Expert Systems in Business

BUTLER COX
FOUNDATION

Research Report 60, October 1987

BUTLER COX FOUNDATION

Expert Systems in Business

Research Report 60, October 1987

Butler Cox & Partners Limited

LONDON
AMSTERDAM MUNICH NEW YORK PARIS

Published by Butler Cox & Partners Limited
Butler Cox House
12 Bloomsbury Square
London WC1A 2LL
England

Availability of reports
Members of the Butler Cox Foundation receive three copies of each report upon publication; additional copies and copies of earlier reports may be purchased by members from Butler Cox.

Photoset and printed in Great Britain by Flexiprint Ltd., Lancing, Sussex.

BUTLER COX FOUNDATION

Expert Systems in Business

Research Report 60, October 1987

Contents

BUTLER COX FOUNDATION

Management Summary

A Management Summary of this report has been published separately and distributed to all Foundation members. Additional copies of the Management Summary are available from Butler Cox

Chapter 1

Expert systems: an opportunity for business now

The purpose of this report is to alert senior business managers and information systems managers to the threats and opportunities presented by expert systems. We also prescribe the actions that you should take now if you are to gain full advantage of this new development in information systems technology.

Our earlier report on expert systems — Number 37, published in 1983 — correctly predicted that expert systems would *not* revolutionise data processing for the next five years. We found then that there were only a handful of commercial projects in ordinary business environments; that few areas of expertise provided demonstrable applications; that the benefits were hard to identify; that the development timescale was not only long but somewhat uncertain; and that available software products were primitive and costly. We also found that most of the applications in use could be understood and used only by other specialists in the given field and not by general office or administrative staff.

Hence, in 1983 we recommended that for the next two or three years most organisations should restrict their activities to small experimental expert systems, whose value would be mainly educational. Only pioneering companies should undertake sharply focused, expensive, and high-risk expert system applications.

Today, though, the situation has changed. It is now clear that expert systems are not just another new technology. Rather, they are a major step forward in applications software, and, as such, create opportunities for the aware and threats to others. Furthermore, the advances made in the last four years mean that expert systems have moved beyond the experimental stage and are now in use in many business sectors. In the near future, there will be a limited window of opportunity for many organisations to exploit the competitive advantage that expert systems can provide.

The conclusions and recommendations set out in this report are based on an extensive programme of research carried out during the first half of 1987, and led by Charles Chang, a principal consultant with Butler Cox in London. He was assisted by David Flint, a principal consultant with Butler Cox in London, and Simon Forge, a senior consultant based in Butler Cox's Paris office. The research base and findings are detailed in Appendix 1 but in summary the research programme comprised:

- Interviews with 58 organisations, most of which were currently using (or were intending to use) expert systems.
- Interviews with 24 suppliers of expert system products, including leading US suppliers.
- Focus group discussions involving 16 organisations (users and suppliers).
- Desk research to evaluate 50 case histories of the use of expert systems and to review the findings of other expert system researchers.
- An analysis of the responses received from 104 Foundation members to the short questionnaire that accompanied the document distributed at the beginning of the research project.
- A questionnaire survey sent to 80 suppliers of expert system products in Europe (of whom 20 replied).
- Discussions with industry experts in the United States and Europe.
- An analysis of the responses to two surveys carried out by Japanese organisations (JIPDEC and ICOT).

The research was international, with data being gathered from Belgium, France, Italy, the Netherlands, the United Kingdom, the United States, West Germany, Japan, and elsewhere. The bibliography at the end of the report lists the publications and documents that considerably influenced our thinking, and that we believe would provide Foundation members with more detailed information about specific aspects of expert systems. For ease of reference, the bibliography is arranged by topic.

CONFUSION OVER THE STATUS AND SIGNIFICANCE OF EXPERT SYSTEMS

Some senior managers — both in systems departments and in business at large — are sceptical about expert systems. They think that the claims for expert systems are nothing more than marketing 'hype', and that they will be of no practical use for the next several years. Other senior managers — often in the same organisations — seem to hold the naive belief that expert systems can solve most of their computer application (and business) problems.

During the past few years, there has been a tremendous growth in information about expert systems. During the course of this study we were inundated with brochures for conferences, research reports, learned journals and other magazines, and for products and consultancy services — all specialising in expert systems. This overabundance of information, however, has led to confusion, rather than understanding, about the importance and status of expert systems. Furthermore, much of the information is provided by organisations with a vested interest; many of the 'consultants' involved have hardware or software to sell. There has been an excess of marketing hype, conflicting opinions on the potential market, and new jargon to learn. Terms like 'fuzzy logic', 'search space', 'combinatorial explosion', 'inheritance classes', and so forth, are not part of the everyday language of commercial systems staff. (We provide a glossary of the most commonly used artificial intelligence and expert system terms at the end of the report.)

Although several market researchers — including Dataquest, Gartner, Frost & Sullivan, Input, and Ovum — have come up with different estimates, they all forecast a rapidly expanding market for expert system products. Estimates of 1986 spending on expert systems in Europe range between $80 million and $160 million, and between $400 million and $1,000 million in the United States — excluding investment by national and regional governments. The general prognosis is that spending on expert system products by European organisations will grow by between five and ten times by 1992.

But there have also been less optimistic statements about the growth of the market. For instance, in December 1986 the highly regarded *New Scientist* magazine reported under the headline "Expert systems: the bubble bursts" that researchers and commercial companies were facing an uncertain future. It quoted an American supplier as believing that the industry is about to enter an "AI winter". Then in May 1987, Brian Oakley, the head of the UK government-sponsored Alvey Research Programme, whose focus is on AI and expert systems research, said that he had to reduce his predictions for the expert systems market in 1990. And in the middle of 1987, one respected market researcher was quoted as saying "slower take-up predicted for AI".

One of the main difficulties is that the subject of expert systems sounds complicated and appears to be very technical and esoteric, and indeed it is. Owing to its intrinsic complexity, many systems staff are almost as unclear about the topic as are the potential users and senior managers. With all the contradictory reporting and often dogmatic views of what expert systems are and are not, and what they can and cannot do, it is not surprising that the key question Foundation members expect this study to answer is:

> *"Are expert systems now ready for serious exploitation and, if so, how can they be fully exploited?"*

From the research we have carried out, there is no doubt that the answer to the first part of that question is 'Yes'. Expert systems can have a significant impact on the organisation, and the technology is now ready for exploitation.

THE BUSINESS SIGNIFICANCE OF EXPERT SYSTEMS

Expert systems represent a discontinuity in the application of computer technology to business problems, akin to other crucial developments such as the introduction of online systems and disc files in the 1960s, or the emergence of personal computers and end-user computing in the early 1980s. They permit new kinds of applications, with widespread usage, to be computerised. They apply computing techniques and power to *knowledge*, the next step beyond the application of computers to data and text processing, thereby enabling computer systems to solve problems that previously could be tackled only by people. For this reason, the term 'knowledge-based systems' is often preferred to 'expert systems', particularly in the United Kingdom. One expert on expert systems (Alex d'Agapeyeff) uses the term 'know-how systems' to describe the majority of so-called expert systems.

'Expert systems' cover a wide range of computer applications that are based on the knowledge or know-how of an expert, a specialist, or a technician. In the case of a doctor (an expert), the expert system might help with medical diagnosis. In the case of a bank-loan officer (a specialist), the expert system might provide advice about whether to grant a loan. And in the case of a PABX engineer (a technician), the expert system might assist in diagnosing faults.

Our definition of an expert system, repeated from Foundation Report 37, is as follows:

> "An expert system is a computer system containing organised knowledge, both factual and heuristic, that concerns some specific area of human expertise; and that is able to produce inferences for the user."

Expert systems are quite different from previous software tools. Figure 1.1 shows the structure of an expert system. The two main components are the knowledge base and the inference engine. The knowledge base is a computer file that contains the facts about a specific domain of human knowledge and the rules for using the facts to solve problems. The inference engine is a program that takes input from the user and infers or interprets the facts in the knowledge base, using the rules stored there. An expert system also contains a user interface through which the knowledge engineer enters new facts and rules into the system, and through which the user accesses and uses the knowledge base. Usually, the user interface contains an explanation facility that allows the expert system to explain its line of reasoning. Sometimes, the user interface is used by applications development staff to develop new applications, which often require interfaces to other systems.

Figure 1.2 overleaf compares conventional applications software with expert systems software. In conventional software, the computer program both controls the application and contains the applications logic. Today's expert systems decouple the applications logic from the control, by storing the facts (data) and logic (rules) in a knowledge base. In the future, there will be a further decoupling, with the knowledge base, data dictionary, and database all being separate entities. Just as conventional software developers hold data definitions separately from data in the form of a

Figure 1.1 Expert systems structure

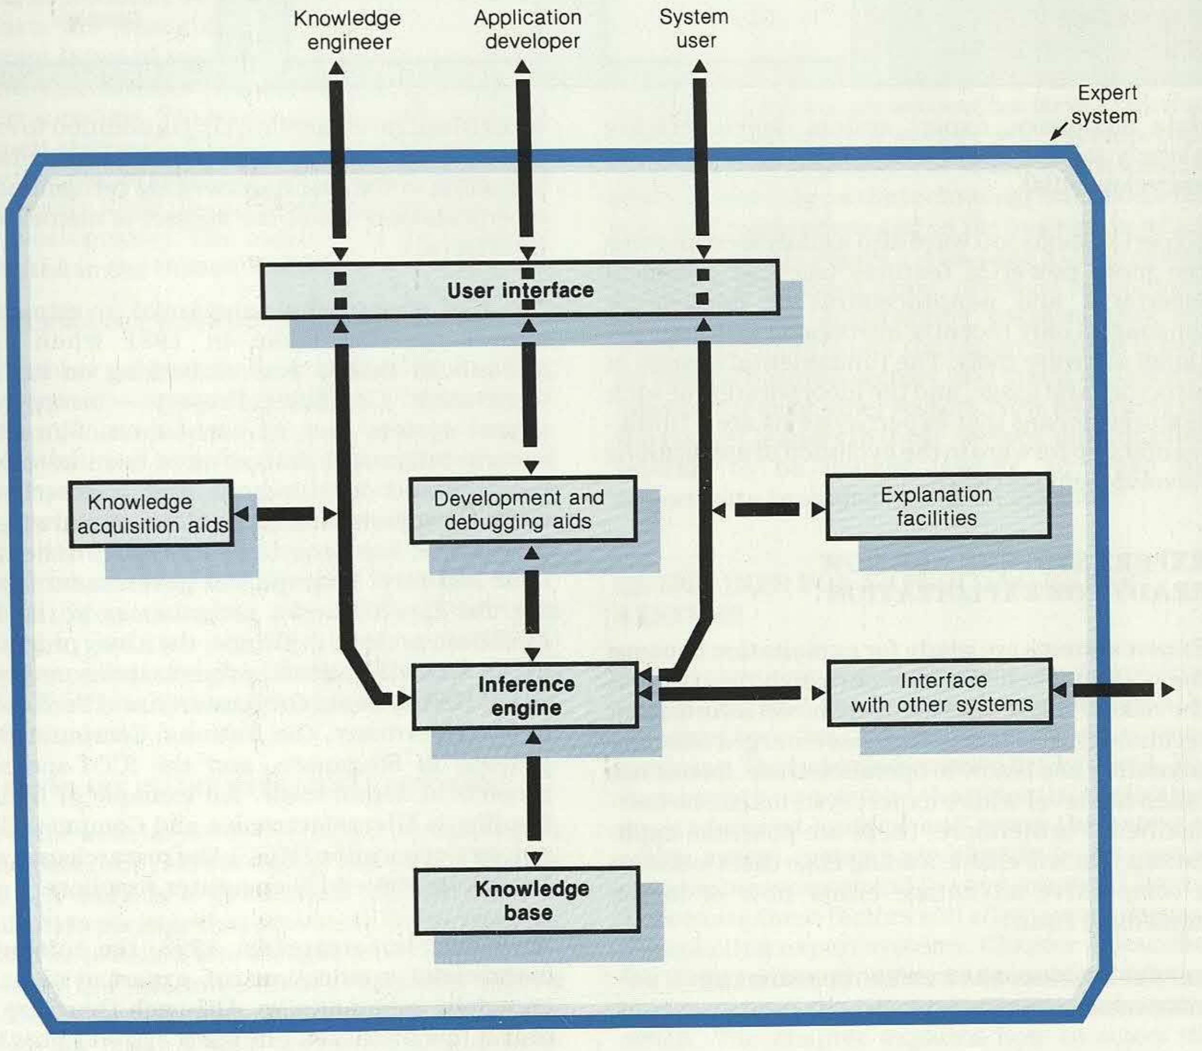

Figure 1.2 Comparison between conventional software and expert systems software

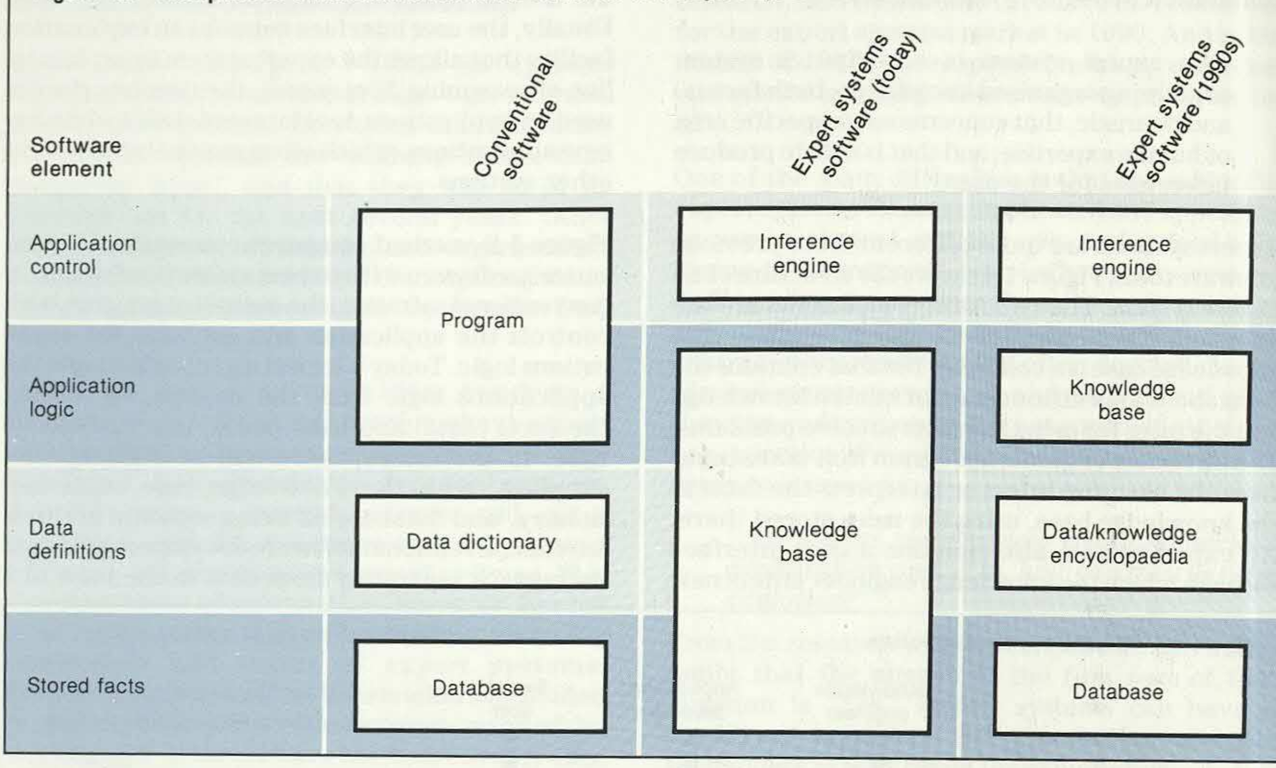

data dictionary, expert system developers are already talking about knowledge dictionaries or encyclopaedias.

Expert systems software also usually incorporates the more powerful features (such as graphical interfaces and nonprocedural or declarative languages) only recently introduced into conventional software tools. The fundamental change in structure and usage, and the incorporation of such features, means that expert systems are a fundamental step forward in the evolution of applications development software.

EXPERT SYSTEMS ARE NOW READY FOR EXPLOITATION

Expert systems are ready for exploitation because the potential benefits now outweigh the cost and the risks of using this relatively novel information technology. The technology has emerged from the laboratory and is now in operational use. Its cost has fallen to a level where expert systems can be cost-justified. Furthermore, there are potential applications that will enable leading-edge users to secure a competitive advantage either now or in the immediate future.

EXPERT SYSTEMS HAVE EMERGED FROM THE RESEARCH LABORATORIES

Expert systems have existed in research laboratories since the late 1960s, alongside other aspects of artificial intelligence (AI). (In addition to expert systems, AI encompasses pattern recognition — including voice and handwriting recognition and interpretation — and the subject of man-machine interfaces.)

The real impetus for substantial investments in expert systems came in 1981 when Japan announced that it was embarking on its 'Fifth Generation' Computer Project — incorporating expert system and AI techniques. Since then, several billions of dollars have been invested in research and development, and in experimental and live projects, in Europe alone. Worldwide, the investment has been both from government and from industry. Examples of government funding are the Esprit/Eureka programmes in the EEC, INRIA-led projects in France, the Alvey programme in the United Kingdom, projects sponsored by the BMFT (Department for Research and Technology) in West Germany, the National Computer Board projects in Singapore, and the ICOT-sponsored projects in Japan itself. An example of industry funding is Microelectronics and Computer Technology Corporation (MCC), the research operation funded by eleven US computer suppliers.

Since our last report in 1983, the number of commercial applications of expert systems has increased substantially. Although there are still only a few examples of expert systems providing substantial business benefits, there are hundreds, if not thousands, of practical applications that are beginning to provide some benefits. Figure 1.3 (on pages 6 and 7) provides a representative sample of the more than 200 practical applications we examined or heard about during our research. These applications cover practically all industry sectors and all the countries where there are Foundation members. Of the 104 members that responded to our initial questionnaire, 13 per cent have operational expert systems, 62 per cent have systems at the pilot, experimental, or development stage, and 15 per cent are considering developing expert systems. Only 10 per cent have done nothing about expert systems.

Interestingly, many of the practical applications are now used by clerks or administrators — not by experts. And there are more examples of systems based on the knowledge of 'specialists' and 'technicians', rather than of 'experts'. This, however, does not detract from, but rather enhances, the utility of expert systems to ordinary businesses.

The application of expert systems is based on a large variety of products (we estimate that more than 400 products are available in Europe) from many different types of supplier — hardware vendors, systems and software houses, and specialist expert system suppliers. There are at least 100 suppliers of expert systems in Western Europe and well over twice that number in the United States (Figure A1.5 in Appendix 1 lists the suppliers that responded to our questionnaire). The majority of the products originated in the United States.

THE TECHNOLOGY WORKS IN PRACTICE

The number of practical examples of working expert systems is growing rapidly. During our research we interviewed dozens of business organisations that now have operational applications of expert systems or are moving from experiments to live operations. The system suppliers that responded to our questionnaire survey claimed that their products were being used in 4,400 applications. Even excluding the 3,500 applications claimed by just two suppliers, there are still 900 discrete applications. And the applications are not trivial. Nearly a third contain at least 5,000 rules (or the equivalent).

THE COSTS ARE IN LINE WITH LIKELY BENEFITS

One of the most important technical developments in the past four years is that of the expert system shell. (A shell is a generalised expert systems application package that provides a framework for a designer to build a knowledge base. The shell also provides the inference engine, together with a predetermined control strategy for using the knowledge base. A shell therefore provides the means for developing an expert system without expert systems expertise.) These products are more powerful and usable than the first-generation expert system environments and the low-level expert system languages (Lisp and Prolog) that were available in 1983.

Also, because it is a new technology, improved versions of shell products tend to be released at least once a year, and sometimes twice. Some expert system shells cost only a few hundred or thousand dollars, plus the cost of a microcomputer or a share of a minicomputer or mainframe; the previous enviroments used to cost tens of thousands of dollars and required the use of specialised hardware, costing similar amounts. (Figure A1.8 in Appendix 1 analyses the pricing information provided by respondents to our supplier survey). There is evidence that microcomputer-based software tools are selling extremely well, whereas software products based on specialised hardware are not selling as well as originally expected either by their suppliers or by market researchers.

