Expert Systems Clubs: Design Methods

H. J. ROBERTS
Manchester Business School, Manchester University, Booth Street West, Manchester M15 6PB

For the implementation of successful expert systems, knowledge engineering techniques should be combined with a participative approach to design. This study is based on the work of three Alvey community clubs which were set up to facilitate the transfer of expert systems technology into industry. They involved both industrial and academic organisations. The aim was to produce expert systems in the areas of planning, insurance and data processing, and to increase awareness of artificial intelligence techniques. The clubs were successful in translating AI theory into worthwhile applications, and also helped establish a community of firms interested in applying AI to genuine business problems.

Received December 1988

1. INTRODUCTION

The theme of this paper is technology transfer: the process by which theory is translated into worthwhile industrial applications. Governments in several countries have attempted to facilitate this process. In the UK, this has included setting up collaborative research projects between academia and industry, funded partly by government research councils. This research was co-ordinated by the Alvey Directorate and covered several fields such as Systems Architectures, Software Engineering, CAD, Man-machine Interface, VLSI, and Intelligent Knowledge Based Systems. This paper is based on the work of three Alvey Expert Systems Community Clubs.

The expert systems clubs were part of the Alvey initiative and were set up to stimulate the application of expert systems techniques to industry. The aim was to encourage firms with a common interest to work jointly to apply the new technology to problems in their areas. The goal of each club was to develop one or more demonstrator expert systems in their field of interest. Costs were met jointly by member firms and the government. Thus members could share the experience of building a system for a fraction of the cost of building one independently.

Expert systems clubs included industrial and academic members, as well as software houses who were contracted to build the systems. Industrial members paid to join the club, whilst academic members contributed expertise to the development effort. Delegates from industrial member firms also contributed greatly in terms of time and expertise in organising and running the clubs. The clubs were therefore a truly collaborative effort.

In setting up the clubs, it was hoped that member firms could gain experience in developing expert systems and awareness of the techniques and methods used. They would also become familiar with the available hardware and software used for AI applications. Furthermore, software houses would be encouraged to formulate development methodologies. Through their involvement with the clubs, all members could gain an appreciation of the difficulties associated with implementing expert systems. Finally, it was hoped that members would be in a position to decide whether or not expert systems would be relevant to their own organisations.

The three clubs studied here were Aries, the Insurance club, Dapes, the Data Processing club and Planit, the Planning club. The aim of the present study was to evaluate the work of the clubs, with particular emphasis on the design methods used for building the systems. The results should improve the theory and practice of expert systems design as it relates to industrial applications. The findings of the study are reported in this paper.

2. METHODS

Information in this study was collected from a variety of sources. Interviews were conducted with steering committee members, member firm delegates and with representatives of the contractors. Further information was obtained from club meetings and working group presentations.

Since the main function of the clubs was awareness, the contractors took great trouble to explain their results, especially in the later stages of the projects. Another useful source of information was the clubs' documentation. Here again the standard was probably higher than it would have been in a normal computer project, as a result of the emphasis on education and awareness.

3. RESULTS

3.1 Expert System Design Methods

The design methods used in the three Alvey clubs studied provide valuable guidelines for the introduction of expert systems into industry. They are described below. However, the differences between designing for a club as opposed to a business organisation are important and will be discussed later in the paper

The methods used by the clubs were basically an extension of conventional systems techniques, with an emphasis on prototyping. An iterative life cycle (Fig. 1), as opposed to the more linear cycle of conventional systems (Fig. 2), was essential. Formal methods of project control also proved to be worthwhile.

