AD-A192 490 KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS IN CONSTRUCTION US ARMY R AND D G..(U) LOUGHBOROUGH UNIV OF TECHNOLOGY (ENGLAND) DEPT OF CIVIL ENGIN.. 1/1

UNCLASSIFIED E G TRIMBLE ET AL. 06 AUG 87 F/G 12/5 NL

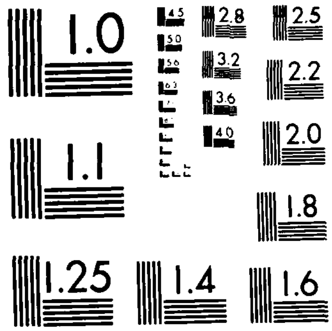

MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

(2)

LOUGHBOROUGH UNIVERSITY OF TECHNOLOGY
DEPARTMENT OF CIVIL ENGINEERING

KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS IN CONSTRUCTION

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

E G Trimble
R J Allwood
A E Bryman

Fifth interim report
Contract DAJA 45-85-C-0033

August 1987

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 Bryman

FIFTH INTERIM REPORT

(Covering the period 1 Jan 87 to 30 June 87)

6 August 1987

1. Introduction

During the period under review our principal effort has been devoted to the development of an expert system model to select materials for very large boilers. This is "the possible new domain" to which reference is made in our last report. This system is called MATSEL and is described below. The work funded by SERC is proceeding and brief reference is also included. The full list of the remaining sections is as follows.

Section 2	CRANES and BID/NO BID
Section 3	PARC
Section 4	MATSEL
Section 5	Evaluation of BREDAMP
Section 6	Work funded by SERC
Section 7	U.S. visit April 1987
Section 8	Work funded by the Alvey directorate
Section 9	SERC Workshop on knowledge acquisition
Section 10	Future work
Appendix 1	PARC critique
Appendix 2	Paper by E G Trimble and C N Cooper presented at the SERC workshop

2. CRANES AND BID/NO BID

The objective of our work on these systems has been to explore knowledge acquisition problems on the basis of our own direct involvement. This objective has largely been achieved. Both systems can be described as practical working tools and it is now up to the two host organizations to exploit them. We are hopeful that BID/NO BID may be put to practical use but are less optimistic about

Overall the evaluation sessions described were very successful and those interviewed were in favour of BID/NO BID. However in order to proceed IDC must purchase and become familiar with a shell program so that they can have their own copies of the system. We are still awaiting their decision on this.

The question of implementation is of crucial importance. We are monitoring other people's experience to see if any factors can be found that tip the balance in favour of actual use. The favourable experience of Stone and Webster inc. is reported elsewhere.

We still have hopes of persuading IDC Ltd to install BID/NO BID and we maintain some pressure in this respect. If we succeed we shall of course monitor their experience of its use.

3. PARC

This is a community club for Project and Resource Control. There are two developers involved namely UNIBIT of Bradford, Yorks and BRAMEUR of Fleet, Hampshire. There have been management problems at UNIBIT which have resulted in the resignation of the managing director. In these circumstances we can expect little from the UNIBIT initiatives except some possible applications of their own software called PARYS which appears to be a large data base with some good search routines. In the formation of PARC some good though over-ambitious objectives were formulated. Professor Trimble's critique of these objectives was discussed with Frank Kearney in April. A copy is attached as appendix 1.

BRAMEUR has formed an informal working party on multi-project scheduling. There are tentative proposals to develop an approach known as Resource Oriented Scheduling; a subject we explored in some depth in the period 1980-83.

It is perhaps worth noting that no staff time has been devoted to PARC. All discussions etc have been conducted personally by Professor Trimble.

4. MATSEL

In our fourth interim report we described discussions we had conducted with Foster Wheeler Power Products Limited with the purpose of identifying an appropriate domain for the development of an expert system. One possibility was a system to diagnose faults during the commissioning of large boiler systems (20 MW to 100 MW capacity) but this domain was rejected because the developments effort required would have far exceeded the available resources and also the precise use of the finished system could not be clearly defined.

Subsequently Foster Wheeler proposed the domain of selection of appropriate steels for use in boiler pressure parts. In this case a number of sub-domains could be readily identified on which initial developments could be centred. Also the task was not as complex as the commissioning domain and an expert who was very keen to develop a system was readily available. This domain was therefore selected define a number of recommended enhancements to the system; these are described in our report to BRE.

Provided that these enhancements are made we believe that the system would find a satisfactory market particularly among local authorities. (In the UK local authorities such as the Nottingham City Council own, manage and maintain a very large number of dwellings that are occupied by subsidised tenants). We believe also that Building Surveyors may wish to buy the system.

We understand informally that we are likely to be commissioned by BRE to undertake the proposed enhancements under a new contract.

6. Work funded by SERC

Mr A E Bryman and Dr J D Cullen are primarily responsible for this work. Professor Trimble is playing a secondary role.

The objective of the work is to examine the experience of others in the development and installation of expert systems. We consider it significant that access to the information has proved difficult to obtain. Various explanations have been postulated and are being explored.