THE WINDOW OF OPPORTUNITY IS HERE

In the evolution of most new technologies, there is a 'window of opportunity' during which leading-edge users may use the technology to gain a competitive advantage without too large a risk that the technology will not work or will cost more than the benefits to be gained. The timing of this window depends not only on the technology but also on the nature of applications and on the business in which it is applied (see Figure 1.4 on page 8). The combination of affordable price, practical technology, and successful applications suggests that the window of opportunity for expert systems is here for several business sectors, and that it will soon be reached for most other sectors. Now is the time for the market leaders in those sectors, and for those aspiring to be market leaders, to grasp the opportunity provided by expert systems.

GUIDELINES FOR EXPLOITING EXPERT SYSTEMS

So how should an organisation set about exploiting expert systems? The first step is to understand the potential benefits that the use of expert systems can bring. This topic is discussed in Chapter 2. You then need to understand the potential application areas of expert systems and assess the extent to which expert systems are already being used in your business sector and in your country (Chapter 3), because these factors will affect your approach to exploiting expert systems. Chapter 4 describes the three main approaches and provides advice on how to choose the best aproach for your organisation. The chapter explains how to select the possible opportunities for exploiting expert systems in each of the main application areas and how

Figure 1.3 Representative sample of expert system applications

Business sector and country	Organisation	Application	Comments
Discrete manufacturing			
France	Bull	Kool — Toolkit designed for own use on several projects.	May lead to commercial product.
	Renault	Sitere — maintenance of automatic transmission.	Sales force required after-sales backup.
Germany	Krupp Atlas Electronik	Kritic — Knowledge Representation & Inference Techniques in Industrial Control — control of a power distribution network.	May lead to a commercial product.
Italy	Pirelli	Gestra — managing transport.	Has developed own methodology for expert systems development.
Switzerland	Baumann AG	Der Rollandenexperte — develops parts lists and specifications for roller-blind manufacture.	
UK	Rolls-Royce	Developing market and product strategy — selecting potential customers and forecasting market penetration.	Two-thirds of systems areas potentially suitable for knowledge-based systems.
UK/US	Digital Equipment	XCON — Configuring Vax computers.	Pioneering success.
US	Cambell's Soups	Repair of soup-cooking plant.	Captures expertise of near-to-retirement repairman.
Banking and finance			
Australia	Australia and New Zealand Banking Group	Configuring and managing bank branch.	
Belgium	Générale Banque	Le Courtier — advisory system for stock portfolio management.	Developed to meet increased demand for investment advice.
Italy	Fideuram	Evaluation of investment revenue for clients.	
UK	TSB Trustcard	Diagnosing and handling communication network problems.	
US	American Express	Authoriser's Assistant — checks credit card use against fraud.	Expanded to other Amex sites.
	Evensky & Brown	Plan Power — advise and manage unique financial situations and long-term goals.	
Public utilities			
France	Eléctricité de France	Diagnosis and control for safety systems for nuclear plants.	Could prevent unnecessary shutdowns without any risk.
Netherlands	Rijkswaterstaat	IVON — forecasts conditions of highways for two-year investment cycle.	
UK	British Nuclear Fuels	Consultative systems for design engineers.	
	British Gas Head Office	Assessing the insurance coverage of computer centres.	New knowledge developed by experts from different disciplines working together.
Transport			
France	RATP	Maintenance, predicting breakdown patterns, and management decisions.	Concerned about immaturity of development tools.
Scandinavia	SAS	Redirecting aircraft when Copenhagen Airport is fogbound.	Developed with Stockholm University.

Figure 1.3 Representative sample of expert system applications (continued)

Business sector and country	Organisation	Application	Comments
Transport			
Singapore	Port of Singapore Authority	Plan loading and unloading of container ships.	A pioneering application.
Process manufacturing			
Netherlands/ UK	Royal Dutch Shell	Ressix — assist in running oil reservoir simulation programs.	Operational at one site, remote access from several others.
UK	Blue Circle Industries	Linkmann — control system for cement-manufacturing plant.	Estimate six months to payback.
	BP Ltd	Gasoil — design aid for gas and oil separators	Has 2,500 rules using Extran induction shell; outputs design drawings.
Insurance			
UK	Legal & General	Aries Fire — risk assessment.	Continuing development with Aries Club.
UK	American International & five others	The Underwriting Advisor — preliminary analysis to decide if further action required.	Designed for 50 workstations on mainframe with simultaneous access.
Mining			
UK	British Coal	UFEL — planning action for methane threats.	Not sure of need for knowledge engineer. Sees need to use mainstream computers.
Professional services			
UK	Touche Ross	Appraising capital investment per project.	Used as tool to increase productivity and effectiveness of auditor.
US	Coopers & Lybrand	Expertax — advising on tax.	Successful commercial application.
Public administration			
UK	DHSS	Performance Indicator Analyst — assists researchers to measure performance of health authorities.	Derived new knowledge or expertise.
Retail and distribution			
UK	W H Smith	Stock Control — forecasting demand on warehouse.	Reveal package uses live data and calls Cobol programs.
Communi-cations			
UK	British Telecom	BT Tracker — diagnosing PABX faults.	A pioneering application.
Construction			
UK	Wimpey Group Services	Checking quality of detailing and workmanship done by contractors.	Believe knowledge engineers will eventually be bypassed.
Education/ research			
UK	South Wales Polytechnic	House price prediction.	Pilot 85% accurate; therefore extending project.
Health care			
UK	North Warwickshire District Health Authority	Answering queries on Aids virus from doctors, the public, and researchers.	

Figure 1.4 Window of opportunity for expert systems

Organisations wishing to use expert systems for competitive advantage should exploit the window of opportunity. The single curve shown in the diagram is an over simplification because different business sectors have reached different stages in accepting expert systems.

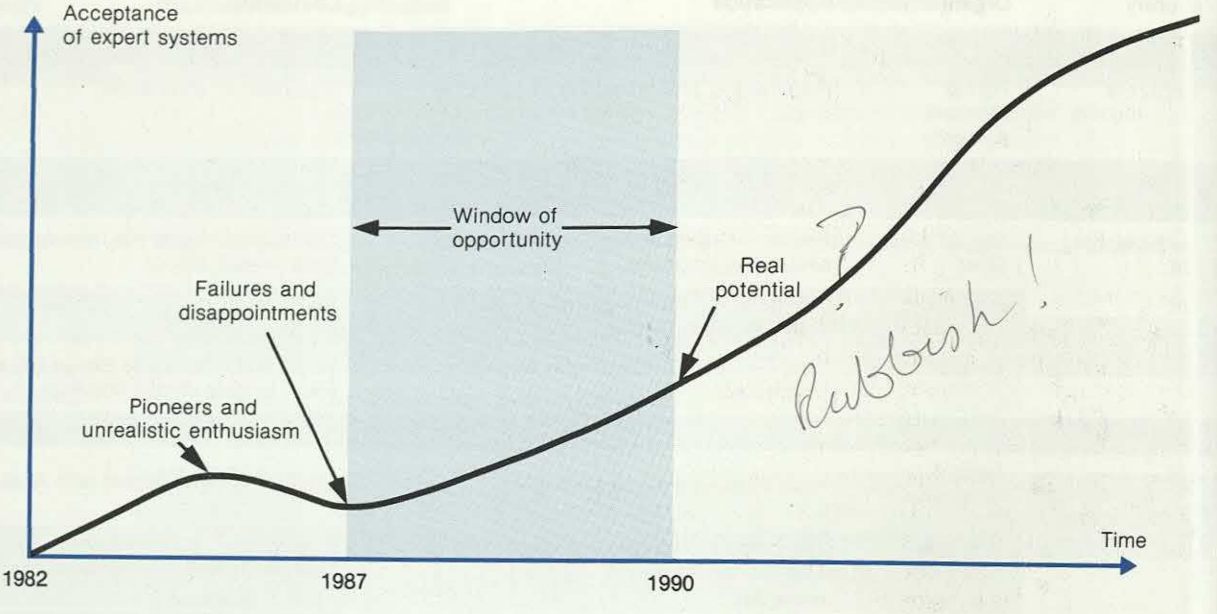

to prioritise the surprisingly large number of opportunities that will be identified.

The introduction of expert systems then needs to be managed. The key is to set up an expert systems support unit and locate it in the most relevant part of the organisation. Chapter 5 describes the role of the unit, and the various organisational models that can be used. The chapter also identifies the skills required for introducing expert systems and provides advice on how to manage the implementation of the first application.

Finally, exploiting expert systems depends on choosing appropriate products and tools. Chapter 6 explains how to do this.

Chapter 2
The benefits of using expert systems

Although expert systems may seem mysterious to the uninitiated, the benefits to be achieved are real. Some expert system applications could probably have been implemented with conventional development tools. However, the emergence of expert system techniques has allowed many applications to be computerised for the first time. Other applications have been re-implemented as an expert system, and this has made them easier to use and more cost-effective. The benefits arising from the computerisation of these applications are diverse. As one might expect, the main benefits of using expert systems stem from replicating or enhancing human expertise and from increasing the productivity of experts, and from solving problems too complex to be solved by traditional computing techniques. Some organisations are also benefiting from expert systems in an indirect way. They have begun to use expert system products as a tool for developing conventional applications because of the better facilities the tools provide for development staff. Finally, many organisations are finding that their use of expert systems is helping them to gain a competitive advantage.

We illustrate each kind of benefit by describing relevant applications chosen from the many already in existence. As you will see, many of the examples have multiple benefits and could be used to illustrate the achievement of other kinds of benefits, but this is not unique to expert systems — it is in the nature of most computer applications.

REPLICATING OR ENHANCING EXPERTISE

Where human expertise is in short supply (thereby creating a bottleneck), expert systems can be used either to replicate the expertise or to allow less expert staff to undertake the work that previously required an expert's personal attention. The ability to enhance the capabilities of non-experts was cited as a major benefit of expert systems by nearly 50 per cent of the respondents in the JIPDEC survey. Specific examples of expert systems being used to replicate or enhance expertise are detailed below.

COMPLIANCE WITH GOOD PROFESSIONAL PRACTICE

From October 1988, financial-services institutions in the United Kingdom will be legally required to ensure that their employees are complying with good professional practice. The purpose of the legislation is to ensure, for example, that a client is advised about the best products on offer, and not the ones that provide the highest profit for the financial-services company. This means that the institutions will need to monitor all the business transactions made by their employees. Such monitoring is almost impossible to carry out by manual methods alone, and even with automated indexing and searching tools, the task is difficult. However, some financial institutions (such as the National Westminster Investment Bank) are developing an automatic monitoring system using an expert system shell called Vanilla Flavor. The details of the transactions executed by each salesman or trader are matched against a set of rules, and any breaches are brought to the attention of the compliance officer, who can then examine the records in more detail. The Vanilla Flavor software and the accompanying hardware costs more than $500,000. However, any serious breach of the law, if proven to be due to negligence on the part of the company, could result in substantial fines and possible revocation of the licence to operate.

A secondary benefit is that it will be easier to incorporate any changes in the compliance rules than it would be with a conventional system.

INTERPRETATION OF TAX LAWS

Coopers & Lybrand, the international audit and accountancy practice, found that its US tax experts were spending the majority of their time analysing and advising on simple and straightforward tax problems. The firm decided to implement an expert system that would enable less qualified personnel to handle the simpler tax problems, freeing the experts to concentrate on the difficult problems where their specialist knowledge is of most value. The result was Expertax — the expert tax advisor system — which has cost some $1 million to develop.

Expertax is now being used in most of Coopers' offices in the United States by general accounting and auditing staff. Everyone gains from this arrangement: the tax expert can apply his or her expertise to the difficult and high-value problems; the non-expert can advise on the simpler tax problems — which form the majority of the cases anyway — without having to refer them to the tax expert, thereby saving both time and effort; the client benefits in reduced time and cost; and Coopers gains an important competitive advantage.

There is also another important benefit. Because of the complexity of the tax laws, different branch offices could easily interpret the laws differently and give different and potentially incorrect advice. The expert system enables the interpretation to be consistent throughout the organisation. Furthermore, in order to develop Expertax, Coopers had to develop a specialist expert system shell (Q-Shell) to handle questionnaires easily and simply. The firm may offer Q-Shell as a product to its clients.

AUTHORISATION OF PERSONAL LOANS IN RETAIL BANKS

Retail banks in several European countries have pilot expert system projects that allow decisions on personal loans to be devolved to individual branch offices. The local staff are assisted by an expert-system-based application in three ways: first, the system assists them to evaluate the risk; second, it assists them with some of the paperwork involved; and third, it refers to head-office specialists those applications that fall outside the authority and conditions that the local office is empowered to deal with. Thus, the same system helps the user, enhances the capabilities of non-experts, and ensures that the current head-office policies and rules are interpreted correctly. Moreover, in today's competitive climate, the ability to provide an immediate answer to a loan request provides the banks with a distinct business advantage. In addition, the system is readily accepted by the local staff because it is seen to be helpful rather than controlling, or even interfering with, their work.

AUTOMATIC TRANSMISSION FAULT DIAGNOSIS BY RENAULT

Renault has built a prototype expert system (called Sitere) in conjunction with Cap Sogeti Innovation to assist in diagnosing faults in automatic transmissions. Such a system is necessary because there are few experts who can correctly diagnose the faults. The aim of the system is to make their expertise more widely available. The dealers benefit by being able to diagnose faults quickly and accurately; customers benefit because the faults are repaired correctly; and Renault benefits because its customers are more satisfied.

PROVIDING SCARCE OR RARELY USED EXPERTISE

Expert systems have been used in situations where expertise is in danger of being lost. Two examples are:

- To capture the knowledge of the expert who administered the pension payments for a scheme that applied to former employees of a company that had been bought out many years ago.
- To capture the expertise of the operator and maintenance staff of a piece of equipment that is no longer being sold or supported by the original manufacturer.

Expert systems have also been used to determine the foreign-exchange remittance rules that apply when a developing country imports certain kinds of capital goods that require special authorisation or licences, and where the particular country or the kinds of goods concerned are not regular features in the company's normal line of business. Hence, the company does not need to employ a person with the relevant expertise, which would be rarely used, nor does it need to employ an expensive consultant when it does require such expertise. However, such an application needs to be used with care. Import/export and foreign-currency regulations change, and the knowledge base must be kept up to date.

PROVIDING BETTER MANAGEMENT CONTROL

Expert systems can be used to ensure the consistent application of head-office policies where the responsibility and authority for implementing the policies is devolved. For example, several organisations (from the public and private sectors) are using expert systems as part of their recruitment programme. Staff can be recruited locally, but can be selected according to central directives and standards. The expert system acts as a filter to eliminate those applicants who are obviously unsuitable. In effect, it acts as a recruitment agency whose task is to identify a shortlist of good prospective candidates, rather than to make the decision on behalf of the hiring organisation.

INCREASING THE PRODUCTIVITY OF EXPERTS

In addition to allowing work previously done by experts to be carried out by lower-level, and less costly, staff, expert systems may also be used to improve the capability of the expert. This was the benefit cited most often (by more than 60 per cent of the respondents) in the JIPDEC survey. One example is the use of expert systems by American Express to authorise credit.

CREDIT AUTHORISATION AT AMERICAN EXPRESS

American Express uses an expert system called Authoriser's Assistant to support its policy of not placing a limit on the amount of credit its cardholders may have at any given time — as long as they clear their account in full at the next billing cycle. The no-limit policy differentiates American Express from the more popular Visa and MasterCard credit cards, but it often causes difficulties for the administrative staff. Three hundred credit authorisers at four locations often need to access up to 13 different databases in order to make the correct decision. Previously, many of the requests for credit were granted by using a mainframe system that used statistical criteria. Even so, the volume of transactions requiring human intervention made the process relatively slow and a potential bottleneck.

After demonstrating the technical feasibility of applying expert systems to the credit-authorisation process (using a shell), a major project was initiated in mid-1985. By April 1986, a standalone prototype had been developed with the aid of Inference Corporation, using its software, ART, which runs on Symbolics specialist hardware. Since then, the Symbolics system has been linked with the mainframe via interfaces with a Sun microcomputer. Not surprisingly, the most difficult part of the entire exercise was not the process of building the knowledge base — difficult as that was — but rather the technical problems of interconnecting incompatible hardware and software. (Figure 2.1 schematically describes the system.)

The system went live in January 1987 and in tests has performed 11 per cent better than the credit authorisers, providing the correct decision in 96 per cent of cases rather than 85 per cent. It is well known that a large part of the potential profits of credit-card and charge-card firms is lost through bad debts. Controlling them or, even better, reducing them is one way for a credit-card company to gain a competitive advantage over its rivals.

DIRECT FINANCIAL PAYOFF

Sometimes, the increased productivity resulting from the use of an expert system can provide a direct financial payoff. One example is provided by a company taxation-optimisation program developed by a major multinational systems vendor for its own use. This application enables the company to pay the minimum amount of tax required by the laws of the countries in which it operates. Because the tax authorities do not yet have an equivalent system, they are hard pressed to prove

Figure 2.1 American Express's Authoriser's Assistant

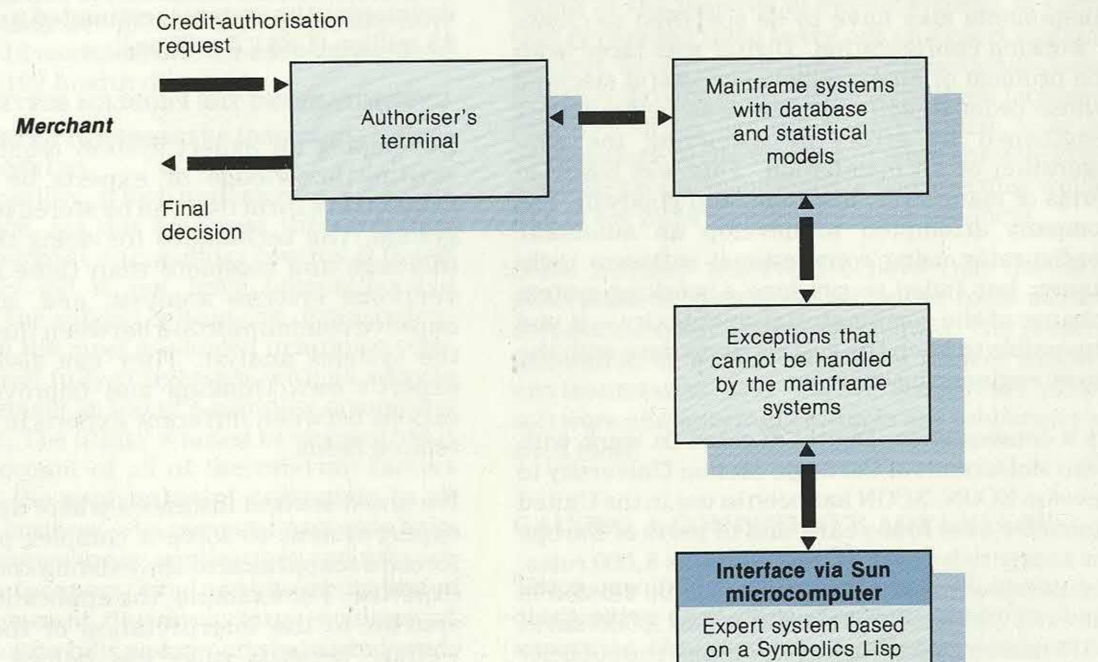

that more tax is due than is declared by the company. However, we hear that one tax authority is busy developing its own expert system to check the validity of these declarations.

SOLVING AND UNDERSTANDING COMPLEX PROBLEMS

Certain types of complex problem have not been capable of solution by conventional computing techniques. Often, these problems require the choice of a solution from an extremely large number of possible answers (the so-called combinatorial problems). In such problems, there may be a few simple separate choices that can be made. The complexity comes from the vast number of ways in which the simple choices may be combined and sequenced. Examples are the configuration of complex computer systems and the optimisation of a port's operations. Sometimes, using expert systems to solve a complex problem provides another benefit. The very process of trying to build an expert system brings about a better understanding of the nature of the problem.

