comparison of three Methods for Personal Information Retrieval. In *Proceedings of the Second Annual Symposium on Document Analysis and Information Retrieval*, Las Vegas, Nevada: pp. 379-398.

Godin, R., Gecsei, J. & Pichet, C. (1989). Design of a Browsing Interface for Information Retrieval. In *Proceedings of the 12th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, N. J. Belkin & C. J. van Rijsbergen (Ed.), Cambridge, MA: ACM SIGIR Forum, pp. 32-39.

Godin, R. & Mili, H. (1993). Building and Maintaining Analysis-Level Class Hierarchies Using Galois Lattices. In *Proceedings of the ACM Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'93)*, A. Paepcke (Ed.), Washington, DC: ACM Press, pp. 394-410.

Godin, R., Mili, H., Arfi, A., Mineau, G. W. & Missaoui, R. (1994). A Tool for Building and Evaluating Class Hierarchies Based on a Concept Formation Approach. In *Proceedings of the OOPSLA 94 Workshop on Artificial Intelligence for Object-Oriented Software Engineering*, Portland, Oregon:

Godin, R., Mili, H., Arfi, A., Mineau, G. W. & Missaoui, R. (1995a). Design of Class Hierarchies based on Concept (Galois) Lattices. *Submitted for publication.*

Godin, R., Mineau, G., Missaoui, R., St-Germain, M. & Faraj, N. (1995b). Applying Concept Formation Methods to Software Reuse. *International Journal of Knowledge Engineering and Software Engineering*, **5**(1), 119-142.

Godin, R., Mineau, G. W. & Missaoui, R. (1995c). Incremental structuring of knowledge bases. In *Proceedings of the International Knowledge Retrieval, Use, and Storage for Efficiency Symposium (KRUSE'95)*, Santa Cruz: pp. Accepted.

Godin, R., Mineau, G. W., Missaoui, R. & Mili, H. (1995d). Méthodes de Classification Conceptuelle Basées sur les Treillis de Galois et Applications. *Revue d'Intelligence Artificielle*, À paraître.

Godin, R. & Missaoui, R. (1994). An Incremental Concept Formation Approach for Learning from Databases. *Theoretical Computer Science, Special Issue on Formal Methods in Databases and Software Engineering*, **133**, 387-419.

Godin, R., Missaoui, R. & Alaoui, H. (1991). Learning Algorithms Using a Galois Lattice Structure. In *Proceedings of the Third International Conference on Tools for Artificial Intelligence*, S. Lee, B. Wah, N. G. Bourbakis, & W. T. Tsai (Ed.), San Jose, Calif.: IEEE Computer Society Press, pp. 22-29.

Godin, R., Missaoui, R. & Alaoui, H. (1995). Incremental Concept Formation Algorithms Based on Galois (Concept) Lattices. *Computational Intelligence*, **11**(2), 246-267.

Godin, R., Missaoui, R. & April, A. (1993). Experimental Comparison of Navigation in a Galois Lattice with Conventional Information Retrieval Methods. *International Journal of Man-Machine Studies*, **38**, 747-767.

Godin, R., Saunders, E. & Gecsei, J. (1986). Lattice Model of Browsable Data Spaces. *Information Sciences*, **40**, 89-116.

Guénoche, A. (1990). Construction du Treillis de Galois d'une Relation Binaire. *Mathématiques et Sciences Humaines*, **109**, 41-53.

Guénoche, A. & Van Mechelen, I. (1993). Galois Approach to the Induction of Concepts. In I. Van Mechelen,J. Hampton,R. Michalski, & P. Theüns (Eds.), *Categories and Concepts: Theoretical Views and Inductive Data Analysis,* (pp. 287-308). Academic Press.

Guigues, J. L. & Duquenne, V. (1986). Familles Minimales d'Implications Informatives Résultant d'un Tableau de Données Binaires. *Mathématiques et Sciences Humaines*, **95**, 5-18.

Krone, M. & Snelting, G. (1994). On The Inference of Configuration Structures from Source Code. In *Proceedings of the 16th International Conference on Software Engineering*, IEEE Computer Society Press, pp. 49-57.

Lebowitz, M. (1987). Experiments with Incremental Concept Formation: UNIMEM. *Machine Learning*, **2**, 103-138.

Malgrange, Y. (1962). Recherche des Sous-Matrices Premières d'une Matrice à Coefficients Binaires; Applications à Certains Problème de Graphes. In *Proceedings of the Deuxième Congrès de l'AFCALTI*, Gauthier-Villars, pp. 231-242.

Mephu Nguifo, E. (1993). *Concevoir une Abstraction à Partir de Ressemblances.* Doctorat Thesis, Montpellier II - Sciences et Techniques du Languedoc.

Michalski, R. S., Stepp, R. E. & Diday, E. (1981). A Recent Advance in Data Analysis: Clustering Objects into Classes Characterized by Conjunctive Concepts. In L. N. Kanal & A. Rosenfeld (Eds.), *Progress in Pattern Recognition,* (pp. 33-56). North-Holland Publishing Company.

Mineau, G., Gecsei, J. & Godin, R. (1990). Structuring Knowledge Bases using Automatic Learning. In *Proceedings of the IEEE Sixth Int'l Conf. on Data Engineering*, Los Angeles, CA: IEEE Computer Society Press, pp. 274-280.

Mineau, G. & Godin, R. (1992). Automatic Knowledge Structuring for Browsing Retrieval. In *Proceedings of the International Conference on Information and Knowledge Management*, Y. Yesha (Ed.), Baltimore, MD: The International Society for Mini and Microcomputers - ISMM, pp. 273-281.

Mineau, G., Godin, R. & Missaoui, R. (1993). Induction of Generic Data Models by Conceptual Clustering. In *Proceedings of the The Fifth International Conference on Software Engineering & Knowledge Engineering*, San Francisco: pp. 554-564.

Mineau, G. W. & Godin, R. (1994). Automatic Structuring of Knowledge Bases by Conceptual Clustering. *IEEE Transactions on Knowledge and Data Engineering*, To appear.

Norris, E. M. (1978). An Algorithm for Computing the Maximal Rectangles in a Binary Relation. *Revue Roumaine de Mathématiques Pures et Appliquées*, **23**(2), 243-250.

Oosthuizen, G. D., Bekker, C. & Avenant, C. (1992). Managing Classes in Very Large Class Repositories. In *Proceedings of the Tools*, Paris: pp. 625-633.

Prieto-Diaz, R. & Arango, G. (Ed.). (1991). *Domain Analysis and Software Systems Modeling*. IEEE Computer Society Press.

Wille, R. (1982). Restructuring Lattice Theory: an Approach Based on Hierarchies of Concepts. In I. Rival (Eds.), *Ordered Sets*, (pp. 445-470). Dordrecht-Boston: Reidel.

Wille, R. (1989). Knowledge Acquisition by Methods of Formal Concept Analysis. In E. Diday (Eds.), *Data Analysis, Learning Symbolic and Numeric Knowledge*, (pp. 365-380). New York: Nova Science Pub.

Wille, R. (1992). Concept Lattices and Conceptual Knowledge Systems. *Computers Mth. Applic.*, **23**(6-9), 493-515.

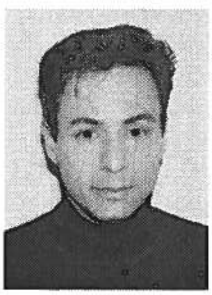

***Guy Mineau** is currently an associate professor at Université Laval, Dept. of Computer Science, where he has conducted his research activities since 1990. His research interests stem mainly from two areas: machine learning (conceptual clustering) and knowledge modeling (using conceptual graphs). Prof. Mineau has been supported by NSERC (Natural Science and Engineering Research Council of Canada) and by FCAR (Fonds pour la formation de Chercheurs et l'Aide à la Recherche) since 1990. He has been and still is involved in major R&D research projects involving academic institutions, business partners, and research centers such as CRIM (Centre de Recherche Informatique de Montréal).*

***Rokia Missaoui** received a Ph.D degree in computer science from Université de Montréal in 1988. In 1987, she joined the Department of Computer Science at Université du Québec à Montréal (UQAM) where she is currently an associate professor. Her current research interests concern: semantic query optimization in deductive databases, object-oriented database modelling and management, conceptual clustering, and knowledge discovery from databases. She has been involved in three important industrial projects during the last five years. Rokia Missaoui is funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and Fonds pour la Formation de Chercheurs et l'Aide à la Recherche (FCAR). She is member of the ACM and the IEEE Computer Society.*

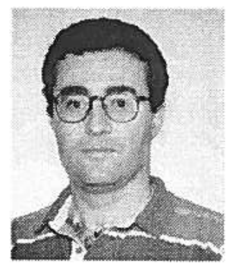

***Hafedh Mili** is an associate professor of computer science at Université du Québec à Montréal (UQAM). He earned an engineering diploma from Ecole Centrale de Paris (1984), a Ph.D. in Computer Science from the George Washington University (1988), and has been at UQAM since. He is interested in OO software engineering, software reuse, AI, and information retrieval, and has published close to 40 referreed journal and conference papers in these areas. He has participated and/or led a number of R&D projects on OO and software reuse funded jointly by the government and various companies (BNR, DMR, Tandem, DEC, CAE, IBM, etc.).*

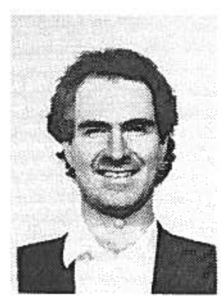

***Robert Godin** is currently a Computer Science professor at Université du Québec à Montréal (UQAM). He earned a Ph.D. degree in Computer Science from Université de Montréal in 1986. His research interests include machine learning, databases, information retrieval and OO software engineering. He has published more than thirty papers in refereed journals and conferences in these areas. His work is supported by the Natural Science and Engineering Research Council of Canada (NSERC) since 1987. He also has been and still is involved in major R&D research projects involving academic institutions, business partners, and research centers such as the Centre de Recherche en Informatique de Montréal (CRIM) and the Centre de Recherche en Analyse de Texte (ATO) of UQAM. He was member of the scientific committee of the CRIM from 1988 to 1993.*

Of Mushrooms and Machine Learning: Identifying Algorithms in a PDP Network

Michael R. W. Dawson & David A. Medler

Introduction

Les chercheurs connexionnistes admettent librement qu'il est extrêmement difficile de déterminer comment les réseaux connexionnistes accomplissent la tâche qui leur a été apprise. "Une chose que les réseaux connexionnistes ont en commun avec le cerveau est que si vous les ouvrez et que vous jetez un coup d'oeil à l'intérieur, tout ce que vous voyez est un gros tas de bouillie" (Mozer et Smolensky, 1989, p.3). De façon similaire, Seidenberg (1993, p. 229) dit: "Si le but de la modélisation par simulation est de clarifier les constructions théoriques existantes, le connexionnisme semble être exactement la mauvaise voie à suivre. Les modèles connexionnistes ne clarifient pas les idées théoriques, elles les obscurcissent." Il en résulte que les connexionnistes ne sont pas très disposés à analyser la structure interne de leurs réseaux pour déterminer l'algorithme qui les sous-tend.

Malheureusement, ce manque d'enthousiasme soulèvent de sérieux doutes quant à l'aptitude des connexionnistes à produire des théories fructueuses sur le traitement cognitif. Parce que les chercheurs comprennent rarement le fonctionnement interne des modèles à traitement parallèle des données, McCloskey (1991) suggère que "les réseaux connexionnistes ne devraient pas être vus comme des théories des fonctions cognitives humaines, ni comme des simulations de théories, ni même comme des démonstrations de points théoriques spécifiques" (p. 387).

La plus grande partie de la recherche dans le projet "Biological Computation Project" (calcul informatique biologique, NdT) à l'Université d'Alberta vise à corriger cette situation. Nous concentrons nos efforts de recherche sur l'interprétation d'un type particulier d'architecture à traitement parallèle des données, les réseaux d'unités de valeur, dans le but de maximiser leur contribution à la science cognitive. Un réseau d'unités de valeur est une extension du réseau à propagation arrière standard qui utilise une fonction d'activation non-monotonique (i.e. gaussienne) dans ses couches cachées et de sortie plutôt qu'une fonction monotonique (par exemple logistique) (voir Dawson et Schopflocher, 1992). Les réseaux d'unités de valeur possèdent plusieurs avantages par rapport aux réseaux à propagation arrière standard, dont un apprentissage plus rapide des classes linéairement inséparables, une meilleure généralisation, et on peut plus facilement augmenter leur taille (voir Dawson et Schopflocher, 1992; Dawson et Shamanski, 1994; Medler et Dawson, 1994). De plus, les unités de valeur ont la caractéristique de pouvoir être interprétés très directement, sans détours.

Introduction

Connectionist researchers freely admit that it is extremely difficult to determine how connectionist networks accomplish the task that they have been taught. "One thing that connectionist networks have in common with brains is that if you open them up and peer inside, all you can see is a big pile of goo" (Mozer & Smolensky, 1989, p.3). Similarly, Seidenberg (1993, p. 229) states that "if the purpose of simulation modeling is to clarify existing theoretical constructs, connectionism looks like exactly the wrong way to go. Connectionist models do not clarify theoretical ideas, they obscure them." As a result, connectionists are reluctant to analyze the internal structure of their networks in an attempt to determine what algorithm lies within.

Unfortunately, this reluctance has raised serious doubts concerning the ability of connectionists to provide fruitful theories about cognitive processing. Because researchers rarely understand the internal workings of PDP models, McCloskey (1991) suggested that "connectionist networks should not be viewed as theories of human cognitive functions, or as simulations of theories, or even as demonstrations of specific theoretical points" (p. 387).