Despite the difficulties we have obtained access to more than 25 systems and have interviewed initiators, knowledge engineers and prospective users. The systems are in various stages of implementation. The facts about the systems are often at variance with the relevant publicly available knowledge. For example some systems that seem to be relevant to construction prove, on examination, to have quite different orientation. More significantly the public claims are often substantially exaggerated.

We expect eventually to be able to report some important findings. We do not anticipate difficulty from SERC in releasing the information.

7. U.S. visit April 1987

Jointly CERL and the National Bureau of Standards agreed to fund a visit to the US by Professor Trimble in April 1987. The purpose of this visit was

- to exchange information on our work with Frank Kearney and his colleagues.
- to attend a meeting of a newly formed committee of RILEM which is to co-ordinate information on expert systems in construction.

The opportunity was taken to visit Stone & Webster inc. who have developed and installed several expert systems.

The visit to CERL provided a most valuable opportunity for exchanging views and for each side to demonstrate recently developed software.

The RILEM meeting was valuable in a similar way. Although a careful record of the meeting was kept by its secretary this had not been circulated at the time of writing this report. The ultimate value of the meeting will of course depend on the follow-up actions.

- Design and procurement of tubes are carried out by different departments. As a result the design engineers only have a general idea of the relative costs of different materials. MATSEL displays the costs of a number of alternatives so that cost implications can be taken into account at the design stage. One senior engineer had taken a design decision which he considered would increase the cost of some tubing by between 25% and 50%. MATSEL showed that the increase was 100%.
- If tubes are going to be bent and attachments are then going to be welded on the choice of minimum allowable tube thickness is complex and mistakes have led to serious problems in the past. Decisions taken by MATSEL are always logically correct.
- Because some aspects of the task were complex some engineers used safe rules of thumb such as "always add 10%". This approach could cost the company money.
- In the proposals department materials selection is always done under pressure. As a result it is common practice to select the most commonly used material for a given part and check that this is acceptable. MATSEL can quickly consider a number of materials and hence indicate when the commonly used material is not the cheapest.

Potential users are enthusiastic about MATSEL because it will enable them to select materials more accurately and quickly and to look at several materials instead of the single material usually considered using hand calculations. IBM have also been very complimentary about MATSEL because they see it as a very good and practical use of their shell product. At the present time Foster Wheeler are seriously considering whether to purchase ESE which at £40,000 to £60,000 is seen as an expensive item of software. We have attended a company seminar at which MATSEL was demonstrated to a committee which advises on computer facilities expenditure. Reactions to the system were favourable.

5. Evaluation of BREDAMP

As previously reported we were commissioned by BRE in October 1986 to evaluate the system called BREDAMP which we had developed under an earlier contract. BREDAMP is a system to diagnose the cause of dampness in buildings. Our evaluation of it commenced with the assumption that it could become a "public domain" system and thus provide a vehicle for disseminating BRE expertise to the general public. This characteristic imposes special constraints on the user interface in the sense that standard definitions of terms and ease of use are essential.

In carrying out our assessment we visited 18 organizations, taking the system with us mounted on a COMPAQ portable computer. We conducted 37 tests using examples provided by the respondents. These tests, and other considerations, helped us to define a number of recommended enhancements to the system; these are described in our report to BRE.

Provided that these enhancements are made we believe that the system would find a satisfactory market particularly among local authorities. (In the UK local authorities such as the Nottingham City Council own, manage and maintain a very large number of dwellings that are occupied by subsidised tenants). We believe also that Building Surveyors may wish to buy the system.

We understand informally that we are likely to be commissioned by BRE to undertake the proposed enhancements under a new contract.

6. Work funded by SERC

Mr A E Bryman and Dr J D Cullen are primarily responsible for this work. Professor Trimble is playing a secondary role.

The objective of the work is to examine the experience of others in the development and installation of expert systems. We consider it significant that access to the information has proved difficult to obtain. Various explanations have been postulated and are being explored.

Despite the difficulties we have obtained access to more than 25 systems and have interviewed initiators, knowledge engineers and prospective users. The systems are in various stages of implementation. The facts about the systems are often at variance with the relevant publicly available knowledge. For example some systems that seem to be relevant to construction prove, on examination, to have quite different orientation. More significantly the public claims are often substantially exaggerated.

We expect eventually to be able to report some important findings. We do not anticipate difficulty from SERC in releasing the information.

7. U.S. visit April 1987

Jointly CERL and the National Bureau of Standards agreed to fund a visit to the US by Professor Trimble in April 1987. The purpose of this visit was

- to exchange information on our work with Frank Kearney and his colleagues.
- to attend a meeting of a newly formed committee of RILEM which is to co-ordinate information on expert systems in construction.

The opportunity was taken to visit Stone & Webster inc. who have developed and installed several expert systems.

The visit to CERL provided a most valuable opportunity for exchanging views and for each side to demonstrate recently developed software.

The RILEM meeting was valuable in a similar way. Although a careful record of the meeting was kept by its secretary this had not been circulated at the time of writing this report. The ultimate value of the meeting will of course depend on the follow-up actions.