CONFIGURING COMPUTER SYSTEMS

Digital Equipment has used an expert system (XCON) to help it solve the problem of configuring Vax installations. The Vax architecture now spans the whole range from microcomputers to mainframes — all with a common operating system and compatible interfaces, such as Ethernet. In the early days, Vax computers were sold mainly through dealers and system integrators. As a result, any combination from about 60,000 discrete components may have to be specified to create a working configuration. Digital was faced with the problem of an extremely successful machine whose order-to-delivery time was being unduly lengthened by errors in specifying the configuration of an installation. This was costly in terms of manpower, lost time, and goodwill. The company attempted to develop an automatic configurator using conventional software techniques, but failed to produce a working system because of the combinatorial complexity — it was impossible to keep the system up to date with the latest engineering changes.

As a consequence, Digital decided to work with John McDermott of Carnegie-Mellon University to develop XCON. XCON has been in use in the United States for over four years, and in parts of Europe for nearly two years. It now contains 8,000 rules, compared with only 400 in the original version in the early 1980s. Digital estimates that XCON saves it $15 million a year. However, despite the publicity given to how easy it is to maintain and change the rules in an expert system, XCON costs Digital in the region of $1.5 million a year. This is a lot of money, but is only one-tenth of the annual savings generated by the system. Only 20 per cent of the annual maintenance effort is concerned with the expert system itself, however; the rest concerns all the activities around the system — such as linking with mainstream systems and data.

XCON has formed the basis for a family of expert systems: XSEL, to help the salesman to configure the order in the first place, and XSITE, which is used to advise on computer facilities and site configuration.

OPTIMISING PORT OPERATIONS

The Singapore Port Authority is developing an expert system to help it decide on the optimum sequence for loading and unloading container ships. At present, the port is able to move 500 containers in eight hours. It believes the use of the expert system will increase this to 640, providing faster turnaround for ships using the port — benefiting the shippers, consumers, and the authority.

The system is being developed in conjunction with the ITI (Information Technology Institute) of the National Computer Board of Singapore. The concept of the system has been proved by a prototype system running on a Lisp machine. The full operational system will run on networked Sun microcomputers linked to the authority's IBM 3081. There will be two phases of development. By the end of 1987, computer-aided manual scheduling will be possible. Fully automatic scheduling will be available two years later. The total cost of developing the system is estimated to be Singapore $3 million (US$1.5 million).

UNDERSTANDING THE PROBLEM BETTER

Developing an expert system requires that the existing knowledge of experts be pooled and codified in a form that can be stored in a computer system. The techniques for doing this are more thorough and complete than those used in conventional systems analysis, and, usually, they improve communication between the experts and the systems analyst. They can also clarify the expert's own thinking and improve communications between different experts in the same or related fields.

We noted several instances where developing an expert system to solve a complex problem had forced a reappraisal of the existing knowledge and expertise. For example, the application of expert systems to the interpretation of the UK social welfare benefits rules has helped to identify anomalies in the statutes and in the detailed interpretation of the rules. Two other examples of the use of expert systems leading to a better understanding of the problem are the identification of the indicators required to measure the performance of a public service, and the identification of the factors required to assess the insurance risk at a computer centre.

Indicators required to evaluate a public service
Under the Conservative government, public administration in the United Kingdom is having to prove that it is delivering value for money, and not just managing on a cost- or budget-control basis. A few years ago, the National Health Service was asked to identify output-performance measurement indicators. Each of the major care groups — doctors, nurses, pharmacologists, maternity, casualty, and so on — identified its own set of indicators; the total was 450. On a pilot basis, some health-service units collected data in order to measure these indicators. The operational research division of the Department of Health and Social Security was then asked to analyse the data being collected. Very quickly, the division realised that it was far from clear how the indicators related to each other and how they could be used to measure the actual outputs achieved and relate them to the expenditure and inputs of the health service.

The operational research division has now loaded some of the data into an expert system shell and has used this as the basis for a dialogue with practitioners in the field. The result is that the practitioners are beginning to identify the key indicators amongst the 450, and to see how they relate to each other and how they reflect the achievement of actual outputs. In the process, the experts are beginning to gain a much better understanding of the business they are in. The indicators are now grouped into 12 modules, one for each care group, covering the 192 health districts.

Factors required to assess the insurance risk at a computer centre
A large public utility, which underwrites most of its own insurance, set out to assess the insurance coverage required at its computer centres. It found that not only did it not have comprehensive expertise on the subject within the organisation, but that even the most advanced insurance companies were not fully conversant with all aspects of risk assessment in today's complex computing environment. The utility wanted to ensure that it had taken account of all of the relevant factors arising from the application of computers to all aspects of its business — its computers provide links with its trading partners; applications and data are used not only for operational activities but also for major management decisions; large volumes of output (customer bills and reports) are produced. Thus, the business is dependent on the machines operating throughout the working day, and sometimes at night or during the weekend as well. When using an expert system to help evaluate the risks, the utility found that it was necessary to bring together specialists from half a dozen different disciplines, and to pool their expertise. The new knowledge that resulted was used to evaluate the adequacy of its computer centre insurance coverage. The utility is now considering whether to make the expertise available commercially.

USING EXPERT SYSTEMS AS DEVELOPMENT TOOLS

Expert systems are designed to support rapid prototyping and easier maintenance, and they provide advanced graphical interfaces (using the now common WIMPs — windows, icons, mouse, and pictures), often in conjunction with an explicit nonprocedural or declarative form of language akin to the more sophisticated versions of fourth-generation languages. One of the key features of expert systems is their ability to allow rapid prototyping without the need for a formal specification. Hence, some users have been encouraged to use expert system tools to develop other conventional systems that did not need an expert systems approach. Although we would not recommend this approach for developing major systems, it can be extremely useful in helping users to formulate their initial requirements and in enabling the systems analyst to understand the requirements better.

Several of the organisations we interviewed told us that it was often the availability of advanced interfaces, rather than the hitherto unproven claims of the benefits of expert systems as such, that persuaded them to try out an expert systems product. In retrospect, these organisations believe that they have benefited from using these facilities, and they often use the expert system tools to demonstrate to the vendors of more conventional software products the kinds of features they would like to see incorporated in future products.

Some industry observers believe that the next generation of system development tools will incorporate expert system techniques. The present generation of IPSEs (integrated project support environments) and CASE (computer-aided software engineering) products are precursors of such tools.

GAINING A COMPETITIVE ADVANTAGE

There are many different ways of gaining a competitive advantage. For some organisations, improving the production process will lead to a competitive advantage, and there are many examples of expert systems being used to improve the productivity of process-control plants in the chemicals, petroleum, other mineral-extraction, cementation, and food-processing sectors.

For others, improved customer service will be the key to gaining a competitive advantage. More than a third of the respondents to the JIPDEC survey said that improved customer service is an important benefit of expert system applications. One example is the Le Courtier system used at the Générale Banque in Belgium to advise customers about stock portfolios. Customers can key in the data and questions themselves through a natural-language interface. Hence, Le Courtier has improved customer service and also improved the productivity of the bank's experts.

Expert system applications can also be used to provide an improved maintenance service. There are several fault-diagnostic applications that speed up the maintenance process by helping to pinpoint the fault faster and more accurately than the average repairman does.

Digital's XCON, American Express's Authoriser's Assistant, and Coopers & Lybrand's Expertax are all examples of companies that have gained, or are gaining, a competitive advantage by exploiting expert systems. Another example is Northwest Airlines' seat allocation system, developed in collaboration with Sperry (now Unisys), where the numbers of seats allocated to different classes and ticketing options are adjusted dynamically so that overall revenue is optimised.

Most of these examples are from the United States. We believe that the reason for this is that many European and other non-American organisations are more reluctant to exploit IT aggressively for competitive advantage. Also, American organisations are seldom reluctant to proclaim their success stories, even if they relate to applications that provide them with a competitive advantage. Nevertheless, we believe that, despite the high-profile examples, American users are no further ahead of the rest of the world in expert systems than they are in other areas of IT. We discuss this further in the next chapter.

Chapter 3

Take-up of expert systems depends on application area, business sector, and country

The extent to which expert systems can be used to improve efficiency or effectiveness or to provide a competitive advantage depends on the type of application, business sector, and the take-up of expert systems in different geographic regions. It is therefore important to understand the main application areas for expert systems and the rate at which they are being adopted by different business sectors and in different countries.

MAIN APPLICATION AREAS

The examples cited in the previous chapter show that expert systems are now being used for a wide range of applications. Researchers in the field of expert systems have produced various classifications of expert system applications, the best-known of which is that of Hayes-Roth. This scheme classifies expert systems by the type of problem addressed: interpretation, prediction, diagnosis, design, instruction, control, and so forth.

The Hayes-Roth classification is very useful once the problem to be solved by an expert system has been identified because it is based on how human knowledge is used, and because it implies a developmental progression in addressing an overall problem. For instance, a particular solution can be categorised as solving a classification problem, or a diagnostic problem, or a repair problem — depending on whether the expert system has been designed just to classify the symptoms of a fault, or to diagnose the cause of the fault as well, or to provide advice on the repair to be carried out once the cause has been diagnosed.

The drawback of the Hayes-Roth classification for business users of expert systems is that it assumes that the application area has already been identified; it does not help to identify the potential opportunities for using expert systems. We prefer to use a classification based on four generic application areas that are familiar to most business systems staff, and that can be used for identifying opportunities. This classification scheme is summarised in Figure 3.1 overleaf, which describes the characteristics of each application area, indicates the level of investment required and the level of payoff that can be expected, and provides examples of typical applications for each area. The type of application area will also determine the most appropriate development tools and methods to use. This topic is discussed in Chapter 6.

In our classification scheme, the universe of possible applications is divided into those of a commercial or administrative nature and those of a scientific or technical nature. The commercial/administrative category is then subdivided into two: information systems and decision-support systems or end-user computing. The scientific/technical category is also subdivided into two: online or realtime systems and offline or standalone systems.

For example, airline seat-allocation applications or front-office applications such as a tax advisor are commercial/administrative information systems, whereas a mathematical model of the economy or of the future of an organisation is a commercial/administrative decision-support application. Process-control-related applications are scientific/technical, online/realtime applications, and systems used to design discrete items of machinery (such as PABXs) or to diagnose faults and repair them, are scientific/technical offline/standalone applications. Although each of the four application areas will have some relevance for most organisations, some application areas will be more important than others for each industry sector.

INFORMATION SYSTEMS APPLICATIONS IN THE COMMERCIAL/ADMINISTRATIVE AREA

Applications in this area are relevant to most sectors, but in particular to financial services, retail and distribution, and public administration, especially where a specialist is needed at the customer interface. All too often, that specialist becomes the bottleneck in the process.

DECISION-SUPPORT OR END-USER APPLICATIONS IN THE COMMERCIAL/ADMINISTRATIVE AREA

Decision-support systems and end-user computing using expert system tools are applicable to all sectors that employ a large number of managers and professionals.

Figure 3.1 Classification of expert system application areas

Application area	Characteristics	Investment required	Potential payoff	Examples
Commercial/ administrative				
Information systems	Complements or extends mainstream systems. systems. Used by general staff.	Variable	Variable	Loan approval; stock control and ordering; airline seat allocation; social welfare advice.
Decision support or end-user computing	Extends decision support or end-user applications. Used by expert or professional staff.	Low	Variable	Assessment of the company's business health; audit planning; risk assessment.
Scientific/technical				
Online/realtime	Part of process equipment. Used by equipment operator.	High	Very high	Sinter plant operation; coal cutter control; nuclear plant instrumentation.
Offline/standalone	Part of repair or design process. Used by repairer or designer.	Modest to high	High	Fault diagnosis; equipment configuration; engineering design.

ONLINE/REALTIME APPLICATIONS IN THE SCIENTIFIC/TECHNICAL AREA

This application area is most relevant to sectors that have process-control requirements and other highly technical or scientific functions. Hence, it is important to the process-manufacturing sectors, such as energy and mineral extraction.

OFFLINE/STANDALONE APPLICATIONS IN THE SCIENTIFIC/TECHNICAL AREA

This application area is most relevant to sectors that have discrete technical or scientific activities, such as after-sales customer engineering and engineering design. Thus, it is important to the discrete-manufacturing sectors, such as engineering, computer manufacturing, and electronics.

RATE OF PROGRESS IS DIFFERENT IN EACH APPLICATION AREA

The extent to which expert systems are already being used in each of the four application areas is different. Applications in the scientific/technical area (online/realtime and offline/standalone) are the most advanced and will develop fastest. Today, there are many examples of standalone diagnostic expert systems, and advanced users are beginning to experiment with expert systems concerned with integrated machinery and integrated instrumentation.

The next most advanced application area is in the use of expert systems for decision-support and end-user computing. Today, expert systems are being used as adjuncts to spreadsheets and for simple modelling, and there are experimental systems for complex modelling and for providing high-quality advice. The use of expert systems in this area is also developing quickly, but not as quickly as in the scientific/technical area.

The use of expert systems for commercial/administrative information systems is the least well-developed applications area. Today, expert systems are used to ensure that rules and regulations are adhered to, and there are experimental expert systems for the application of complex knowledge and expertise. This application area is developing less rapidly than the other three, although the rate of progress will increase as tools become available for interfacing (and subsequently integrating) expert systems with mainstream data processing.

TAKE-UP IS DIFFERENT IN DIFFERENT SECTORS

In order to take advantage of expert systems (or of any other new technology), each organisation must seek to answer the following questions:

— How will expert systems technology affect my business sector?

- What are my competitors and peers doing about expert systems?
- What has my organisation done about expert systems so far?
- What is it planning to do in the future?

Without answers to these questions, the organisation is vulnerable to competitive activities. At best, it may miss the opportunity either for gaining a competitive advantage or for improving its effectiveness and efficiency.

The take-up is different for different business sectors, though, interestingly, it is not very different in the three main geographic areas: the United States, Japan, and Europe. Figure 3.2 shows the take-up (and projected take-up) of expert systems by several business sectors in Western Europe in terms of progress through the well-known product life-cycle stages (as explained in the key to the figure). The surveys carried out by ICOT and JIPDEC indicate similar progress in Japan: about 125 expert systems were identified, only two of which started before 1984; more than 40 were started during 1984 and 1985, and around 100 were started in 1986; another 200 projects were expected to begin in 1987.

THE DEFENCE INDUSTRY IS THE MOST ADVANCED

In Europe, the defence industry (which relies heavily on applied research) is the most advanced in exploiting expert systems. This is not surprising, given the research nature of AI, from which expert systems originated. (Our comments about the defence industry should be treated as an 'informed estimate', because the secret nature of much defence work makes it impossible to gather direct evidence.) During 1987/88, the use of expert systems in this sector will be moving into the upper part of the 'pacing-technology' stage. This means that if you are in the defence business and regard yourself as a market leader, or you aspire to be one, you have to be using expert systems today. This sector is also the most advanced in the United States and Japan.

MANUFACTURING

After defence, European manufacturing industries are making good progress in exploiting expert systems, with process manufacturing leading discrete manufacturing. We predict that expert systems will remain a pacing technology in the manufacturing industry sector until well into 1989. The implication is that if you are a manufacturing company and have not yet begun to use expert systems, you still have time to catch up. The position is similar in the United States and Japan, with many more applications in the electronics and electricals subsector of manufacturing than in any other subsector.

FINANCIAL SERVICES

Our research shows that the next most advanced sector in the use of expert systems is financial services, with banking leading insurance. However, there is some evidence to suggest that the

Figure 3.2 Take-up of expert systems in Western Europe by business sector

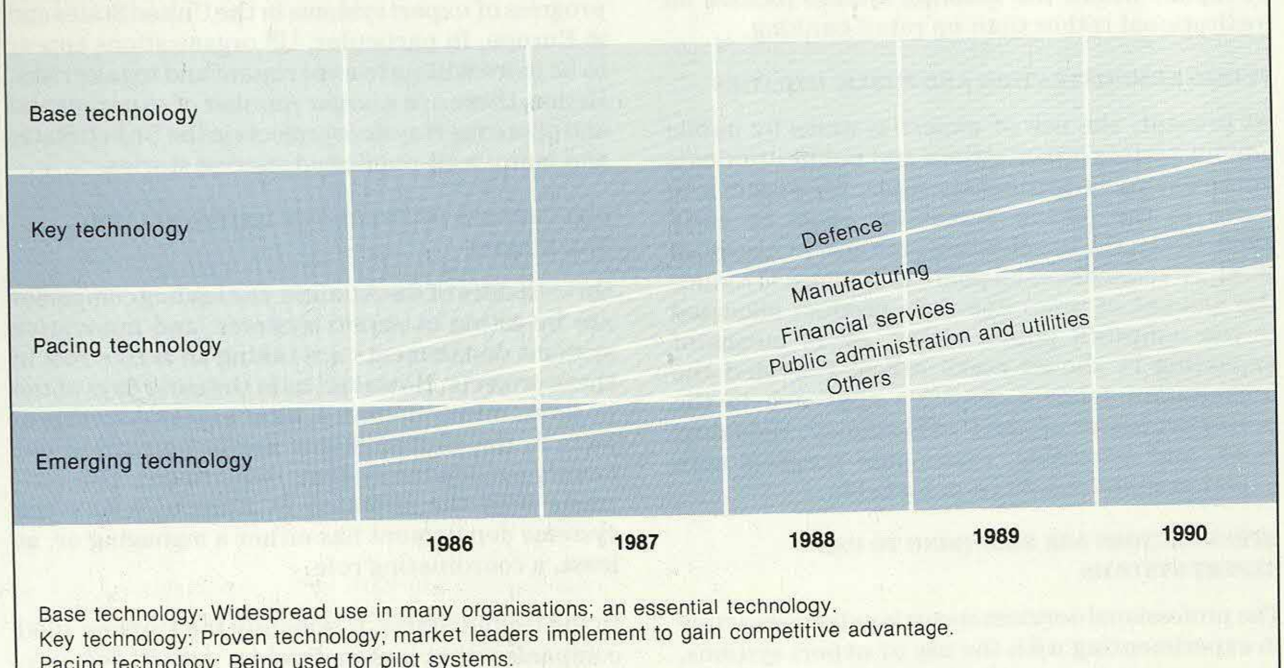

Base technology: Widespread use in many organisations; an essential technology.
Key technology: Proven technology; market leaders implement to gain competitive advantage.
Pacing technology: Being used for pilot systems.
Emerging technology: Still in the laboratory; few prototypes in use.

pace of advance might be faster in this sector than in the manufacturing sector. For instance, an Arthur D Little study in 1986 found that 32 per cent of American insurance firms were active in expert systems; and a Coopers & Lybrand study of the top 100 insurance firms in the United States showed that by 1987 only 21 of them had not begun work on expert systems (see Figure 3.3).

Figure 3.3 The top 100 US insurance firms are embracing expert systems technology

State of development	*1986*	*1987*
Operational application	2	22
Under development	33	57
Not yet started	65	21

(Source: Coopers & Lybrand)

In our view, manufacturing industries will stay ahead of the financial-services industry in their use of expert systems. In particular, manufacturing industries will have to find new ways of competing with the developing and newly industrialised nations. Their motivation for using expert systems will therefore be based as much on survival as on growth in the market. The explosive expansion of the market for financial services will mean that this sector will not be under such great pressure to find new ways of competing in the marketplace. We therefore believe that expert systems will remain a pacing technology in the financial-services sector until the early 1990s. This means that if you are a financial-services company and have not yet begun to use expert systems, you have a breathing space of two or three years.

However, if you want to gain an advantage over your existing and prospective competitors, then expert systems — whether as a front-end to the foreign-exchange or commodities-dealing workstations, or as advisors to branch-office staff dealing with loan applications or investment portfolios — may well enable you to do so. Although the take-up of expert systems in the financial-services sector is advanced in the United States, this is not the case in Japan, where the banking system focuses on institutional rather than on retail banking.

PUBLIC ADMINISTRATION AND PUBLIC UTILITIES

At present, the use of expert systems by public administration organisations and public utilities is at the emerging-technology stage. We expect it to move to the pacing-technology stage by early 1988. Given the unrelenting pressure in nearly all Western countries of all political leanings to reduce the amount of gross domestic product allocated to the nonprofit public sector, every means of exploiting IT should make sense. Provided the organisation takes a sensitive approach to the personnel issues of staff redundancy, redeployment, and retraining, reasonable progress with expert systems should be achievable.

OTHER SECTORS ARE BEGINNING TO USE EXPERT SYSTEMS

The professional-services sector is extremely active in experimenting with the use of expert systems, as are certain parts of the transport and the mineral-extraction sectors. And expert systems can now be regarded as an emerging technology in many other business sectors. In Japan, for example, there are many examples of expert systems in the construction industry.