Much of the research at the Biological Computation Project at the University of Alberta aims to rectify this situation. We are focusing our research efforts on the interpretation of a particular type of PDP architecture—networks of value units—in an attempt to maximize their contribution to cognitive science. A value unit network is an extension of the standard back propagation network that uses a nonmonotonic (i.e., Gaussian) activation function in its hidden and output layers instead of a monotonic (e.g., logistic) function (see Dawson & Schopflocher, 1992). Value unit networks have several advantages over standard back propagation networks, including faster learning of linearly inseparable classes, better generalization, and better "scaling up" properties (e.g., Dawson & Schopflocher, 1992; Dawson, & Shamanski, 1994; Medler & Dawson, 1994). Furthermore, value units have a characteristic that makes them very straightforward to interpret.

Characterizing Hidden Units with Jittered Density Plots

Consider using a set of patterns to train a network of value units. After training, one could again present each pattern to the network and record the activity that each pattern produced in each hidden unit. This amounts to "wiretapping" each hidden unit while the stimulus set is being presented. One could use this information to create a jittered density plot for each hidden unit (e.g., Figure 1).

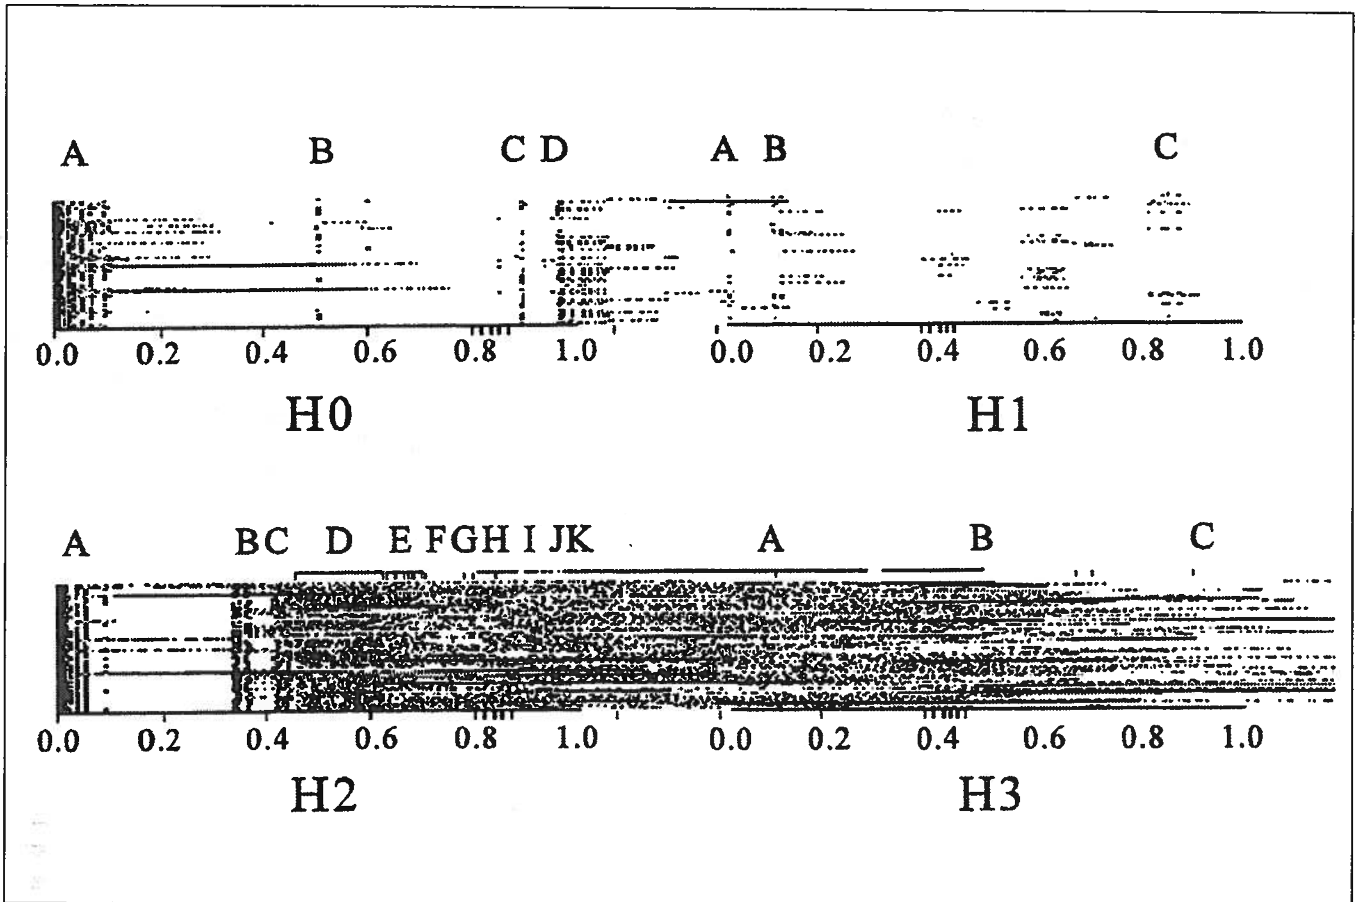

Figure 1. Jittered density plots for a network with four hidden units trained on the mushroom problem that is described later in the text. Each hidden unit is given its own density plot. Letters are used to label the bands that were analyzed for definite features, as described in the text below.

The jittered density plots in Figure 1 illustrate the result of wiretapping a network that consisted of four different hidden units. In each of the density plots, a single dot is used to represent the activity in that hidden unit produced by presenting one pattern to the network. The density plots in Figure 1 were produced in a network trained on 8124 different input patterns; consequently, each density plot is composed of 8124 different dots. The horizontal position of each plotted point along the x-axis represents the activation produced by one of the training patterns, whereas the y-axis represents a random jittering introduced to prevent points from overlapping (Chambers, Cleveland, Kleiner, & Tukey, 1983, pp. 19-21).

Berkeley, Dawson, Medler, Schopflocher, and Hornsby (1995) found that in many cases, the density plots for hidden value units are highly structured. Density plots of value unit activations typically reveal distinct "bands," as is the case for the four density plots in Figure 1. (Berkeley et al., 1995, did not find that this structure was typical of density plots for standard back propagation networks using the logistic activation function.)

This "banding" provides an important method for interpreting the kinds of features to which a hidden unit is sensitive. Each band in such a density plot supports a coherent interpretation: each pattern that falls into a band is characterized by definite features—a specific feature or set of features. These definite features can be quickly identified by calculating simple descriptive statistics. To illustrate the utility of identifying definite features, let us consider an example problem.

The Mushroom Problem

Providing a biological classification of a mushroom (i.e., correctly identifying a mushroom's genus and species) requires paying attention to a wide variety of features, some that are apparent from a casual inspection, others that are apparent only after microscopic examination. Box 1 provides an algorithm for correctly classifying all mushrooms within Schlimmer's (1987) data set.

Schlimmer's data set consisted of the hypothetical description of 23 different mushrooms in the Agaricus and Lepiota family (see Lincoff, 1981, pp. 500-525). Each mushroom was described as a set of 21 different features. Multiple featural descriptions of one species of mushroom were possible because one species might be found in several different habitats, have more than one possible odour, etc. The total data set consisted of 8124 different instances. 4208 of these patterns corresponded to edible mushrooms;

the remaining 3916 training patterns corresponded to inedible mushrooms (i.e., mushrooms that were definitely poisonous, or were of unknown edibility and therefore not recommended).

The purpose of the current experiment was to determine if an artificial neural network could learn to identify correctly a mushroom as edible or not. In particular, we were interested in seeing whether—after the network converged—we could determine the rules that it used to classify mushrooms. The network that we trained had four hidden value units, and one output value unit.

The network was trained to generate a response of "1" to an edible mushroom, and a response of "0" to a poisonous mushroom. Initial connection weights for the network were randomly selected from the range [-1.0,1.0]. Biases for each hidden unit and for the output unit were set to 0, and were not modified during training. The network was trained using Dawson and Schopflocher's (1992) modified generalized delta rule, with a learning rate of 0.01 and no momentum. Network convergence (i.e., a hit on every pattern) was achieved after 108 sweeps through the training set.

Box 1. Algorithm for identifying mushrooms

Step 1 *What is the odor of the mushroom?*

```
If it is almond or anise then it is edible.
If it is creosote or fishy or foul or musty or pungent or
spicy
  then it is poisonous.
If it has no odor then proceed to Step 2.
```

Step 2 *Obtain the spore print of the mushroom.*

```
If the spore print is black or brown or buff or chocolate or
  orange or yellow then it is edible.
If the spore print is green or purple then it is poisonous.
If the spore print is white then proceed to Step 3.
```

Step 3 *Examine the gill size of the mushroom.*

```
If the gill size is broad, then it is edible.
If the gill size is narrow, then proceed to Step 4.
```

Step 4 *Examine the stalk surface above the mushroom's ring.*

```
If the surface is fibrous then it is edible.
If the surface is silky or scaly then it is poisonous.
If the surface is smooth the proceed to Step 5.
```

Step 5 *Examine the mushroom for bruises.*

```
If it has no bruises then it is edible.
If it has bruises then it is poisonous.
```

Interpretation of the Network

The density plots for the four hidden units in the trained network are shown in Figure 1. As can be seen, all four hidden units reveal distinct banding. But, are these bands interpretable?

Following the practice of Berkeley et al. (1995), descriptive statistics (i.e., means, standard deviations, and correlations) were computed for the set of features that defined the mushrooms that fell into each band identified in Figure 1. Almost all of these bands revealed a rich set of definite features (for a detailed list, see Dawson, Medler, & Berkeley, 1995).

For example, consider Band H of Hidden Unit 2. The following description is true of all the mushrooms that fall into this band: *Free gill attachment, narrow gill size, tapering stalk shape, stalk surface above ring silky or smooth, stalk surface below ring silky or smooth, stalk colour above ring is pink or white, stalk colour below ring is pink or white, white veil colour, one ring, several population. If it has bruises, then odour is almond or anise, crowded gill spacing, and pendant ring type. If it does not have bruises, then odour is foul or musty, close gill spacing, and evanescent ring type. If odour is almond or foul, then spore print colour is brown or green or orange. If odour is anise or musty, then spore print colour is buff or chocolate or purple or white. Leaves or path or urban or woods habitat.*

This level of detail is typical of the definite features for all of the bands identified in Figure 1.

An examination of the definite features that were identified indicated that Hidden Unit 0 sends a strong signal when a member of a specific class of edible mushrooms is encountered. All of the mushrooms that fall into Band C have almond or anise odour, which is characteristic of an edible mushroom. Some of the mushrooms that fall into Band D of Hidden Unit 0 are edible in virtue of having almond or anise odour. The others that fall into this band have no odour, but are edible because the pass all five questions in the algorithm detailed in Box 1. With respect to this latter set of mushrooms, mushrooms that fall into Band B are nearly identical to them with one exception—the Band B mushrooms have bruises. As a result, they pass the first four steps of the algorithm, but fail (i.e., are identified as poisonous) in the fifth step.

Hidden Unit 1 detects inedible mushrooms based on odour. In particular, all of the mushrooms in Band B and C of Hidden Unit 1 have foul odour, which is diagnostic of a poisonous mushroom.

Hidden Unit 2 captures different classes of mushrooms related to each other in terms of bruises, gill attachment, and number of rings (almost all the mushrooms that produce high activity in this unit have no bruises, free gill attachment, and one ring). The classes diverge from one another with respect to combinations of other features, primarily odour and ring type. As a result, this unit does not pick out features that are only definitive of poisonous (or edible) mushrooms. For example, Band J in Hidden Unit 2 picks out edible mushrooms, while band K picks out some mushrooms that are poisonous (i.e., they have foul or musty odours) and some mushrooms that are edible (i.e., they have almond or anise odours).

Hidden Unit 3 generates strong activity when it detects the presence of a combination of features that are characteristic of inedible mushrooms. Specifically all of the mushrooms that fall in Bands B and C are inedible, and have free gill attachment, close gill spacing, stalk colour above the ring that is not yellow, a white veil colour, a spore print colour that is not yellow, and a pendant ring type if it has no bruises, or a nonpendant ring type if it has bruises.

How does the network use the features encoded in the hidden unit activity bands to determine whether a mushroom is edible or not? For the most part, three of the hidden units in the network (hidden units 1, 2, and 3) detect features that are characteristic of poisonous mushrooms. As a result, the network defines a very large class of edible mushrooms as those that fail to produce activity in its hidden units; 4064 (96.58%) of the edible mushrooms produce bands of [0-A, 1-A, 2-A, 3-A], where 0-A is read "Band A of Hidden Unit 0". These mushrooms are classified as edible because the network fails to find any poisonous features in them.

While this approach works for most of the mushrooms, it fails for a minority of them. A small percentage of edible mushrooms fall into the same (high activity) bands in Hidden Unit 2 as do a number of poisonous mushrooms. As a result, a different approach must be used to identify the edibility of these mushrooms. The network pools the activation of hidden units 0 and 2 to solve this problem, essentially by intersecting the features detected by the two units with an AND operation. The high activity bands of Hidden Unit 2 capture groups of related mushrooms, but include both poisonous and edible mushrooms. The intersection of the members of the high activity bands from both units define a set of edible mushrooms—this intersection picks out the edible mushrooms from the bands in Hidden Unit 2 that also contain poisonous mushrooms. Consequently, the three other "rules" needed to identify the remaining 144 edible mushrooms are: [0-C, 1-A, 2-H, 3-A], [0-D, 1-A, 2-J, 3-A], and [0-D, 1-A, 2-K, 3-A]—each of these rules uniquely identifies 48 or 1.14% of the remaining edible mushrooms.

