AD-A192 468 KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS IN CONSTRUCTION(U) LOUGHBOROUGH UNIV OF TECHNOLOGY (ENGLAND) DEPT OF CIVIL ENGINEERING E G TRIMBLE ET AL. 1/1

UNCLASSIFIED FEB 88 DAJA45-85-C-0033 F/G 12/5 NL

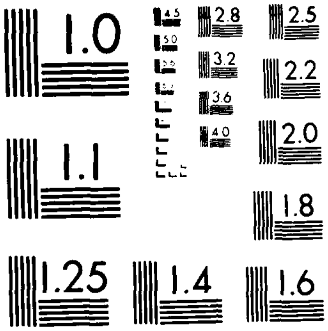

MICROCOPY RESOLUTION TEST CHART

NATIONAL BUREAU OF STANDARDS 1963-A

LOUGHBOROUGH UNIVERSITY OF TECHNOLOGY

DEPARTMENT OF CIVIL ENGINEERING

KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS IN CONSTRUCTION

U.S. ARMY R 7 D GROUP (UK)

E G Trimble

R J Allwood

A E Bryman

C N Cooper

Sixth interim report

Contract DAJA 45-85-C-0033

February 1988

THE DESKTOP GURUS

Expert systems are now being used on real engineering problems, and appear to be fulfilling their early promise. Nigel Atkinson reports on one company's progress.

Science fiction has for decades described computers which, in an instant, can assess immense problems to help their hopelessly inadequate human masters to a solution. As a result there is something of a mystique surrounding attempts at artificial intelligence which has led to high expectations for expert systems. Many apparently level-headed people seem to assume that rules are simply fed into such a system and, miraculously, a previously unobtainable solution will appear. In reality, they are only another software tool, which process data according to preprogrammed rules.

There has been incredible industrial interest in this area over recent years according to Peter Andow — an expert in the AI field, who has recently moved from Loughborough University to champion expert systems at KBC Process Automation in Southampton. Five years ago, all the major companies wanted to send someone on a course, says Andow, but virtually none could supply a lecturer. Now nearly every company of a reasonable size seems to have one or more engineering applications, though Andow feels many others have 'toys', distanced from reality. However, most have given positive, though sometimes unexpected, benefits to their authors.

Andow has seen many companies rush into diagnostic expert systems which the intended users, plant operators, never see. 'In developing a system, engineers find out more about the plant and add equipment which makes it easier to operate anyway. Alternatively, having found out what the rules are for diagnosis, it is found that the operators can be trained to apply them without the help of a computerised expert system.'

Successful applications of expert systems in the process industries have tended to concentrate on plant operation. Valuable fault diagnosis systems have been developed for operators, and there have also been interesting developments in the computer evaluation of trips and alarms. Plant maintenance functions have also benefitted from application of artificial intelligence. Attention is now turning to design, with process contractors beginning to adopt the technology. One of these, Foster Wheeler, has developed a system for material selection of boiler tubes which is currently used on industrial contracts.