The one sector that surprised us is retail and distribution. We found very few examples of expert systems in firms in this sector, and most of these examples were in the area of optimisation, such as vehicle routeing or stock control — the focus of traditional operational research. We believe that there ought to be many opportunities for using expert systems in retailing — for intelligent market forecasting, or targeted marketing, for example. The expert system suppliers we surveyed believe that retail and distribution companies are active in developing plans for using expert systems — behind manufacturing, but ahead of public utilities, and transport and communications.

PROGRESS IS DIFFERENT IN DIFFERENT COUNTRIES

In our research we found several significant differences and some interesting similarities between the progress of expert systems in the United States and in Europe. In particular, US organisations appear to be more willing to experiment and to take risks. Hence, there are a larger number of experimental and pilot expert system projects in the United States and many well-publicised success stories.

DIFFERENCES BETWEEN THE UNITED STATES AND EUROPE

On both sides of the Atlantic, the leading companies are investing in expert systems, and innovative systems departments are taking an active role in these projects. However, as in the early days of the business microcomputer, most expert system projects in the United States are being carried out largely outside the systems department. This contrasts with the situation in Europe, where the systems department has either a managing or, at least, a coordinating role.

In the United States, it is mainly the Fortune 1000 companies that are active in expert systems, whereas in Europe much smaller firms as well as the major organisations are very active in applying expert system techniques. Overall, we believe that US organisations are not significantly ahead of organisations in other parts of the world in making effective use of expert systems. At first sight, this may seem surprising but, unfortunately, many consultants and market researchers have confused the maturity of US expert system suppliers and products (which are several years ahead of those in Europe) with the state of users and applications there (which is, at best, only several months to a year or two ahead of those in Europe — though, of course, the situation varies from country to country in Europe). This view was endorsed by a technical manager of Pactel's Intelligent Systems Centre, who said at the KBS87 Conference in London in June 1987 that the United Kingdom, in many cases, is as advanced as the United States in applying expert systems, particularly in petrochemicals, utilities, and financial applications.

There is one key difference between expert system projects in the United States and Europe, however. There are several multimillion-dollar commercial/administrative information systems projects in the United States; in Europe, apart from government-subsidised projects, we do not know of anything on the same scale. Only where expert systems are being integrated with process-control systems do projects approach this scale. Although this difference reflects the differences in national and regional business cultures, it could make a significant difference in the exploitation of expert systems to gain a competitive advantage. The recent thrust of both American Airlines' Sabre and United Airlines' Apollo reservations systems into Europe, and the belated attempt by two groups of European airlines to develop an alternative, is an example of the risks that European organisations run by not investing heavily and quickly in new information technology.

PROGRESS IN JAPAN

We visited ICOT in Japan and also had access to the results of a survey carried out in April and May 1986 by the JIPDEC Institution in Japan, which analysed the responses of 203 organisations. The main findings of our Japanese research are that:

— Japanese organisations are farther behind than European organisations, in terms both of products and of using expert systems.

— Although most of the suppliers are Japanese, more than half of their products originated in the United States (this is true in Europe as well).

— In terms of the role and involvement of the systems department in expert system projects, Japanese organisations are closer to the United States (where the majority of initiatives come from users) than to Europe (where the systems department usually leads or coordinates expert system projects).

— More use is made of languages than of shells, and of minicomputers than of microcomputers. As in the United States, more use is made of specialist AI hardware than in Europe.

— The majority of applications are in the diagnostic area, with product design a close second.

Chapter 4

Selecting the best opportunities for using expert systems

Before beginning the task of identifying the best opportunities for exploiting expert systems, it is necessary to select the best overall approach for using expert systems in your organisation. The approach chosen will determine how successful you will be in exploiting expert system applications.

You will then need to identify and prioritise the opportunities for using expert systems. Much has been written about how to select the best potential application areas for expert systems. Although this material is useful for evaluating opportunities (and we incorporate many of the ideas later in this chapter), it is not particularly helpful in identifying the application areas on which to concentrate. Instead, it provides advice for discarding unsuitable opportunities. We therefore describe how to identify and prioritise the possible opportunities in each of the four application areas in our classification.

SELECTING THE BEST APPROACH

There are three different approaches you can adopt for exploiting expert systems. You can accept expert systems as just another technical tool; you can use them as a new pragmatic solution to certain applications needs; or you can use them as a strategic weapon — as a means of radically changing your business.

USE EXPERT SYSTEMS AS A TECHNICAL TOOL

One approach you can adopt is to treat expert systems as just another technical tool. They may be new, clever, and, perhaps, more useful or more cost-effective than any existing software development tool. But as far as you are concerned, they are merely a technical tool. With this approach, the 'technicians' in the systems department should be left to evaluate, select, and use expert systems as they think best. Users and senior management will not be involved in directing the organisation's expert system activities.

The technical-tool approach is likely to be the least-cost and least-risk option, but it is also likely to provide the least benefit. Left to itself, the systems department is unlikely to place a high priority on expert systems. It already has enough new technology to come to terms with. Hence, by the time your organisation uses expert system tools for operational applications, the products will be mature, the suppliers stable, and the lessons well learnt by other pioneering organisations.

There is nothing wrong with this approach — if it is the most appropriate for you for the next two to three years. In this case, the technical aspects of this report, particularly Appendix 2, will be most relevant.

USE EXPERT SYSTEMS AS A NEW PRAGMATIC SOLUTION TO BUSINESS PROBLEMS

The second approach you can adopt is to involve the users from the start, perhaps with a user champion leading the expert systems initiative. Some users already regard expert systems as a pragmatic solution to many business problems that could not be tackled easily before. Their information about expert systems and their motivation for applying them comes from professional peers, journals, conferences, and directly from suppliers at trade shows. In many ways, this approach is similar to the way that personal computers were introduced in the late 1970s and early 1980s. Then, the microcomputer was promoted in all kinds of business and domestic periodicals and was demonstrated at all sorts of business events. Business managers and professionals began to use personal computers to solve problems that were not being tackled by the systems department. Today, expert systems are also seen by some users as the means of solving problems that have not been tackled by the systems department.

The pragmatic-solution approach to using expert systems is likely to result in a large number of discrete applications, each of which has been installed to meet a particular requirement. This may well be the most suitable approach for some organisations for the next two or three years. The drawback of this approach, however, is that there is no overall strategy for exploiting expert systems and no policy regarding products or suppliers. This is all very well if there are adequate financial and technical resources, but exploiting expert systems in this way is likely to be inefficient. The drawback of the pragmatic approach can be avoided if the systems department takes a lead in establishing priorities and standards.

For organisations adopting the pragmatic-solution approach, the practical example-based sections in Chapters 2 and 3 will be of most use. However, the technical aspects of selecting the right software (Chapter 6), and setting up the right kind of support team (Chapter 5) will also be important.

USE EXPERT SYSTEMS AS A STRATEGIC WEAPON

The third approach you can adopt to exploiting expert systems is to recognise that you can use them to change the way you do business and, as a result, to gain a competitive advantage or to substantially improve your efficiency or effectiveness. Expert systems can do this because they represent a major technological discontinuity. They have the potential to be a new strategic business weapon.

This approach means that strategic applications need to be identified; that business objectives, threats, and opportunities, as well as your organisation's strengths and weaknesses, need to be assessed; and that the potential and the limitations of expert systems must be understood. It is then a relatively easy task to select the business functions, the domains of expertise, the applications, and the users that offer the best opportunities for exploiting expert systems. This approach is the most costly, and has the highest risk, but it provides the greatest potential for making a long-term difference to the business itself. For it to succeed, it clearly needs the commitment of senior user management.

Organisations that decide on the strategic-weapon approach must not only take heed of all the technical issues and the application lessons in this report, they also need to be concerned with exploiting technology to gain a competitive advantage. This topic forms the subject of the next Foundation Report.

DECIDING ON THE BEST APPROACH FOR YOUR ORGANISATION

At first sight, the choice between the three possible approaches may not seem an easy one to make. In fact, it is straightforward. It should be based on your own organisation's business objectives and strategy.

First, if your organisation is, or aspires to be, a leader in a strongly competitive and fast-moving market, then it is essential to assess the major business threats and opportunities posed by expert systems. You therefore need to use the strategic-weapon approach.

However, if the competitive pressures in your marketplace are less and your business strategy is to follow closely behind the innovators in your sector, then you will need to seek out and be ready to use expert systems as tactics dictate. You should therefore adopt the pragmatic approach.

Finally, if your organisation operates in a less competitive market (or in a nonprofit sector) and your business strategy is to minimise the risks of technical innovation, then the best policy is to wait until the need for the new technology arises and until the products are fully proven. You should therefore treat expert systems as just another technical tool.

Of course, many Foundation members have multiple business streams that span several markets and have a mix of business strategies. You should examine each separately to identify the best approach in each area of business. You may well find that you need to use a different approach for each business stream.

The choice will also depend to some extent on the rate at which expert systems are being adopted by your particular sector. If expert systems are about to become a pacing technology in your sector, you probably cannot afford to adopt the technical-tool approach. You will need to adopt the strategic-weapon or pragmatic approach and time your expert system initiatives according to the activities of other organisations in your industry. Timing will be most critical in applying expert systems as a competitive weapon, where being first (or at least the first company to apply expert systems well) is how the advantage is often gained.

IDENTIFYING THE POSSIBLE OPPORTUNITIES

Although there are differences in detail for each of the four application areas, there is one common guideline for identifying expert system opportunities, which is illustrated in Figure 4.1 overleaf. The figure shows that opportunities for expert systems exist where an expert or specialist applies his or her skills at the interface either between two computer systems or between a computer system and another person (a customer, supplier, senior manager, fellow worker, and so on). In this situation, the expert or specialist often becomes the bottleneck for information flow or decision making.

We now discuss the different guidelines for identifying possible expert system opportunities in each of the four application areas.

Figure 4.1 Expert system opportunities exist when an expert or specialist forms the bottleneck at the boundary with a computer system

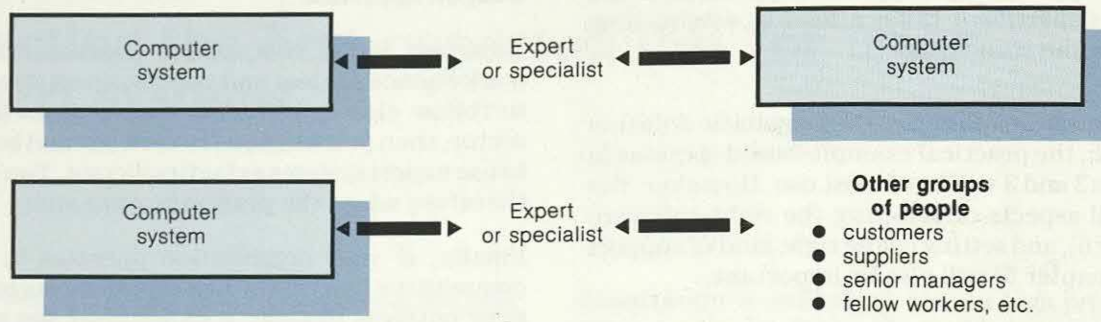

INFORMATION SYSTEMS IN THE COMMERCIAL/ ADMINISTRATIVE AREA

Possible opportunities for expert systems in the information systems application area exist when:

- There is a significant (even critical) business function where a bottleneck is being caused by the unique skills of a specialist. For example, Lend Lease in Australia develops and manages large property sites. It found that one of its project planners consistently produced more accurate estimates for multimillion-dollar (and multiyear) projects than his colleagues. This planner's skill is being captured in an expert system.
- There would be a direct commercial benefit from distributing the expertise normally found only at the head office. One example is an expert system that allows branch-office staff to use well-tried underwriting practices for export guarantees, thereby relieving the head-office experts of the more mundane cases and allowing them to concentrate on the small number of difficult and time-consuming cases.
- There are experts who act as the interface between different computer systems. For example, the purchasing officer of a manufacturing company might take the data from the company's stock-control and sales-order-analysis systems and combine it with his or her own knowledge of the business to decide whether to initiate an order via the purchasing system. There is an opportunity to exploit expert systems if assisting these experts will improve their productivity. Using expert systems to help the experts who act as the interfaces between computer systems has an important implication, however. The expert system will need to interface with existing mainstream software and data. We see this as being a growing requirement, with expert systems becoming one element of an integrated solution to many business problems.

DECISION-SUPPORT SYSTEMS OR END-USER COMPUTING IN THE COMMERCIAL/ ADMINISTRATIVE AREA

The characteristics to look for when seeking possible expert systems in the decision-support systems or end-user computing areas are:

- A staff function where the impact of the decision is significant, even crucial, to the success of the business, but where the experts are overwhelmed by the need to manipulate a mass of basic data that can be combined and permutated in many ways, which prevents them from exploiting their expertise to the full. A typical example would be marketing analysts in an entrepreneurial consumer-goods firm.
- Situations where experts use unstructured but definable methods for their analysis, and where providing the experts with expert systems assistance would result in better decisions. At Robson Rhodes in the United Kingdom, for example, an expert system has been developed to take account of qualitative, rather than quantitative issues. Robson Rhodes is a management consultancy firm specialising in providing advice about corporate finance. The expert system is designed to assist bank managers and corporate-lending advisory staff in making corporate-lending decisions. The expertise provided by Robson Rhodes' corporate finance unit (and modelled in the expert system) is aimed at improving borrowing decisions and lending decisions by taking account of borrowers' and lenders' perceptions of the likely risks and returns. Robson Rhodes has tried to capture the unit's subjective knowledge in the knowledge base.

ONLINE/REALTIME SYSTEMS IN THE SCIENTIFIC/TECHNICAL AREA

Possible opportunities for expert systems in online and realtime applications exist when:

- There is a significant, probably crucial, automated production process where successful operation of the process depends on the skill of an expert who optimises the quality of the product by 'fine-tuning' the process. Broken Hill in Australia, for example, operates a sinter processing plant. It is using a Macintosh-based expert system to provide less-skilled operators with the type of correlated data that the experienced operator uses in making fine-tuning decisions. Blue Circle Industries has a similar application for its clinker kilns. It estimates that a rule-based expert system will save 10 per cent of the energy costs and improve the quality of the product. In addition, because the system will allow the process to run at lower temperatures, the refractory lining of the kilns will last 30 per cent longer — resulting in an anticipated saving of $450,000 a year per kiln.

— The experts operating the process have to work under time pressure and have to cope with large amounts of data that must be cross-related. Both Electricité de France and British Nuclear Fuels have pilot expert systems for instrumentation displays in nuclear power plants. The idea is to display the relevant data only, and allow the operator to call up the detailed raw data when it is needed. The aim of the system is to reduce the number of times a nuclear power plant is unnecessarily shut down. According to British Nuclear Fuels, it costs £2 million per day in lost production and restart costs to shut down a nuclear power station.

OFFLINE/STANDALONE SYSTEMS IN THE SCIENTIFIC/TECHNICAL AREA

There are potential expert system opportunities in the offline/standalone scientific/technical application area when:

— There is an important design or maintenance function in which the role of the expert is critical. Northrop in the United States uses expert systems for production planning and design of military aircraft involving 10,000 parts. Before the introduction of the system, the process-planning exercise took between 8 and 12 hours. It now takes between one and two hours, and there are almost no errors.

— The experts in the design or manufacturing function use permutations and combinations of large amounts of data, some of which may be of suspect accuracy. Several electricity utilities are experimenting with applying expert systems to the design of power distribution networks. The design of such networks involves an extremely large number of possible combinations of hundreds of factors. Existing systems impose limits on the combinations that can be considered and therefore compromise the validity of the results.

— There is an incomplete understanding of the behaviour of the physical environment. British Coal has an expert systems application that predicts the likelihood of encountering methane when opening up a new coal face or mining an existing face in a different direction.

— Direct and tangible benefits, such as improved customer service, would result from providing the experts with expert systems assistance. British Telecom has a fault-diagnostic expert system that enables new engineers to service older-generation PABXs. This has resulted in improved customer service and has reduced the costs of training engineers to maintain obsolete equipment. It has also prolonged the useful life of the equipment and increased the revenue from it.

PRIORITISING THE OPPORTUNITIES

Using these guidelines will usually result in more candidates for potential expert system opportunities than you will be able to cope with, even with a reasonably sized expert systems team. First, you should eliminate any projects that are not suitable candidates for today's expert systems technology:

— If conventional systems technology can be used to solve the problem, then do not use expert system tools — they are still relatively immature and should be used only when existing tools are inadequate.

— Problems that take an expert more than a few days to solve are not yet candidates for expert systems because they are beyond the capabilities of today's expert systems technology.

You will then need to prioritise the remaining opportunities by considering factors relating to the business requirements of the opportunities, to the availability of experts and expertise, and to the system development requirements. The factors are listed in their order of importance for prioritising the opportunities.

The business requirements to consider are:

— The relevance of the opportunity to the systems strategy and business objectives of your organisation.

— The potential payoff of the opportunity and the significance of the problem. Too many of the early projects we heard about during our research were in the 'interesting-exercise' category. Hence, even though they were technically successful, nothing ever happened after they were completed. Worse, the organisation had moved no further toward creating an understanding of the real potential of expert systems, much less toward generating a belief that they should be taken seriously now. In other words, there must be a business reason for proceeding with the project. The best way to make sure that there is a payoff is to identify a user sponsor or 'champion'. The existence of a user champion will prevent sceptics from saying that the project is a technical exercise to satisfy the curiosity of the systems department.

- The extent to which the consistency or timeliness of decision making can be improved.
- The likelihood of obtaining management commitment. Major projects could take several man-years of development effort and substantial amounts of the experts' time — both for the initial development and for any subsequent knowledge-base maintenance and enhancement.
- The extent to which the intended users of the system (who often are not the expert or specialist whose knowledge or know-how is being modelled) are enthusiastic and committed to the success of the project.
- The feasibility of integrating the expert system with the existing organisational procedures. Sometimes, a successful pilot system has been abandoned because the scale of the changes in the organisational structure or procedures required to implement a fully operational system has been too great to contemplate.

Next, factors concerning the availability of experts and expertise must be considered:

- The solution of the problem should not require general knowledge or common sense. Such problems cannot yet be tackled by computer technology.
- There should be a vital need to replicate or distribute the expertise or specialism.
- The experts or specialists constitute a bottleneck.
- The expertise or specialism could be lost if it is not captured in an expert system — either because the expert is nearing retirement or because he or she is likely to change jobs in the foreseeable future.
- The domain of expertise should be well bounded and narrow, with commonly accepted standards or methods. Current expert system tools and techniques cannot handle broad domains.
- An expert or specialist should be willing to cooperate in the project. Failing this, there must be a body of recognised and accepted test cases to start from. In particular, the expert must not feel that his or her job is threatened by the expert system.

Finally, the system development requirements must be considered:

- There should be a reasonable chance that a working prototype can be produced in less than a year. Users and senior managers despair if it takes longer.
- It should be feasible to interface the expert system with appropriate existing systems. More than 25 per cent of the applications reported by the vendors that responded to our survey are linked, or even integrated, with a mainstream system.
- The cost of implementing a full operational version of the system should be considered. We heard about a surprisingly large number of apparently successful experimental and pilot projects that had been abandoned once the cost of fully implementing the system had been calculated. A small-scale pilot can be implemented for a relatively moderate cost, whereas the cost of providing hardware and software for use throughout the organisation can be very significant — either in terms of the cost of central-processing capacity or in terms of the cost of the large number of terminals or workstations required.

In March 1987, *EDP Analyzer* published a checklist of questions that could be asked to determine if a potential expert system opportunity is worth pursuing. Some Foundation members may find the questionnaire a useful adjunct to the factors listed above, and it is set out in Figure 4.2.