Design included the following stages:

Choice of problem area
Knowledge acquisition
Intermediate representation (not used for all systems)
Implementation
Testing

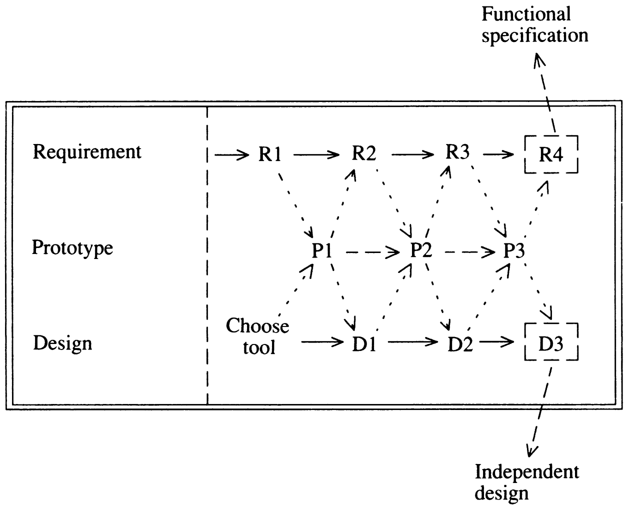

Figure 1. Planit Club prototyping methodology.

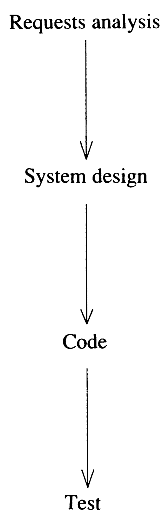

Figure 2. Conventional methodology.

All three clubs produced prototypes before going on to produce their final demonstrator systems. This approach has several advantages. Knowledge engineers may have a limited understanding of the problem domain at the start of the project. Building a prototype allows them to become familiar with the domain and to experiment with different knowledge structures. The expert can be tested to make sure he has the appropriate knowledge and is both willing and able to communicate it to the knowledge engineer.

In business applications, prototypes may be used to demonstrate the strengths of expert systems techniques to senior management as well as to the expert. In this case, more than one prototype may be needed, to demonstrate different aspects of the proposed system. Several club members stressed the importance of gaining management support at an early stage in the project: a prototype can help to ensure such support is forthcoming.

Methods of prototyping varied between the clubs: both rapid prototyping and structured prototyping methods were used. In rapid prototyping the knowledge engineer elicits knowledge from the expert and builds it directly into the system. The expert then tries the system and points out any faults or omissions, which are then corrected (Fig. 3). The process is repeated until the prototype reaches the desired stage. One advantage of this method is the high speed of development. It was found to be a successful technique for the development of small systems.

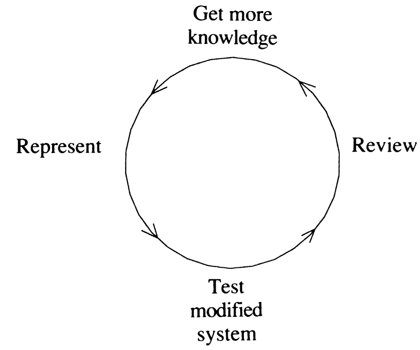

Figure 3. Rapid prototyping.

Structured prototyping, which involved the use of some form of intermediate representation, was the method most commonly preferred in the clubs. With this method, the prototype is built initially in the form of a 'paper model' which can be tested before proceeding to the machine based implementation.

The use of an intermediate representation has several advantages. The most appropriate knowledge structure for the domain can be investigated because the paper model does not impose artificial restrictions: the main limitation is the knowledge engineer's imagination. This contrasts with the use of rapid, machine-based prototyping, where the knowledge has to fit into available structures such as production rules or frames. When an intermediate representation is used, hardware and software decisions can be left until the knowledge structure becomes clear, thus minimising the chance of costly mistakes. Furthermore, the same paper model can be used to implement the system in several different environments, for example to compare the merits of different shells or tools as was done in the Dapes club. An added advantage is that the model serves as a continually updated functional specification, providing documentation during the course of the project.