Conclusion

Connectionists must interpret the internal structure of their networks if PDP models are to contribute to cognitive science (e.g., McCloskey, 1991). In this paper, we have provided a very brief overview of a technique for interpreting value unit networks (for more detailed accounts, see Berkeley et al., 1995; Dawson, Medler, & Berkeley, 1995). Our hope is that this kind of technique will allow cognitive scientists to generate detailed algorithmic accounts of PDP models of cognitive processes, and thus allow these models to inform other empirical investigations of cognition (e.g., by suggesting experiments to be performed by cognitive psychology).

References

Berkeley, I. S. N., Dawson, M. R. W., Medler, D. A., Schopflocher, D. P., & Hornsby, L. (1995). Density plots of hidden value unit activations reveal interpretable bands. *Connection Science*, 7, 167-186.

Chambers, J. M., Cleveland, W. S., Kleiner, B., & Tukey, P. A. (1983). *Graphical methods for data analysis.* Belmont, CA: Wadsworth International Group.

Dawson, M. R. W., Medler, D. A., & Berkeley, I. S. N. (1995). The secret meanings of hidden units: Jumping on the banding wagon. *Cognition*, under editorial review.

Dawson, M. R. W., & Schopflocher, D. P. (1992). Modifying the generalized delta rule to train networks of nonmonotonic processors for pattern classification. *Connection Science*, 4, 19-31.

Dawson, M. R. W., & Shamanski, K. S. (1994). Connectionism, confusion, and cognitive science. *Journal of Intelligent Systems*, 4, 215-262.

Lincoff, G. H. (1981). *National Audubon Society field guide to North American mushrooms.* New York: Alfred A. Knopf Publishers.

McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. *Psychological Science*, 2, 387-395.

Medler, D. A., & Dawson, M. R. W. (1994). Training redundant artificial neural networks: Imposing biology on technology. *Psychological Research*, 57, 54-62.

Mozer, M. C., & Smolensky, P. (1989). Using relevance to reduce network size automatically. *Connection Science*, 1, 3-16.

Schlimmer, J. S. (1987). *Concept acquisition through representational adjustment* (Technical Report 87-19). Doctoral dissertation, Department of Information and Computer Science, University of California Irvine.

Seidenberg, M. (1993). Connectionist models and cognitive theory. *Psychological Science*, 4, 228-235.

Author's Notes

This research was supported by an NSERC research grant awarded to MRWD, and by a Province of Alberta Graduate Fellowship awarded to DAM. Address correspondence to Dr. Michael Dawson, Biological Computation Project, Department of Psychology, University of Alberta, Edmonton, AB, T6G 2E9. Phone: 403-492-5175. E-mail: mike@psych.ualberta.ca.

The Inductive Learning Model: A Foundation for AI

Lev Goldfarb

Depuis 18 ans, Lev Goldfarb s'efforce de trouver un modèle formel (mathématique) qui clarifierait la nature du processus d'apprentissage inductif. L'objectif principal du programme de recherche dans lequel il est actuellement impliqué est d'explorer les forces et les limites théoriques et pratiques d'un nouveau modèle d'apprentissage inductif.

Some of us, partly because of a genetic predisposition, become attracted to science in order to seek some unified abstract paradigms. It goes without saying that to recognize such a general model, even when it is staring one in the face, the researcher should also be genetically endowed with an appropriate aesthetic sense.

Over the last 18 years I have been preoccupied with finding a formal (mathematical) model that would clarify the nature of the inductive learning processes. To my great surprise, this issue was not perceived by the leading (and founding) schools of AI to be even *one* of the central issues confronting AI, a circumstance which, in my opinion, has had a very negative effect on the development of the entire field up to now (see also N. Cercone and G. McCalla comment in [1], p.18: "Surprisingly, until recently, learning has not been a major concern for AI").

In the late 80's, however, once the political events intervened (in the form of the recent rise of connectionism, which properly belongs to the field of pattern recognition), the situation has begun to change, as can be seen from the content of the recent leading AI textbooks [2,3]. As a result, the old issue of incompatibility of the symbolic and the numeric (including connectionist) models again came to the fore. The "again" refers to the fact that from the very beginning AI had severed practically all ties with pattern recognition based solely on the incompatibility of the two paradigms. Thus, ironically but fortunately, the "politically" successful connectionists unwittingly brought the extremely divisive but fruitful old issue back to the front burner.

In this respect, it is important to note that within the area of pattern recognition proper the symbolic-numeric dichotomy has also played a very prominent role: vector space based models and syntactic (structural) models are the only two main competing paradigms.

As I entered the field of pattern recognition in 1977, it became clear to me that the general concept of distance measure should become the key concept around which a satisfactory recognition model can be built. In my doctoral thesis [4], a new mathematical framework for pattern recognition was proposed, in which the recognition problem formulated in a general "distance space" (i.e., a pseudo-metric space) is "translated" in a proposed standard manner into one in a low-dimensional pseudo-Euclidean vector space. The distance is typically that between the structured object representations. Thus, since various "symbolic" distances on strings, trees, etc. are becoming widely available, more efficient and reliable vector space *classification* algorithms can now be applied to patterns that are represented by strings, trees, etc. The basic idea of this analytical framework is to approximate the distributions in the metric space by those in the efficiently constructed pseudo-Euclidean vector space. We have applied the framework to shape recognition [5,6], library information retrieval [7], speech recognition [8]. Some applications by others are, for example, to texture classification [9] and waveform recognition [10,11]. Although authors of the last paper mention such potential application areas as radar target identification, signal processing and other forms of sensory communications processing, it is not an exaggeration to say that the framework applies to almost all domains of pattern recognition.

Initially, it did not occur to me, as to practically all participants in the development of pattern recognition methods, including connectionists, that the problem of inductive learning has a somewhat different flavor than that of pattern recognition. The basic problem of pattern recognition is to develop techniques for *recognizing* new objects based on the information provided by the finite training set of objects. Thus, the focus of pattern recognition is on the decision algorithms. It became clear to me much later that the problem of inductive learning is that of learning the *class structure*. To be sure, the two problems are closely related, but it also became clear that the process of inductive learning must precede the process of recognition and this fact has profound consequences. It has also become gradually clear to me that the vector space model was not developed to (and simply cannot) capture the structure of the learning class in practically all real-world problems.

The corresponding argument about the inappropriateness of the classical vector space based learning models is so simple that one can capture it in one (the following) sentence. Since the object representation space is linear, the only mathematically *legitimate* class structures that can be (inductively) constructed based on the finite training set of vectors are linear. In reality, because the corresponding linear subspaces induced by the training vector data representation hardly ever capture the structure of the class from which the training set was extracted, practically all vector space based models introduce in a (necessarily) arbitrary manner, i.e., in a manner not related to the mathematical structure of the underlying vector space, some class of non-linear functions that is used to approximate "class" boundaries based on the vector representation of the training data. The latter "class" has little to do with the class from which the training data was collected, unless, of course, the parameters of the training data has been chosen by the designer in a very particular manner that, in essence, already presumes the structure of the class to be known. Moreover, there are uncountably many classes of non-linear functions, each of which is equally (non-)eligible to be chosen as the corresponding approximating class of functions.

What is an appropriate formulation of the inductive learning problem? The basic problem can be stated as follows: Given a finite set C+ of positive training objects that belong to possibly infinite set C, the concept to be learned, and a finite set C- of negative training objects that do not belong to the concept C, find an analytical model that would allow one to capture the *structure* of C, i.e., to recognize if a new element belongs to the concept C. In other words, on the basis of the finite training set (C+,C-) the model must allow to form the "idea" of *inductive* generalization corresponding to the concept C.

I now believe that it is meaningless to speak of recognition without the concept of (inductive) generalization, where generalization is simply the state of the agent after the inductive learning process.

The reason why a formal model of inductive learning, must be a fundamentally new mathematical model can be intuited from the following fact. To construct such a model, one must find a (new) way of encoding, or capturing, a possibly infinite set C of positive objects based on the finite training set. Mathematicians, so far, have not addressed this question, so no scheme to answer it has existed within the present mathematics. The potential model should be able to extract the structure of the class C from its finite subset C+. By the structure of the class I mean 1) the symbolic features that make the objects of the class similar to each other and/ or different from other objects, and 2) the emergent interrelationship among these features. The inductive learning process would then involve the discovery of the structure of the class.

The main objective of our present research program is to explore theoretical and practical strengths and limitations of the new model of inductive learning originally proposed in [12] (see also [13]-[15]). The new model - evolving transformation systems (ETS) - is the first formal model that unifies the symbolic and numeric learning paradigms. The model has emerged as a result of an attempt to unify two fundamentally different approaches to pattern recognition - syntactic approach, based on Chomsky's formal language theory, and the vector space approach, based on the approximation of the decision surfaces in the Euclidean space. It appears now that ETS offers a very natural model of a central/driving cognitive process - inductive learning - around which one can build all other AI processes. As above, inductive learning should be understood as a process by means of which an agent can learn/encode the appropriate class structure from some finite training set. ETS makes it quite clear that it is a change in the metric geometry of the object representation space that is a critical step in the learning process. This is accomplished naturally only under a symbolic representation space by adding new macro-operations to the existing set of operations, where the distance between two object representations is defined as the "shortest" weighted sequence of operations that can transform one of them into the other. Moreover, ETS offers a very satisfactory formal (axiomatic) definition of a symbolic system, which has been the object of interest in AI since its inception.

What makes ETS fundamentally different from all present mathematical models is the new way in which the concept of continuity is introduced into a discrete computational model. It is the introduction of continuity that allows the system to efficiently expand during the learning process its finite set of symbolic operations, by means of which the distinguishing "features" of the learning class are captured.

We have applied the new model to analogical reasoning [16], handwritten digit and character recognition [17,18], inductive learning of a subclass of regular languages [19,20], and protein sequence analysis and classification [21]. Though ETS model suggests literally scores of major theoretical and applied research projects, we have been able to pursue very few of them (in part due to the very poor funding of the project). Currently, we are working on two projects. In one project (with Sanjay Deshpande), we attempt to develop a general approach to low-level vision processes suggested by ETS, i.e., we are exploring a systematic approach to the very initial low-level symbolic representations that are "emerging" during the inductive learning processes. In another project (with John Abela), we intend to show systematically the total inadequacy of the classical vector space based inductive learning (including connectionist) models.

References

[1] N. Cercone and G. McCalla (1984). Artificial intelligence: Underlying assumptions and basic objectives, *J. Amer. Soc. for Information Science*, 35(5), 280-290.

[2] E. Rich and K. Knight (1991). *Artificial Intelligence*, second edition , McGraw-Hill, New York.

[3] P. H. Winston (1992). *Artificial Intelligence*, third edition, Addison-Wesley, Reading, Massachusetts.

[4] L. Goldfarb (1985). A new approach to pattern recognition. In L. N. Kanal and A. Rosenfeld (eds.), *Progress in Pattern Recognition* 2, North-Holland, Amsterdam.

[5] L. Goldfarb, (1984). A unified approach to pattern recognition, *Pattern Recognition* 17, 575-582.

[6] L. Goldfarb and R. Verma (1988). Hybrid associative memories and metric data models, Digital and Optical Shape Representation and Pattern Recognition, ed. R.D. Juday, *Proc. SPIE* 938, 302-316.

[7] R. D. Smith (1986). *An Associative Information Retrieval System*, Master's Thesis, University of New Brunswick.

[8] R. Verma and L. Goldfarb (1991). A metric approach to isolated word recognition, *Proc. 4th U.N.B. AI symposium*, U.N.B. Press, Fredericton, N.B., 169-182.

[9] A.K.C. Wong, H.C. Shen and P.W. Wong (1990). Search-effective multi-class texture classification, *Intern J. of Pattern Recognition and Artificial Intelligence* 4(4), 527-552.

[10]V.H. Clarson and J.J. Liang (1989). Mathematical classification of evoked waveforms, *IEEE Trans.* SMC-19, 68-73.

[11] J.J. Liang and V.H. Clarson (1989). A new approach to classification of brainwaves, *Pattern Recognition* 22, 767-774.

[12] L. Goldfarb (1990). On the foundations of intelligent processes I: An evolving model for pattern learning, *Pattern Recognition* 23, 595-616.

[13] L. Goldfarb (1992). What is distance and why do we metric model for pattern learning?, *Pattern Recognition* 25, 431-438.

[14] L. Goldfarb and S. Nigam (1994). The unified learning paradigm: a foundation for AI. In V. Honavar and L. Uhr (eds.), *Artificial Intelligence and Neural Networks: Steps toward Principled Integration*, Academic Press, Boston, MA.

[15] L. Goldfarb, J. Abela, V.C. Bhavsar and V.N. Kamat (1995). Can vector space based learning model discover inductive class generalization in a symbolic environment?, *Pattern Recognition* Letters 16, 719-726.

[16] A. Satriawan (1990). *A Tree Transformation System for Pattern Learning*, Master's Thesis, U.N.B.

[17] S. Dewi (1991). *Dynamic Selection of Primitives in Metric Model for Pattern Learning*, Master's Thesis, U.N.B.

[18] V.N. Kamat (1995). *Inductive Learning with the Evolving Tree Transformation System*, Ph.D. Thesis, U.N.B.

[19] W.B. Santoso (1992). *Learning Algorithm for the Reconfigurable Learning Machine*, Master's Thesis, U.N.B.

[20] S. Nigam (1993). *Metric Model Based Generalization and Generalization Capabilities of Connectionist Models*, Master's Thesis, U.N.B.

[21] J.M. Abela (1994). *Topics in Evolving Transformation Systems*, Master's Thesis, U.N.B.

***Lev Goldfarb** received his Diploma in Mathematics and Computer Science from Leningrad (now St. Petersberg) University in 1973 and Ph.D. in Systems Design from the University of Waterloo in 1980. He has been with the Faculty of Computer Science at the University of New Brunswick since 1992. His main research interests are in the foundations of basic cognitive processes, foundations of artificial intelligence, machine learning, pattern recognition and computer vision. His paper [12] won an Annual Pattern Recognition Society Award. He is an associate editor of* Pattern Recognition.