Figure 4.2 Expert system project-selection criteria

Criteria	*Score* D	*P*	*M*	*N*	*N/A*
Economic justification					
Is there a need to make the knowledge of a specially trained or talented individual more widely available?	6	4	2	0	0
Will the expertise be lost if not captured by an expert system?	6	4	2	0	0
Will there be significant savings or payoffs from an expert system?	6	4	2	0	0
Is there a need to increase the efficiency of the decision-making process through improved consistency and timeliness?	5	3	1	0	0
Is there a need to record the decision and the reasons?	3	2	1	0	0
Is it expensive to train individuals to deal with the problem area?	5	3	1	0	0

(Continued on next page)

Figure 4.2 Expert system project-selection criteria (continued)

Criteria	*Score* D	P	M	N	N/A
Economic justification (continued)					
Will building an expert system help future developments?	4	2	1	0	0
Can the proposed expert system be integrated into existing service or product lines to increase their value?	5	3	1	0	0
Will an improved understanding of the problem, gained through expert system development, be valuable to the organisation?	4	2	1	0	0
Total score: More than 30 = Excellent opportunity 20 to 30 = Good opportunity if cost is not too great Less than 20 = May not provide sufficient payback					
Availability of expertise					
Is there an expert available who solves problems significantly better than the majority of the intended users of the expert systems?	6	4	2	0	0
Is the expertise to be used accurate and correct?	5	4	1	0	0
Are there a few key people with specialised knowledge or expertise spending excessive time helping many others?	5	3	1	0	0
Are the experts articulate enough to explain their methodology?	5	3	1	0	0
Are the experts willing to work with a knowledge engineer?	5	3	1	0	0
Will the experts have time to complete the development process?	5	3	1	0	0
If multiple experts contribute, is one the final authority?	4	2	1	0	0
Total score: More than 20 = Adequate expertise available 15 to 20 = Expertise may not be adequate Less than 15 = Insufficient expertise available					
Suitability for expert systems					
Does the problem require mainly experience-based reasoning?	5	3	1	0	0
Is the problem solution dependent on common-sense reasoning?	0	1	2	5	0
Does the problem require small amounts of time for the expert to solve or explain (less than two hours), or can it be subdivided?	5	3	1	0	0
Do the users require all possibilities to be known in advance?	0	1	2	5	0
Is the problem domain well bounded?	3	2	1	0	0
Are the interfaces with users or other 'systems' easy to achieve?	3	1	0	0	0
Total score: More than 18 = Suitable for expert systems 11 to 18 = Some difficulties may be experienced Less than 11 = Unsuitable for expert systems					
Acceptance by user community					
Is there adequate managerial commitment for the effort?	5	3	1	0	0
Are the users committed to using the expert system?	5	3	1	0	0
Will the introduction of an expert system cause political or control repercussions either from its use, contents, or recommendations?	0	1	2	5	0
Can the expert system be integrated into existing procedures without upheaval?	3	2	1	0	0
Will the system handle a real and necessary business need?	5	3	1	0	0
Is it acceptable to complete the system in phases?	5	3	1	0	0
Total score: More than 21 = Organisation will support the project 16 to 20 = Success may be difficult to achieve Less than 16 = There may be obstacles that will hinder the project					

Each question can be answered in one of five ways:
D = Definitely, P = Probably, M = Maybe, N = No, N/A = Not applicable

(Adapted from EDP Analyzer, March 1987, based on work done by Mind Path Technologies, Dallas, Texas)

Chapter 5

Managing the introduction of expert systems

Once the overall approach for using expert systems has been selected, and the most promising opportunities have been identified and prioritised, it is necessary to manage the introduction of expert systems. We believe that the introduction of expert systems should be coordinated by an expert systems support unit, and in this chapter we discuss the role of such a unit and the most appropriate location for it.

It is also necessary to ensure that the right skills are available for the introduction of expert systems or to provide appropriate training where the skills are lacking.

Careful attention must then be given to managing the implementation of the first expert system application. This chapter provides appropriate guidelines. The final management activity identified in the chapter is the need to monitor future developments in the fast-moving field of expert systems. We provide guidance on what to look out for.

ESTABLISH AN EXPERT SYSTEMS SUPPORT UNIT

Before an expert systems support unit is set up, an organisation needs to define the role of the unit, both for the immediate future during the initial stages of exploiting expert systems and for the following two to three years. Planning any further ahead is likely to be premature, both because the area of expert systems is developing rapidly and because the organisation's understanding and exploitation of expert systems will change considerably in that time. Some organisations have formed an expert systems unit, but have restricted it to a research role, without any real support from, or contact with, either senior systems management or user management. Unless the sole purpose of working with expert systems in your organisation is to experiment, a unit such as this is likely to be a complete waste of valuable and scarce resources.

ROLE OF THE EXPERT SYSTEMS SUPPORT UNIT

The expert systems support unit has four different roles:

— To promote the use of expert systems and provide appropriate education.

— To plan and coordinate the implementation of expert systems.

— To provide support and consultancy services.

— To develop expert systems.

Some units will perform all four roles; others a combination of one or more, depending on whether the responsibility for exploiting expert systems lies with the user community, or the systems department, or is shared, and on the level of investment that will be made in expert systems (see Figure 5.1).

Promotional and educational role
If you want the user community (or the organisation's research and development department) to take the bulk of the responsibility for exploiting expert systems and to provide most of the effort, the main role of the expert systems support unit will be to promote the use of expert systems and to provide expert systems education. Although this type of support unit will be acting as the catalyst for exploiting expert systems, it will need to acquire actual experience of real products, ideally by using them on in-house projects. This can be achieved by working with the users on their projects.

Figure 5.1 Roles for the expert systems support unit depend on where the responsibility for expert systems lies

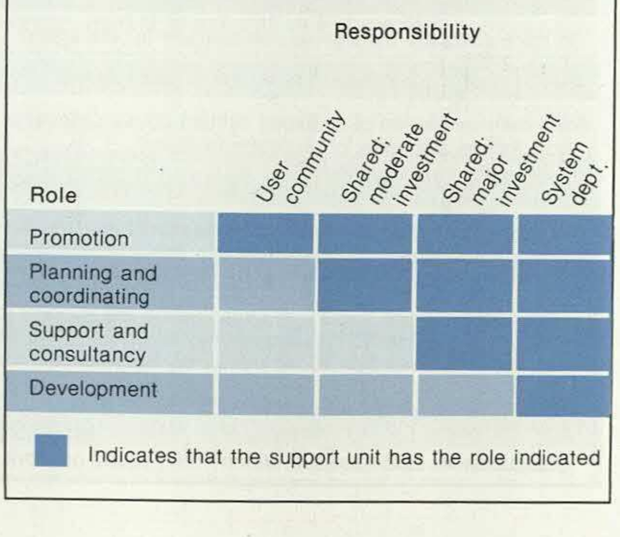

In many respects the promotional role of the expert systems support unit is similar to that of the end-user computing support unit (or information centre). Successful ways of promoting end-user computing throughout the organisation could be adopted for the promotion of expert systems. The expert systems support unit might thus publish a regular newsletter giving news about a network of user 'experts' that can provide local support for everyday queries. It might also organise in-house seminars involving users from your own organisation, speakers from organisations in the same sector, and speakers from suppliers and consultants. Such events can be an excellent vehicle for spreading awareness about expert systems, and for enabling users to share their experiences, in a controlled environment.

In addition to promoting the use of expert systems in the organisation, the expert systems support unit will also need to educate managers and users about their potential and to support the technical training activities in the systems department. The main education and training needs of different groups of staff are summarised in Figure 5.2.

Managers need to be able to identify the high-payoff opportunities for using expert systems. This ability is particularly important if expert systems is already a pacing technology in your sector, but your organisation has not yet started to use them. The purpose of the management-education activity is, first, to alert senior managers to the threats and opportunities and, second, to gain their backing and commitment for some experiments as precursors to operational projects. They need to be aware that, even though benefits can come from modest investment in expert systems, major benefits are most likely to require substantial investment of money, people (experts, systems staff, and users), and management time and attention. Finally, they need to be aware that significant expert systems — like any other significant computer system — will impact on people and on the organisation. If the approach to be taken is to use expert systems as a strategic weapon, the impact will be on the way business is carried out.

Figure 5.2 Education and training in expert systems for different groups of staff

Group	Education and training
Users and managers	Policy and direction Threats and opportunities Potential and limitations
Systems department	Prepare for future role
Business analysts	Matching problems and solutions
Knowledge engineers	Knowledge acquisition and documentation
Designers and programmers	Development tools and techniques

The potential of expert systems also needs to be demonstrated to the user community. Many users, in spite of all the publicity, will still be bemused (if not confused) by expert systems and will not be in a position to identify the opportunities for using expert systems in their field of activity. The education effort should be directed at creating an enthusiastic user community that is aware of the potential, as well as the limitations, of expert systems. One particular concern in the user community may be that expert systems will lead to job losses and de-skilling. These fears should be allayed wherever possible.

For experts and other specialists, the potential of expert systems to enhance their skills and improve their personal productivity needs to be promoted. They too may fear that expert systems will de-skill their role or make them redundant. These fears also need to be allayed.

There will also be a need to promote the concept of expert systems within the systems department. Even though widespread support and acceptance of expert systems by the systems department is not necessary in the initial stages, we believe that, eventually, the expert systems activity will become an integral part of the systems department's responsibilities. Many systems departments now have the reputation of being slow to embrace a new and different technology (witness the introduction of the business microcomputer), and it will do no harm to alert the systems department at an early stage to the concepts of expert systems and the potential for using them.

Planning and coordinating role

When the organisation is ready to make a substantial investment in expert systems, the expert systems support unit should take on a planning and coordinating role. This role requires a wider range of skills than those required for a promotional role. The team members need to be able to interpret the users' ideas so they can be translated into actual requirements for the suppliers. They need to be good at listening both to prospective users and to suppliers and external consultants, and to be able to draw their own conclusions about the match between requirements and the proposed solutions. They also need to be able to advise the users about suitable products, helping them to distinguish between the reality and the marketing hype in the information they will be receiving.

The team members will also need to be able to persuade some or all of the users to accept compromises in their requirements and timescales and to forgo the use of their favourite products or consultants, for the common or longer-term good of the majority. In addition, members of the unit will need to be familiar with the business requirements of the organisation and with the need to align systems planning with business planning.

Part of the coordination activity will be to evaluate and shortlist preferred suppliers and products. Chapter 6 provides guidelines for selecting products and tools. The establishment and maintenance of lists of preferred suppliers and products can also help to create an environment where users from different parts of a large organisation can share their experience and learning. Such a list also prevents inexperienced users or experts from wasting their time in performing product-evaluation and selection exercises. It should also enable the users to employ the best products for the given requirements.

In time, as an organisation develops many expert system applications, and as the tools evolve to a state where multiple knowledge bases can be used, the role of this type of support unit might be expanded to include knowledge-base management. The role of a knowledge-base manager would be similar to that of the database manager. But for almost all organisations, such a development will be several years in the future.

Support and consultancy role
Where the organisation expects to make a major investment in expert systems and the systems department has an active role in helping users to exploit expert systems, the main role of the support unit should be to provide support and consultancy services. This role requires the unit's staff to be skilled in all aspects of expert systems. It is most likely that staff with some of the skills will have to be recruited from outside the organisation, and that use will have to be made of external consultants and suppliers. However, the ultimate goal of this role is to encourage and enable the users to become self-sufficient. Whilst it will be perfectly in order for the unit to assist and to teach by carrying out projects, active steps should to be taken to transfer the unit's skills and experience to the user community.

Development role
Where the systems department has the responsibility for expert systems, the role of the support unit will also include the development of expert systems. Although early projects will invariably need support and resources from suppliers and/or consultants, major in-house projects should eventually be developed by the organisation's own staff. Staff skilled in developing expert systems will therefore be required in the support unit. It may be possible to recruit people with the right skills, but because expert systems development is a new discipline, skilled staff are rare and in great demand. Many organisations will have no option but to retrain existing development staff.

LOCATION OF THE SUPPORT UNIT

The most appropriate organisational location for the expert systems support unit depends on the type of application (see our classification in Figure 3.1), and the stage of development your organisation has reached in exploiting expert systems. We identify three stages:

— Experimental stage, where the aim is to learn about expert systems, rather than to use them to solve a particular problem.

— Pilot-project stage, where the aim is to test out expert systems in a restricted but, nevertheless real environment; if the prototype is successful, it will lead to a fuller implementation.

— Operational stage, which may begin in a restricted way with the application being confined to a geographic region, for example, or to a product division or particular market sector.

Figure 5.3 shows the most suitable location of the expert systems support unit for each combination of application type and stage of development.

It is important to realise that as the use of expert systems progresses through the experimental, pilot-project, and operational stages, an ever-widening set of people and factors will have to be taken into consideration. During the experimental stage, management responsibility for expert systems is with the group running the experiments (research and development, operational research, or the systems department). At the pilot-project stage, it will be necessary to involve user representatives and their management as well. For applications in the scientific/technical area that are to interface with process-control systems, it will be necessary to involve staff from the engineering and production functions, and, usually, technical representatives from the suppliers of the process-control systems. Operational expert systems in the information systems area will often need to interface with mainstream data processing systems, which means that systems development and data centre staff will have to be involved. And because operational expert systems will have a large impact on the organisation, user management, and sometimes, senior management, must also be involved once this stage is reached.

During the early days of exploiting expert systems (experimental and pilot-project stages), the expert

Figure 5.3 Location of the expert systems support unit changes as the use of expert systems evolves

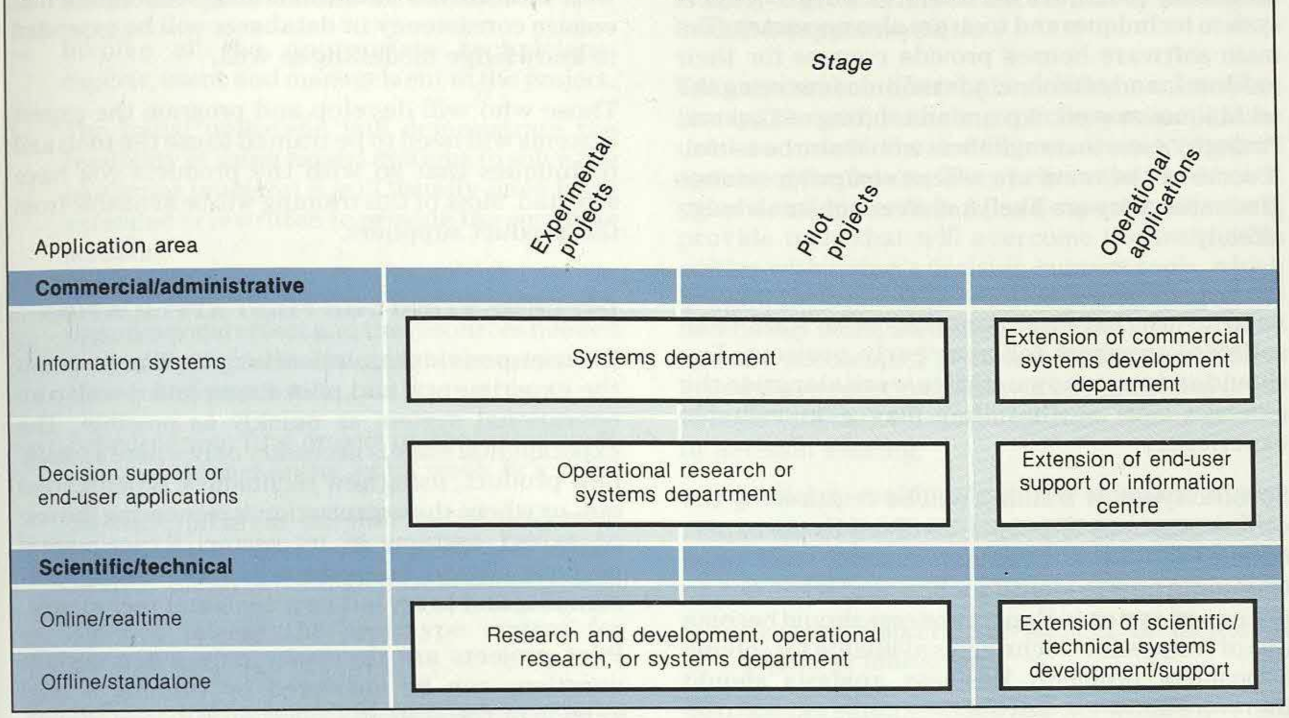

systems support unit should therefore be located centrally, regardless of the way you normally distribute technical specialities in your organisation. The possible locations are the research and development department, the operational research department, or the systems department. Any one of these would be a candidate if the focus of your expert system activities is to be on applications in the scientific/technical area. For information system applications in the commercial/administrative area, the unit should be located in the systems department. For decision-support or end-user computing applications, the unit can be located within the operational research department or the systems department. If the unit is not located in the systems department, there should be a formal relationship between the two.

Once the use of expert systems has reached stage 3 (operational applications), the expert systems support unit should be integrated with the more established system-support functions of the organisation. Thus, for applications in the technical and scientific application area, we recommend that the unit should be included within the technical and scientific systems support function. For the commercial and administrative information systems area, it should be located within the normal commercial systems development function. And for the decision-support or end-user computing area, it should be located within the end-user support or information-centre function.

However, if your organisation has different systems units for technical and scientific development and for commercial and administrative development, it may be sensible to have a small central technical-support unit for expert systems within your normal systems support function to coordinate the two types of expert systems activity.

PROVIDE THE RIGHT SKILLS

In addition to the skills referred to earlier, many organisations are seeking advice about whether the expert systems support unit requires specialist knowledge-engineering skills. They are unclear about whether knowledge engineering is an entirely new discipline, combining skills in cognitive psychology with business analysis skills. Or is it possible that most experienced systems analysts could add knowledge engineering to their skills portfolio? We support the view of Professor Feigenbaum of Stanford University, who said in his address to the UK members of the Foundation at their annual conference in June 1987: "Your most competent experts and your most competent computer specialists can do the work. You do not need specialists trained in artificial intelligence. . . . "

The skills supposedly required of a knowledge engineer are so demanding that it is no wonder there are so few of them. The role requires psychological skills, some technical skills, and 'political' skills. It is a rare individual indeed who combines all of the required skills, and we believe that a team approach to knowledge engineering is the only practical approach.

In addition to knowledge-engineering skills, technical skills specific to the different kinds of expert system techniques and tools are also necessary. The main software houses provide courses for their products, and often include modules covering the techniques as well. An understanding of Lisp and Prolog by some team members would also be useful. If some of the team are recent computer science graduates, they are likely to have such knowledge already.

We have already mentioned that you will probably need to use external consultants or staff from software suppliers for your early projects. You should assign your own staff to work alongside the external staff so that they may gain valuable experience.

Specific types of training will be required by the different groups of people involved in the expert system activities. Business analysts will need training in how to identify when, and when not, to use expert systems. Expert systems should become one of the possible techniques available for solving a business problem. Business analysts should therefore know enough about expert systems to be able to propose them as the best technical solution to a specific business problem.

Those who will perform the role of knowledge engineers will need training in knowledge-acquisition methods and techniques. In most countries where there are Foundation members, in-depth knowledge-engineering training courses are provided by suppliers, consultants, and academic institutions. The suppliers that replied to our survey derive 20 per cent of their income from consulting, which includes a considerable number of training assignments. In the UK there are excellent 'journeyman' courses run by Imperial College in London and by the Turing Institute in Scotland. (The term 'journeyman' means a skilled artisan who works with a master craftsman to improve his or her skill.) The Imperial College course is spread over six months and forms part of a postgraduate MSc course. The Turing Institute course is residential and lasts up to three months. Both courses include work on a problem of the student's own choosing — which is usually nominated by the sponsoring company. The journeyman returns to his or her company with practical experience of how to set up and develop an expert system. By April 1987, 30 journeymen from 20 companies had completed their training successfully.

Knowledge engineers also need to be trained to document their work in a standard way so that others in the same organisation can take over their projects with the minimum of disruption. In time, we believe that relational database technology will evolve to encompass knowledge bases and that the well-established data-modelling techniques that ensure consistency in databases will be extended to knowledge modelling as well.

Those who will develop and program the expert systems will need to be trained to use the tools and techniques that go with the products you have selected. Most of this training will be available from the product suppliers.

IMPLEMENTING THE FIRST APPLICATION

We recommend that, wherever possible, you omit the experimental and pilot stages and develop an operational system as quickly as possible. The experimental stage is necessary only where a totally new product, using new techniques, is being tried out, or where the organisation is pioneering the use of expert systems in its sector. Experimental projects should be designed to increase understanding and to try out new tools and techniques.

Pilot projects are necessary only when certain questions can be answered by running a trial version of the system. Everyone involved in such a project should be aware that if it meets certain predefined criteria, a full operational system will be developed.