In the discussions with Stone & Webster details were obtained of their commercially-based services in expert systems work. Their brochure states that their applications include

- Centrifugal-pump diagnostics
- Rotating-equipment-vibration diagnostics
- Unit commitment
- Technical Specifications tracking
- Weld-defect diagnostics
- Welder-qualification-test selection
- Welding-procedure selection

Details of these systems can be obtained from their offices at 245 Summer Street, Boston Mass 02210. In view of our plans to develop systems for vibration diagnosis we were given a demonstration of one of the Stone & Webster systems developed for this domain. It is understood that it has been installed at about 500 sites. Most of their other systems are also in day to day use.

We were impressed by the extent that Stone & Webster have exploited expert systems in real situations. There was insufficient time to explore the reasons for their success in any depth but one factor emerged namely their policy to concentrate on the solution of defined problems without being concerned about whether expert system technology was being used. Solutions with any proportion of expert system content appear to be equally acceptable.

This problem-orientation may provide lessons for those of us who wish to see expert system solutions exploited more extensively.

8. Work funded by the Alvey directorate

We have been appointed by the Alvey directorate (a British government body) to contribute to the development of 2 (and possibly 3) expert systems relating to vibration analysis. One system will be concerned with the diagnosis of faults in aero-engines; another with the balancing of alternator rotors used in major power stations. For each system raw data will be taken from transducers and analysed in existing electronic equipment. Output from the latter will form an important part of the input to the expert systems that will eventually assist the engineers in diagnosis. Thus the systems will be different from those we have been concerned with so far in that they will have mixed input; partly manual and partly electronic. We believe that this experience will enrich our insight into knowledge acquisition technology and have proposed to Dr J Zavada that we should temporarily re-assign Mr C N Cooper to this work. This will enable relevant work to proceed despite the funding problems that have arisen from the major movement of the $/£ exchange rate since the award of our CERL contract. We shall be free to report on the knowledge acquisition work though not on the details of the domain knowledge. (This parallels the conditions we accepted for our work with Foster Wheeler on MATSEL). Dr Zavada accepted our proposal in his letter dated 13 July 1987.

9. SERC workshop on knowledge acquisition

In June Professor Trimble and Chris Cooper were invited to present a paper at a two day workshop on knowledge acquisition organised by the Science and Engineering Research Council. A copy of the paper presented "Experience of knowledge acquisition for expert systems in construction" is attached. Among the key conclusions in this paper are the following:

- System development should be client led.
- It is unlikely that a single method of knowledge acquisition can be adopted for development of an application - rather a number of methods will be required.
- Using a computer program to induce rules from cases may provide some enlightenment but is unlikely to provide working rules except perhaps for simple systems.
- Our experience has been that a prototype system should be demonstrated to the domain experts as soon as it starts to give plausible and recognisable advice. This dispels misconceptions at an early stage and in general provokes further knowledge elicitation and promotes enthusiasm on the part of the experts.
- We have recently had good experience using a paper model as an intermediate representation for the knowledge elicited. This helps the knowledge engineer cope with the mass of information gathered, and with the right experts can itself be used as a knowledge elicitation tool. We see the model as a two-way communication tool. It helps directly in obtaining clear confirmation of the expert's statements and acts as an unambiguous specification to be followed by the person who codes the system.

A few salient points made by other contributors to the workshop will now be described:

Margaret Welbank is the author of an authoritative literature survey on knowledge acquisition techniques published in 1983. She has just completed a second survey and presented some of the new findings. She commented that in her experience analysis and structuring the knowledge is the most difficult part of building a knowledge-based system. We have also found that knowledge analysis and structuring may be just as difficult as knowledge acquisition. Welbank cited a number of references to the use of paper models (which she refers to as "intermediate representations"). We are following these up in order to compare them with our own work in this area.

Anna Hart presented some findings on rule induction. Her theme was that inductive methods are useful tools in some circumstances but that they must be used with great care. In particular skill and care is needed to select examples for a training set and to select the attributes to be used by the induced rules. The rules produced must be inspected with care and will usually need to be adjusted by hand before incorporation with a knowledge-base. Her conclusions are reproduced below and are very much in line with our own findings on rule induction:

"As experts find it difficult to be objective about their knowledge or to articulate it then there is a role for some tools to help in knowledge elicitation.. However, any general purpose tool, while having good features, will have restrictions and limitations. They are really modelling tools enabling people to explore representations, and suggest models. The results are suggestive rather than proven, and need very careful evaluation and justification. It may well be that the results are the means to an end, and not the end in itself. The methods may raise questions which point the expert in the direction of further ideas or examples. In some cases a traditional statistical method may be more effective. The output is rather 'clinical' and needs augmenting with the explanations and justifications which form such an important part of a KBS. The main restriction with general purpose tools is that they take little account of the context and content of a problem and do presuppose that a general method of inference is appropriate. They must therefore be used with care."

There appeared to be a division between researchers oriented towards experimental behavioural science and those who were engaged in knowledge acquisition for real expert systems. Much of the work of the former group involved experimental assessment of the effectiveness of social science techniques such as protocol analysis (analysis of a commentary given by an expert as he performs a task), laddered grid techniques and card sorting techniques. Practitioners such as Ellman and Welbank commented that keeping a rapport with the expert was of paramount importance and they felt that the use of techniques such as these could well lose this rapport. Therefore they relied exclusively on "safer" interview techniques and demonstrations of prototype systems. Experimental work such as that presented by Burton and Shadbolt which involved elicitation of mineral classification knowledge from geology undergraduates should ultimately provide an objective assessment of the relative effectiveness of interview and other techniques. It should then be possible to assess whether techniques other than interviewing are justified on the basis of faster or more complete knowledge acquisition.