The experience of the clubs suggests that for larger industrial applications, particularly where several experts are involved, the use of structured prototyping and intermediate representations is likely to be essential. However, rapid prototyping may be the quickest and easiest way to get started on smaller and more straightforward systems. It could also be used to get a quick feel of part of a larger problem domain. Both methods of prototyping were used successfully in the clubs, and both are likely to be used in future applications of expert systems techniques.

Knowledge acquisition is the process of getting domain knowledge out of the expert and into a form suitable for use in the system. The clubs found that different problems required different methods of knowledge acquisition. Structured interviews formed an important part of the process for the Aries club. Their systems were for assessment of fire risk on clothing premises and for equity selection. Knowledge engineers did background reading in the equity selection area in particular, because the expert was busy and could not afford to waste time on basic explanations. Two knowledge engineers with complementary skills carried out each interview, which was also tape recorded. They discussed the interview as soon as possible afterwards in order to make the best possible use of the gathered information. The knowledge was then organised and added to the intermediate representation (Fig. 4).

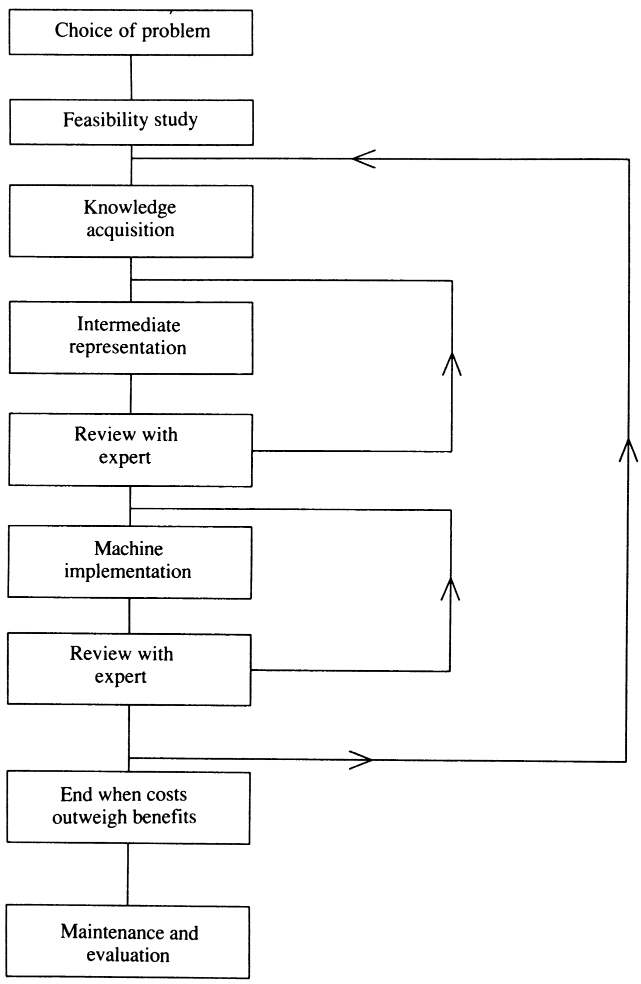

Figure 4. Expert system life cycle.

The Dapes club used rapid prototyping as a method of knowledge acquisition for their early systems in the domain of computer network control. Knowledge was elicited as the expert reacted to the prototype versions of the system, pointing out faults or omissions. This approach was also used for one of their final demonstrator systems, a help desk system for reporting and advising on faults in a computer network.

Dapes also investigated more formal methods, such as the use of intermediate representations, for their other demonstrator, a system for computer network fault diagnosis. The methodology used was influenced by a ESPRIT project (KADS: P1098).[1] The knowledge engineer made a list of all the physical objects in the problem domain, and found as many ways as possible to categorise them. This process helped to give clues as to how the knowledge should be structured. In identifying concepts and problem solving strategies he would get the expert to 'walk through' imaginary cases. He also watched the expert on the job solving real problems. Different information was gathered by each of these methods. The expert tended to be more logical with imaginary cases than when he was doing the job for real. Also, the knowledge engineer could not ask for explanations when the expert was busy working, he could merely write down what the expert had done. The expert was available for knowledge elicitation for about half a day per week, although he could be contacted at other times if needed.