For the small percentage of Foundation members that have not yet started on expert systems, we paraphrase the advice for beginning your first project given in 'A Strategic Guide to Implementing an Expert System' (published by systems dynamics limited, a UK consultancy):

- Your first expert system should not be a large, high-risk project.
- Do not raise your organisation's expectations too high for the first project.
- Expect the project to take longer than you originally thought.
- Do not invest in specialised machines or expensive software tools until you have sufficient experience with expert systems to know what they can do for your organisation.

We also commend the advice regarding the development of initial expert systems given by Peter Sell of Digital at a meeting of UK Foundation members:

- Management, not technicians, should be responsible for the project.
- Start with the problem, not the solution.
- Find a real business problem to solve, rather than one that is of interest to the systems department.
- Start with the least expensive usable software.

- Start by developing a prototype rather than a functional specification.
- Involve all the appropriate technicians, experts, users, and management in the project.
- The initial prototype will demonstrate the feasibility of using expert systems to solve the particular problem; it will usually have to be extended or rewritten to provide the complete solution.
- After the initial prototype project, assess the organisational effect and the resources needed to develop the full solution, and then enhance the prototype solution.
- Schedule your first project to last for between one and three months, not a week or a year.
- Measure progress on your initial project in terms of the functionality modelled, rather than by progress through formal checkpoints.
- Do not release the prototype system for general use.

The need to involve everyone who has something to contribute, or has a vested interest, was underlined in a presentation at the annual UK Foundation Conference in 1987 by Anthony Butler, chairman of the UK Alvey expert systems project for insurance applications. He said that the development group had concentrated on the needs of the experts and had only belatedly realised that the users of the system would be branch-office staff, not the experts, and that the requirements of branch-office staff had not been taken into account.

It is also important to ensure that the intellectual property rights of the knowledge base and inferencing rules of any expert system developed are protected so they cannot be used by other organisations (perhaps competitors) without your permission. We know of several instances where a software house has been employed to develop an expert system and has, quite legitimately, claimed the rights to the knowledge base and inference rules and has included them in an expert systems application package. You should therefore ensure that the intellectual property rights are protected in the contract with the software house.

MONITOR FUTURE DEVELOPMENTS

The field of expert systems is evolving rapidly, and the techniques and tools available are continually changing. To ensure that you are not left with yesterday's best solution, your expert systems support unit should monitor the developments that are occurring not only in the tools, but also in the application of expert systems. One application that is likely to grow in importance is that of intelligent decision-support systems.

Two researchers (Remus and Kotteman) are beginning to consider how expert systems might be used to create the intelligent decision-support system of the future. They believe that within ten years developments in AI and expert systems will provide tools that will overcome the main difficulties with today's decision-support tools, which require the user to have an understanding of the technology being used and are not designed to solve specific problems. Furthermore, the researchers argue that expert systems will be required in order to cope with the two most important characteristics of decision making:

- Decision making involves selecting appropriate decision-making tools and approaches, and obtaining and filtering the information available.
- Decision makers are subject to numerous cognitive limitations.

We believe that Remus and Kotteman are too pessimistic in their view that it will be ten years before such tools are available. Our view is that intelligent decision-support systems, based on expert systems, will be available within five years.

The most significant technical developments that are likely to occur in the immediate future are described in Appendix 2, but the key ones to watch for are:

- The availability of expert system tools that integrate with ordinary software, including transaction-processing monitors, database software, and system development tools.
- The availability of automated knowledge-engineering methods and tools.
- The emergence of compatible software that can be used across the whole range of hardware — microcomputers, minicomputers, and mainframes.

In the scientific/technical area, the integration of expert system tools into process-control equipment will come about as much because of users experimenting as because of the original equipment supplier agreeing to accept and incorporate such tools in its equipment. The exploitation of expert systems in this way parallels that of the microchip, which is now used extensively in electrical appliances, motor cars, and industrial equipment. In a similar way, the use of expert system tools in the engineering-design activity and in product diagnosis and repair will need the direct and active commitment of the respective professional and line functions.

In the decision-support and end-user computing area, the IT suppliers will be instrumental in providing integrated 'smart' tools — such as intelligent modelling software. However, the key developments described above will have the most impact in the general information systems area — the area in which the systems function has a leading role to play in the exploitation of expert systems.

Chapter 6

Selecting and using expert system products and tools

As explained in Chapter 1, one of the reasons for the recent expansion in the application of expert systems has been the emergence of development products and tools. They have enabled applications to be developed economically and quickly without the need for specialist skills. The technology is still immature, however, and it is evolving rapidly. Over half the products offered by the suppliers in our survey are less than three years old — which is not surprising for a new market segment.

Some of the users we interviewed were reluctant to invest in the existing products because of their short expected lifespan. But if it is worth developing expert system applications, then the use of tools is essential to contain the development cost and timescale. Short product lifespans of IT products is not a problem confined to expert systems.

It is necessary to select the kinds of tool appropriate to the application and to the skills of the development team. You also need to bear in mind the limitations of the current products and tools and the period of time before those limitations are likely to be overcome. To help in selecting the most appropriate tools, we describe below the characteristics of the various kinds of products and tools available, summarise the limitations of the current technology, and show how the different kinds of tool are suited to the various application areas. Appendix 2 contains more technical discussions of the state of the art in expert system products and tools.

THE PRODUCTS AND TOOLS AVAILABLE

The software products and tools available to the expert systems developer range from programming languages through to specific application packages. They differ in the kinds of skill (and hence the type of developer) required to use them and in the range of applications for which they may be used. Figure 6.1 overleaf shows the six main kinds of expert systems software available, classified according to these two parameters.

The most basic types of tool are *programming languages* such as Lisp and Prolog. They can be used to solve any kind of problem, but are suitable for use only by specialists.

Expert system toolkits or *development environments,* such as ART, KEE, and Knowledge Craft, can also solve a wide range of problems, but can be used by a wider range of developers. Products like these require special-purpose hardware but are very sophisticated and powerful. Increasingly, new toolkits are becoming available for use with conventional hardware (many on microcomputers, others on minicomputers, and some on mainframes). Many of these products are just as sophisticated as those requiring special-purpose hardware, but their operational performance is not as good.

With languages or the simpler toolkits, the developer has to construct the expert systems framework. The need to do this has been removed by the emergence of *general-purpose expert system shells*, which provide a ready-made framework that enables applications to be developed relatively quickly and with less effort. It was the advent of microcomputer-based shells that gave expert systems its great boost between 1984 and today. Many of the second-generation shells available now are more powerful than the first-generation application environments available three years ago. With the increasing use of Intel 80386 chips in microcomputers, the performance of shells will continue to improve. Shells have a narrower range of application than languages or toolkits, but the more sophisticated toolkits also incorporate shells.

Specialised shells are designed to cater for a specific class of applications or for a specific domain of expertise. An example is Q-Shell, developed by Coopers & Lybrand as an adjunct to Expertax. Q-Shell's user interface is based on a questionnaire technique that is used both for developing and for using the subsequent expert systems application.

Induction-logic tools work on the premise that the problem-solving rules can be induced from the existing case material, instead of having to go through the extremely difficult and tedious process of extracting knowledge from the domain experts.

Figure 6.1 Types of expert systems software

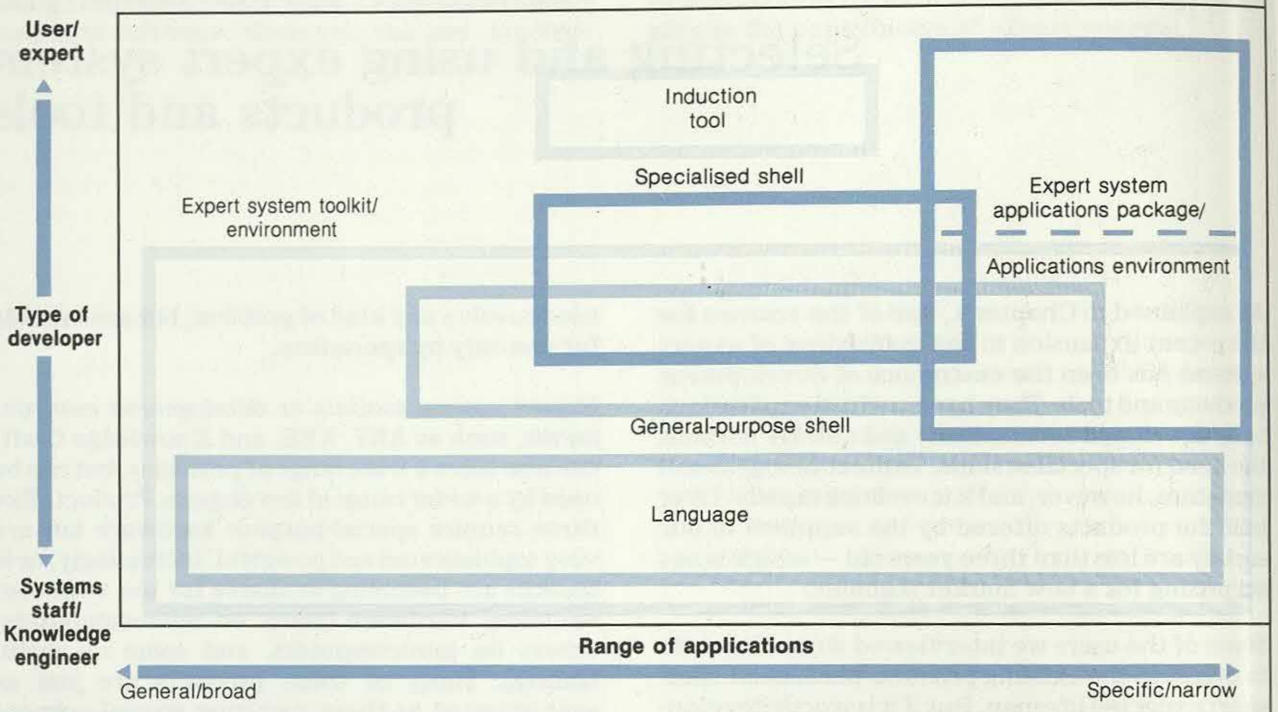

The only limitations are that there must be a sufficient number of cases to represent all the conditions and that the case material must contain all the relevant factors in order to allow the correct and complete rules to be induced. Where this case material is inadequate, many developers use induction tools as part of the knowledge-acquisition process rather than as the sole basis of the application system. The most prominent example of an induction tool is Expert-Ease (and its successor Extran) developed by Professor Michie and marketed by Intelligent Terminals.

The final kind of expert system products are *expert system application packages* and *application environments*. Expert system application packages include a basic knowledge base for the subject domain, and, possibly, a set of inference rules already developed for another organisation. These products are usually based on an expert systems tool that provides all the user-friendly features (such as WIMPS) that are now an essential part of expert system products.

An application environment is a cross between a development environment and an application package. An example of this type of product is Parys from Business Information Techniques in the United Kingdom. Parys was designed originally for the development of applications in the personnel and administration areas. It now provides a series of applications-oriented shells that can be used for drafting contracts, for analysing customer requirements, and for administering personnel. Parys also provides:

- A database package.
- A relational query language.
- A word processing package.
- A report generator.

Application packages and environments have emerged during the last two years. The suppliers who responded to our survey gave details of ten application-package products. We estimate that in the middle of 1987 there were some 50 such products based on expert systems software available in Europe, with more becoming available all the time. They are typically priced at between $10,000 and $100,000. These prices reflect market value rather than the cost of producing the packages.

LIMITATIONS OF CURRENT PRODUCTS AND TOOLS

Although today's expert system products and tools can be used to produce worthwhile applications, they do have some limitations. However, the field of expert systems is developing rapidly, and many of the limitations will be removed within the next three years or so. There are other limitations that will take longer to remove because they will require research breakthroughs or because AI researchers do not take account of the way in which expert systems will be used in a business-computing environment. The latter situation means that initial products and tools are aimed at specialists and specialist standalone applications, rather than at general-purpose computing. One particularly important limitation that will not be removed in the short term concerns the process of acquiring and maintaining the knowledge base on which the expert system is based.

KNOWLEDGE ACQUISITION AND REPRESENTATION DIFFICULTIES

Expert systems are often most needed where there is a scarcity of experts, which means that it is difficult for them to find the time for their expertise to be acquired in a form that can be used by an expert system. Moreover, different experts on the same subject may disagree with each other. And sometimes there are disagreements between head-office experts and practitioners in the field. Some expert system products now provide facilities such as blackboarding that allow experts with different points of view to share their knowledge. (Blackboarding is described in more detail in Appendix 2.)

Another difficulty can be caused because experts are unwilling to cooperate in the development of an expert system because they fear that the resulting system will degrade, or even eliminate, their jobs. Even if an expert is willing to contribute to the development of an expert system, he or she may be reluctant to assist in maintaining and modifying the system.

Another factor to consider when selecting expert system tools is the way in which they represent the knowledge acquired from the experts. Different knowledge-representation techniques are suitable for different types of problem. The two most common methods are *production rules* and *frames*. In the production-rules method, the domain knowledge is represented as a set of rules, typically in the form:

If A or B and C . . . then D. . . .

This format is similar to that used in many fourth-generation languages. Whilst the knowledge base is more easily read and understood in this form than in the frames method, it is very tedious to prepare the knowledge base using production rules. It is akin to writing programs before subroutines, macros, and data divisions were invented. Production rules are more suitable for expert systems where the knowledge domain is narrow and where the facts are largely discrete and do not fall into a hierarchy of classes.

By contrast, the frames method is a more structured and useful way of describing the knowledge; it enables the properties both of knowledge entities and of logic procedures to be held with those entities. The frames method of representing knowledge is more suitable where the facts fall into one or more hierarchies of classes, such as in the scientific classification schemes. The method can be thought of as providing a frame of reference that allows the properties of each logical entity to be stored in 'slots' in the frame. It also allows properties to be transferred or inherited from a higher member of a class of entities. For example, cars may inherit the properties common to vehicles (a means of propulsion, a steering mechanism, a load-carrying capacity, and so on). When a new type of vehicle is added to the database it will automatically inherit the properties common to vehicles, but not those specific to cars.

Logic procedures can also be stored in slots of the frame. These procedures will be invoked when certain conditions are triggered. Other slots may contain pointers or relationships to other frames. As anyone with data-modelling or entity-modelling backgrounds can see, the ideas behind the frame method are very similar to the ideas behind entity modelling.

Indeed, knowledge engineering and knowledge bases are sometimes perceived as a natural extension of data modelling and relational databases. However, the expert systems field is at the same stage of development that the database field was in the early 1970s. Today, the database field is changing rapidly from being a craft to being based on science and engineering. There are now well-established theories and methodologies for data modelling and databases. Equivalent theories and methodologies for knowledge modelling and knowledge bases do not yet exist, although several research projects are currently investigating these areas. The most notable of these is an Esprit project (known as KADS) that involves STC and Scicon from the United Kingdom, SCS from Germany, and the University of Amsterdam. The project is concerned with knowledge acquisition and documentation and is described in Appendix 2 on pages 49 to 50.

In general, the slow but steady convergence between expert systems and relational database technology may speed up developments in acquiring and storing knowledge. We also cover this topic in more detail in Appendix 2.

LIMITED EXPLANATION CAPABILITY

Many of today's products have a limited explanation facility — even though this is often claimed to be a characteristic feature of expert systems. In most of today's expert systems, the explanation facility displays the logic rules (or equivalent) that were used to reach the result. If the rules were written in a 'natural language' then the explanation is relatively clear. But if the rules comprise a string of codes and symbols, a typical user would not be able to make much sense of the explanation.

MOST LIMITATIONS WILL BE REMOVED DURING THE NEXT THREE YEARS

Most of today's expert system products and tools cannot link to existing data, software, and applications. Although full integration of expert systems with mainstream data processing will take longer than three years to achieve, gateway interfaces are already available in some products and tools, and should be widely available by the end of the 1980s. We discuss the prospects for integrating expert systems with other types of information processing systems more fully in Appendix 2.

The following limitations are likely to be overcome within the next three years:

- The processing power of today's microcomputers limits the size and scope of microcomputer-based expert systems, and therefore makes it impossible to upgrade some experimental systems to operational systems. The increasing availability of hardware based on the new Intel 80386 chip will ease this limitation, allowing larger and more difficult problems to be tackled with microcomputer-based systems.
- Some applications are useful only if multiple users in the same office can share them on the same microcomputer, which is not currently possible with most of the systems. Also, most of the shells and toolkits allow an application to access only one knowledge base.
- Many of the advanced development tools can be used only on specialised AI hardware, as can the operational versions of the resulting expert systems. Often, it is not practical to use expensive specialised hardware — for example, where the system is to be distributed widely throughout the organisation. In this situation, it should be possible to run the expert systems application on a dumb terminal linked to a standard mainframe or minicomputer, or on a standard business microcomputer.
- There will be an increasing number of products and tools that can be used over the full range of hardware (microcomputers to mainframes). These products will be particularly useful where large-scale and small-scale versions of the same basic expert system are required in different parts of the organisation. An example of a product that works both on IBM mainframes and on the PC is AION. It was developed by ex-IBM employees who became impatient with the lack of progress of their work within IBM. The work they carried out in IBM is now being marketed by IBM as ESE.

SELECT PRODUCTS AND TOOLS TO MATCH THE TYPE OF APPLICATION

Earlier in the report, we identified four expert system applications areas (see Figure 3.1). Applications in each area are likely to be developed by different categories of staff, who will require different types of products and tools. The resulting applications will be used in different environments and in different ways, and the importance of the operational system's user interface will therefore vary, depending on the type of application. To a large extent, the quality of the user interface (or the ergonomics of the operational system) will be determined by the products and tools used to develop the system. It is therefore vital to select the most appropriate products and tools for the particular application type, for the staff who will develop the system, and for the quality of the user interface that will be required. Figure 6.2 summarises these parameters for each of the four application areas. They are discussed in more detail below.

INFORMATION SYSTEM APPLICATIONS IN THE COMMERCIAL/ADMINISTRATIVE AREA

Information system applications in the commercial/administrative area require development tools that can at least interface, and ideally integrate, with conventional software. You should therefore select products that are capable of interfacing or integrating with your teleprocessing monitor, data dictionary, and database management systems and with programs written in ordinary languages such as Cobol and PL/1. The products should be part of a comprehensive toolkit so that, for example, a prototype developed using an expert systems shell can easily be rewritten in a lower-level language such as C (in a Unix environment) or Lisp or Prolog. In some cases, the operational application will contain a mixture of shell-based code and lower-level code, similar to the way fourth-generation languages are used in many installations today. Eventually, as expert systems are integrated more with mainstream data processing, a complete system could include expert system subsystems, as well as others written in fourth-generation and older languages.

Many of your more important expert system applications will be run on standard data processing equipment, not on specialist AI hardware. Some organisations will be able to justify using AI hardware for experimental purposes, or even as development tools. But few organisations will be able to justify using specialist AI hardware for systems, especially if the application is to be distributed throughout the organisation. The expert systems software should therefore be able to run on the complete range of standard hardware.

The staff involved in developing this type of expert system include professional developers from the

Figure 6.2 Expert system tools and development staff for each application area

Application area	Tools required	Development staff	Importance of user interface*
Commercial/ administrative			
Information systems	Tools that interface/integrate with mainstream tools. Toolkit/environment/required. Micro/mini/mainframe portability required.	Expert/specialist User Systems staff	Important
Decision support or end-user applications	Microcomputer-based shells. Uncertainty, fuzzy logic tools. Induction tools	Expert Systems staff (as advisor)	Useful
Scientific/technical			
Online/realtime	Specialist tools.	Plant supplier Plant operator In-house expert Systems staff	Crucial
Offline/standalone	Microcomputer-based tools with fault diagnosis. Specialist workstations for design applications.	Repairer or equipment designer In-house expert Systems staff	Important

*The importance of the user interface to some extent determines the tools required.

systems department and the experts whose knowledge is being captured by the expert system. The eventual user of the expert system also needs to be involved at the requirements-definition stage because the experts are seldom good judges of what the end users really need. Many organisations that have developed expert systems in this area are looking for a complete methodology that starts with defining the requirements, continues through acquiring the knowledge about the domain, and finishes by supporting the design process. Such a methodology is needed, but it is premature to specify one at the current stage of development of expert systems. However, the KADS project (described in Appendix 2) is aiming to develop a knowledge-engineering methodology.

The user-interface requirements of expert systems in the commercial/administrative area are no different from those for ordinary mainstream information systems.

