





British Computer Society Workshop Series

# Expert Systems 85

Proceedings of the  
Fifth Technical Conference of the  
British Computer Society  
Specialist Group on Expert Systems

Edited by Martin Merry





# Expert Systems 85

Proceedings of the Fifth Technical  
Conference of the British Computer Society  
Specialist Group on Expert Systems

University of Warwick, 17-19 December 1985

MARTIN MERRY

*Hewlett Packard Laboratories*



*The right of the  
University of Cambridge  
to print and sell  
all manner of books  
was granted by  
Henry VIII in 1534.  
The University has printed  
and published continuously  
since 1584.*

CAMBRIDGE UNIVERSITY PRESS

*Cambridge*

*London New York New Rochelle*

*Melbourne Sydney*

Published by the Press Syndicate of the University of Cambridge  
The Pitt Building, Trumpington Street, Cambridge CB2 1RP  
32 East 57th Street, New York, NY 10022, USA  
10 Stamford Road, Oakleigh, Melbourne 3166, Australia

© British Informatics Society Ltd 1985

First published 1985

Printed in Great Britain at the University Press, Cambridge

Library of Congress cataloguing data available

*British Library cataloguing in publication data*

British Computer Society: *Specialist Group on  
Expert Systems. Technical Conference (5th: 1985:  
University of Warwick)*

Expert systems 85: proceedings of the Fifth  
Technical Conference of the British Computer  
Society Specialist Group on Expert Systems,  
University of Warwick, 17-19 December 1985.—  
(The British Computer Society Workshop series)

1. Expert Systems (Computer science)

I. Title II. Merry, Martin III. Series

006.3'3 QA76.9.E96

ISBN 0-521-32596-X

## **Expert Systems 85**

## THE BRITISH COMPUTER SOCIETY WORKSHOP SERIES

EDITOR: P. HAMMERSLEY

The BCS Workshop Series aims to report developments of an advanced technical standard undertaken by members of The British Computer Society through the Society's study groups and conference organisation. The Series should be compulsive reading for all whose work or interest involves computing technology and for both undergraduate and post-graduate students. Volumes in this Series will mirror the quality of papers published in the BCS's technical periodical *The Computer Journal* and range widely across topics in computer hardware, software, applications and management.

Some current titles:

Data Bases: Proceedings of the International Conference 1980

Ed. S. M. Deen and P. Hammersley

Minis, Micros and Terminals for Libraries and Information Services

Ed. Alan Gilchrist

Information Technology for the Eighties BCS '81

Ed. R. D. Parslow

Second International Conference on Databases 1983

Ed. S. M. Deen and P. Hammersley

Research and Development in Information Retrieval

Ed. C. J. van Rijsbergen

Proceedings of the Third British National Conference on Databases (BNCOD 3)

Ed. J. Longstaff

Research and Development in Expert Systems

Ed. M. A. Bramer

Proceedings of the Fourth British National Conference on Databases  
(BNCOD 4)

Ed. A. F. Grundy

People and Computers: Designing the Interface

Ed. P. Johnson and S. Cook

Expert Systems 85

Ed. M. Merry

## *Organising Committee*

*General Chairman:* Bernard Kelly

*Deputy chairman:* Tom Addis

### *Programme committee*

Martin Merry (Chairman)  
Richard Young (Deputy Chairman)  
Max Bramer  
Julie Gadsden  
Alison Kidd  
Tim O'Shea  
Peter Ross

John Lumley (Applications Stream)  
Keith Clark  
Jim Hunter  
Ken MacCallum  
Kevin Poulter  
Robert Worden

### *Referees*

Tom Addis  
Max Bramer  
Bill Clocksin  
Robert Corlett  
Peter Hammond  
Jim Hunter  
Caroline Knight  
Ken MacCallum  
Tim O'Shea  
Allan Ramsay  
Bill Sharpe  
Sam Steel  
Robert Worden

Peter Alvey  
Keith Clark  
Derek Coleman  
Julie Gadsden  
Graham Higgins  
Alison Kidd  
John Lumley  
Jo Marks  
Kevin Poulter  
Peter Ross  
Karen Sparck Jones  
Steve Todd

## *Preface*

This volume is the Proceedings of the Fifth Annual Conference of the British Computer Society Specialist Group on Expert Systems, held at the University of Warwick in December 1985. Following the precedent set in 1984, it includes an introductory paper written by the programme chairman.