Machine Learning at the National Research Council's Institute for Information Technology (IIT)

Peter Turney

Introduction

L'apprentissage machine est depuis longtemps un domaine important de technologie à l'Institut de technologie de l'information (ITI) du Conseil national de recherches. L'apprentissage machine a joué dans le passé, et joue toujours aujourd'hui un rôle fondamental dans bon nombre de projets. Cet article décrit brièvement quelques-unes de nos activités dans ce domaine de recherche. Comme le travail en rapport avec l'apprentissage machine à l'ITI a été distribué dans plusieurs projets, laboratoires et groupes, nous avons décidé d'adopter une approche orientée sur la personne dans cet article, plutôt qu'une approche orientée sur le projet ou orientée sur le groupe. Pour obtenir plus d'information sur l'ITI, visitez notre site sur le World Wide Web, à l'adresse http://www.iit.nrc.ca/.

Introduction

Machine Learning (ML) has long been an important area of technology at the National Research Council's Institute for Information Technology (IIT). ML plays a fundamental role in a number of projects, both past and present. This article briefly describes some of our activities in this area of research. Since ML work at IIT has been distributed across many projects, laboratories, and groups, we have decided to take a person-oriented approach in this article, instead of a project-oriented or group-oriented approach. For more general information about IIT, visit our World Wide Web site, http://www.iit.nrc.ca/.

Pierre Boulanger

Dr. Pierre Boulanger has been a research officer at IIT for the past ten years. He is currently with IIT's Visual Information Technology Group, where the focus of his research is on computer vision and environment modelling.

Dr. Boulanger's work in computer vision has involved the development of advanced clustering techniques using Bayesian decision trees and robust statistics. He is also working on fuzzy graph matching methods to find the correspondence between scene models and information produced by range and intensity sensors. These matching methods are applied to integrate multi-view information, to produce a coherent and multi-resolution representation of indoor and outdoor environments. His recent publications are in early vision processes, such as Bayesian sensor fusion, graph matching, and segmentation.

Martin Brooks

Dr. Martin Brooks leads the Interactive Information Group at IIT. His current research interests include empirical modelling and approximation theory. He is an associate editor of IEEE Expert.

In his recent research, Dr. Brooks has developed the concept of scale dependent monotonicity, which is a generalization of the usual mathematical concept of monotonicity. It is useful for generating empirical models of noisy processes generating time series data. Any time series may be analyzed in terms of its alternating rising and falling behaviour, segmenting the data at each local minimum and maximum. However, noise in the data will cause there to be many short segments. Scale dependent monotonicity provides a natural smoothing operator, resulting in a coarser segmentation that does not track the noise. By analyzing multiple time series generated by the same process, one can determine the natural scale of that process and generate piecewise monotone models.

Peter Clark

Dr. Peter Clark has been involved in research in machine learning and knowledge-based systems at the Turing Institute, UK (1985-91), Ottawa University (1991-93) and at IIT (1993 to present). He is currently on leave of absence to University of Texas at Austin.

Dr. Clark has been involved in collaborative research with University of Ottawa on using background knowledge to guide inductive learning. Most standard inductive tools are knowledge-poor, and consequently often induce rules which are nonsensical to domain experts (based on their background knowledge). The goal of this research was to represent and use an expert's knowledge to eliminate obviously nonsensical rules from the search space. The resulting learning algorithm searches only the space of rules which are plausible to the domain expert, rather than the entire rule space. The expert's knowledge was represented using a qualitative model, and experiments demonstrated that the accuracy, efficiency and (most importantly) explainability of the induced rules could be increased using this technique.

Alain Desilets

Mr. Alain Desilets holds a Masters in Management Sciences from the University of Waterloo. He has been at IIT since 1992, where he has built expert decision support systems for iceberg management and for financial engineering. Before joining IIT, he worked at GIRO, a firm specializing in systems for optimal scheduling of bus personnel. His current research interest is in combining machine learning and community-based navigation for doing information retrieval.

Mr. Desilets is currently involved in two projects that use machine learning for information retrieval on the World Wide Web. In both projects, machine learning is used to leverage information contained in human-generated topic hierarchies. The rationale behind this approach is that humans are still better than machines at interpreting the meaning of documents.

The first project will build a web robot that recognizes topic hierarchy web pages, by employing rules generated by decision tree induction. The robot will create a database of topic hierarchy pages, which can then be used to help people find information.

The second project deals with the merging of distinct topic hierarchies. This is a difficult problem, since different hierarchies may have conflicting views of the same documents. This problem will be approached as a form of unsupervised learning, where a general structure must be learned from particular cases.

Fazel Famili

Dr. Fazel Famili has been involved in research and development projects applying machine learning since 1988. His primary concern in his work at IIT has been to use the vast quantity of data collected from industrial processes to provide support for intelligent process management. Among his achievements (with his colleague Dr. Peter Turney) has been development of an intelligent data analysis tool called IMAFO.

IMAFO applies machine learning techniques (decision tree induction) to analyze data collected from an industrial process (e.g., semiconductor production) and assist engineers and operators to understand why operations or processes fail. IMAFO discovers relationships between parameters and problems of a process.

IMAFO was first applied to process optimization for electrochemical milling of aircraft turbine blades. IMAFO was then used in a joint project, with Mitel and Gennum, for a number of applications in process optimization and yield control of semiconductor manufacturing. A number of enhancements were made to IMAFO during this joint project, including the development of a User Assistant. IMAFO and parts of the User Assistant have been commercialized and are sold as Q-YIELD (trade mark) by Quadrillion Corporation of Ottawa.

Dr. Famili's current research focuses on data pre-processing and use of machine learning techniques to analyze data collected from the operation of commercial aircraft. His goal is to develop an algorithm that can find patterns in the data that provide explanations for problems and failures of varying degrees of severity.

Innes Ferguson

Dr. Innes Ferguson joined IIT in 1994 after obtaining his Ph.D. from the University of Cambridge (UK) while on a Bell-Northern Research Postgraduate Scholarship. His research interests include adaptive agents, multimodal multimedia communications, and information navigation and filtering.

In his current research, Dr. Ferguson has been investigating, jointly with Grigoris Karakoulas of Carleton University (Ottawa), information filtering (IF) mechanisms which can take advantage of the inherent distribution and dynamics of today's information networks. These domain characteristics suggest an IF system design based on a coordinated collection of robust, decentralized, adaptive agents.

The objective of the work is to develop a multi-agent framework in which the various agents involved in the IF task are able to adapt to changes in both the information space and the users' interests. The first step has been the definition of a market model, called SIGMA. SIGMA is being applied to the task of personalized filtering of Usenet netnews. The agents are based on CALVIN, IIT's in-house agent framework. The CALVIN framework makes it possible to integrate a number of different learning and adaptation techniques. A major focus of future work will be evaluating the system, both from the point of view of measuring the performance of the computational economy and from the point of view of improved access to Usenet.

Grigoris Karakoulas

Dr. Grigoris Karakoulas joined IIT in 1993 as a visiting research fellow after obtaining a Ph.D. in Computer Science from Athens University in Greece. He is currently a research associate in the Department of Systems and Computer Engineering at University of Carleton and an affiliate member of the Interactive Information Group (II Group) at IIT. His research interests include reinforcement learning, intelligent agents, reasoning under uncertainty and applications of machine learning in medicine and finance.

Dr. Karakoulas's current research focuses on adaptation and learning in multi-agent environments. He is exploring the problem of developing a computational market model, SIGMA, for distributed information filtering tasks. He is involved in collaborative work with Dr. Innes Ferguson (also a member of II Group), applying the SIGMA market model to the task of personalized information filtering of Usenet Netnews. Various learning and adaptation techniques are embedded in SIGMA: reinforcement learning, relevance feedback, and bidding.

In his recent research, Dr. Karakoulas has investigated pragmatic issues associated with the scaleability of reinforcement learning with respect to real-world learning-to-control tasks. These issues include stability of control policies, generalization, partial observability, exploration and cost of experimentation. He has developed an algorithm for learning robust policies in stochastic non-linear systems and an algorithm for cost-effective classification in tasks such as medical diagnosis, repair planning, and credit decision-making.

Joel Martin

Dr. Joel Martin recently joined IIT from the University of Pittsburgh. His research focuses on machine learning and reasoning under uncertainty. His recent work has included publications in unsupervised learning and bayesian networks.

Dr. Martin's current research applies unsupervised machine learning to information retrieval, filtering, and organization. He is addressing the problem of automatically organizing a large set of electronic documents and providing that organization to various communities of users. A probabilistic, unsupervised method is used to cluster related documents together and to provide a general hierarchical structure to those documents. This resulting structure promotes browsing through the documents and can provide decision-theoretic judgments for information filtering.

Some of Dr. Martin's recent research has investigated the interaction between supervised learning (learning to predict one aspect or feature of a domain) and one form of unsupervised learning (learning to predict any feature of a domain). He has developed a learning algorithm, DP1, that permits a user to provide different utilities (values) of predictive success for each aspect of a domain. In this way, the user can decide which aspects are worth predicting and can weight more important aspects more heavily than less important ones.

Salvatore Morgera

Dr. Sal Morgera is a Professor at McGill University in the Department of Electrical Engineering, the Director of the Information Networks and Systems Laboratory (INSL) at McGill, a Participant in the Canadian Institute for Telecommunications Research (CITR), and an IIT Guest Worker. His research is principally in the areas of wireless information networks, adaptive signal processing, and pattern analysis and visual systems design. Results on the performance of certain structured estimation and classification methods under realistic, finite sample size conditions led to his being named a Fellow of the IEEE in 1990 and his innovative high capacity multiservice wireless network architectures were a key element in his being named an IEEE Distinguished Lecturer in 1993.

Dr. Morgera's current IIT-related research focuses on two areas: (1) early sensory coding and scene representation and the role that a new area, sparse signal processing, may play in the development of a general information processing strategy and (2) cooperative processing as a means for approaching the capacity limits promised by information theory for practical wireless and broadband networks. Other research efforts ongoing at INSL McGill include the development of a next generation global positioning system, the design of a vision system for automated manufacturing and inspection applications, the development of new algorithms for broadband interference cancellation for wireless networks, and the design of data compression algorithms. Much of the research is funded by industrial partners.

Dale Schuurmans

Mr. Dale Schuurmans is completing a Ph.D. in Computer Science at the University of Toronto, and will be joining IIT this September. His general research interests include: clarifying the different types of learning tasks and different applications for learning systems, developing effective tools and techniques for solving learning problems, and understanding the inherent limitations (and possibilities) of learning systems. An underlying theme of his research is improving the pragmatic performance of machine learning systems through the mathematical analysis of their behaviour.

Most of Mr. Schuurmans' recent research focuses on developing efficient techniques for learning classification rules from data. He has developed a number of new techniques for reducing the amount of training data needed to learn with statistically guaranteed accuracy and reliability.

Mr. Schuurmans has also actively investigated several generalizations of classification learning that frequently occur in practice. One deviation from standard classification learning occurs when the descriptions to be classified are incomplete or corrupted in some way. This requires the system to make classifications on the basis of incomplete information.

Recently, Mr. Schuurmans has also developed a precise characterization of the conditions that permit fast (exponential) learning, and contrasted these with the conditions where slow (rational) learning is the best that can be guaranteed. This resolves a question posed by neural-network researchers to the computational learning theory community.

Peter Turney

Dr. Peter Turney joined IIT in 1989, after obtaining a Ph.D. from the University of Toronto. His research is mainly concerned with machine learning. A unifying theme in his publications is pragmatic constraints on inductive concept learning, such as cost, context, stability, and complexity. He is Co-editor of *Canadian Artificial Intelligence* magazine and Advanced Information Systems Co-ordinator for the *Journal of Artificial Intelligence Research.*

Dr. Turney's current research is in machine learning applied to information retrieval and filtering. He is exploring the problem of extracting a small set of keywords and keyphrases from a document, to concisely summarize the content of the document. The learning algorithm will be a form of supervised learning from examples, where the training data will consist of pairs of documents and their keywords. Information retrieval and filtering are core problems for the Interactive Information Group, which is addressing the problems of Digital Libraries.

For the past two years, Dr. Turney has been investigating the problem of cost-sensitive classification. The prototypical example of the problem of cost-sensitive classification is medical diagnosis, where a doctor would like to balance the costs of various possible medical tests with the expected benefits of the tests for the patient. He has developed a learning algorithm that attempts to find the optimal balance.

Recent ML Publications at IIT

P. Boulanger. Hierarchical segmentation of range and colour images based on Bayesian decision theory. *Proceedings of the Fifteenth International Workshop on Maximum Entropy and Bayesian Methods,* Santa-Fe, New Mexico. In press.

P. Boulanger and J.A. Beraldin. Three dimensional motion estimation using range and intensity flow. *Proceedings of the ISPRS Intercommission Workshop: From Pixels to Sequences,* Zurich, Switzerland. pp. 330-337. 1995. NRC 38365.

M. Brooks. Approximation complexity for piecewise monotone functions and real data. *Comput. Math. Appl.* 27(8): 47-58. 1994. NRC 35087.

P. Clark and S. Matwin. Using qualitative models to guide inductive learning. In *Proc. Tenth Int. ML Conference* (ML-93), Ed: P. Utgoff, 49-56. 1993. NRC 37079.

V. Conan, P. Boulanger, and S. Elgazzar. Robust position estimation of a mobile vehicle. *Proceedings of the Europto Conference on Sensors and Control for Automation,* Frankfurt, Germany, pp. 37-46. 1994. NRC 37130.