DECISION-SUPPORT OR END-USER COMPUTING APPLICATIONS IN THE COMMERCIAL/ ADMINISTRATIVE AREA

Expert system applications in the decision-support or end-user computing area are the least demanding in terms of the products and tools required to develop the system. Usually, microcomputer-based expert system shells are adequate. However, if the subject area and available case material are suitable, tools for inducing the knowledge-base rules can be useful, as can tools for dealing with uncertainty factors or fuzzy logic. (Fuzzy logic is the application of mathematical techniques to quantify the qualitative aspects of human judgement. Fuzzy logic converts relative factors to quantifiable ones — for example, converting 'very high', 'high', 'moderate', 'low', and 'very low' interest rates to actual percentage points.) Facilities for downloading data from mainstream systems for use by the expert system can also be valuable.

One example of an expert systems product designed for use in the decision-support area is ICL's Reveal. This product combines financial and other modelling techniques with fuzzy logic. One of the pioneers in the implementation of fuzzy logic is Peter Jones, the developer of Reveal, which was originally available in the United Kingdom on IBM mainframes through Tymshare. Reveal is now also available on ICL mainframes. Peter Jones has formed his own company, Creative Logic, and has developed the Leonardo family of products. Leonardo 3 incorporates fuzzy logic and other uncertainty factors. (A fuller explanation of these techniques can be found in the earlier Foundation Report, No 37, on expert systems.)

Expert systems in the decision-support area are often developed (and used) by the expert. Alternatively, the system may be developed by a departmental colleague of the expert, in much the same way as many other decision-support or end-user computing applications are developed by the 'local' end-user computing expert. However, someone from the systems department may be involved in the development in an advisory capacity.

Again, the user interface requirements are no different from other decision-support or end-user computing applications.

ONLINE/REALTIME APPLICATIONS IN THE SCIENTIFIC/TECHNICAL AREA

Online/realtime expert systems in the scientific/technical area require special development tools, often associated with specialist data-acquisition devices (such as sensors and transducers) to link the expert system to process-control machinery. More often than not, the operating environment is hostile, with high temperatures, humidity, and dust or other contaminants in the air. The staff involved in developing expert systems in this area include in-house experts, developers from the systems department, staff who operate the process-control equipment, and representatives of the equipment suppliers.

The design of the human interface is a crucial aspect of this type of expert systems application. In the example of the sinter plants at Broken Hill, the initial design was extremely sophisticated, with an Apple Macintosh being used to display only the data that was relevant to the particular situation. However, the operators found the presentation difficult to assimilate and use. They felt that vital information was missing, even though in fact they were being presented with all the relevant data and with no extraneous data. The system was redesigned to mimic the strip-charts previously used, but with additional pointers to indicate critical combinations of threshold values. The company says that it will probably reincorporate some of the more sophisticated features in the next version of the system, once the operators have become accustomed to the new method of presentation.

OFFLINE/STANDALONE APPLICATIONS IN THE SCIENTIFIC/TECHNICAL AREA

Offline/standalone expert systems in the scientific/technical area will require the same kind of tools as online/realtime applications if the main focus of the application is to design a product and to link the automatic-design system with an automated production system.

However, if the application is concerned with fault diagnosis and field maintenance, the requirements are less stringent. The main requirement is to develop expert systems that are portable and robust and that can run on inexpensive hardware. The staff involved in the development of expert systems in this area include in-house experts, staff from the systems department, and the designer of the equipment being designed or maintained.

A good user interface is obviously important for design applications. It is also important for maintenance applications, where a good user interface will facilitate fault diagnosis and maintenance.

IN CONCLUSION

In our 1983 report on expert systems, our advice was that the technology was not yet ready for widespread commercial implementation. That situation has now changed. In the past four years the technology has matured: expert systems are now an essential technology in some business sectors, and will soon become so in others.

In this report we have shown how expert systems are being used for a wide range of applications by most kinds of business. We have identified the benefits that can be gained from expert systems and provided advice on how to select and prioritise potential opportunities for exploiting expert systems. We have shown how an expert systems support unit can be set up, and have explained what its role is. The final chapter of the report has provided guidelines for selecting expert system techniques and tools.

The underlying message of the report is that expert systems are fast becoming just another facet of business data processing. Increasingly, expert systems will not require specialised and esoteric hardware; they will run on conventional computing equipment ranging from microcomputers to mainframes. Furthermore, expert system techniques are converging rapidly with mainstream data processing, particularly with relational database techniques.

Expert systems are set to become an essential part of mainstream business data processing. If you have not yet started to use expert systems for operational applications, now is the time to do so.

Appendix 1
Research base and findings

The research carried out for this report covered users and suppliers of expert systems in Europe, in the United States, and in the Asia-Pacific region, including Japan. In addition to interviewing users and industry experts, we also surveyed suppliers of expert system products in Europe. Finally, we made use of several other researchers' reports and reviewed a significant portion of the considerable volume of published literature. (The bibliography lists those publications and articles that particularly influenced our thinking.)

Figure A1.1 summarises the user and supplier research work carried out. The focus group discussions were held in France, in the Netherlands (attended also by representatives of Belgian organisations), and in the United Kingdom. The results of the supplier survey are set out later in this appendix.

NON-EUROPEAN RESEARCH

Although the bulk of the user research was carried out among European organisations, we also thoroughly researched the situation in the United States, particularly amongst suppliers and industry experts. We met with Professor Feigenbaum of Stanford University, who started work on expert systems nearly 20 years ago and who is regarded by many as the 'father' of expert systems, and with representatives from two other well-known expert systems research establishments — the Rand Corporation and Carnegie-Mellon University (where XCON was developed). We also talked with Microelectronics and Computer Technology Corporation (MCC), the research consortium formed by major US vendors.

We sought the views of US vendors that specialise in AI and expert system products, including Symbolics, Inference Corporation, AI Corporation (who developed Intellect, the AI-based natural-language interface to mainframe databases), Teknowledge, Neuron Data (whose products run on the Apple Macintosh and IBM PC), and Carnegie Group, where John McDermott, the developer of XCON, now works. We also talked with Cigna and Athena, both of whom have developed successful commercial products; with Xerox Parc, where much of the early work on windowing techniques and the use of a mouse was carried out and where Smalltalk and the InterLisp specialist hardware were developed; with Relational Technology, to seek that company's views on the convergence of database and knowledge-base technology; and with Cullinet, where we explored the convergence of expert systems software with conventional software.

We also visited ICOT in Japan to hear at first hand about developments in the fifth-generation

Figure A1.1 Summary of user and supplier research carried out

Type of research	European users	Non-European users	Total users	European suppliers/experts*	Non-European suppliers/experts	Total suppliers/experts*
Face-to-face interview	23	2	25	3	12	15
Telephone interviews	32	1	33	6	3	9
Focus-group discussions	14	—	14	2	—	2
Desk research cases	30	20	50	—	—	—
Total	**99**	**23**	**122**	**11**	**15**	**26**

*Excluding the 20 suppliers that responded to our supplier questionnaire survey

Figure A1.2 Results of three surveys of the state of development of expert systems

Butler Cox survey of 104 Foundation members, covering 155 applications (January/February 1987)

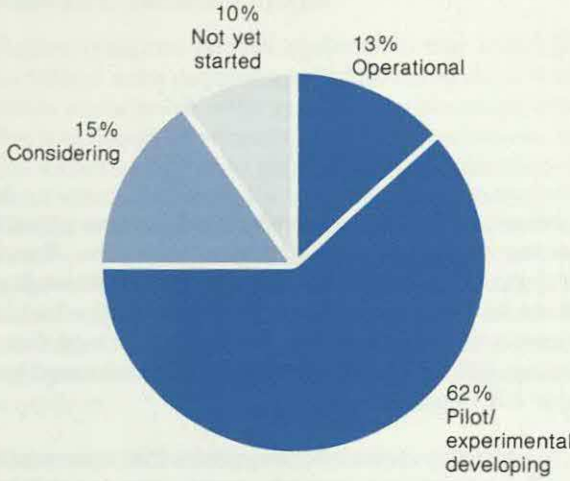

Pactel survey of 257 UK companies (March 1987)

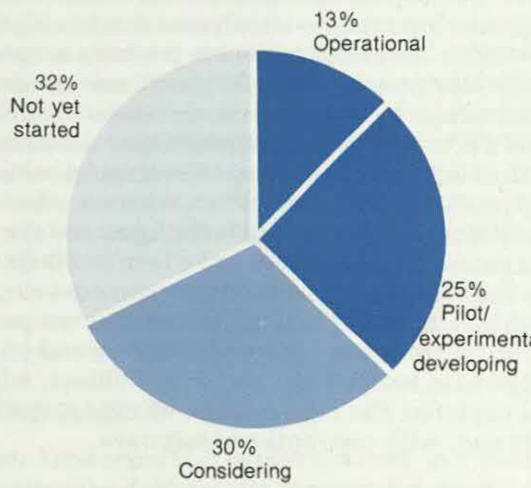

JIPDEC survey of 203 Japanese companies, covering 235 applications (April 1986)

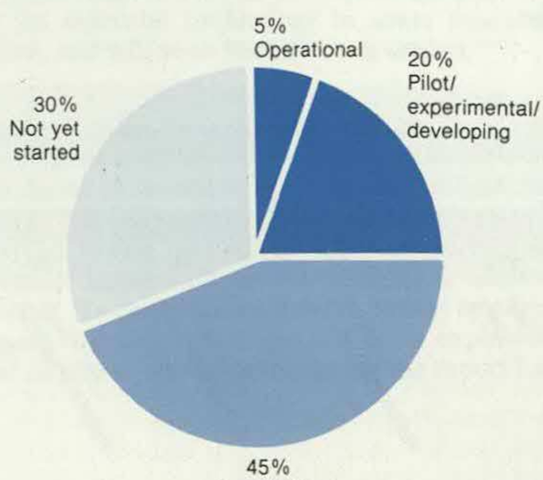

research programme. In Japan, we were given access to two Japanese surveys (one carried out by ICOT, the other by JIPDEC) about the use of, and attitudes towards, expert systems.

MEMBERS' RESPONSES TO INITIAL QUESTIONNAIRE

At the beginning of the research, we sent a short questionnaire to all Foundation members, seeking their views about expert systems and their involvement with them. We received 104 replies, several of which included copies of in-house material that provided very useful input to our research.

Figure A1.2 shows the status of expert system activities amongst the 104 members that responded. The figure shows that, early in 1987, 75 per cent were already actively involved with expert systems. For comparison purposes, Figure A1.2 also shows the results of two other surveys — one carried out in March 1987 in the United Kingdom by Pactel and *Business Computing and Communications* magazine, and the other in April 1986 in Japan by JIPDEC. The results of both of these other surveys are broadly similar, but they show that the companies surveyed are not as advanced in using expert systems as are Foundation members (only 38 per cent in the Pactel survey were actively

Figure A1.3 Expert systems hardware and software used by 104 Foundation members

Hardware

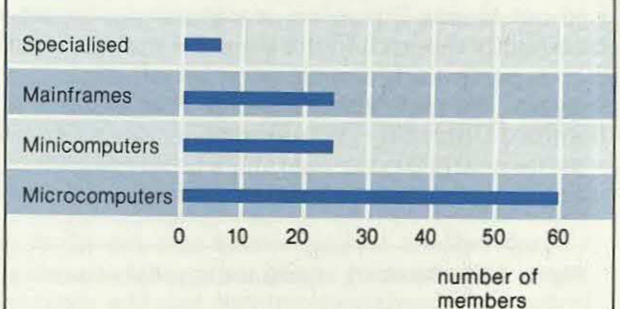

Software

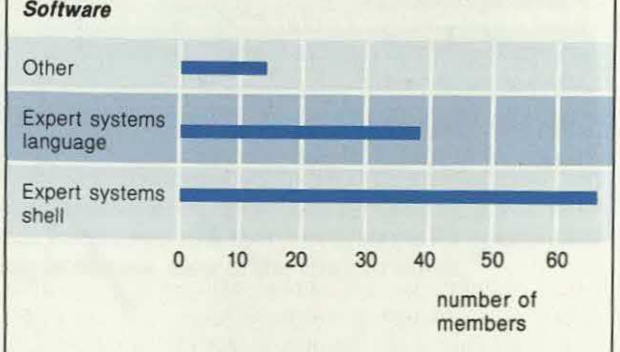

Note: Some members reported the use of more than one type of hardware and/or software

Figure A1.4 Expected impact of expert systems on the business in the next three years

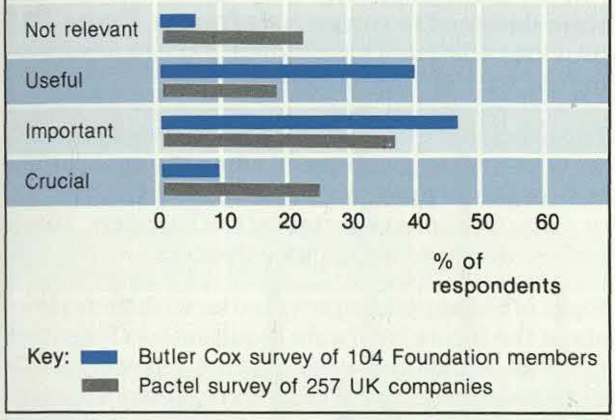

involved). However, both surveys included a higher proportion of smaller companies than did our Foundation member survey. Also, the Japanese survey was carried out nearly a year before our survey, and we would not be at all surprised to find that the situation has since changed in Japan.

Figure A1.3 shows the type of hardware and software being used for expert systems by Foundation members. More than half of those with expert system projects are using expert system shells and microcomputers. Specialised AI hardware is being used by less than six per cent of the respondents.

Figure A1.4 shows the perceived impact of expert systems reported by the respondents both to our survey and to the Pactel survey. We were surprised to see that only nine per cent of Foundation respondents thought that expert systems would be crucial to the business in the next three years, and that five per cent thought they would not be relevant at all. By contrast, 24 per cent of the respondents to the Pactel survey thought that expert systems would be vital to their business, although 22 per cent thought they would have no impact at all. It is interesting to compare these findings with the results of our supplier survey (see below), where 16 per cent of the respondents believed that expert systems would be crucial to their customers' businesses and only six per cent said that expert systems would have no impact.

EXPERT SYSTEM SUPPLIERS AND PRODUCTS IN EUROPE

Between February and April 1987, we sent a questionnaire to 80 organisations that supply expert system products in Europe. The 20 organisations that replied are listed in Figure A1.5, together with their addresses and telephone numbers. The majority of them (11) had been in business for

Figure A1.5 Expert system suppliers that responded to our survey

France

Cognitech 167 Rue de Chevaleret, 75013 Paris 1/45.83.73.00

Neuron Data 97 Rue d'Areau, 75014 Paris 1/45.89.81.43

Germany

Brainware GmbH Voltastrasse 5, 1000 Berlin 65 06121-372011

IntelliCorp GmbH Rosenheimerstrasse 143a, D-8000 München 80 089-414361

InterFace Computer GmbH Garmischerstrasse 4/V, D-8000 München 2 08951-0860

The Netherlands

Lithp Systems BV PO Box 65, 1200 AB Landsmeer 029-084623

United Kingdom

Artifical Intelligence Ltd Intelligence House, 62-78 Merton Road, Watford WD1 7BY 0923-47707

Business Information Techniques 20-26 Campus Road, Bradford BD7 1HR 0274-736766

Computer Research Systems Ltd 5 Bridge Street, Bishop's Stortford, Herts CM23 2JU 0279-506717

Creative Logic Ltd Brunel Science Park, Kingston Lane, Uxbridge, Middlesex UB8 3PQ 0895-70244

Expert Systems International 9 Westway, Oxford OX2 0JB 0865-242206

ICL Arndale House, Arndale Centre, Manchester M4 3AR 061-833-9111

Intelligent Applications Ltd Kirkton Business Centre, Kirk Lane, Livingston Village, West Lothian EH54 7AY 0506-410242

Intelligent Environments Ltd Northumberland House, 15-19 Petersham Road, Richmond, Surrey TW10 6TP 01-940-6333

Intelligent Terminals Ltd George House, 36 North Hanover Street, Glasgow G1 2AD 041-552-1353

ISI 11 Oakdene Road, Redhill, Surrey RH1 6BT 0737-71327

Software Architecture and Engineering Sussex Suite, City Gates, 2-4 Southgate, Chichester, West Sussex PO19 2DJ 0243-789310

Systems Designers plc Pembroke House, Pembroke Broadway, Camberley, Surrey GU15 3XD 0276-686200

Telecomputing 244 Barners Road , Oxford OX4 3RW 0865-777755

The Vanilla Flavor Company 6 St Clements Street, Winchester SO23 8HN 0962-68428 between two and five years, and five of them had been in business for five years or more. Twelve of them employed less than 20 professional staff and only one employed more than 80. Thus, they are small businesses by general standards, but about the average size for software houses.

Excluding one supplier that also provided hardware, just over half of their aggregate revenue came from software sales. The overall breakdown of revenue sources was:

Software sales	55 per cent
Applications development	25 per cent
Consultancy, training, and other	20 per cent

Thus, users are spending almost as much on services as on products, underlining the shortage of expert system skills in the user community.

PRODUCTS SUPPLIED

Between them, the 20 suppliers provided 79 different products, 35 of which were expert system shells (see Figure A1.6). Expert system application packages have been introduced only recently, and most of the products are less than a year old. They are usually introduced by generalising a bespoke system developed for a particular customer. Of the remaining products, nearly three-quarters are no more than three years old.

HARDWARE REQUIRED TO RUN PRODUCTS

Thirty-four of the 79 software products are available for microcomputers, mainly the IBM PC or compatibles. The next largest category of products were those designed to run on specialist workstations, including Unix-based technical workstations such as the Sun, although there were only slightly more of these than products designed to run on minicomputers. Only six of the products were designed to run on mainframes. Figure A1.7 provides a detailed breakdown of the hardware required for the 79 products.

Figure A1.6 Types of expert system products provided by suppliers responding to our survey

Type of product	Number of products	% of total
Expert system shell	35	44
AI/expert system languages	15	19
Expert system development environments	13	16
Expert system application packages	11	14
Expert system toolkits	3	4
Expert system utilities	2	3
Total	**79**	**100%**

The dominance of microcomputer-based products was even more pronounced when the applications base of the 79 products was analysed. Of the 4,000 or so applications reported by the suppliers, 80 per cent were based on microcomputers.

Eight of the suppliers provided us with their views about the future hardware requirements for expert systems. All of them expected more powerful PCs to be used, together with advanced workstations (like the Sun), rather than specialist Lisp machines. Products aimed at Lisp machines are expected to be phased out as the capability of Unix-based workstations increases. As with other PC-based tools, the suppliers expected there to be a need to link expert system products with systems running on mainframes and superminicomputers.

PRODUCT PRICES

It is not surprising that PC-based products are the most popular — they cost significantly less than the

Figure A1.7 Hardware required to run expert system products reported in our supplier survey

Type of hardware required		Number of products	% of total
Microcomputers		34	43
IBM/PC	31		
Apple Macintosh	2		
BBC	1		
Special workstation		17	21
Lisp	8		
Other	9		
Minicomputers		15	19
Vax	7		
Other	8		
Mainframes		6	8
ICL	3		
Other	3		
Other		7	9
Total		**79**	**100%**

Figure A1.8 Prices of 79 expert system products reported in our user survey

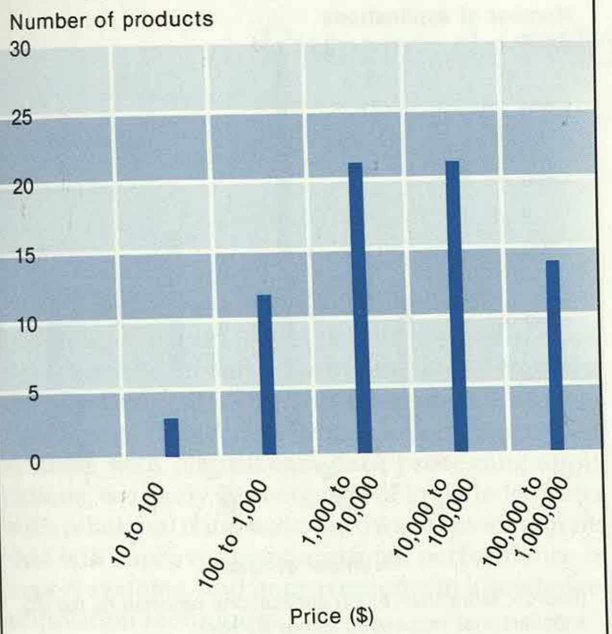

other types of product. Figure A1.8 shows the distribution of the prices for the 79 products reported to us. The products cover a wide range of prices, with most being in the range $5,000 to $100,000. Analysing the price of products by the type of hardware required to run them shows that:

- Forty per cent of PC-based products cost less than $1,000; none of the other products were that inexpensive.
- Lisp-based products cost $10,000 or more, as do half the mainframe products; only 10 per cent of PC-based products cost that much.