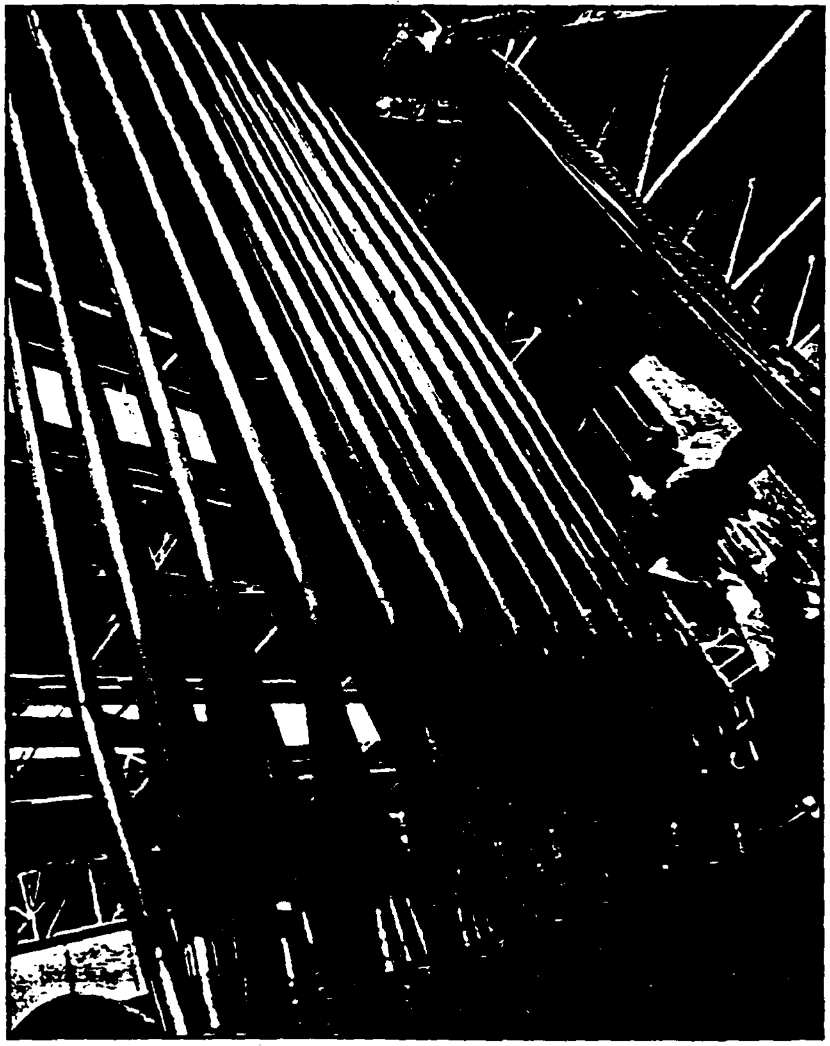

Material selection for boiler tubes like these is now done using an expert system.

Involvement in expert systems at Foster Wheeler Power Products, in London's Mornington Crescent, began earlier this year. Martin Allen, engineering computing systems manager, had wanted to try the technology for some time, and his opportunity came with two surprise offers: a free six month software trial from IBM, and the time of Chris Cooper, an experienced knowledge engineer researching at Loughborough University Civil Engineering Department. The main question remaining was the choice of application. As Allen explains, 'We wanted to choose an engineering subject, and one where we had the knowledge at hand so that we did not depend on one engineer who may suddenly become very busy on a contract.'

Allen's department designs large boilers, and several expert system applications were assessed. ► 34

33 ◄ Material selection for boiler tubes was a subject already earmarked for a conventional program. It was seen as an ideal subject required within every contract, including the proposal stage, the rules involved are based on codes and the engineers concerned were readily available.

The next task was to acquire the knowledge. This step is described by many in the AI field as a major bottleneck, though Cooper considers the identification of structures within the gathered knowledge as equally difficult. Allen felt an important factor would be the reaction of the engineers. 'We were a little worried they would feel we were replacing them with a computer. The people aged under 50 were very open and thought the work we were doing was very worthwhile. Those about to retire were a little more skeptical but once they saw the system working they tended to be much more enthusiastic. Material selection is not one of the most enjoyable subjects — searching through codes to make sure you get the right part is always as bit tedious. The system is seen as a great help — particularly in the proposals department.'

Two engineers were selected for the initial interviews, and each went through manual examples referring to BS1113 and the company design manual. From the interviews and references, Cooper derived a series of rules, known as the paper model, which formed the basic structure of the system.

Hour-long interviews continued on a weekly basis, slowly becoming more focused on difficult aspects such as tube-bending and costing. The paper model was continually updated and circulated around the department for comments. As all the rules were in English, it was readily understood, leading to faster development and checking than conventional systems. 'We have found that it does not require much expert input for us to gather the knowledge required to build a system', said Cooper. 'The total amount of interaction with experts was about 30hours, with the total mantime for the whole project around 500hours.' Around 12-15 people were consulted for the system, ranging from quality assurance and design engineers to the manager of the procurement department to discuss costing.