The proceedings include all the refereed papers which were presented at the conference, together with the invited papers by Austin Tate and Abe Mamdani; papers from the other invited speakers were not available at the time of going to press.

I would like to thank all those concerned for their work in putting together the programme for this conference: in particular the members of the Programme Committee, and all those who refereed papers.

Martin Merry  
Programme Chairman

# Contents

	<i>Organising committee</i>	<i>page</i> vii
	<i>Preface</i>	ix
1	Expert systems – some problems and opportunities <i>M. Merry</i>	1
2	What do users ask? – Some thoughts on diagnostic advice <i>A. L. Kidd</i>	9
3	Relevant criteria for choosing an inference engine in expert systems <i>H. Reichgelt and F. van Harmelen</i>	21
4	Generalised alpha-beta pruning as a guide to expert system question selection <i>C. S. Mellish</i>	31
5	Expert system techniques: an application in statistics <i>L. HaKong and F. R. Hickman</i>	43
6	A model based expert system for HW troubleshooting driven by compiled control knowledge <i>F. Giovannini and F. Malabocchia</i>	65
7	ESCORT: The application of causal knowledge to real-time process control <i>A. Paterson, P. Sachs and M. Turner</i>	79
8	A review of knowledge-based planning techniques <i>A. Tate</i>	89
9	Allocating abilities to actors <i>C. Trayner</i>	113
10	Choice making in planning systems <i>D. Croft</i>	125
11	The ECO browser <i>D. Robertson et al.</i>	143
12	A CAD/CAPP expert system shell <i>V. I. Kirov</i>	157
13	Symbolic uncertain inference: a study of possible modalities <i>J. G. Ganascia and Y. Kodratoff</i>	169



14	Inference under uncertainty <i>A. Mamdani, J. Efstathiou and D. Pang</i>	181
15	SOJA: A daily workshop scheduling system, SOJA's system and inference engine <i>C. Lepape</i>	195
16	Real-time multiple-motive expert systems <i>A. Sloman</i>	213
17	O-plan – Control in the open planning architecture <i>K. Currie and A. Tate</i>	225
18	An expert system for efficient office object management <i>K. L. Mannoek and C. H. C. Leung</i>	241
19	Comprehension by model-building as a basis for an expert system <i>J. Cunningham</i>	259
20	Acquisition of control and domain knowledge by watching in a blackboard environment <i>C. D. B. Boyle</i>	273
21	Altering the description space for focussing <i>J. Wielemaker and A. Bundy</i>	287
22	Deep knowledge representation techniques <i>A. G. Cohn</i>	299
23	An application of knowledge based techniques to VLSI design <i>H. J. Kahn and N. P. Filer</i>	307
24	Where's the expertise? Expert systems as a medium of knowledge transfer <i>H. M. Collins, R. H. Green and R. C. Draper</i>	323

## EXPERT SYSTEMS - SOME PROBLEMS AND OPPORTUNITIES

Martin Merry  
Hewlett-Packard Research Laboratories  
Bristol

### INTRODUCTION

We are constantly being told that we are living in an electronic age; that we are undergoing the second industrial revolution; and that the information technology era is now upon us. Journalists describe the latest advances in computing in tones usually reserved for slow motion pictures of the Wonders of Nature or close up views of Halley's Comet.

Despite all this reverence, however, computer science is very young. It has roots in a number of different disciplines (engineering, mathematics, psychology, neurophysiology,...) and still forms a rather uneasy synthesis of ideas from these areas. Unsurprisingly, therefore, progress has been much greater in some areas than in others.