A. Famili and P. Turney. Intelligently helping human planner in industrial process planning. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* (AIEDAM), 5(2), 109-124, 1991. NRC 31771.

A. Famili. Use of decision-tree induction for process optimization and knowledge refinement of an industrial process. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* (AIEDAM), 8(1), 63-75, 1994. NRC 35070.

A. Famili. The role of data pre-processing in intelligent data analysis. *Proceedings of the International Symposium on Intelligent Data Analysis (IDA-95),* Baden-Baden, Germany. In press. NRC 38373.

G. Karakoulas. Probabilistic exploration in planning while learning. In *Proceedings of the Eleventh International Conference on Uncertainty in Artificial Intelligence,* Montreal, Canada. 1995. NRC 38386.

G. Karakoulas. Cost-effective classification for credit decision-making. In *Proceedings of the Third International Conference on Artificial Intelligence Applications on Wall Street,* New York City, USA. 1995. NRC 38385.

G. Karakoulas and I. Ferguson. A computational market for information filtering in multi-dimensional spaces. In *Proceedings of the AAAI Fall Symposium on AI Applications in Knowledge Navigation and Retrieval,* Cambridge, MA, November 10-12, 1995. In press. NRC 38387.

J. Martin and K. VanLehn. Student assessment using Bayesian nets. *International Journal of Human-Computer Studies,* 42. In press.

J. Martin and D. Billman. Learning overlapping concepts. *Machine Learning.* 1994.

J. D. Martin. DP1: Supervised and unsupervised clustering. In *Proceedings of the 1994 European Conference on Machine Learning,* Catania, Italy. 1994.

J. D. Martin. Goal directed clustering. In *Proceedings of the AAAI Spring Symposium on Goal-Directed Learning,* Palo Alto, CA. 1994.

D. Schuurmans and R. Greiner. Practical PAC learning. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* (IJCAI-95). In press.

D. Schuurmans and R. Greiner. Sequential PAC learning. *Proceedings of the Eighth Annual Workshop on Computational Learning Theory (COLT-95).* pp. 377-384. ACM. 1995.

D. Schuurmans. Characterizing rational versus exponential learning curve. *Proceedings of the Second European Conference on Computational Learning Theory (EuroCOLT-95).* Ed: P. Vitanyi. pp. 272-286. Berlin: Springer. 1995.

D. Schuurmans and R. Greiner. Learning default concepts. *Proceedings of the Tenth Canadian Conference on Artificial Intelligence (CSCSI-94).* Ed: R. Elio. pp. 99-106.

P. Turney. Cost-sensitive classification: empirical evaluation of a hybrid genetic decision tree induction algorithm. *Journal of Artificial Intelligence Research (JAIR),* 2: 369-409. 1995. NRC 38369.

P. Turney.Technical note: Bias and the quantification of stability. *Machine Learning.* In press. NRC 38313.

P. Turney. Low size-complexity inductive logic programming: The East-West Challenge considered as a problem in cost-sensitive classification. *Proceedings of the Fifth International Inductive Logic Programming Workshop (ILP-95).* In press. NRC 39164.

The Ottawa Machine Learning Group

Robert C. Holte, Stan Matwin, Mario Marchand

Le Groupe d'apprentissage machine d'Ottawa est un groupe de recherche de l'Université d'Ottawa. Les activités du groupe vont de la recherche pure aux applications industrielles. Les sujets d'intérêt principaux sont l'apprentissage empirique, le raisonnement basé sur des cas, la programmation logique inductive et le changement de représentation. On présente dans cet article un survol du travail de Robert C. Holte, de Stan Matwin et de Mario Marchand, tous trois du Groupe d'apprentissage machine d'Ottawa.

The Ottawa Machine Learning Group (OMLG) is a research group that functions at the University of Ottawa. The group consists of three professors (Robert C. Holte, Stan Matwin, Mario Marchand), and a number of graduate students and associates. The focal point of the group is the weekly seminar, in which members of the group present their own work, discuss the current literature, or invite speakers from the outside. Activities of the Machine Learning group range from pure research to industrial applications. The ML group's primary topics of interest are empirical learning, case-based reasoning, inductive logic programming, and change of representation. Examples of past and current application areas include software engineering (an NSERC Strategic Grant "Machine Learning Applications in Software Reuse"), as well as managing and interpreting the data on natural resources. Several systems/projects that arise from this work are: CAESAR [Fouque, Matwin 93] (CAse-basEd SoftwAre Reuse), CABARESS [Clark et al. 94] (CAse-BAsed REmote Sensing Shell), GUIDAR (Graphical User Interface Design And Reuse), LOPSTER [Lapointe, Matwin 92] (inductive LOgic Programming with Subunification of TERms). The work of our group has been sponsored by NSERC, Canada Centre for Remote Sensing, Natural Resources Canada, Canadian Space Agency, and Cognos, Inc, and MacDonald Dettwiler Associates.

We will list different research projects and research areas according to the three investigators involved in OMLG. The work of Stan Matwin belongs in three areas: Inductive Logic Programming, induction in the presence of knowledge, and learning from remotely sensed data.

In ILP, the training data and the learned concept description are both formulas in first order logic. ILP is often perceived as learning definitions of Prolog programs from the examples of the desired I/O behaviour of these programs. Our group's work in ILP on learning recursive definitions is often quoted in the ILP literature [Aha et al. 94], [Ling et al. 93], [Lapointe Matwin 92].

Learning of fully explained rules (joint work with P. Clark, University of Texas) investigates the influence of domain knowledge available to an inductive learner on the quality of the learned result. Suppose we have a system or a device whose approximate model is available, plus a task-specific set of examples characterizing the behaviour of the device system. Can the two be converted into high quality task-specific rules describing the behaviour of the system? Our initial work on this problem [Clark, Matwin 93a, 93b], produced encouraging and well received results.

We have also initiated a number of applications-driven projects. Machine Learning and Software Reuse (NSERC Strategic Grant with R. C. Holte and F. Oppacher, Carleton U.) investigated the application of ML techniques for useful Software Reuse tools and for the enhancement of reuse in general. Machine Learning from Text (NSERC Strategic Grant with S. Szpakowicz) investigated an integrated ML/NLP system that would acquire useful symbolic rules directly from technical texts in English. S. Matwin has also worked for several years on applications of Machine Learning in Remote Sensing, with Canada Centre for Remote Sensing and with the Pacific Forestry Centre. The work produced a system that generates agents for the use of a complex, heterogeneous computing environment in an information system for forestry [Matwin et al. 95]. Another result is a system for crop recognition from radar imagery [Clark et al. 94]. We are currently starting a new project with MacDonald Dettwiler Associates in which machine learning will be applied to the problem of detection of oil spills from radar images of the sea.

In neural networks, M. Marchand has recently made several important contributions to the understanding of the computational complexity of neural network learning. Both analytical and numerical methods to design and analyze learning algorithms for neural networks were used. Marchand has obtained Probably Approximately Correct (PAC) learnability results for certain classes of neural nets ([Hancock et al. 94a], [Golea Marchand 93a], [Golea, Marchand 93b], [Marchand Hadjifaradji 95], [Golea Marchand 94b, Golea et al. 93]). We plan to continue this work using analytical and numerical methods to design and analyze learning algorithms for neural networks. More specifically, we want to extend the work to study the learnability of stochastic non-overlapping neural networks under k-blocking distributions (see [Marchand,Hadjifaradji 95] for a preliminary study in this direction). Because of both the noisy aspect of these neural network classes and the generality of the input distribution, it is quite reasonable to expect that the corresponding learning algorithms will be useful for solving practical learning problems.

The goal of R. Holte's work in learning and search is the development of learning techniques that speed up various kinds of search, including generate- and-test search

[Holte Drummond 93], state-space search [Porter et al. 90, Drummond et al. 94] and human-controlled search. R. Holte's current research on state-space search investigates both analytical techniques (abstraction, caching) and inductive ones akin to those used in reinforcement learning. Because information collected on one search might, or might not, result in speedup of subsequent searches, the standard "utility" problem (and the related exploitation-versus-exploration issue) is a central issue in this research.

R. Holte's current research on human-controlled search uses as its testbed the task of searching in a library of object-oriented code for a particular class ("browsing"). The focal issue is how to quickly and accurately infer the human's search goal from his/her browsing actions. This inference necessarily involves inductive learning of some sort. To be useful, learning must succeed in real-time; this requires the learning system to have a strong, accurate bias. Much of this research is in the development of this bias, in the form of (a) a weak domain theory, (b) task-specific generalization techniques, and (c) a carefully engineered set of "features." As such, this work constitutes a major case study in bias engineering. To evaluate and compare learning techniques, we have developed a novel empirical method, in which the human searcher is replaced by an automated browsing agent. This permits us to automate the evaluation process and do extensive testing which would be impossible using human subjects.

R. Holte is also continuing to investigate basic scientific issues in machine learning. Some of the issues of interest are: the problem of learning small disjuncts [Holte et al. 89], the induction and use of very simple classification rules [Holte 1993; Holte Maass 1995], and data engineering [Turney 95].

In the past three years, OMLG has been a host to several post-doctoral fellows that have worked with us on different projects. We were lucky to have the cooperation of David Aha (NSERC International Research Scholar, currently at Naval Research Labs, Washington, DC), Peter Clark (currently at University of Texas, Austin), Cao Feng (currently at Fulcrum), Jean-Francois Delannoy (currently at RES), Gilles Fouque (currently at IntelliCorp), Denys Duchier (currently at Simon Fraser University), Xiaobin Li (currently at Concordia University), and Kazumi Saito (currently at NEC).

The description of OMLG would be incomplete without mentioning graduate students involved in the group and their research topics: Daniel Charlebois (machine learning and planning with an application to remote sensing), Johanne Morin (Inductive Logic Programming), Riverson Rios (Inductive Logic programming and The Civic Hospital Project), Messaouda Ouerd (Learning from the Use of Distributed Databases), Thierry Rouget (Inductive Learning in the Presence of Knowledge), Guiming Chen (Inductive Logic Programming in Fuzzy Logic), John Ng Yuen Yan (learning applied to browsing), Chris Drummond (reinforcement learning).

References

[Matwin et al. 95] Matwin, S. Charlebois, D., Goodenough, D., 1995, Machine Learning and Planning for Data Management in Forestry, to appear, *IEEE EXPRT and Intelligent Systems.*

[Fouque, Matwin 93] G. Fouque, S. Matwin, 1993, A Case-based Approach to Software Reuse, *International Journal of Information Systems* (Kluwer) 1, pp. 165-197.

[Clark et al. 94] Clark, C. Feng, S. Matwin, K. Fung, 1994, Improving Image Classification by Combining Statistical, Case-based, and Model-based Prediction Methods, to appear, *Fundamenta Informaticae.*

[Aha et al. 94] D.W. Aha, S. Lapointe, C.X. Ling, S. Matwin, 1994, Learning Recursive Relations with Randomly Selected Small Training Sets, *11th International Conference on Machine Learning*, pp. 12-18.

[Ling et al. 93] X. Ling, S. Lapointe, S. Matwin, 1993, Constructive Inductive Logic Programming, *Procs. of the 10th Int'l Joint Conf on AI*, Chambry, France, pp. 1030-1035.

[Clark, Matwin 93a] Clark, S. Matwin, 1993, Using Qualitative Models to Guide Inductive Learning, *10th Int'l. Conf. on Machine Learning*, Amherst, MA, pp. 49-56.

[Clark, Matwin 93b] P. Clark, S. Matwin, 1993, Learning Decision Theories from Abstract Backgound \ Knowledge, *Procs. of the European Conf. on Machine Learning*, pp. 360-365.

[Lapointe, Matwin 92] Lapointe, S., Matwin, S., Subunification: A Tool for Efficient Induction of Recursive Programs, *Procs. of the 8th International Machine Learning Conference*, pp. 273-280.

[Porter et al. 90] B.W. Porter, E.R. Bareiss, and R.C. Holte, (1990) Concept Learning and Heuristic Classification in Weak-Theory Domains, *Artificial Intelligence* 45 (1-2):229-263.

[Drummond et al. 94] C. Drummond, D. Ionescu, and R.C. Holte (1994), Accelerating the Searching of Software Libraries through Goal Inference", submitted to IEEE Transactions on Software Engineering. (39 pages)

[Golea,Marchand 93a] Mostefa Golea and Mario Marchand, On Learning Perceptrons with Binary Weights, *Neural Computation* vol. 5, pp. 767 - 782 (1993).

[Golea, Marchand 93b] Mostefa Golea and Mario Marchand. Polynomial Time Algorithms for Learning Neural Nets of Nonoverlapping Perceptrons, *Computational Intelligence* 9, pp. 155-170 (1993).

[Holte Drummond 93] R.C. Holte and C. Drummond (1993), "Speeding Up Human Problem-Solving by Machine Learning, *ML'93 workshop on Knowledge Compilation and Speedup Learning*, June, 1993. (4 pages).

[Hancock et al. 94] Thomas Hancock, Mostefa Golea and Mario March nd, 1994. Learning Nonoverlapping Perceptron Networks from Examples and Membership Queries, *Machine Learning*, vol. 16, pp. 161-183.

[Marchand, Hadjifaradji 95] Mario Marchand and Saeed Hadjifaradji, Learning Stochastic Perceptrons under K-Blocking Distributions, to be published in *Advances in Neural Information Processing Systems* 7, Morgan Kaufmann Publishers, (1995).

[Golea-Marchand 94b] Mostefa Golea and Mario Marchand, On Learning Simple Deterministic and Probabilistic Neural Concepts in Shawe-Talor J., Anthony M. (eds.), *Computational Learning Theory: Euro-COLT'93*, pp. 47-60, Oxford University Press, (1994).