Once the paper model appeared to offer an effective solution, Foster Wheeler programmer Andy Jones began to use the IBM ESE software. Until this stage the software to be used was irrelevant, though low-cost shells require rules to be very rigidly constructed. Programming ESE from the paper model proved relatively simple, and considerably quicker than a conventional language like Fortran. In addition to the direct correlation between the rules in the model and those in the system, features of the software such as the automatic generation of screens and questions, interactive compiling and de-bugging were found to be very helpful.

With the initial system working, engineers again went through examples — this time using the system to see if it gave the same answers. Some discrepancies were found to be errors by the engineer, who had failed to take relevant aspects of the code into account. The final system contains 145 rules, most of which represent actual knowledge rather than software manipulations. Knowledge is drawn from BS1113, the Power Products design manual and engineers — with some of the human knowledge previously undocumented.

A feature of the final system, named Matsel, is the assessment of material costs. 'We felt that the early consideration of costs was important,' said Allen. 'The numbers given at the moment are based on crude prices, but ultimately they could be linked to a database maintained by the procurement department so that the numbers given to the proposals engineers would always be up to date.

Reaction to Matsel within Foster Wheeler has been enthusiastic according to Allen, who was himself impressed by the speed it was developed. He sees several advantages in the further use of expert systems in design. First, they provide a consistent and correct approach to design which can easily be up-dated in accordance with new standards. During the development of Matsel one of the company rules was changed, leaving the manual out of date for several weeks. The expert system was corrected in five minutes, and at a much lower cost. The change was also much cheaper than reprinting pages, which may not fully replace the original versions leaving the possibility of design being carried out incorrectly. Indeed, Allen sees expert systems making paper documents, such as manuals, obsolete.

Second, expert systems can be developed by the end user without early involvement of computer personnel. They do not involve complicated computer jargon, and are easily maintained. Ease of checking is also an advantage according to Cooper. 'Expert systems solve problems using logic, and can therefore explain the reasoning they are using. Computer systems which are able to explain their reasoning are very flexible and very user-friendly. I think that is the most important reason why we should want to use an expert system approach rather than a procedural language.'

Information can also be collected in an expert system which would otherwise be lost to the company, either through retirement, promotion or people leaving work elsewhere. This is particularly valuable in areas of novel technology, where few employees are involved in detailed design, and in these areas they would also allow a more rapid expansion of expertise once commercial viability has been proven.

Suggestions abound at Foster Wheeler for future applications of expert systems. In addition to other material selection problems, Allen sees a definite need for advice in complicated areas such as corrosion. 'Erosion/corrosion was the only area we had real difficulty with during the development of Matsel. We talked to a lot of people, but the knowledge seemed very woolly.' In its current form, Matsel gives information on possible difficulties and company information sources.

Perhaps a surprising application seen by Allen is easing the use of other design software. 'Some of our software is quite complicated to use, and with a tendency to buy more generic software the problem is going to get worse. What density of mesh might you choose in a finite element model? Do you need a high concentration around a particular nozzle or not? An expert system could help the user choose the method of application.'

Artificial intelligence will undoubtedly be widely used in the next few years, and its capabilities will increase as a result. Matsel is one of the first generation of true engineering applications and shows the potential benefits to engineers. Tedious duties can be performed quickly without loss of accuracy, while guidance can be sought on technically difficult areas. Widespread use of expert systems does have ramifications for the training of engineers and their ability to change established systems. However, if management appreciate the pitfalls as well as the advantages, the new systems will be welcomed by most engineers. □

Report for the US Army research development and standardization group (UK).
Contract no. DAJA 45-85-C-0033
Requisition no. R & D 5133

Subject	Knowledge acquisition for expert systems in Construction
Principal investigator	Professor E G Trimble
Associated investigators	Dr R J Allwood. Mr A E Brymar
	Mr C N Cooper

SIXTH INTERIM REPORT

(Covering the period 1 July 87 to 31 December 87)

18 February 1988

1. Introduction

As indicated in section 8 of the fifth interim report Mr C N Cooper was temporarily assigned to a project funded by the (British) Alvey Directorate early in July 1987. We shall be free to report to CERL on the methods of knowledge acquisition in this work though not on the details of the domain knowledge.

In these circumstances our scope for reporting is limited. The items covered are as follows.