The third club, Planit, developed a new method of knowledge acquisition for domain areas in which expertise is diffused among several people. This was the case with their problem domain, planning. The club had no specialists in this particular area, although many of the members did have planning expertise in addition to their skills in other areas such as engineering. The club therefore set up working groups to find out how club members went about planning in three areas: project planning, process planning and job scheduling. The groups defined the planning process and developed a new and original understanding of it. These findings were extremely valuable in their own right, and also formed the basis for the club's interactive planners' assistant.

Two categories of tool were used predominantly in the implementation of the systems described here.[5] Most commonly used were powerful expert system development toolkits, such as KEE (Intellicorp) or ART (Inference Corporation). These hybrid toolkits offer several forms of knowledge representation and inferencing mechanism, along with sophisticated user interface facilities. They are flexible to use and allow access to the underlying programming language so that desired features can be hand crafted if necessary. KEE and ART were confirmed to offer excellent development environments. However, both Aries and Planit found KEE to be rather slow for operational use. Extra memory helped alleviate this problem but did not solve it completely.

Expert shells were also used for development purposes by the Dapes club: the shells chosen were Expertech's XI and XI-plus, and KES 2 (Software Architecture and Engineering). If the problem domain can be handled adequately by a shell, then development will be quicker and easier than if a large toolkit were used. This is because shells generally provide a limited set of facilities in a highly accessible fashion. They are also cheaper than toolkits and run on readily available hardware such as IBM PCs. Dapes compared ART and KES by implementing the same system, based on an intermediate representation, on both tools. Although the ART version incorporated more advanced features, both environments were capable of supporting useful systems.

Aries and Planit delivered their systems on smaller, PC-based shells rather than on large toolkits. This was because the shells and the hardware to run them were adequate and less expensive than toolkits. For an expert system to gain widespread industrial acceptance, it must run on hardware that is available to industry. Porting systems from toolkits to PC-based shells was done from the intermediate representations rather than by re-coding, and was found to be quicker and easier than the initial development.

Testing is an important aspect of expert system design. Because of the heuristic nature of the knowledge, it is important to test it under as many and as varied conditions as possible. Test cases were used by the clubs during development. These could be run automatically after changes had been made to the system, to make sure that the changes did not cause the system to give incorrect answers. Systems were also tested by allowing the expert to interrogate them, comparing the resulted generated by the system with those expected. Exhaustive testing, as would be required for fully operational systems, was not carried out on these demonstrators.

3.2 Implementing Expert Systems

Alvey clubs gave firms the chance to experience the development and implementation of one or more expert systems in their field of interest. They were therefore able to gain an awareness of the problems involved, and how they might be overcome. The following section covers some of these issues.

3.2.1 Choice of Problem Area

At the time the Alvey clubs were formed, expert systems were a solution in search of a problem. This situation is the reverse of what would generally be expected. However, some of the methods used in selecting suitable applications would be equally useful in determining whether or not a given problem could be solved using expert systems techniques.

Guidelines on the choice of suitable problems were provided by the contracted software houses, and were similar to those in general use.[2] These guidelines were successful in allowing the choice of several problems which formed the basis for the systems built by the clubs.

Additional criteria were also used which were specific to systems built in precompetitive research environments like the Alvey clubs. The problem should be interesting to club members, it should not be commercially sensitive, it should be in an area familiar to them and should be of medium complexity. Outside the context of an awareness club, perhaps the most important criteria for choosing a problem is that its solution should meet a genuine business need. However for the clubs, the solution of genuine business problems was a bonus.

One result of the clubs' emphasis on awareness was that because the systems were not primarily intended to solve genuine business problems, the environment in which they were developed was artificial. Potential users had little, if any, involvement in the decision to build an expert system. The importance of user participation in the design of computer systems, and expert systems in particular, has been stressed.[4] Without such participation, maximum benefit from the new systems may not be obtained.