As far as hardware performance goes, exaggeration is scarcely needed. The last 30 years have seen 6 orders of magnitude improvement in hardware performance/cost. Expected lifetimes for new hardware products decrease constantly (as anyone who has bought a micro knows - if only you'd waited 6 months you'd have been able to buy something substantially better and cheaper).

In other areas of computing, however, progress has been rather less meteoric. This is due in part to the difficulty of the many problems that need to be solved, but is also due to delays in technology transfer from laboratories to general use. For example, think how long it took Pascal to emerge from universities into widespread use: even now, FORTRAN and COBOL still have a strangle-hold over many sectors of computing. Other examples now reaching a wider audience include: object oriented programming; logic programming; and the "overlapping windows" user interface paradigm. The typical gestation period for new ideas, languages, etc. appears to be between 10 and 15 years.

Expert systems are no exception. The well-known early expert systems, DENDRAL and MYCIN date from the late 60's and early 70's (Buchanan et al 69; Shortliffe 74), and yet the expert systems "boom" has really only appeared over the last two or three years. Most current applications work involves very few substantive new ideas over these early systems. Arguably, the most noticeable advances in current application systems are inherited directly from advances in the underlying hardware - systems can now be developed on reasonably sized hardware and run relatively quickly.

While all this is rather depressing from the point of view of wanting to see new ideas taken up quickly, it should at least make it somewhat easier to predict what's likely to happen in the next few years. In this paper we have a quick glance at three particular topics currently being explored in research laboratories: knowledge-based planning, new architectures for expert systems, and qualitative reasoning. This will hopefully show us what is likely to come into public view in the near future.

#### KNOWLEDGE-BASED PLANNING SYSTEMS

Most existing expert systems work in *analytic* domains, where problem-solving consists of identifying the correct solution from a pre-specified finite list of potential answers e.g. fault diagnosis is concerned with identifying which potential fault is actually present. Many possible application domains, however, do not have this restriction: these domains are *synthetic* - problem-solving involves actually synthesizing a new solution. This is substantially more complex.

Over the past few years there has been increasing interest in building expert systems in these sort of domains, drawing on techniques from AI planning and expert systems. AI planning systems actually pre-date expert systems - significant work was being done in the mid 60's (e.g. Doran & Michie 1966) - but the convergence of the two streams is relatively recent. Expert systems applied to synthetic domains necessarily use planning techniques; modern planning systems use significant amounts of knowledge from the application domain to help formulate their plans. A survey of knowledge-based planning techniques will be found in the review paper by Austin Tate (Tate 86) in this volume.

It is not only in synthetic domains that the techniques of knowledge-based planning are important for expert systems. The control of expert systems themselves is in itself a real-time planning task (Lesser 1984). This will arguably be the dominant use of these techniques as larger and larger expert systems are built.

In my own work on expert systems I have found that the large majority of problems that people have brought to me, looking for expert system solutions, have required knowledge-based planning techniques; I have regretfully sent most of these people away again because the necessary techniques just did not exist or were not sufficiently robust. This is slowly changing. More and more experimental systems are being built to tackle these sort of problems; the number of people working in these areas is steadily increasing. For example, the Special Interest Group on Planning, formed as part of the Alvey Expert Systems Research Theme, has gone from strength to strength over the last two years and now holds regular workshops. In this volume there are more papers on these sorts of systems than in any of the earlier proceedings of BCS Expert Systems conferences; we expect this number to increase again next year.

#### NEW ARCHITECTURES FOR EXPERT SYSTEMS

The traditional expert system architecture, consisting of a single *knowledge base* and an *inference engine*, is well-known. Within this framework there are many variations - different flavours of knowledge representation, uncertain reasoning, control strategies etc. - for a good analysis of a number of traditional expert systems see Johnson and Keravnou 1985.

Whilst a lot has been achieved within this framework, as more complex problems have been tackled, it has been found to be inadequate, and a number of variations have emerged. Probably the best known is the *blackboard* architecture, where the expert system has a number of distinct knowledge bases, each of which contribute to problem-solving by writing hypotheses etc. about the problem to be solved on a common "blackboard". This architecture was originally developed for the speech understanding system HEARSAY-II (Erman et al 1980).