References

"Experience with Knowledge Acquisition for Expert Systems in Construction" E G Trimble and C N Cooper.

"Perspectives on Knowledge Acquisition" M Welbank.

"A Formal Evaluation of Knowledge Elicitation Techniques for Expert Systems: Domain 1" M. Burton and N. Shadbolt.

"Automatic Knowledge Generation: Possibilities and Restrictions" A. Hart.
"An example of Knowledge Elicitation: Reality versus theory" J. Ellman.
All references: SERC Workshop on Knowledge Acquisition for Engineering Applications, Abingdon, England, 1987.

Section 10 Future work

In the immediate future we shall concentrate on the work funded by the Alvey directorate (see section 8). However our experience so far in this project has shown that opportunities for relevant minor studies quite often arise. We shall continue to report on these along with the main study.

PARC critique

The following four sheets provide comment on a document entitled "Functional priorities for the PARC system based on members' opinions". The document itself is restricted and is not included. Thus the comments are in several instances incomplete. However they are included on the suggestion of Frank Kearney and it is hoped that they may provide an aide memoire regarding some of the possible applications of expert systems in the field of project management.

PARC

Thoughts on product intentions

Geoffrey Trimble, December 1986

The following notes are based on the intentions mooted in pages 1 to 3 of the document "Functional priorities for the PARC system based on members opinions. (Steering group meeting no 1 17 October 1986)". They are personal observations offered as an aid to clarifying what is being suggested by the developers.

While most of us have been thinking in terms of an expert system approach this should not prejudice the end product. We should instead address the objectives and be prepared to accept any computer methodology that can achieve them. However the concept of defined goals and the rules that point to the selection of one (or more) from a set of such goals may be helpful in exploring the possibilities. This concept is used as the basis of the comments that now follow on each of the proposed tasks.

1. Project definition and parameterisation

My view is that variables such as project size and complexity will help in assembling the rules that select appropriate goals elsewhere. I know however that Horace Mitchel has discussed "parameterisation" with some potential users who appear to see it as an object in its own right.

2. Project manager selection

I assume that there will be a limited number of candidates (as goals) from whom we are to select a best candidate or a short list. No quarrel with this but we may have to provide for some quite complex rules.

3. Baseplan/Framework selection

When we know what these words mean we shall need to agree on some alternatives from which to select. (More exactly, each user will need to do this and to produce his own rules).

4. Project specification

This implies further variables that will be used to select goals elsewhere (see item 1).

5. Selection of planning methods/tools

If by "methods" is meant scheduling techniques, I have produced a demonstration expert system that can do this. For the real world it would be necessary to extend and refine the rule base quite extensively.

the CRANES system. This is because CRANES is still only able to address a small part of the large quantity of complex domain knowledge elicited.

The technical development of these systems has been largely quiescent during the period under review. A few minor improvements have been made to the graphics facilities in CRANES but apart from this no other development work has taken place. A paper entitled "CRANES - a rule-based assistant with graphics for construction planning engineers" is to be presented by Chris Cooper at the Third International Conference on Civil and Structural Engineering Computing in London in September.

An evaluation of BID/NO BID was carried out in February in cooperation with the host organization, IDC of Stratford-upon-Avon. Six of the company negotiators (who are the target users of the system) were asked to consult the system with regard to current or recent projects. In the course of the evaluation ten consultations were carried out. The results of this evaluation were very encouraging. The advice given by the system and the scores for individual projects were in general similar to user expectations. However several examples occurred where users and BID/NO BID arrived at similar scores for different reasons and there is clearly scope for further refinements of the knowledge base. A number of detailed comments were made about the wording of system questions and this is clearly an area which needs careful attention during system development.

In general the users considered that the system would provide an unbiased assessment of a project and that it would be valuable to have this available to compare with their own more subjective views. Also the system forced them to think about all aspects of a bid opportunity, some of which they would tend to overlook. The young inexperienced negotiators were particularly enthusiastic about BID/NO BID and one was very keen to have a copy on his own computer so that he could have free access to it. The more senior negotiator was more critical but saw a role for the system as a training tool. He was sympathetic to the idea of BID/NO BID and had contributed to it at an early stage. However he made some perceptive comments about its use in practice. In particular the assumption behind the system was that a negotiator would obtain the details of a project and then decide whether to bid or not. In practice project opportunities tended to evolve over a number of months with the project details emerging very slowly. The negotiator would have to interact with the client throughout this period and could not wait until most of the information was available before running BID/NO BID and hence deciding whether to proceed. Furthermore the negotiator should at all times encourage the client to appoint IDC without recourse to competitive bidding. These comments are valuable because they reinforce our findings in other domains that correct identification of the role and users of a system is of extreme importance.

If tools are software tools this implies a wide-ranging evaluation of available software and the development of a complex rule-base to control the selection.