[Golea et al. 93] Mostefa Golea, Mario Marchand and Thomas Hancock. On Learning *-Perceptron Networks with Binary Weights, *Proceedings of the NIPS'92 conference*, Nov 30-Dec 3 1992, Denver Colorado. In Giles C.L., Hanson S.J. and Cowan J.D. (eds.), *Advances in Neural Information Processing Systems* 5, pp. 591—598, San Mateo CA, Morgan Kaufmann Publishers, (1993).

[Holte et al. 89] Robert C. Holte, Liane Acker, and Bruce W. Porter (1989), Concept Learning and the Problem of Small Disjuncts, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, Morgan-Kaufmann, pp. 813-818.

[Holte 93] Robert C. Holte (1993), Very Simple Classification Rules Perform Well on Most Commonly-Used Datasets, *Machine Learning*, vol. 11, pp. 63-91.

[Auer et al. 95] P. Auer, R.C. Holte, W. Maass (1995), Theory and Applications of Agnostic PAC-Learning with Small Decision Trees, *Proceedings of the Twelfth International Conference on Machine Learning (ML'95)*, pp. 21-29.

[Turney 95] P. Turney (1995), Working Notes of the IJCAI'95 workshop on Data Engineering for Inductive Learning.

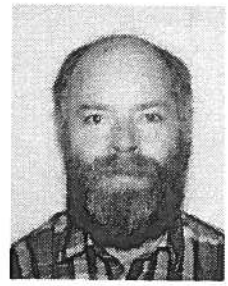

***Robert Holte** is an associate professor of Computer Science at the University of Ottawa. His research area is Machine Learning Artificial Intelligence.*

***Stan Matwin** is a professor of Computer Science at the University of Ottawa. He is the current President of the Canadian Society for Computational Studies of Intelligence (CSCSI) and the Head of IFIF WG 12.2 (Machine Learning). His research interests are in machine learning and its applications, and in knowledge-based systems, with a special focus on combining data and knowledge-driven learning. He has worked for several years with Forestry Canada and the Canada Centre for Remote Sensing on applications of learning and planning in the natural resources sector. He is currently Research Director of a project with MDA, Inc. of Vancouver, Canada, in which learning is applied to detection of oil spills on sea from satellite radar images.*

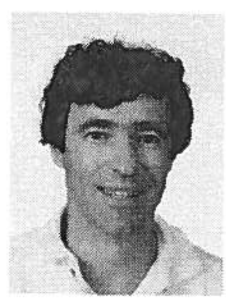

***Mario Marchand** currently works in an area of machine learning research known as computational learning theory. He uses rigorous mathematical analysis for the design of computationally efficient learning algorithms. Recently, he designed noise-tolerant statistical methods for efficient learning of neural networks. Mario Marchand is currently an associate professor of Computer Science at the University of Ottawa. He received his Ph D. in theoretical physics from the University of Sherbrooke in 1987.*

Become a Member of CSCSI/SCEIO

For more details, see page 36.

Machine Learning at The University of Western Ontario

Charles X. Ling

The long-term goal of Charles Ling's research is to study computational models of human learning, and to design and implement machine learning programs to model human learning. His current research can be roughly divided into three areas: machine learning algorithms and their applications; computational cognitive modeling of human learning tasks; and, learning in first-order logic. In the area of machine learning, he has studied a variety of machine learning algorithms and their applications in real world problems. He has worked on improving the traditional decision tree learning algorithm [3], nearest-neighbour learning algorithms [10], knowledge-base refinement with uncertainly [9], and reinforcement learning algorithm (on the balance scale problem) [6]. He has applied decision tree algorithm on a market problem for a major Canadian insurance company on predicting potential customers for agents to approach. As for computational modeling, he has been doing research on both symbolic [3] and connectionist [5] learning algorithms, studying and comparing their strengths and weaknesses, and applying them to various cognitive learning tasks. It is his belief that symbolic learning algorithms can serve as ideal competence models, which underlie many high-level cognitive processes. Indeed, he has shown that many landmark cognitive learning tasks, some of which are traditionally believed to be amenable only to connectionist modeling, are better modeled by symbolic algorithms. These cognitive tasks include: learning the past tense of English verbs [7, 3]; non-conscious information acquisition [8]; English word pronunciation (submitted); and, developmental model of the balance scale [11]. Learning in propositional logic has limitations; it cannot represent compactly concepts in first-order. Charles Ling and other researchers (Stan Matwin, David Aha, Stephane Lapointe) have studied several important problems in ILP, such as the lack of sufficient instances for new predicates [2] and learning singly recursive programs efficiently [1]. He has also made significant effort in applying techniques of new predicate introduction to model scientific revolution [4].

References

[1] D. W. Aha, S. Lapointe, C. X. Ling, and S. Matwin. Learning recursive relations with randomly selected small training sets. In *Proceedings of the 1994 International Conference on Machine Learning*, pages 12 -18, 1994.

[2] S. Lapointe, C. X. Ling, and S. Matwin. Constructive inductive logic programming. In *Proceedings of IJCAI-93 (Thirteenth International Conference on Artificial Intelligence)*, pages 1130-1136. Morgan Kaufmann Publishers, 1993.

[3] C. X. Ling. Learning the past tense of English verbs: the Symbolic Pattern Associator vs. connectionist models. *Journal of Artificial Intelligence Research*, 1:209 - 229, 1994.

[4] C. X. Ling. Introducing new predicates to model scientific revolution. *International Studies in the Philosophy of Science*, 9(1):19 - 36, 1995.

[5] C. X. Ling. Overfitting and generalization in learning discrete patterns. *Neurocomputing*, 8:341-347, 1995.

[6] C. X. Ling and R. Buchal. Learning to control dynamic systems: A progressive quantization approach. *Adaptive Behavior*, 3(1):29-49, 1994.

[7] C. X. Ling and M. Marinov. Answering the connectionist challenge: a symbolic model of learning the past tense of English verbs. *Cognition*, 49(3):235-290, December 1993.

[8] C. X. Ling and M. Marinov. A symbolic model of the nonconscious acquisition of information. *Cognitive Science*, 18(4):595 - 621, 1994.

[9] C. X. Ling and M. Valtorta. Refinement of uncertain rule bases via reduction. *International Journal of Approximate Reasoning*, 13(2):95-126,1995.

[10] C. X. Ling and H. Wang. Deciding optimal weights for IBL algorithms. *Artificial Intelligence Review*, accepted, 1995.

[11] William C. Schmidt and Charles X. Ling. A decision-tree model of balance scale development. To appear in *Machine Learning*, 1995.

***Charles X. Ling** obtained his BSc. in the Department of Computer Science of Shanghai Jiao Tong University in China in 1985, and PhD. in the Department of Computer and Information Science, University of Pennsylvania in 1989. Since then, he has been an Assistant Professor in the Department of Computer Science, The University of Western Ontario. He will research-visit in University of Hong Kong in 1996. His WWW home page is http://www.csd.uwo.ca/faculty/ling from which you can access some of his publications.*

Machine Learning at McMaster University

Ivan Bruha

My research is devoted to machine learning, learning from examples, and involves various topics, particularly for attribute-based approach. I have implemented a completely new extended version of the CN2 covering learning algorithm (called CN4), and am designing and testing various topics and comparing it with other well-known ML algorithms, e.g., C4.5.

In more detail:

(i) Economic learning concerning attribute cost.

(ii) Discretization of numerical attributes. Two existing approaches, off-line and on-line discretization, have been designed for CN4 and Kex, another rule-based machine learning algorithm. Besides the discretization, learning concepts with continuous classes is also studied; therefore, I am incorporating regression statistical algorithms to CN4. The off-line processing of continuous classes for Kex is also being investigated, and both approaches will be compared.

(iii) Unknown attribute value processing. Five different routines have been embedded into CN4 and tested. Now, I am designing an algorithm for processing unknown values which are "irrelevant" or "not-applicable" for certain observations. Such a routine will utilize external domain-specific knowledge provided by an expert-designer of a training database of observations (examples).

(iv) Quality of rules. An unordered-mode classification requires a numerical evaluation to be attached to each decision rule. I studied and designed various formulas for the rule quality. Currently, I am also investigating how these qualities can be combined in order to solve conflict situations when classifying a given unseen object.

(v) Incremental learning. I am convinced that the incrementality of learning should be studied in additional ways. I am investigating new approaches to this topic, particularly, how a learning system can forget and eventually retrieve some 'pieces' of knowledge. This incremental manner of learning requires embedding powerful statistical measures and techniques into the traditional symbolic learning systems.

(vi) Hybrid representation and systems. Attempts exist to combine various knowledge representation tools as well as machine learning paradigms. Such multistrategy or hybrid systems may solve the complexity of learning. I am particularly interested in combining symbolic learning with neural nets. Each paradigm has its advantages and drawbacks, but only its sophisticated merging may enhance the believability of such learning systems. Besides the above hybrid systems, neural nets seem to be very useful in emulating probability distribution functions and belief functions in statistical applications.

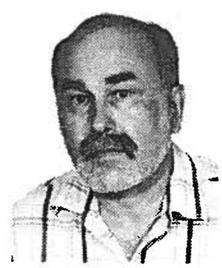

__Ivan Bruha__ was born in Prague, Czechoslovakia (since 1993: Czech Republic). He received the degrees Dipl. Ing. at Czech Technical University, Prague, Faculty of Electrical Engineering (1969), RNDr. at Charles University, Prague, Faculty of Mathematics and Physics (1974), and PhD. (major artificial intelligence) at Czech Technical University, Prague (1973). He taught at Czech Technical University, Faculty of Electrical Engineering for 12 years. He is presently an associate professor at McMaster University, Ontario, Department of Computer Science and Systems. He has published three monographs and about 70 papers in peer reviewed journals and proceedings. He collaborates with a group at Prague University of Economics in machine learning (project Kex and others).

Harry Rogers

Fifth annual IRIS/PRECARN conference Keeping up with the changes

The times they are a-changin'. Bob Dylan would have been right at home at the fifth annual IRIS/PRECARN conference in Vancouver this June. It seemed as if there was something or someone new everywhere — the latest results from the second phase of IRIS, the announcement of the new projects in PRECARN's Phase 2 program, a chance to meet Harry Rogers, PRECARN's new president, and much more.

All this "newness" fit in well with the theme of the conference: *Intelligent Systems and the New Economy*. The two-day meeting began with a plenary session that focused on the evolution of intelligent systems technologies, the realities of the new economy, and how IRIS and PRECARN are adapting to the challenges ahead.

The plenary opened with a warm welcome from Ron McCullough, PRECARN's interim president. He urged the over 400 conference participants to keep up with the changes. "Yes, it's a game," he said. "However, we have the talent. We can win."

In his keynote address, University of British Columbia computer science professor Alan Mackworth gave a personal view of the evolution of intelligent systems. His thesis: "In the past ten years, there's been a revolution in intelligent systems. If we don't understand how the science is changing, we will build applications on the 'old' science — and they won't work."

Mackworth outlined some of the limitations of GOFAIR (Good Old Fashioned Artificial Intelligence and Robotics), which include assumptions of a static environment, deterministic world, and a robot with complete knowledge of the world. "We've got to rethink, recast the concepts of AI," he said. "A key element is the coupling between the robot and the world."

Jean-Claude Gavrel, former PRECARN VP, who recently moved to the Centre de recherche informatique de Montréal, reviewed some of the changes in the structure of PRECARN projects. He noted that the initial "pyramid" model — basic research at the bottom, precompetitive R&D in the middle, and product development at the top — has been modified using concepts from concurrent engineering. Rather than compartmentalizing the research effort, this model has teams of researchers, technology developers, suppliers, and end users working simultaneously on projects

"Of course, having a model is no guarantee of success," Gavrel said. "But not following one certainly increases risk of failure. In the PRECARN/IRIS network, there are 24 end users, 23 suppliers, 27 technology developers, 12 research organizations, and 22 universities. The challenge is now yours to put the network to work."

Joanne Curry, Director of the SFU University and Industry Liaison Office, gave a view from the "trenches of technology transfer." She outlined the motivations for university/industry linkages and reasons for the success or failure of partnerships.

She said that one of the primary causes of failure is that the deliverables and outcomes are not clearly communicated and managed. Success factors include a company with a long-term perspective and realistic expectations, and a project with goals that are compatible with the university researchers' personal interests.

"In our office, we have a supply and demand problem — we have a lot more researchers than companies," Curry said. "Companies need to see that being involved with universities is more than good corporate citizenship, it is part of a good R&D and business strategy."

Michael Brown, president of venture capital company Ventures West Management Inc., talked about the availability of capital in Canada. The picture he painted was not rosy.

He noted that a huge shift in where venture capital comes from has taken place. Funds controlled by labour unions now account for almost 50 per cent of the total amount of capital. Because these funds are so large and their investments must be paced, they are usually only interested in large deals.

"If you have a small project, you have to go and search out the odd labour pool that might be interested in your project or go to government sponsored funding organizations," Brown said. "All this is happening at a time when organizations like NSERC and Western Economic Diversification are cutting back. The funding problem is not getting solved by these initiatives."

PRECARN president Harry Rogers had the happy task of announcing the 15 projects that received the "go-ahead" in PRECARN's Phase 2 program. He thanked Industry Canada for its support and noted that the IRIS and PRECARN network is one of the few programs that continues to be funded by Industry Canada.

"We must all work to improve the competitiveness of Canadian companies," Rogers said. "The best way to do this is to combine the best and brightest of our human resources in training, development, and research networks."

STOP WHINING: Hugh Wynne-Edwards delivers a strong message to S&T community

Hugh Wynne-Edwards calls it his "best adventure yet." He's referring to his job as president and CEO of B.C. Research Inc., a privately-owned research corporation built from the remains of B.C.'s defunct provincial research organization. A description of his latest foray into the business world provided the starting point for his speech at the closing banquet of the IRIS/PRECARN conference.