Section 2	Work funded by the Alvey directorate
Section 3	Work funded by SERC
Section 4	MATSEL
Section 5	Machine assisted knowledge acquisition
Section 6	Plans for the completion of the project

2.0 WORK FUNDED BY THE UK ALVEY DIRECTORATE

In July 1987 we started work on a research contract funded by the Alvey Directorate, a funding body for information technology research formed by the British Government The project we are involved with comprises research into the applicability of expert systems to vibration analysis. The project consortium is as follows:

Stewart Hughes Limited - a company specialising in systems for measurement and analysis of vibration.
Rolls Royce plc - manufacturers of aeroengines.
GEC - a major engineering company with wide ranging interests including the design and manufacture of power station alternators.
Loughborough University - Department of Civil Engineering.
Loughborough University - Human Science and Advanced Technology Research Centre (HUSAT).

The objective is to examine two possible expert system applications which relate to the particular interests of Rolls Royce and GEC respectively. Rolls Royce are interested in vibrations in gas turbine engines. Tests are carried out by attaching transducers to the engine which is then accelerated from zero to full speed. The recorded vibrations are examined at discrete speed intervals using a fourier transform analysis so that the results from a given transducer can be displayed on a diagram which is known as a "Z-MOD". The Z-MOD has axes of running speed and frequency, with amplitudes of vibration shown by the use of colour. A typical plot appears as a complex of straight and curved lines. A skilled vibration engineer is able to deduce the phenomena which give rise to these lines - which include shaft vibration, blade vibration, electromagnetic interference, flutter, bearing faults etc. The objective of the project is an expert system which can perform this interpretation task. Rolls Royce are developing Fortran routines which can detect the lines on the Z-MOD and describe these in terms of line of best fit, amplitude, width, cross-sectional shape etc. The expert system will be able to call these Fortran routines, and will then use rules to deduce the phenomenon which has given rise to each line.

An expert system could in the long term have two potential uses. The first would be to carry out preliminary checks on Z-MOD diagrams. Because engine tests are very expensive large numbers of transducers are used in order to maximise data collection Hence a single engine test may result in the production of 300 Z-MOD diagrams. It would be helpful to have an automated system which could categorise these diagrams as either (a) normal or (b) unusual and requiring detailed attention from a human expert.

The second use would be as a vibration workstation which a vibration engineer could interact with as a means of interpreting and understanding Z-MOD diagrams. For example a user would be able to place a cursor over a particular phenomenon shown on a screen display of a Z-MOD diagram, and would then receive an explanation of that phenomenon with details of the reasoning process used by the system.

GEC wish to explore the use of an expert system in the balancing of large alternator rotors. The only source of vibration considered in this project is that due to an out-of-balance shaft (in contrast to the Rolls Royce work in which the all types of vibration phenomena are considered). The objective of the GEC work is a system which will help a user to correct the out-of-balance in a shaft so that vibrations are brought within predefined acceptable levels.

Alternator shafts are balanced by changing the numbers and locations of small balance weights. Many constraints are imposed on the balancing procedure by the construction of the alternator - the shaft is generally only accessible at bearings and balance weights can only be positioned in a limited number of places along the shaft shown as balance planes. Furthermore the number of transducers available for taking measurements on the rotor is generally very limited.

Balancing is to some extent a trial and error process. A number of test runs of the alternator are carried out with a different combination of balance weights used each time. An algorithm is used to process the combined test results and deduce an optimum disposition of balancing weights.

A single test involves running an alternator up to full speed, maintaining full speed for several hours so that "heat soak" can occur, and then reducing speed to zero. As a result the balancing process can take several days, and if the alternator concerned is a 660 MW set this "outage" can represent a major cost to the operator in terms of lost power generation. The ultimate objective of the expert system is to assist a vibration engineer in balancing the alternator with the aim of minimising the number of test runs required before balance can be achieved. The system will be interfaced to Fortran routines developed by GEC for complex numerical work, but is expected to contain rule sets to advise on where transducers should be located, the modes of vibration seen, and the best configurations of balance weights for tests and for final balance etc.