6. Controls and measures planning

Even if we concern ourselves with only one scheduling technique and one tool there are many different ways in which control can be established. Similar comments apply to "resource requirements analysis" which appears in the right hand column.

7. Resource and team specification and selection

There is a lot of expertise in team building and a lot of current research. Are we to be offered a limited range of alternative procedures for these processes and thus risk a trivial solution or shall we sponsor some extensive research to enable real expertise to be embodied in the PARC system?

8. Sensitivity appraisal

The developers have rightly distinguished between "critical" activities and "sensitive" activities. For a computer system to work we must arrive at an acceptable definition of "sensitive" activities. The proposed facility broadens this concept to include sensitive aspects of the project, the team, and the resources. Much more explicit definitions are needed if this interesting concept is to be properly captured and exploited.

9. Estimating/guesstimation

Many man years have been devoted to computer aided estimating in construction. More man-years have gone into debating whether the products offered are safe to use.

Much effort too has gone into parametric estimating based on statistical methods such as regression analysis. What, actually, are we being offered?

10. Tool set up

A facility to enable the PARC system to generate input to existing commercial software is much to be applauded. It contrasts with the "start-from-scratch" approach that has been adopted by the PLANIT/Project planning group. At some stage we need to make a selection from the wide range of software available as, presumably, we cannot expect a "universal coupling" to suit any package.

11. Contracting

If this is restricted to a choice between "lump sum", "target", and "reimbursable" contracts this could be very useful. A realistic extension to select from a wide range of contract clauses would be very demanding.

12. Resource and team orientation

A brief look at the PARYS system suggests that some interesting work has been done by BIT to identify training needs. (BIT has now been re-named UNIBIT Ltd).

13. Tasking and activity management

This sounds intriguing but needs closer definition.

14. Monitoring and trouble shooting

It may be that someone has defined a limited range of ways in which these tasks can be performed. If so, we need only the rules to point us to the right one. (I have no personal knowledge of any such limited range of options).

15. Project "driving" and contingency response

Similar comments to item 14.

16. Completion checks

Each general type of project eg R & D, admin, construction, will have its own completion requirements and these will almost certainly vary within type. For this facility to work it seems that each user will have to develop his own checklists.

17. Resource release

Somewhat similar comments to item 16.

18. Post-mortem analysis. Human resource management, Physical resource management

Somewhat similar comments again.

19. Large scale projects

The explanation needs to be expanded.

20. Multiple projects and resource allocation across projects

There are well-established techniques to deal with the scheduling of multiple projects represented as networks. It is not clear what further work is proposed.

21. Bid/no bid

I have developed an expert system for this domain for one design/construct contractor. Each user is likely to want different rules but the underlying concepts may be very similar.

22. Front end to "conventional" products

This seems to duplicate item 10.

23. Multiple, everyday, projects.

Another intriguing suggestion that needs amplification!

OBSERVATION

Many of the foregoing observations suggest that more knowledge is needed about the various "tasks" before a satisfactory system can be produced. This may reflect a mis-interpretation on my part and I have to accept that I have been looking for a product that contains a fair amount of domain knowledge. If the basic intention is to produce only an empty shell into which each user slots his own knowledge then some of my comments are not relevant. However there remains the need to define the general structure of each domain so that each user will be free to insert his own knowledge in a way that suits his needs. On the face of it this could be very demanding as many of the proposed "tasks" have wide variations of definition.

EXPERIENCE OF KNOWLEDGE ACQUISITION FOR EXPERT SYSTEMS IN CONSTRUCTION

E G Trimble, Professor of Construction Management, Department of Civil Engineering, Loughborough University of Technology.

C N Cooper, Research Assistant, Department of Civil Engineering, Loughborough University of Technology.

INTRODUCTION

The authors are currently engaged in research into methods of knowledge acquisition for expert systems. Although their original focus was on construction industry applications the findings so far suggest that the choice of industry does not of itself affect the knowledge acquisition method selected. Rather it is the situation that determines the appropriate method. By situation we mean such factors as:

- o Available time of the domain expert
- o Whether he holds his expertise in explicit or intuitive form
- o Whether he is motivated to help the process or hinder it.

The authors are associated with two research projects into knowledge acquisition methods. The first involves the development of a number of expert systems in conjunction with industrial clients as summarized below:

CONPLANT - selection of materials handling equipment for multi-storey construction sites.

BREDAMP - diagnosis of the cause of dampness in buildings.

CRANES - selection of tower cranes for multi-storey construction sites (incorporating a graphic interface).

BIDDER - decision of a design and construct contractor on whether to bid for a project.

NETWORK - diagnosis of faults in a national computer network.

MATSEL - material selection for boiler pressure parts.

The first five systems utilise the SAVOIR commercial shell program and run on an IBM PC XT. The last named uses the IBM shell ESE and runs on an IBM mainframe computer. Different client organizations and experts have been involved in each case and this has enabled the authors to gain a broad spectrum of experience of knowledge acquisition.

The second area of research is being undertaken by Mr Alan Bryman and Dr Joe Cullen of the Department of Social Sciences at Loughborough University in association with Professor Trimble. The aim of this work is to look at the experiences of industry in developing in-house expert systems. The work involves questionnaire surveys and interviews with staff from a broad range of companies.