"I've spent my life in S&T—in academia, in government, in big companies, and in smaller companies," he said. "I've given so much S&T advice, I'm sick of it. Now I get to walk the talk. I've never enjoyed anything so much."

For Dr. Wynne-Edwards, part of "walking the talk" is making sure B.C. Research makes money out of the fruits of science. And the key to that, he says, is the management of science. He identified two main determinants to success: the environmental sustainability of the way products are made and the way people use them, and the capacity to accelerate innovation and rapidly introduce products.

"S&T lies at the root of both these skill sets. Not science alone, but the ability to convert science to technology and carry it to successful commercialization."

Dr. Wynne-Edwards sounded the alarm about Canada's ability to do this. "Canada has to be a good place to raise capital, do research and commercialize products. It's not now.

"Our staggering investment in postsecondary education should show up in economic performance. But in Canada, we haven't coupled the two. We've made a monstrous economic mistake."

The solution? Replace exports of low value with exports of high added value, which means concentrating on commercializing science.

And if you don't like that message? Dr. Wynne-Edwards has a two-word reply: "Stop whining." He has even developed a copyrighted logo—an "unhappy" face with a red line through it—as part of his "no whining" message.

"I'm sick of hearing members of the S&T community say that science and discovery can't be predicted. Their message is 'Get out of my face. Give me more money.'

"We've got this culture of entitlement and the people in the S&T community are card-carrying members. The entitlement industry promotes whining and its head office is in Ottawa. Well, the party's over. The S&T community should stop whining first.

"We must commit to making the right connections, to making sure that the path we're on will lead to products.

"I urge you to look in your labs, at the tools you have invented. Think about what they could be in others' hands. Above all, stop whining. You've had the unquestioning support of Canada. Now is the time for change.

"The great thing about change is —although it may destroy and disrupt—it also renews. I know this from personal experience. Here I stand, with all this grey hair, in the best adventure of my life. I couldn't wish anything better for you."

IRIS Researcher Wins Award

When PRECARN researcher Tom Calvert won the 1995 Xerox Canada-Forum award from the Corporate Higher Education Forum (CHEF), the $5000 he received went directly to one of his research projects. "We really needed the money," he said.

Dr. Calvert knows how tight funding can be. His work at Simon Fraser University spans computer science, kinesiology, and engineering science. One of his projects resulted in a spin-off company that markets *Life Forms*, software that choreographers use to compose and preview dance sequences using computer-animated figures. And he is involved with PRECARN's Intelligent Graphic Interface project as well as projects in IRIS 2.

Dr. Calvert's work also reaches beyond the research lab —to university administration and the development of collaborative research activities. The CHEF award recognizes these efforts, citing him for "remarkable achievement in facilitating corporate-university research collaboration."

As well as founding the University/Industry Liaison Office (UILO) at SFU, Dr. Calvert was president of the Science Council of B.C. from 1990-92. He also served as chair of kinesiology, dean of interdisciplinary studies, director of engineering, and vice-president research and information systems. He is currently director of SFU's centre for systems science.

"Looking back over the years since the establishment of the UILO in 1985, it's amazing to see how dramatically university culture has changed," says Dr. Calvert. "Today there's a willingness on the part of universities to find links with industry and to work with companies. Ten years ago a spin-off company from a university was rare. Now it's commonplace."

But there's still some distance to go to find the best way for universities and industry to collaborate, Dr. Calvert adds. He says that smaller companies are the most challenging interface.

"Smaller companies tend to work on a much shorter time frame," he says. "This puts a lot of pressure on the university team to do something new, implement it, and deliver in one year.

"Smaller companies are where the bulk of the work is. So it's very important for us to work on ways to enhance these collaborations. PRECARN's consortium approach is one mechanism. We need more."

__Harry Rogers__ is President and CEO of PRECARN and Director of IRIS.

To obtain more information on PRECARN, please contact __Lise McCourt__ at PRECARN, Tel: (613) 727-9576, Fax: (613) 727-5672 or E-Mail: mccourt@precarn.ca.

BOOK REVIEWS

Simply logical: Intelligent reasoning by example
Peter Flach Tilburg University, The Netherlands, John Wiley & Sons, 1994, 240 pp; ISBN 0-471-94152-2, Cdn$47.95

Reviewed by
Messaouda Ouerd
University of Ottawa

Books taking a theoretical approach to logic programming are often inaccessible because they are overloaded with detailed mathematical proofs of theorems and lemmas. The practical approach to artificial intelligence programming often ignores important theoretical aspects of logic. Peter Flach has written a practical book that successfully gives the flavor of a combination of logic, artificial intelligence, and computer programming, which are usually treated as distinct subjects in other books. He shows how to implement the theory in runnable Prolog programs and provides the user with a disk with the Prolog programs.

The book is subdivided into three parts. Part I (chapters 1-3) provides an introduction to the main concepts of the programming language Prolog. It addresses the fundamentals of resolution theorem proving in logic and discusses some notions of proof procedure used in Prolog. Part II (chapters 4-6) demonstrates the power of Prolog applied in some central areas of artificial intelligence, including graphs, heuristic search, and knowledge representation. Part III (chapters 7-9) presents a number of recent extensions of logic programming, mainly inductive logic programming. Most chapters also include some paper-and-pencil exercises. The book concludes with an appendix covering a discussion of some other useful programs in logic: a program for logical conversion from predicate logic to clausal logic and a program for predicate completion handling negative information as well as answers to selected exercises.

Chapter 1, "A brief introduction to clausal logic," begins by describing the main concepts in logic programming, such as program clauses, structured terms, recursive data structures, and answering queries. The computation process required to answer a query positively or negatively is illustrated with an example of the London Underground circuit. These discussions of logic programming in Prolog are of benefit to inexperienced programmers as they presuppose no knowledge of the Prolog language.

Chapter 2 is entitled "Clausal logic and resolution: theoretical background." The author introduces the formalism of clausal logic and shows how it can be used in Prolog to perform logical inferences. He starts by describing concepts of resolution theorem proving such as Herbrand models and resolution refutation, and goes on to describe variants of clausal logic and defines for each of them the syntax and the semantics. This chapter relates logic programming theory and Prolog programming practice in a sound manner. Personally, I would like to have seen a little more synthesis and analysis of the logic definitions and concepts.

Chapter 3, "Logic programming and Prolog," describes Prolog itself as a logic programming language. The practical aspects of Prolog such as SLD-trees, negation as failure, arithmetic expression, accumulators, second-order predicates, and meta-programs are discussed. Using the notion of SLD-tree, the author explains the difference between "green" cuts which affect the efficiency but not the meaning of the program, and "red" cuts, which alter the meaning of the program. Although red cuts can improve efficiency, they should be replaced as much as possible by such higher-level constructs as NOT and IF-THEN-ELSE. I would give the title of this chapter to all of Part I, since it explains the concepts of logic programming in terms of the Prolog language.

In Part I, examples of structured knowledge including Prolog terms, proof trees, and SLD-trees were described. SLD-trees play an important role in Prolog, since its proof procedure is based on the principle of SLD-resolution. Those SLD-trees are graphs and searching the space of possible solutions to a query Q results in finding all paths in the SLD-tree from the root of the tree to the empty clause. In Part II—chapters 4, 5, and 6—fundamental artificial intelligence techniques such as searching graphs and forward-chaining are introduced and developed towards their implementation in Prolog, resulting in complete Prolog programs. These programs can be used as building blocks for sophisticated applications. What is nice about these programs is that each one is well described in Prolog and run with some data.

In Part III, the author has nicely assembled several ideas and techniques (Prolog programs) from reasoning with knowledge expressed in natural language, reasoning with incomplete information, and inductive logic programming.

Chapter 7, "Reasoning with natural language," gives a basic introduction to some language concepts, and to definite clause grammars and parsing in Prolog. An example is given of converting quantified English sentences to Prolog rules.

So far, techniques for reasoning with a complete, consistent, and unchanging model of the world have been described. Unfortunately, in many problem domains it is not possible to create such models. In Chapter 8, the author describes some Prolog program techniques for solving problems with incomplete models, such as default reasoning, the closed-world assumption, abduction, and diagnostic reasoning. The goal of inductive reasoning is to infer predicate relations from a set of positive and negative instances of the concept to be learned and a background knowledge. Two main approaches are usually used: bottom-up approaches (also called generalization techniques) and top-down approaches (or specialization techniques). The last chapter of this book describes two Prolog programs inducing relations using these techniques.

Chapters 8 and 9 assume a strong background in the field of reasoning with incomplete information and inductive reasoning "or are written with a graduate student audience in mind." It would be more useful if the author had referred the reader to conference proceedings and journal articles discussing the concepts used.

I found the programs to be quite satisfactory. They are well documented and very helpful to the user. I would recommend this book mostly because of the strength of its excellent examples.

Messaouda Ouerd is a doctoral candidate in Computer Science at the University of Ottawa. She is a member of the Ottawa Machine Learning Group at the University of Ottawa. She is interested in machine learning and databases.

BRIEFLY NOTED

Expertise and technology: Cognition and human-computer cooperation *Jean-Michel Hoc, Pietro C. Cacciabue, and Erik Hollnagel (editors)* (CNRS—University of Valenciennes; CEC Joint Research Centre, Ispra, Italy; and Human Reliability Associates) Hillsdale, NJ: Lawrence Erlbaum Associates (Expertise: Research and applications series, edited by Robert R. Hoffman, Michelene T.H. Chi, K. Anders Ericsson, and Gary A. Klein), 1995, xvi+289 pp; hardbound, ISBN 0-8058-1511-2, US$45.00

This book is the first in a new series on *Expertise: Research and Applications* which will address topics and issues related to the central concept of "expertise." It is largely human expertise and its support through information technology that is addressed, so that both the series and the book under review are primarily concerned with 'human-computer interaction' rather than 'artificial intelligence.' The book is the product of an interdisciplinary working group including psychologists, ergonomists, computer scientists, and control engineers, collaborating through a series of bi-annual meetings since 1987. The 26 authors of the 17 papers in the volume are all well-known and respected contributors to the human-factors literature, and they present research from a wide range of insitutions in seven countries in the EC and USA. Taken overall, the book is an excellent presentation of the state of the art in the discipline that has come to be termed "cognitive ergonomics". It encompasses the applied psychology of the human operators of complex systems, the types of cognitive model that can provide insights into human behavior in controlling such systems, the development of systems that support human operators effectively, and the training of operators to be effective with particular systems. —***Brian R. Gaines**, Knowledge Science Institute, University of Calgary*

Techniques d'analyse et de génération pour la langue naturelle *Claire Gardent and Karine Baschung* (Université d'Amsterdam and Université Paris X—Nanterre) Paris: Editions Adosa (Langue naturelle et traitement de l'information, no. 4, dirigée par Gabriel G. Bès), 270 pages; livre broché, ISBN 2-86639-005-9, FF 250

L'objectif de cet ouvrage est de mettre à la disposition du public français un manuel de référence répertoriant un certain nombre de techniques fondamentales pour le traitement automatique du langage naturel, tout en caractérisant les problèmes évoqués et leurs solutions indépendamment de théories particulières et de langages de programmation particuliers.

Dans cette perspective sont traités les problèmes relatifs au développement d'analyseurs et de générateurs, ainsi qu à celui des grammaires d'unification — problèmes à nos yeux centraux pour la linguistique informatique contemporaine. Une implémentation possible de ces techniques de base est également proposée, en l'occurrence au moyen du langage PROLOG.

Ce livre s'adresse principalement aux étudiants et aux chercheurs en linguistique informatique — tant universitaires qu'industriels — aux étudiants en linguistique, aux étudiants en informatique, ainsi qu'aux enseignants désireux d'appuyer leur cours sur un manuel de référence en français. Dans la mesure où les développements techniques peuvent être sautés aisément (les programmes et les exercices étant matérialisés de manière typographiquement claire), il est également susceptible d'intéresser les néophytes soucieux, à titres divers, des questions de langue, et curieux de se documenter sur les problèmes et les perspectives du traitement du langage humain par l'ordinateur. *—De la notice publicitaire de l'éditeur*

BOOKS RECEIVED

Books marked with a + in the list below are scheduled for review in a future issue. Reviewers are still sought for books marked with a *. Readers who wish to review books for *Canadian Artificial Intelligence* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, or send electronic mail to gh@cs.toronto.edu or gh@cs.utoronto.ca. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

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

***Empirical methods for artificial intelligence** *Paul R. Cohen* (University of Massachusetts) Cambridge, MA: The MIT Press, 1995, xvi+405 pp; hardbound, ISBN 0-262-03225-2, $US55.00

Time-constrained memory: A reader-based approach to text comprehension *Jean-Pierre Corriveau* (Carleton University) Mahwah, NJ: Lawrence Erlbaum Associates, 1995, xx+408 pp; hardbound, ISBN 0-8058-1711-5, $US79.95; paperbound, ISBN 0-8058-1712-3, $US34.50

Idioms: Structural and psychological perspectives *Martin Everaert, Erik-Jan van der Linden, André Schenk, and Rob Schreuder (editors)* (Utrecht University, University of Amsterdam, and University of Nijmegen) Hillsdale, NJ: Lawrence Erlbaum Associates, 1995, vi+329 pp; ISBN 0-8058-1505-6, $US59.95

***Reasoning about knowledge** *Ronald Fagin, Joseph Y. Halperin, Yoram Moses, and Moshe Y. Vardi* (IBM Almaden Research Center, the Weizmann Institute of Science, and Rice University) Cambridge, MA: The MIT Press, 1995, xiii+477 pp; hardbound, ISBN 0-262-06162-7, $US45.00

Artificial minds *Stan Franklin* (University of Memphis) Cambridge, MA: The MIT Press, 1995, xi+449 pp; hardbound, ISBN 0-262-06178-3, $US30.00