The clubs were unable to take a participative approach to the design of their systems as they never intended the systems to be fully operational within the timescale. Thus there were no real users to involve in the design process. Systems were, however, tested out in development sites by the people for whom they would have been intended. Some met with an encouraging reception and were later developed further by individual club members to customise them for operational use. For others the reaction from users was mixed. Although the systems appeared to do everything they were intended to do, doubts were expressed as to whether they would ever be used for real. One contractor was disappointed that the host site rejected the system developed for them. This rejection occurred despite so much time having been spent on the system, and the fact that it was a marked improvement on the pre-existing system, in the contractors' opinion.

These experiences highlight an obvious though sometimes ignored aspect of system design: the chances of producing a successful system are greatly reduced if potential users do not want it. For this reason it is important to determine their requirements right at the start of the project. Had this been done in the case of the systems described here, they may have been built differently, or not at all. However, this was not possible: it was never the intention of the clubs to produce operational systems. Rather, the success of these systems was seen in their contribution to expert systems awareness in the UK.

3.2.2 Finding an Expert

In choosing a host site for their development work, the clubs took into account the availability of suitable experts. There were few problems in this respect. Time was a problem for at least one club; the most useful expert is often also the busiest. There were also problems with commercial sensitivity which were directly related to the nature of the clubs and would be less likely to occur in a normal business environment.

Although single experts were the norm, Planit used working groups to take the place of the expert. The use of several experts is to be recommended since it avoids the risk of producing an idiosyncratic system and reduces dependence on individuals. Despite overheads such as time required with more experts, Planit found their knowledge engineering process to be valuable and enlightening.

3.2.3 Finding Knowledge Engineers

Finding suitably qualified knowledge engineers for business applications is difficult, although this was not a problem for the clubs. Until recently, a large number were involved in the various Alvey projects. The clubs gave firms the opportunity to experience expert system development without the costs of hiring or training specialized personnel. Club members become familiar with the work of their contractors. They also learned enough about expert systems to decide whether they would want to hire knowledge engineers from software houses or develop their own 'in-house' capabilities.

The attributes needed in a knowledge engineer were described by Alvey contractors. An ability to learn and understand new subjects, along with good communication and social skills, is important. It also helps if the knowledge engineer has a reasonable degree of maturity and practical experience: new graduates are not suitable in this respect. Flexibility of outlook and creativity are required in structuring the knowledge. Systems analysts who have become entrenched in a particular design methodology may therefore be unsuitable as knowledge engineers, although in other respects the qualities required are similar. A knowledge engineer should also be familiar with computers, although programmers are often unsuitable because of their procedural style of thought.

Knowledge engineering is acknowledged to be a bottleneck in the process of implementing expert systems. It is both time consuming and expensive. It is therefore wise to have a knowledge engineer who will recognise when expert systems techniques are not appropriate. The costs of large systems are prohibitive unless the problem is important to the business and cannot be solved by conventional methods.

3.2.4 Validation, Update and Maintenance

This is the least developed area of expert systems research: many systems simply have not reached this stage. The systems produced by the clubs also fell into this category although club members were interested in the question and expressed a desire to know more.

The clubs found that several different tests can be carried out in validating the system. Results from the system under development can be compared to those predicted by the intermediate representation to ensure correct coding. Results from the system can be compared to those of an expert without the system, to check that the reasoning and results generated by the system are correct. Another test is to compare results from the system operated by a typical user with those when it is operated by someone familiar with it, such as the implementor. This shows up problems in the user interface and in expression of questions and explanations. Other comparisons are also possible, such as user with system against expert with or without the system, or expert with system against expert without the system.