Some of the author's experience of knowledge acquisition will now be described. Much of this material was recently presented at a conference in Washington D.C. but some adjustments have been made to reflect recent findings.

SITUATIONS THAT AFFECT THE ACQUISITION PROCESS

It is clear that the nature of the situation within which the knowledge is acquired will have a major influence on the knowledge acquisition methods to be selected. The categories so far identified are:

1. The knowledge is held in a largely intuitive undefined format.
2. As category I but some closely similar domains have been examined previously.
3. Cases can be defined that reflect a body of decision-making within the domain.
4. There is published material about the domain.
5. The domain expert has sufficient knowledge about expert systems to enable him to define the knowledge (or at least to play a significant role in its definition).

Superimposed on this list of categories are other dimensions such as:

- o The "depth" of knowledge to be represented, i.e. does it represent fundamental knowledge such as that relating to molecular structure or "heuristic" knowledge which includes a substantial amount of personal opinion.
- o The attitude of individual experts to the system.
- o The extent to which a consensus among experts can be found.

The foregoing categories are now elaborated.

Intuitive knowledge

Some knowledge engineers favour a method which requires the development of a prototype system based very often on the prior knowledge of the knowledge engineer. The prototype is demonstrated to the domain expert who suggests modifications and amplification. The changes are made and the revised system demonstrated again. The iterations of this process continue until the domain expert is satisfied that the model is acceptable. If a good initial model is produced this method can be very productive. However it can have the effect of prejudicing the responses of the expert and thus diverting him from some of the subtle, more intuitive knowledge, that might be of crucial importance in the operation of the system.

An alternative is to start with a blank sheet of paper and ask the domain expert to tell you what he knows. A fairly extensive set of knowledge is then assembled before the initial system is coded. This approach is fundamentally better but its success is critically dependent on the time that the domain expert can devote to the process.

Intuitive knowledge with precedents

Where systems have been produced for very similar domains it may be safe to introduce a short-cut in the form of structured interviews based on the content of the previous systems. The danger of prejudicing responses must always be borne in mind.

Defined cases

There are several computer programs that will induce rules from sets of cases. Of these EXPERT-EASE is probably the simplest and best known. At first sight this approach has much to recommend it. However, extensive trials of the early programs have revealed some disconcerting problems. One of these is that the natural sequence of questioning that is inherent in a domain is not respected. For example a pair of questions might be:

- Is the pipe a drain?
- Have you performed a drain-test?

If the sequence of these questions is reversed as it may be in rule-induction the confidence of the user will quickly evaporate. Another problem is that, like regression analysis, rule-induction works on cases irrespective of any causal connections.

It is not, of course, imperative to use the rule-induction program. Manual examination of sets of cases will often indicate relationships that can be coded on an ad hoc basis. The expert will frequently find it easier to recall his knowledge through the recounting of case histories than through other forms of interview.

Published material

There is a lot of interest in the use of expert systems to guide users in the interpretation of regulations and codes of practice. Clearly, in this situation, there should be no problem of human interaction as the views of the human experts should be fully recorded in the published text. As an aside it should be noted that attempts to "computerize" regulations were made before the recent surge of interest in expert systems. These attempts often revealed inconsistency and vagueness which made full "computerization" difficult. Some investigators have suggested that this should be anticipated since differences between the views of the members of the drafting committee eventually have to be resolved by compromise.

Coding by the expert

When the domain expert is also a reasonably competent user of computers it may be possible for him to produce his own expert system without the use of an intermediary. This approach is only possible where the expert exhibits a high degree of self-knowledge and is likely to be unsuccessful where the knowledge is largely intuitive. The expert may or may not be inclined to use an expert system shell to assist him (and constrain him) in his efforts.

SELECTING THE METHOD

The previous section identifies the following methods of knowledge acquisition:

- *Unstructured interviews*
- Structured interviews
- Case histories
- Prototype system evolved iteratively
- Rule induction

To this list should be added:

- Observational

In the observational method the knowledge engineer observes the domain expert as he performs tasks which require him to draw on his expertise.

- Paper models

We have recently focussed our attention on the use of paper models after seeing the successful use of this technique by the ARIES Alvey Community Club. The knowledge engineer creates a document detailing the rules elicited and develops this as knowledge acquisition progresses. The document records the status of each rule - i.e. finalized, tentative, needs review etc and also records the source from which each rule was obtained. A natural language format is adopted so that the paper model can in many cases be reviewed by the domain experts. We have found that this review process can itself stimulate further knowledge acquisition.

o o o

Our studies started with the view that we should be able to isolate situations (or combinations of factors) that would point to the selection of a single method. Our experience has not borne out this view and for example in one of our applications five different methods were used at different stages. Thus our current advice is to acquire a feel for the alternative methods and to use them flexibly as the position unfolds. The following comments augment those in the preceding section:

Unstructured interviewing has the great merit of not prejudicing the responses of the domain expert. Thus less obvious points emerge that can be very important. The method however is time-consuming and requires patience on the part of the expert.