Computers and English language learning *John Higgins* (University of Stirling) Oxford: Intellect and Norwood, NJ: Ablex Publishing Corp, 1995, viii+140 pp; hardbound, ISBN 1-56750-189-3 and 1-872526-40-4, $US42.50; paperbound, no ISBN or price listed

Constraint grammar: A language-independent system for parsing unrestricted text *Fred Karlsson, Atro Voutilainen, Juha Heikkila, and Arto Anttila (editors)* (University of Helsinki) Berlin: Mouton de Gruyter (Natural language processing 4, edited by Annely Rothkegel), 1995, viii+430 pp; hardbound, ISBN 3-11-014179-5, DM 268

An introduction to computational learning theory *Michael J. Kearns and Umesh V. Vazirani* (AT&T Bell Laboratories and University of California, Berkeley) Cambridge, MA: The MIT Press, 1994, xii+207 pp; hardbound, ISBN 0-262-11193-4, US$32.50

The trouble with computers: Usefulness, usability, and productivity *Thomas K. Landauer* (University of Colorado) Cambridge, MA: The MIT Press, 1995, xiii+423 pp; hardbound, ISBN 0-262-12186-7, US$27.50

Artificial life: An overview *Christopher G. Langton (editor)* (Santa Fe Institute) [A reprinting of the first three issues of the journal in *Artificial Life*] Cambridge, MA: The MIT Press (Complex adaptive systems series, advised by John H. Holland, Christopher Langton, and Stewart W. Wilson), 1995, xi+340 pp; hardbound, ISBN 0-262-12189-1, $US42.00

***Case-based reasoning in design** *Mary Lou Maher, M. Bala Balachandran, and Dong Mei Zhang* (University of Sydney, University of Wollongong, and Commonwealth Scientific and Industrial Research Organization, Sydney) Mahwah, NJ: Lawrence Erlbaum Associates, 1995, xiii+246 pp; hardbound, ISBN 0-8058-1831-6, US$39.95; paperbound, ISBN 0-8058-1832-4, $19.95

Existential cognition: Computational minds in the world *Ron McClamrock* (State University of New York at Albany) Chicago: The University of Chicago Press, 1995, ix+205 pp; hardbound, ISBN 0-226 55641-7, US$28.95

HORATIO: A middle-sized NLP application in Prolog Archibal Michiels (Université de Liège) Liège: Université de Liège, 198 pp and diskette; paperbound, ISBN 2-87233-015-i, no price listed

Participating in explanatory dialogues: Interpreting and responding to questions in context *Johanna D. Moore* (University of Pittsburgh) Cambridge, MA: The MIT Press (The ACL—MIT Press series in natural language processing, edited by Aravind K. Joshi, Karen Sparck Jones, and Mark Y. Liberman), 1995, xiv+352 pp; hardbound, ISBN 0-262-13301-6, US$49.95

Type logical grammar: Categorial logic of signs *Glyn V. Morrill* (Polytechnic University of Catalonia) Dordrecht: Kluwer Academic Publishers, 1994, xiv+307 pp; hardbound, ISBN 0-7923-3095-1, US$99.00

Cooperative computer-aided authoring and learning: A systems approach *Max Muhlhauser (editor)* (University of Karlsruhe) Dordrecht: Kluwer Academic Publishers, 1995, xiii+351 pp; hardbound, ISBN 0-7923-9527-1, US$95.00

***Applications of neural networks** *Alan F. Murray (editor)* (University of Edinburgh) Dordrecht: Kluwer Academic Publishers, 1995, xii+322 pp; hardbound, ISBN 0-7923-9442-9, US$110.00

Computational learning theory and natural learning systems. Volume III: Selecting good models *Thomas Petsche, Stephen José Hanson, and Jude Shavlik (editors)* (Siemens Corporate Research and University of Wisconsin—Madison) Cambridge, MA: The MIT Press, 1995, xxiii+411

Artificial intelligence: A modern approach *Stuart J. Russell and Peter Norvig* (University of California, Berkeley and Harlequin Inc) Englewood Cliffs, NJ: Prentice-Hall (Prentice-Hall series in artificial intelligence), 1995, xviii+932 pp; hardbound, ISBN 0-13-103805-2

Computational lexical semantics *Patrick Saint-Dizier and Evelyne Viegas (editors)* (Institut de Recherche en Informatique de Toulouse, CNRS and Brandeis University) Cambridge University Press (Studies in natural language processing, edited by Branimir K. Boguraev), 1995, ix+447 pp; hardbound, ISBN 0-521-44410-1, US$69.95

Computerized grammars for analysis and machine translation *Bengt Sigurd (editor)* (University of Lund) Lund: Lund University Press (Travaux de l'Institute de Linguistique de Lund 29), 1994, 148 pp; also distributed by Chartwell-Bratt, Bromley, Kent, U.K.; paperbound, ISBN 91-7966-304-4 (Lund) and 0-86238-383-8 (Chartwell-Bratt), US$30.00

Spoken natural language dialog systems: A practical approach *Ronnie W. Smith and D. Richard Hipp* (East Carolina University and Hipp, Wyrick and Company Inc) New York: Oxford University Press, 1994, xiv+299 pp; hardbound, ISBN 0-19-509187-6, US$49.95

Computational architectures integrating neural and symbolic processes: A perspective on the state of the art *Ron Sun and Lawrence A. Bookman (editors)* (University of Alabama and Sun Microsystems Laboratories) Boston: Kluwer Academic Publishers (The Kluwer international series in engineering and computer science), 1995, xviii+475 pp; hardbound, ISBN 0-7923-9517-4, US$120.00

Representing time in natural language: The dynamic interpretation of tense and aspect *Alice G.B. ter Meulen* (Indiana University) Cambridge, MA: The MIT Press, 1995, xii+144 pp; hardbound, ISBN 0-262-20099-6, US$25.00

Conceptual structures: Current practices [Proceedings of the Second International Conference on Conceptual Structures, University of Maryland, August 1994] *William M. Tepfenhart, Judith P. Dick, and John F. Sowa (editors)* (AT\&T Bell Laboratories, University of Maryland, and State University of New York at Binghamton) Berlin: Springer-Verlag (Lecture notes in artificial intelligence, edited by J.G. Carbonell and J. Siekmann, volume 835), 1994, viii+331 pp; paperbound, ISBN 3-540-58328-9 and 0-387-58328-9, no price listed

The creative process: A computer model of storytelling and creativity *Scott R. Turner* (University of California, Los Angeles) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, ix+298 pp; hardbound ISBN 0-8058-1576-7, US$49.95

CSCSI/SCEIO Membership

I wish to join CSCSI/SCEIO and receive Canadian Artificial Intelligence

- [ ] **• Web Access Only ($30.00* Cdn./yr.)**
- [ ] **• Printed Copy and Web Access ($40.00 *Cdn./yr.)**

I am a student

- [ ] **• Web Access Only ($15.00* Cdn./yr.)**
- [ ] **• Printed Copy and Web Access ($15.00* Cdn./yr.)**

I am a member of CIPS

- [ ] **• Web Access Only ($25.00* Cdn./yr.)**
- [ ] **• Printed Copy and Web Access ($30.00* Cdn./yr.)**

Name ______________________________

Mailing Address ______________________________

____

____

Please mail your membership to:

CIPS
430 King Street West, Suite 106
Toronto, Ontario
M5V 1L5

Phone: (416) 593 - 4040
Fax: (416) 593 - 5184

For more information contact CIPS or a member of the executive.
***Includes Applicable G.S.T.**

Advertisers Index

Call for Papers - CP96

Second International Conference on Principles and Practice of Constraint Programming

Cambridge, Massachusetts, USA, August 19-22, 1996

CP96 will be held in conjunction with the first appearance in the USA of the Conference on the Practical Applications of Constraint Technology, PACT-USA 96, August 23, 1996.

Constraints have emerged as the basis of a representational and computational paradigm that draws from many disciplines and can be brought to bear on many problem domains. The conference is concerned with all aspects of computing with constraints including: algorithms, applications, environments, languages, models, systems. Contributions are welcome from any discipline concerned with constraints and from any domain employing constraints. Papers are especially welcome that bridge disciplines or combine theory and practice. Workshops and tutorials can be organized by interested groups on August 19.

The proceedings of the conference will be published by Springer in the Lecture Notes in Computer Science series. Authors of selected papers will be invited to submit revisions to a special issue of the Constraints journal devoted to the conference. An award will be given to the paper that best exemplifies the interdisciplinary spirit of the conference. Submissions are due by February 15, 1996 and should be sent to the Program Chair.

A detailed call for papers and additional conference information is available from the conference web site: http://www.cs.ualberta.ca/~ai/cp96.

Program Chair

Eugene C. Freuder
Department of Computer Science
University of New Hampshire
Kingsbury Hall M208, College Road
Durham, New Hampshire 03824 USA
ecf@cs.unh.edu

Organizing Committee

Alan Borning, Jacques Cohen, Alain Colmerauer, Eugene Freuder, Herve Gallaire, Jean-Pierre Jouannaud, Paris Kanellakis, Jean-Louis Lassez (chair), Ugo Montanari, Anil Nerode, Vijay Saraswat, Ralph Wachter

Sponsored in part by the Canadian Society for Computational Studies of Intelligence / Société canadienne pour l'étude de l'intelligence par ordinateur (CSCSI/SCEIO).

Established in 1983, Applied AI Systems, Inc (AAI) is the oldest Artificial Intelligence company in Canada. The emphasis of our business is on real world applications of cutting edge intelligent systems technology. AAI 's long-term commitment to the "paying respect to science" approach is now reaping benefits. The company is international in approach, with its members constantly traveling the globe to meet other researchers and practitioners, and participating in all major conferences, workshops, and symposia in related fields.

Reseach and Development Areas: --- Intelligent Robots --- Autonomous Systems --- Interface Agents and other "Intelligent Agents or Soft Robots" --- Evolutionary Systems and Artificial Life --- Evolutionary Robotics (ER) --- Neural Network applications--- Real Time Knowledge Based Systems (RKBS) --- Speech Recognition ---

Service Provided to Clients: --- Custom System Development and Implementation --- Technical Research Reports --- Sales of Intelligent Robots and AI Products --- Training ---Consulting ---

Eight Legged High Performance Walking Robot

OCT-I ™

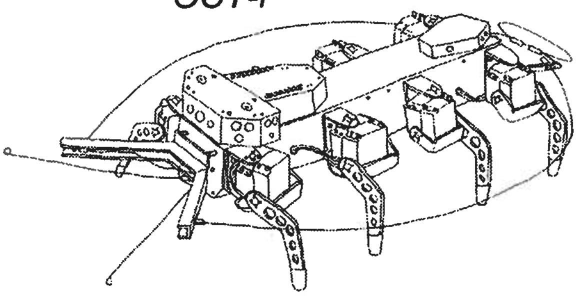

- Eight Legged lobster-shaped robot with 9 whisker sensors (6 front, 1 rear and 2 side).
- Good for walking robot experiments and educational exhibition.
- Processor: Motorola 68332 with 64KByte ROM (expandable to 1 MByte) and 1 MByte RAM.
- Sensors: IR transmitters (2 front, 2 side, 2 back), 4 IR receivers, 11 ambient light sensors (photodiodes), 16 current sensors, 2 on each leg.
- Measures 65 cm long x 28 cm wide (45 cm with whiskers) x 13 cm high.

- Highly affordable and powerful 2-wheel drive mobile robot base with 2 DC motors and encoders.
- Good for research of multiple intelligent behaviors or interactions between a robot and its environment.
- Saphira™ software for navigation and map making, in C, with hooks for custom routines.
- Onboard MC68HC11-based controller and 9-pin RS-232 serial port for external I/O, and internal 10-pin ribbon cable connector to support onboard serial I/O.
- Additional applications can be added in C or C++.
- 7 sonars and 2 range-sensing IR's.
- Measures 40.5 cm long x 33 cm wide x 25 cm high.

REAL WORLD INTERFACE, INC

Price-Performance Mobile Robot Base

Pioneer 1™

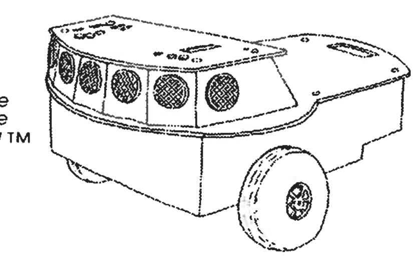

- Small (60 mm diameter and 30 mm height), modular mobile robot.
- Over 200 sold world wide.
- Good for learning in robot control, collective behavior, and Evolutionary Robotics (ER) experiments.
- Laser robot location/orientation measurement system.
- Has 256 KByte RAM and 8 infrared/light sensors to avoid objects.
- Optional modules: linear vision, gripper, additional sensor turret, video camera module, and I/O turret.

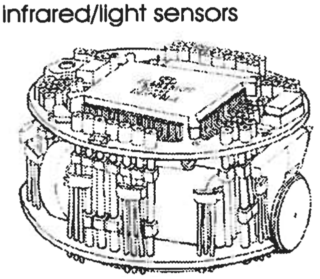

Miniature Mobile Robot

Khepera™

For further information on pricing or ordering, please contact:

Applied AI Machines & Software, Inc.
340 March Road, Suite 601
Kanata, Ontario, Canada K2K 2E4
Phone: (613)592-7729
Fax: (613)592-9762
E Mail: 73051.3521@compuserve.com

For further information on consultation or research and development, please contact:

Applied AI Systems, Inc.
340 March Road, Suite 600
Kanata, Ontario, Canada K2K 2E4
Phone: (613)592-3030 Fax: (613)592-2333
E Mail: aai@AppliedAI.com
http://fox.nstn.ca/~aai/