In general, frameworks like the blackboard architecture are needed when different types of knowledge are used in an expert system,

which may need, for example, different forms of knowledge representation. This leads to partitions in the knowledge base and then to different means of drawing inferences from the different types of knowledge. A number of systems tackling such problems have been built, leading to a host of different architectures (many of these are called blackboard architectures but have many differing features; others, like MOLGEN mentioned earlier, use rather different techniques). Different architectures have also been developed for many other reasons e.g. to cluster related pieces of knowledge (Aikins 83).

I have been involved in a number of feasibility studies for systems with the requirement to handle varying kinds of knowledge. Many problems fall into this category; these are not solved by simply putting knowledge into a conventional shell. Something more flexible is needed - today's toolkits are a short term response to the problem.

Like knowledge-based planning systems, we believe systems based round new architectures will become much more prominent over the next few years. However, many of these architectures are in a state of flux - it will probably be a long time before they stabilize. In many ways, this is because they lack any sort of formally defined semantics - they can be criticized as short term technical fixes to problems that need to be understood much better before longer term solutions can be found. The next topic we consider is one direction for research on such long term problems.

#### QUALITATIVE REASONING

The first two topics we have discussed both appear to be ideas whose time has come; work in both areas has acquired its own momentum and no great predictive skill is needed to see that these topics will become more and more influential. Qualitative reasoning, unfortunately, does not come into this category.

Most current expert systems only represent very shallow expert knowledge - a collection of fragments of compiled expertise. These, in general, form a small fragment of the problem-solving skills an expert brings to bear upon a problem. Typically, in addition to rules of thumb, the expert is able to reason about the deep structure of a problem. If the domain is sufficiently concrete, this reasoning may be completely quantitative e.g. the manipulation of particular mathematical

formalisms. In general, however, much of this reasoning is qualitative, and is concerned, for example, with questions of causality.

We can illustrate this by means of a toy example. An expert system for fault diagnosis in a vending machine might have a rule like "IF you have inserted 10p AND a cup has come out AND there is no drink in the cup THEN hit the machine just to the left of the tea-no-sugar button". This rule is highly specific to a particular (model of) vending machine. If one wished to build an expert system to diagnose faults in an arbitrary vending machine, based round a model of the machine in question, it would need to be able to follow a chain of reasoning similar to "the coin is progressing through the machine...it has passed the mechanism which releases the cup...it has not reached the mechanism which releases the drink...it is therefore stuck at a certain place...this is immediately behind just to the left of the tea-no-sugar button...therefore hit the machine here".

Qualitative reasoning is concerned with capturing this kind of problem-solving skill. Slightly blurring the definition, it is concerned with all kinds of representation and manipulation of "deep" knowledge, and also "common-sense" reasoning. It is a problem that has been looked at throughout the whole of the history of Artificial Intelligence (e.g. see McCarthy 58).

Much work has been done; there are now expert systems based round qualitative models (Weiss et al 78; Mozetic et al 83). However, these systems were largely ad hoc; there is still no principled way of handling qualitative reasoning and many extremely hard problems remain to be solved (some argue that these problems are inherently insoluble e.g. see Dreyfus 1981).

Nevertheless, at least a partial solution to some of these problems needs to be found in order to realize many of the claims currently being made for expert systems. For example, one property of human experts is *graceful degradation* - if you ask an expert a question which is slightly outside his domain of expertise you tend to get an answer which is reasonable, if not completely correct; if you ask a similar question of an expert system you get garbage in response. [This of course is true of other computer systems - if you give a Pascal program to a FORTRAN compiler you don't expect it to produce almost working object code, but at least it produces a string of error messages - current

expert systems don't "know what they know", and are as likely to produce wrong advice as no advice at all].