Structured or focussed interviewing achieves results quickly and is appropriate when the knowledge engineer is fairly confident of his understanding of the domain. This understanding may result from prior knowledge or from the results of earlier acquisition methods. Interviews can be broadly focussed or narrowly focussed. During a broadly focussed interview a knowledge engineer might pose questions such as "describe the parts of a typical boiler system" whereas during a narrowly focussed interview "are there any welding problems associated with carbon steel?" would be more typical.

Prototyping has much to recommend it particularly as each interaction can provide cues to prompt the expert in his thoughts about his intuitive knowledge. As with structured interviewing there is a danger that less obvious points may get overlooked.

The authors have very limited experience of the observational method. However it must be a beneficial approach in providing at least the initial evidence that the knowledge engineer will require in structuring the problem. Furthermore observation can be used to check that the expert performs his tasks in the way he claims to use during interview sessions, i.e. during interview sessions he may be telling the knowledge engineer how he thinks the job ought to be done rather than how it is done in practice.

Rule induction appears to be satisfactory for simple well-defined applications. However for applications even of quite modest levels of complexity we have found that rules prepared by induction are unsatisfactory for direct incorporation in the system. Despite these shortcomings we have found that attempts to apply rule induction to limited modules of a total application can force the expert into considering factors that are not revealed by other methods. This must improve the validity of the knowledge base even if the induced rules are themselves discarded.

SOME HUMAN FACTORS

Knowledge acquisition for expert systems is a human process and several of the human aspects have already been mentioned. The purpose of this section is to itemize the human problems that arise so that readers can be aware of them. This is not to suggest that we can yet offer solutions; the process is likely to remain largely ad hoc for some time. Before proceeding it is worth reminding ourselves that the process of knowledge acquisition is the transfer and transformation of expertise from some source (usually human) to a computer program.

Resistance

The domain expert may fear that, by giving up his knowledge, he will weaken his position within his organization. Unless some incentive can be engineered such an expert is unlikely to provide the basis for a useful system. Organizational resistance may also arise and has been observed in the Community Clubs established in Britain by the Alvey Directorate. For example one club member may provide an expert but then realize that commercially valuable skills could be transmitted via the system to a competitor. It should be noted, on the other hand, that positive motivation may be encountered when an expert is bored with providing personal advice in one subject and would welcome the chance to have this process automated.

Accessibility and prejudicing responses

An expert may have the necessary knowledge and motivation but may have other duties that prevent his spending an adequate amount of time with the knowledge engineer. We have already mentioned the dangers of prejudicing responses by over-structuring interviews and by offering detailed prototypes. However, the methods that prejudice responses are usually quicker so some compromise will often be necessary.

Cues and examples

Experts are often better at doing things than explaining what they are doing and why. So one method of obtaining knowledge is to watch the expert at work and then ask why he did what he did. The problem is that the expert often cannot recall from his subconscious the rules and relationships that have become intuitive. A method that also deals with this problem is to generate artificial examples as cues and to ask the expert what he would do in these circumstances. Our experience in obtaining the Bayes factors for BREDAMP is an illustration of this method (see below).

Rapport and roles

Clearly the knowledge acquisition process will proceed more smoothly and effectively when rapport is established between the knowledge engineer and the domain expert. As a corollary to this, it is usually better to separate the tasks of knowledge acquisition from those of coding the information for the computer. This enables the knowledge engineer to concentrate on the knowledge as perceived by the expert and on establishing a good human relationship with him.

Much of our work has been undertaken in construction industry domains and we believe that our engineering background has enabled us to develop a rapport with the domain experts more easily than would have been the case had we been from an unrelated discipline. However a knowledge engineer in this situation must be careful to avoid distorting the expert's knowledge by introducing his own ideas.

SOME PRACTICAL NUANCES

Development environment

So far the great variety of domains that have been attempted and the approaches adopted have made generalized analysis virtually impossible. However, one conclusion is inescapable, namely that there are at least two distinct kinds of situation namely:

- o A client has identified his own need for an expert system and engaged an employee or contractor to deliver a system to the client's requirements.
- o An enthusiast, typically an academic, has defined an interesing application and has persuaded host organizations to provide relevant knowledge.

The nature of the relationship between the client (or host organisation) and the knowledge engineer is of crucial importance. Where a client has defined his own needs it is likely that the experts will be readily available and that they will be uninhibited by company policy and commercial confidentiality. They may still be inhibited by their own motivations. At one extreme they may be eager to impart the knowledge in order that the expert system will eventually relieve them of routine tasks which have degenerated into boring chores. At the other extreme they may fear that revealing their knowledge will undermine their security and generate a situation where they become redundant.

Our experience is that where an application is undertaken by an enthusiast the objectives and in particular the intended uses to which the system will be put are often not clearly defined. The experts may also be less readily available and inhibited by commercial considerations. As a result such systems are much less likely to succeed than those which are client inspired. However this type of situation may in some cases have its compensations when the enthusiast himself has substantial domain knowledge.