An interesting feature of the GEC work is that the company had previously purchased a computer package for the balancing of rotors at a cost of £100,000. This was a conventional program which instructed an operator on what to do but gave him no insight into how decisions were being made. The balancing engineers reacted strongly against this "black box" solution and refused to use it. A major aim of the current work is to produce a system which gives advice with justification rather than simply imposing actions on the user.

Roles within the project consortium are as follows: Rolls Royce and GEC will provide knowledge in their respective domains and have also developed Fortran algorithms where appropriate. Stewart Hughes Ltd are supplying an expert system shell with interfaces to the Fortran modules. We are eliciting knowledge and coding this into the shell. HUSAT is advising on human-computer interface issues (these are likely to be particularly important for the GEC work in the light of their previous experiences with the "black box" package). We joined the consortium at a late stage after the project had been running for one year and as such are finding the March 1989 completion date to be a tight deadline. The intention was originally that Stewart Hughes would develop their own shell in LISP. However after one year of work this attempt was abandoned and the LISP shell Goldworks was purchased from Gold Hill Computers Inc. We experienced some delay in delivery of the Personal Computer with 8 MB RAM required to run Goldworks and have in the interim developed two demonstration systems using the recently launched Leonardo shell.

In both the Rolls Royce and GEC domains two experts are available. In each case the individual who has been most informative is junior to the second expert. Knowledge acquisition has been carried out by interviewing and so far six sessions have been carried out with Rolls Royce and five with GEC. The work with Rolls Royce is going well and we have built up a good rapport with the experts.

We have found it more difficult to develop a rapport with the GEC experts who are understandably disappointed with lack of progress on their system (caused both by the abandonment of the original LISP shell and by problems with recruitment of personnel prior to our joining the consortium). We are also experiencing problems eliciting the knowledge to deduce the mode of vibration of the rotor from sparse transducer data. The personnel we are dealing with do not undertake this procedure on a regular basis and therefore speak from a knowledge of theory rather than practice. (Direct approaches to the engineers who routinely carry out balancing have been strongly resisted because it is claimed that their experiences with the "black box" software referred to earlier have made them resistant to the concept of a computer aid to balancing). GEC have therefore used a graduate engineer to examine theoretical rotor vibration data and develop rules for deducing vibration modes from sparse transducer data. This has resulted in the knowledge being a moving target in that the knowledge changes markedly, even fundamentally, between knowledge acquisition sessions.

3.0 WORK FUNDED BY SERC

Since the previous interim report, data collection and analysis has been competed and a summary report outlining the results of the research work also completed. A more detailed technical report covering the work is in preparation.

The object of the work was primarily to examine the experiences of initiators, developers, knowledge engineers, experts and users of expert systems, with particular regard to knowledge acquisition. A total of 70 applications had been analysed by the time the survey had been completed. These applications cover a wide range of domains, functions and organisations. Because of a number of factors, not least the limited access made available to the researchers by developers, it did not prove possible to interview all those involved in applications development. Another factor which constrained the collection of data, particularly in relation to users' and potential users' experiences, was the fact that a substantial proportion of the systems studied were not operational. Although most of the systems included in the survey were reputed to be 'real world' systems, around half of them were still under development, had been abandoned, or suspended or were being used purely for demonstration purposes.

Most of the systems that were actually operational were performing routine functions. They were almost entirely selection systems, and all but a few were rule based systems. Most of the operational systems had been developed on PC's using commercial shells. These systems were characterised by the adoption of a knowledge acquisition strategy based on iterative fast prototyping. This strategy usually entailed an initial exploratory phase, involving usually unstructured interviews with experts, followed by the rapid development of a prototype expert system. This system was subsequently used for feedback and further refinement of the knowledge base.

Most of the non-operational systems were developed for complex domains. These domains typically entailed a high degree of procedural knowledge and heuristics or involved domains in which the expertise was widely spread among a number of expert sources or domains which involved underlying 'deep' knowledge models. Many of the systems were directed towards complex functions, such as high level reasoning processes, training routines or the emulation of complex processes like computer assisted design and multi level organisational behaviour. In many cases, these systems were being developed using object-orientated representations, involving development environments like LISP and POP11. These systems were more likely to be developed using an 'evolutionary' method of knowledge acquisition. Apart from a more prolonged knowledge acquisition phase, this type of acquisition strategy chiefly differed from fast prototyping in relation to the way in which 'intermediate' knowledge was represented to the experts in order to generate feedback. Feedback under the evolutionary process was more likely to be undertaken by means of representations like paper models, flow charts and semantic nets, rather than by using a prototype system or shell.