In the clubs more emphasis was placed on methods than on the results of testing. Test examples were provided by the experts and run after the system had been changed. For example, two thirds of the test cases might be classic examples, and the other third unusual in some way. Testing was a time consuming process. One club had automatic testing modules built into their system. However, they found that often it took as much effort to rebuild the automatic test after a change as it did to rebuild the system. For this reason, automatic testing was abandoned.

Update and maintenance are fundamental problems that could not be fully covered by the Alvey clubs. One important question concerns responsibility for updating the system. Should it be done by the expert, the users, the designers or someone else? One suggestion was that the system should be treated in the same way as a database: users could be given permission to update certain facts in the knowledge base, but should not be allowed to make changes which would have far reaching effects within the system.

4. DISCUSSION

This study confirms that the three Alvey clubs have been successful in their aim of deploying expert systems in their fields of interest. Member firms have gained a practical awareness of expert systems which would not otherwise have been possible without their own expert system development projects. Contractors have also benefited from the opportunity to carry out projects in this new area. It can be concluded that the clubs have stimulated a greater awareness of expert systems design in the UK, as was intended.

However, many challenging problems in the area of expert systems remain to be solved. There are important differences between systems built as awareness exercises and those built as operational applications.[6] Moreover, these differences appear to be largely irreconcilable within the context of precompetitive research. As a result, the methods of development used in the Alvey clubs are limited in their approach to certain aspects of design. Areas which need extension are discussed below.

The lack of user participation in the design of these systems has already been mentioned. Adverse consequences which could have been a result of this included the rejection by host sites of systems which had been designed for them. It is possible that if there had been greater user involvement in the design of the systems, as opposed to the largely administrative interactions which sometimes took place, the systems might by now be operational. A model for the successful use of this approach is Digital Equipment Corporation's design of the XSEL system.

Organisational aspects of design were not included in the scope of these projects. However, if expert systems are to become more widespread in industry then these aspects should not be ignored.[7] Introduction of such systems will mean changes in the jobs of the people who are to use them, which could lead to more widespread changes within the organisation. The implications of such changes must be considered and managed. This requires involvement from personnel skilled in the management of change. It should not be simply left in the hands of the computer department who are probably more interested in technical than organisational matters.

5. CONCLUSIONS

The aim of the Alvey clubs was to stimulate the application of expert systems technology in British industry. They have certainly been successful in spreading awareness among their immediate members. Firms have started to customise the systems developed by the clubs to suit their own needs. Others have begun to design their own systems: one or two even have systems in regular use. However, it is likely that further work will be needed before expert systems become a mature technology. In particular, the work of the clubs could be complemented by additional work on user participation and on organisational aspects of design.

Technical difficulties are seldom a cause of failure in expert systems projects for business applications. The club experience confirms this view: they suffered no insoluble technical problems. Some systems were, however, rejected by users, a problem that could also occur in single organisations if attention were not given to user requirements at an early stage in the project. Lack of management support can also be a cause of failure, though the clubs did not suffer from this problem.

The successful design of expert systems for business applications requires at least equal attention to be given to organisational as to technical aspects of design. This can be achieved by combining knowledge engineering techniques, such as those used by the Alvey clubs, with a participative design approach. This approach allows future users to be responsible for the development and testing of their system. User needs can be analysed systematically by a user design group consisting of users, experts and knowledge engineers. The ETHICS methodology used by Digital Equipment Corporation for the development of XSEL has proved to be suitable for this purpose in the design of both expert and conventional computer systems.[5]

Acknowledgements

This research was funded by an Alvey grant from the Science and Engineering Research Council. I am grateful to Enid Mumford for her valuable suggestions, and to Caroline Ball and Steve Hickey for help in preparing the manuscript. The assistance of the Aries, Dapes and Planit clubs and their contractors was essential to this project, and is gratefully acknowledged.