Qualitative reasoning, then, is an area in which much work is needed. We do not expect to see many application papers presented at Expert Systems conferences in the near future which are built round qualitative models - we do hope, however, to see more papers on theoretical work in this area.

### CONCLUSION

In this paper we have briefly looked at some of the current work on expert systems which is likely to make itself widely felt in the near future. The first two topics, knowledge-based planning systems, and systems which use differing types of knowledge already have many exemplars and are likely to become of great commercial significance relatively quickly. The third topic, qualitative reasoning, is already making itself felt: it is conspicuous by its absence.

In the introduction to the paper we pointed out that a slow rate of perceived progress can be attributed to two things: the time taken for technology transfer, and the difficulty of many of the problems to be solved. While most of this paper has been motivated by the first of these reasons one must not underestimate the second. Qualitative reasoning is but one example of an extremely hard problem area where major theoretical breakthroughs are needed before we will see significant progress (others include: reasoning about time, a *good* way of handling uncertainty, reasoning with defaults, learning from mistakes, etc.).

One advantage of the expert systems boom is that these problems are receiving far more publicity than ever before, and that hopefully therefore more effort is being expended on them; a disadvantage is that expectations have been raised which will need some of these problems to be solved before they can be met. We must be careful not to assume that mere publicity is sufficient to solve these problems; mounting large scale projects which require unsolved research questions to be answered before they can be completed is rather unwise. Raising unjustifiable expectations in the early 70's led, amongst other things to the removal of much of the funding for work on AI in the UK - we must not let this happen again.

ACKNOWLEDGEMENTS

I would like to thank John Lumley, Bill Sharpe and Steve Todd for their constructive comments on earlier drafts of this paper.

REFERENCES

- Aikins, J. S. 1983. Prototypical knowledge for expert systems. *Artificial Intelligence* 20 pp 163 - 210.
- Buchanan, B.G., Sutherland, G.L. and Feigenbaum, E.A. 1969. Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry. In Meltzer, B. and Michie, D. (eds) *Machine Intelligence 4*. Edinburgh University Press.
- Doran, J. E. and Michie, D. 1966. Experiments with the graph traverser program. *Proc Royal Society A*, pp 235 - 259.
- Dreyfus, H.L. 1981 From micro-worlds to knowledge representation: AI at an impasse. In Haugeland, J. (ed.) *Mind Design* Bradford Books.
- Erman, L. D., Hayes-Roth, F., Lesser, V. and Reddy, D. 1980. The HEARSAY-II speech-understanding system: Integrating knowledge to resolve uncertainty. *Computing Surveys* 12, pp 213 - 253.
- Johnson, L. and Keravnou, E.T. 1985 *Expert Systems Technology: a guide*. Abacus Press.
- Lesser, V. 1984. Control in complex knowledge-based systems. Tutorial presented at *First Conference on Artificial Intelligence Applications*, Denver.
- McCarthy, J. 1958. Programs with common sense. *Mechanization of thought processes*, *Proc Symp Nat Phys Lab* vol 1 pp 78 - 84 London : HMSO. Reprinted in Minsky, M. (ed.) *Semantic Information Processing* MIT Press 1968.
- Mozetic, I., Bratko, I. and Navrao, L. 1983 An experiment in Automatic Synthesis of Expert Knowledge through Qualitative Modelling. *Proc Logic Programming Workshop 1983* Algarve, Portugal.
- Shortliffe E.A. 1974. MYCIN: a rule-based computer program for advising physicians regarding antimicrobial therapy selection. Ph.D. dissertation, Stanford University.
- Stefik, M. J. 1981a. Planning with constraints. *Artificial Intelligence* 16 pp 111 - 140.
- Stefik, M. J. 1981b. Planning and meta-planning. *Artificial Intelligence* 16 pp 141 - 169.
- Tate, A. 1986. A Review of Knowledge-Based planning techniques. This volume.
- Weiss, S. M., Kulikowski, C.A., Amarel, S. and Safir, A. 1978. A model-based method for computer-aided medical decision making. *Artificial Intelligence* 11 pp 145 - 172.