The main conclusions of the research were that knowledge acquisition was a relatively minor problem in the development of routine systems, those which were normally developed using a commercial shell and an interative fast prototyping strategy. The main constraints affecting the development of this type of system were more likely to be organisational factors, for example a lack of commitment from higher level management; difficulties experienced in gaining access to experts and hostility towards the implementation of the application.

Knowledge acquisition was recognised as a problem in complex domains and in systems performing complex functions, but was not found to be, in itself, the major development problem. The problems encountered in eliciting knowledge in, for example, complex engineering design, are inextricably bound up with problems in developing adequate theoretical models to portray human cognitive behaviour and learning processes and in developing suitable formalisms to represent them.

The research clearly pointed out the need for further research in cognitive modelling and its links with knowledge representation. Other research areas highlighted by the work include the identification of appropriate knowledge domains, particularly in construction engineering, and further research in the investigation and evaluation of automated knowledge acquisition programs like Boeing's 'Aquinas' and U.S. West's 'KRIMB'.

5.0 MACHINE-ASSISTED KNOWLEDGE ACQUISITION

During our work on this project we have become aware of developments in the use of computers to assist in the process of knowledge acquisition. Notable amongst these developments in the AQUINAS system developed by Boeing at Seattle. No provision was made in our budget for the project to study these developments and the budget is already over-spent as a result of adverse changes in the dollar/pound exchange rate.

In these circumstances it seems that there are two alternatives.

1. To seek additional, quite modest, funding to study the developments in machine-assisted knowledge acquisition.
2. To deliver our final report later in 1988 referring to the existence of these developments but without any detailed comment.

We should appreciate advice from CERL on how we should proceed in this matter.

6.0 PREPARATION OF THE FINAL REPORT

Our plan is that the final report should be based on all our findings to date plus any further findings from our work for the ALVEY Directorate (see section 2) up to 30 June 1988. We have allocated July and August for the preparation of the report.

4.0 MATSEL

In our last report we described how we had worked with Foster Wheeler Power Products Ltd to develop an expert system for the selection of an appropriate steel for boiler tube fabrication. We have subsequently made contributions to a Foster Wheeler internal report on this work and shall be able to draw on those contributions when we write our final report to CERL

MATSEL was well received by potential users and by those in Foster Wheeler associated with its development. However the system was developed to enable Foster Wheeler to evaluate the IBM Expert System Environment (ESE) over a six month free trial period and MATSEL could only be used subsequently if the company were to purchase the ESE shell. This shell is available solely for use on IBM mainframe computers and currently costs £30,000 (thirty thousand pounds).

Foster Wheeler have not so far purchased ESE and are therefore unable to continue the work with MATSEL. We understand from the Engineering Computing Manager who was associated with the MATSEL project that this decision is the result of very tight company budgets in a slack market for industrial boilers. Should the market improve then the purchase of ESE would certainly be reconsidered.

A feature of the Foster Wheeler situation is that purchase of an expensive shell such as ESE could only be justified if adequate resources were also made available to tran programming staff and to develop a number of additional worthwhile applications. Hence the cost of the shell is only one of several considerations.

The coding of MATSEL was carried out by a Foster Wheeler programmer who had no previous experience of expert systems. On the basis of this experience he decided to specialise in expert systems work and has since taken a job in this field with Thames Water. one of the UK's biggest water utilities. This loss of expertise will handicap further expert system work in Foster Wheeler.

Currently it seems that large numbers of expert systems are under development both in academic institutions and in industry but only a tiny proportion of these get to be used on a routine basis. Our own experience reflects this and we suggest that getting an expert system into everyday use may be the most difficult aspect of development. We intend to comment further on this in our final report.

An article describing the MATSEL work appeared in the magazine Process Engineering in October 1987. A copy is appended to this report.

END
DATE
FILMED
6-1988
DTIC