REFERENCES

1. A. Anjewierden, Knowledge acquisition tools. AI Communications, 0(1), pp. 29–38 (1987).
2. P. Harmon and D. King, Expert Systems. Pub. John Wiley & Sons, New York (1985).
3. E. Mumford, Designing Human Systems. Pub. Manchester Business School, UK (1985).
4. E. Mumford, Participative design for expert systems. Keynote address to Fourth International Expert Systems Conference, London 7th June (1988).
5. E. Mumford and B. Macdonald, Xsel's Progress: The Continuing Journey of an Expert System. Pub. John Wiley & Sons, New York (1989).
6. H. J. Roberts, Use of tools and shells in three Alvey expert systems clubs. Proceedings of Conference on the Use of Expert Systems in Oil and Gas, London, UK. Pub. IBC Technical Services Ltd., London, UK, 27–42 (1988).
7. H. J. Roberts, Design methods in three Alvey Expert Systems Clubs. AISB Quarterly, No. 65, 18 (1988).
8. H. J. Roberts (1988 c). Expert systems and the personnel department. Personnel Management, 52–55, June (1988).

Announcements

15–18 July 1991

Multi-Media Information: The Second International Information Research Conference

Churchill College, Cambridge

Sponsored by the British Library Research and Development Department and the University of Pittsburgh School of Library and Information Science

The Theme

Many of the IT systems currently in use owe much of their character to earlier research into the handling of purely textual and numeric data. With the advent of technologies such as CD-I (compact disc interactive), DVI (digital video interactive) and those allowing the digitisation of large archives of data, there are now opportunities to meet the pressing need to merge many types of data: text, numeric data, graphics, sound, and still and moving images.

Research in the early 1980s resulted in the 'hypertext' concept, which affords the user the ability to construct and read text in non-linear fashion. By extending the hypertext principle to other modes, 'hypermedia' systems have been developed. The earliest commercial applications stemmed from the television industry, with the increasing sophistication of interactive video, and from the education/training industry with computer-based packages. The variety of products now available on optical disc indicates that publishers are also becoming aware of the potential of electronic publication using multimedia.

Adding temporal data types to a database system raises a number of fundamental technical problems. The traditional mechanisms for storage, retrieval, delivery and presentation of these types are not easily integrated with current computer technology. Significant progress must be made in all of these areas before sound and motion pictures can be processed with the same ease and simplicity as we now handle text.

The Audience

Research into multi-media is of interest to anyone concerned with the creation, storage, retrieval and dissemination of information in science and technology, arts and humanities, education and training, entertainment, creative leisure, the work environment, and the home environment.

This conference is aimed at promoting serious discussion between researchers, providers and users of information including those working in computer, communication and library sciences. The conference will consist of formal papers and syndicates with ample time for discussion. There will also be an exhibition with the facility for poster sessions.

To ensure effective interaction between delegates, the number of participants will be restricted to 100. The conference proceedings will be published as soon as possible after the event.

Are You Interested?

Would you like to give a presentation on:

- the theory or concept of multi-media information
- organising information for multi-media use
- the impact of the technologies on the research process
- future applications in commerce, industry, science and education
- policy implications for decision makers and researchers?

If so, please send a detailed abstract (100–200 words) upon which your 20-minute presentation would be based, by **16 January 1991**. All authors will be notified of acceptance or rejection by **late February 1991**, and those whose papers are accepted will be given a discount of £100 on the delegate fee.

Would you like to take part in the exhibition or present a poster session? If so, please register your interest as soon as possible.

Would you like to attend the conference as a delegate?

The Venue

The conference will be held at Churchill College in the university city of Cambridge, easily accessible by road, rail and air. The College, built in the 1960s as a memorial to Sir Winston Churchill, was the first major work of modern architecture in Cambridge. It has excellent conference facilities with accommodation in single study-bedrooms and is set in grounds of some 42 acres.

Enquiries

For further information on any aspect of the Multi-Media Information Conference, please contact:

Ms Karen Merry, Multi-Media Information Conference, the British Library R & D Department, 2 Sheraton Street, London W1V 4BH, U.K. Tel: 071 323 7050; fax: 071 323 7251.