A possible approach to development

We have found that the approach to knowledge acquisition must be flexible and be specific to the domain under consideration. Our own general approach can be characterised as the following progression:

(i) one or two unstructured sessions

(ii) case histories and broadly focussed interviews

(iii) narrowly focussed interviews

(iv) prototyping and narrowly focussed interviews

In our work the demonstration of a prototype system to the experts has provided a valuable stimulus to further knowledge elicitation. We have found that the point at which the expert first sees this prototype should be selected with care. On the one hand the knowledge engineer may become prematurely over-enthusiastic about the system and as a result demonstrate a piece of code which is trivial, thus losing the confidence of the expert. On the other hand we have encountered an expert who had serious misconceptions about the capabilities of the system being developed, and such misconceptions must be dispelled at an early stage. In the light of experiences such as these our advice is to demonstrate the prototype as soon as it is capable of giving a piece of recognisable and plausible advice.

Following our most recent work we also recommend that a paper model should be developed as a means of documenting the knowledge acquisition process and as a focus for discussion with the domain experts.

Uncertainty

The BREDAMP system generated some useful insights into the problems of uncertain knowledge. This system was commissioned by the Building Research Establishment; its purpose is to diagnose the cause of dampness in buildings. The domain expert was exceptionally cogent and well motivated. However it was necessary to attach probabilities to the goals e.g. to conclude that there was a 90% likelihood that the cause of the dampness was rising damp. For this the dependencies between variables are calculated using Bayes theorem for which affirmative and negative factors must be established. While we could expect the domain expert to describe the behaviour of dampness phenomena it was impractical to obtain from him estimates of the BAYES factors. To overcome this problem we first obtained from him the key *factors* relating to each cause (or goal). We then compiled tables in the following form and asked the domain expert to suggest values to replace the question marks:

Factor	Suggested values	
Evidence	A stain	A stain
Height of stain	9 inches	15 inches
Age of building	8 years	9 years
Component wetter inside than out	Yes	Don't know
Positive salts test	Don't know	Yes
Probability of rising damp	?	?

To derive BAYES factors from these data it was sufficient to use an ad hoc approach i.e. a combination of simultaneous equations and trial and error. A better approach would have been some form of regression analysis although it can often be difficult to elicit a sufficient number of cases to make this approach possible.

This case illustrates a further point namely confidence limits. At present BREDAMP offers only a set of probabilities for each of the defined causes of dampness. For example:

Rising damp	90%
Rain penetration	27%
Others	less than 5%

The rising damp figure may in fact mean that the probability is in the range 89-91% or it may mean that it is in the range 80 - 100%. A user would react differently if he had these ranges available. With a narrow range he is likely to conclude that he has gone as far as the system will allow and he may then decide to take remedial measures to cure the problem on the assumption that the cause is in fact rising damp. If the wider range (80-100) is shown he will probably undertake additional, quite cheap, tests to improve the reliability of his diagnosis. This extension of the information provided by a system has been mentioned in several contexts, but no actual implementation has so far been identified by the authors.

CONCLUSIONS

The following are offered as reminders of the key points in this paper:

- o System development should be client led.
- o Knowledge acquisition methods depend much more on situation than domain.
- o Flexibility of approach is essential to the knowledge acquisition process. Factors which will determine this approach include:
  - the form in which the knowledge is available
  - the depth of knowledge (i.e. fundamental or heuristic)
  - the degree of consensus among experts
  - the attitudes of individual experts to the system.

- It is unlikely that a single method of knowledge acquisition can be adopted for development of an application - rather a number of methods will be required.
- Cues and examples can help an expert to recall intuitively held knowledge.
- Using a computer program to induce rules from cases may provide some enlightenment but is unlikely to provide working rules except perhaps for simple systems.
- Even with a very responsive expert ascertaining Bayes factors is best done by examples.
- Our experience has been that the prototype system should be demonstrated to the experts as soon as it starts to give plausible and recognisable advice. This dispels misconceptions at an early stage and in general provokes further knowledge elicitation and promotes enthusiasm on the part of the expert.
- We have recently had good experience using a paper model as an intermediate representation for the knowledge elicited. This helps the knowledge engineer cope with the mass of information gathered, and with the right experts can itself be used as a knowledge elicitation tool.

FURTHER READING

Welbank, M. A review of knowledge acquisition techniques for expert systems. British Telecom research laboratories, 1983.

Gotts, N.M. Knowledge acquisition for medical systems - a review. Artificial Intelligence in Medicine Group, University of Sussex, 1984.

Buchanan, B.G. Some approaches to knowledge acquisition. Dept. of Computer Science, Stanford University, 1985.

Allwood, R.J., Stewart, D.J., Trimble, E.G., Shaw, M.R. and Smith, J. BREDAMP - an expert system to diagnose the causes of dampness in buildings. Dept. of Civil Engineering, Loughborough University of Technology and Building Research Establishment. To be published shurtly.

Cooper, C.N. CRANES - a rule-based assistant with graphics for construction planning engineers. To be presented at 3rd Int. Conf. on Civil and Structural Engineering Computing, London, september 1987.

Wijesundera, D.A. and Harris, F.C. The integration of an expert system into the construction planning process. Proc. of 2nd Int. Conf. on Civil and Structural Engineering Computing 1985, Vol 2, pp399-405. (Describes CONPLANT).

END
DATE
FILMED
6-1988
DTIC