USER AND APPLICATIONS BASE

In total, the 20 suppliers had provided their products to more than 6,000 user organisations, although half these (mostly in the United States) were reported by one supplier. Only 4,400 application projects were reported, however, which seems to imply that some organisations purchase the products and do not use them. The reason for the discrepancy is that the users do not inform the suppliers of all of the uses being made of their products. Again, one supplier dominated, reporting 2,500 projects (also mostly in the United States).

The results of analysing the user organisations by industry sector are shown in Figure A1.9 overleaf, which shows the breakdown for all organisations and for those in Europe. The European percentage breakdown is very similar to that for all organisations.

We also analysed the applications using the Hayes-Roth classification of expert systems (diagnosis, design, advice, planning, monitoring, repair, prediction, and other). The results are shown in Figure A1.10 overleaf. The largest class is diagnosis systems, followed by design.

Figure A1.11 analyses the applications according to size, measured in terms of the number of rules (or equivalent). The majority of projects have between 500 and 1,000 rules, although there are more than 500 projects with at least 5,000 rules (however, two-thirds of these large projects were reported by just one supplier).

STAGE OF APPLICATIONS DEVELOPMENT

We asked the suppliers in our survey what stage of development applications built with their products had reached. Their responses corresponded with the data we gathered from user organisations and confirmed that most of the projects are at the development or pilot/experimental stage (see Figure A1.12). For the purpose of comparison, the figure sets out the equivalent Japanese data from the JIPDEC survey.

INTEGRATION WITH MAINSTREAM DATA PROCESSING SYSTEMS

Six of the 20 suppliers rated the need to integrate mainstream data with expert systems as crucially important to the successful operational use of expert systems. And seven of them believed that full implementation of major expert systems required significant organisational change. We asked the suppliers to classify the applications for which their products were being used according to their level of integration with mainstream data processing systems. The results are shown in Figure A1.13. A surprisingly high proportion (25 per cent) are already embedded in mainstream systems, and a further third have some type of link with mainstream systems.

Figure A1.9 Organisations using expert systems products, by industry sector

Industry sector	All organisations		European organisations	
	Number	% of total	Number	% of total
Education and research	1,200	20	390	24
Government (including defence)	1,000	17	270	16
Discrete manufacturing	700	12	180	11
Process manufacturing	450	7	120	7
Retail and distribution	340	6	90	5
Utilities	330	5	140	9
Transport and communications	280	5	60	4
Insurance	260	4	75	5
Banking and finance	250	4	80	5
Professional services	190	3	90	5
Health care	100	2	35	2
Other	920	15	110	7
Total	**6,020**	**100%**	**1,640**	**100%**

Figure A1.10 More than half of expert systems are used for diagnosis or design applications

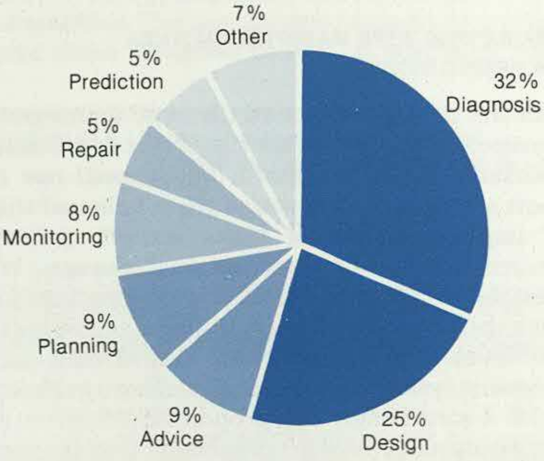

(Source: More than 4,000 applications notified by the 20 suppliers that responded to our survey)

Figure A1.11 Most expert systems have between 500 and 1,000 rules

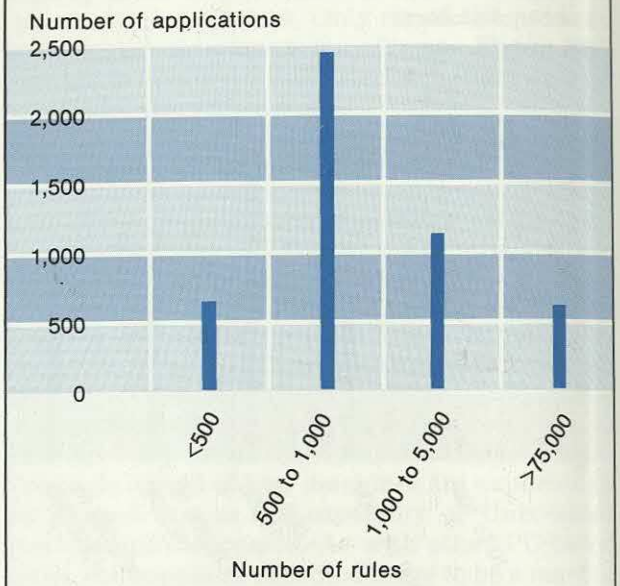

(Source: More than 5,000 applications reported by the 20 suppliers that responded to our survey)

Figure A1.12 Stages of development of expert system applications

Stage of development	Butler Cox survey		JIPDEC survey	
	Number of projects	% of total	Number of projects	% of total
Operational	234	11	27	12
Pilot/experimental	446	22	61	26
Development	1,260	62	144	61
Abandoned	96	5	3	1
Total	**2,036**	**100%**	**235**	**100%**

Figure A1.13 Integration of expert system applications with mainstream systems

Level of integration	Number of projects	% of total
Embedded in mainstream system	1,220	25
Data shared with mainstream system	330	7
Data transferred from/to mainstream system	900	19
Manual link with mainstream system	320	7
Standalone	2,050	42
Total	**4,820**	**100%**

Appendix 2

Future developments in expert systems

In this appendix, we highlight the developments in the field of expert systems that we believe will impact on the business use of expert systems over the next five years. The developments fall into four main areas: the trend towards integrating expert systems with mainstream data processing applications; the likely convergence of knowledge bases with relational databases; hardware developments that will improve the operational performance of expert systems; and improvements in knowledge-acquisition techniques.

INTEGRATION WITH MAINSTREAM DATA PROCESSING

The majority of the expert system products available today are not able to interface with conventional software. However, most of the suppliers that responded to our survey, and most of those we talked with, are aware of the need to integrate their products with mainstream software and applications and are working towards this goal.

With the exception of specialist hardware-based products (such as KEE, ART, and Knowledge Craft), the few products and suppliers that today provide some form of interface to mainstream systems include:

- Top-One, from Telecomputing.
- Guru, from MDBS.
- Parys, from Business Information Techniques.
- Leonardo, a new product from Creative Logic.
- Expert 2000, a recent joint development between Thorn EMI and Helix Expert Systems.

The products based on specialist hardware provide an expert systems environment together with a variety of tools. Their origins are firmly in the AI-research community, but they are being enhanced to provide more conventional software facilities and interfaces. For instance, there is now a version of KEE (called PC Host) that operates on a PC, and IntelliCorp (the suppliers of KEE) recently announced two new products: Connection, for downloading data from mainframe databases for use in an expert system; and Intelliscope, which provides an end-user front end to mainframe databases.

Thus, the general trend is for suppliers of specialist expert system products to evolve their products so they can interwork with mainstream data processing applications and software. This trend is mirrored by the established system and software houses, who realise that there will be a continuing need for expert system products and that expert systems offer them a business opportunity. For example, MSA, the world's biggest software house, recently declared that its future products will incorporate knowledge-based features. In particular, MSA's data dictionary product, called Information Expert (which is *not* an expert system) is being provided with expert system front ends for the built-in fourth-generation language.

Another example is Cullinet's Application Expert product, which is designed to be used for the rapid development of applications. The main features of Application Expert are that:

- It can be integrated with the current systems environment.
- Its applications can access and update conventional databases by using the EXL language.
- It provides access to the most popular database management systems on Vax computers and IBM mainframes.
- It is built to run in a transaction-processing environment.
- It allows expert system procedures to be embedded in conventional software.

Finally, we use Telecomputing's Top-One to illustrate the type of expert system features we believe will become commonplace within the next five years for general information systems applications, such as major transaction-processing and database-oriented systems.

In addition to providing features similar to those of Application Expert, Top-One can handle a

Figure A2.1 Telecomputing's Top-One handles a wide range of concurrent relationships

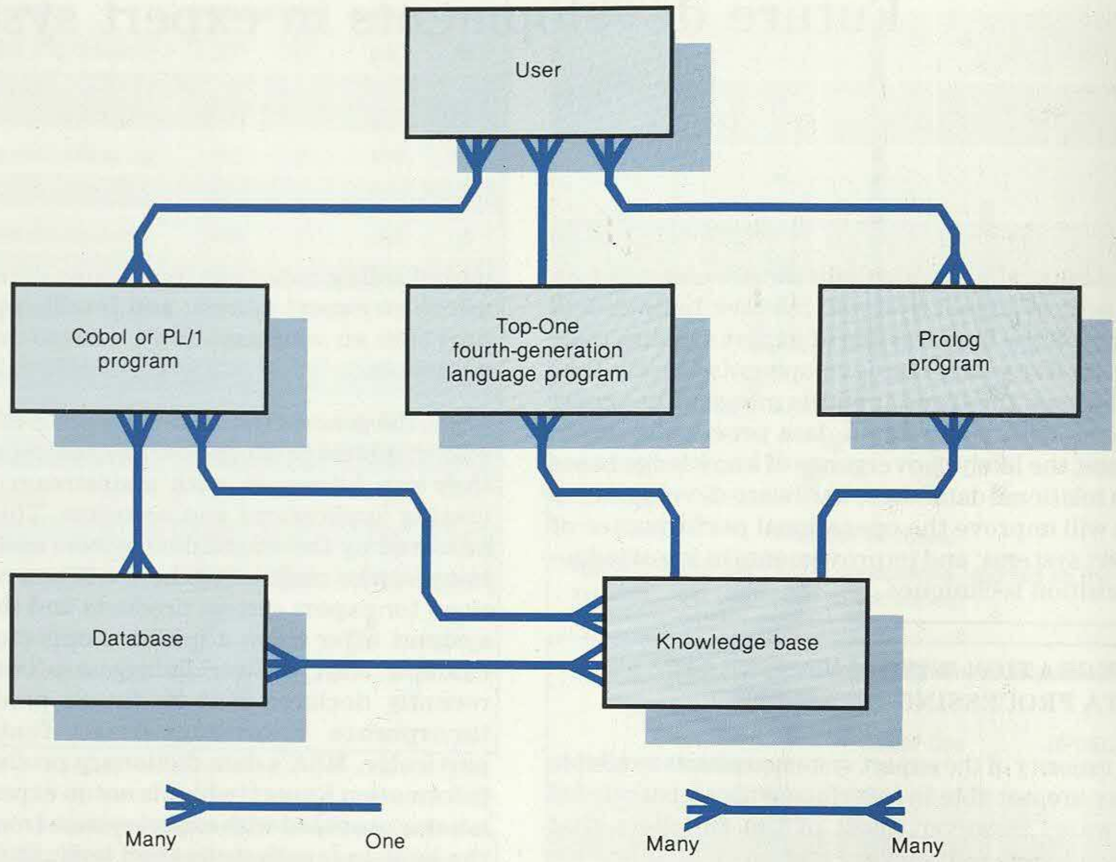

wide range of concurrent relationships (see Figure A2.1).

A user can access, through the same application system, any combination of a Cobol (or PL/1) program, a program written in Top-One's built-in fourth-generation language, or a Prolog program. Each of these programs can itself be accessed by many concurrent users. And each of the programs can access the same knowledge base, and the Cobol (or PL/1) program can concurrently access multiple knowledge bases and multiple conventional databases. Furthermore, a knowledge base can be accessed by more than one program of each of the three types. Finally, there can be many-to-many links between knowledge bases and conventional databases.

Top-One acts as the master control program and runs under CICS on IBM hardware or its own transaction-processing monitor (TPS) on ICL machines. The structure of an application system built with Top-One is shown in Figure A2.2 overleaf. The PACE (Programming and Consulting Environment) element shown in the figure is similar to an expert system shell. Top-One is designed with all the recovery and integrity features associated with modern transaction-processing database software. It also provides measures of the processing and inferencing loads and of the levels of access to the database and knowledge base, so that the developer can easily tune the application system, either by replacing PACE procedures with Prolog routines and fourth-generation routines with Cobol procedures, or by reorganising the database and the knowledge base.

At the time we spoke with Telecomputing (March 1987) there were four main users of Top-One:

— A property company had used Top-One to build an application to bill the tenants in a shopping mall.

— An engineering firm had used Top-One to archive the rarely used skills required for preparing certain international tenders.

— A public utility had developed a quantity-surveying application.

— An equipment-rental firm had built applications for calculating sales commission, for planning, and for modelling the future performance of the business.

Figure A2.2 Structure of an application built with Top-One

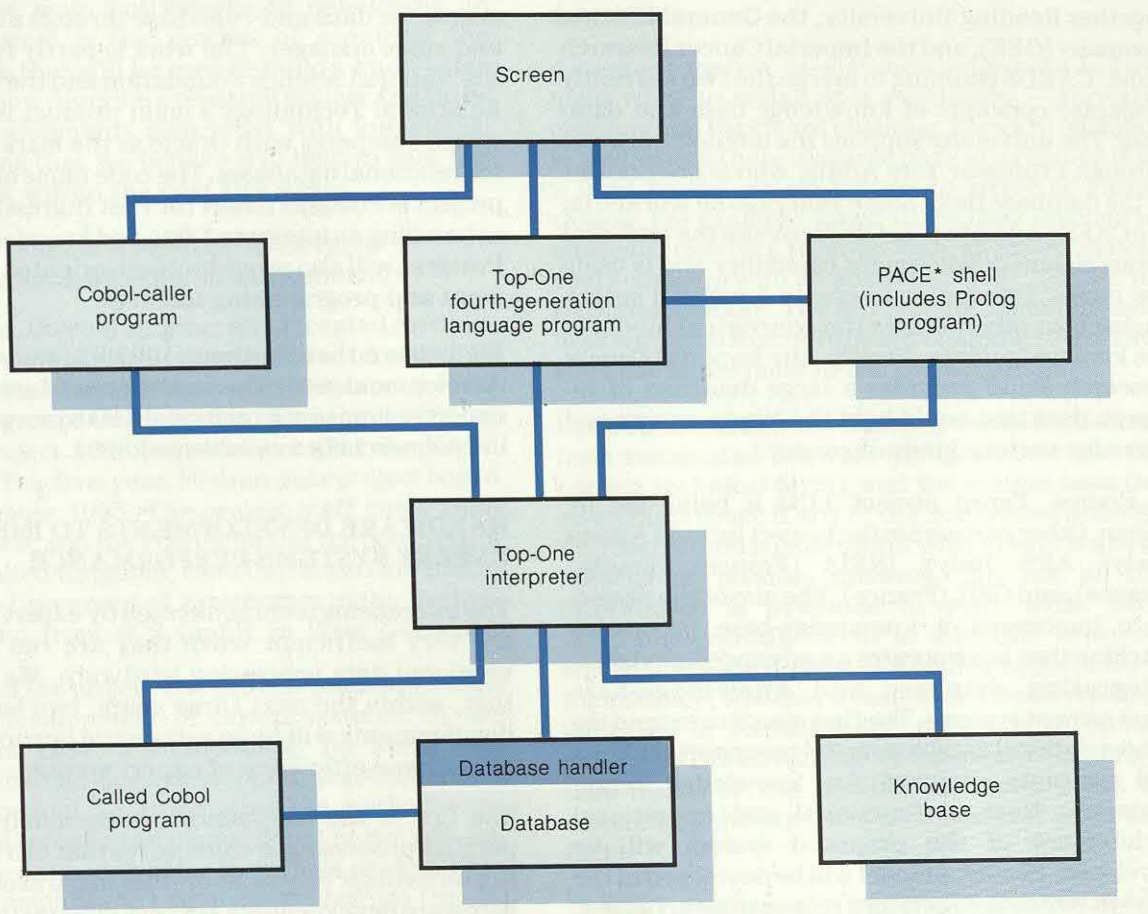

*PACE = Programming and Consulting Environment

We believe that future expert systems software designed for the development of information systems will look remarkably like Top-One.

CONVERGENCE OF RELATIONAL DATABASES AND KNOWLEDGE BASES

There is an obvious similarity between relational databases and knowledge bases. Advanced database systems include a data dictionary that contains a data schema or definitions of the data — comprising definitions of entities, attributes, and relationships. Expert systems include a knowledge base that defines facts (equivalent to entities and attributes) and the rules equivalent to relationships) for using the facts. Advanced dictionary products, such as ICL's DDS, contain a process model that describes the business-transaction process. In the future, 'knowledge-encyclopaedia' products are likely to contain a knowledge-base schema where the 'process' model describes the problem solving-process. Furthermore, the statements in a typical relational database-access language, such as SQL, resemble an expert system's rule.

The similarities are not surprising because both relational database theory and logic programming (the basis of rule-based expert systems) have their roots in predicate calculus, a branch of mathematics and logic. For a more learned and thorough treatment of the similarity between expert systems and relational databases, we refer you to a paper by Professor Nijssen of the University of Queensland, Australia. Professor Nijssen has for many years been deeply involved with defining standards for conceptual schema. In general, the academic world is showing increasing interest in this topic, and a recent issue of the British Computer Society's *Computer Bulletin* described several academic research projects on expert database systems. (The relevant papers are listed in the bibliography at the end of the report.)

There are also several government-sponsored projects investigating the relationships between expert systems and relational databases. In the United Kingdom, an Alvey project called CARDS (conceptual and relational database server) brings together Reading University, the General Electric Company (GEC), and the Imperial Cancer Research Fund. CARDS is aiming to merge the two currently disparate concepts of knowledge base and database. The university supplies the intellectual input through Professor Tom Addis, who is no stranger to the database field. Some years ago he worked on the CAFS project in ICL. GEC provides the technical research and development capability and is using the Inmos Transputer to create a parallel search engine that incorporates the conceptual model of the knowledge/data. Finally, the Imperial Cancer Research Fund provides a large database of research data that could hold the key to causes and cures for various kinds of cancer.

In France, Esprit project 1133 is being led by Sagem. Other partners in the project include Agusta (Italy), ARS (Italy), INRIA (France), Simulgo (France), and CRIL (France). The aim of the project is to implement a knowledge-base processing machine that incorporates an advanced model for integrating database and knowledge-base management systems. The first step is to extend the Sabre relational database model to support the rules and concepts required for knowledge representation. Next, a functional and operational architecture of the proposed system will be developed. Finally, a model will be developed of the ways in which the facts and rules can be accessed. The project is scheduled to last for three years and has a budget of about $4 million.

Several suppliers are also beginning to work on linking relational-database and knowledge-base systems. We have already touched upon Cullinet's Application Expert, which interfaces with its IDMS database, and mentioned IntelliCorp's KEE Connection, which connects an expert system with SQL. Telecomputing's Top-One has a similar link. Recently, Intelligent Terminals and Concurrent Computer announced that they have developed a software tool, code-named Reliance Expert, that integrates Concurrent's Reliance relational database system with Intelligent Terminal's Extran. (Extran induces the rules for an expert system from case material.) One planned application of Reliance Expert is to identify faults in NASA's spacecraft rocket engines from engine-test data, where a single test produces 50M bytes of test data (so far, there have been 1,400 engine tests). The relational database will be used to organise and hold the data in an easy-to-handle form.

In the United States, Relational Technology's Michael Stonebraker and colleagues at the University of California, Berkeley, are working on the integration of expert systems and relational databases. The approach adopted is to combine the data and rules into a single base, and for applications to access the data-and-rules base through a database-and-rules manager. The work is partly funded by the National Science Foundation and the US Navy. Relational Technology's main product is Ingress, which competes with Oracle as the market leader for relational databases. The code name of the new project is Postgres (short for Post Ingress). As well as handling an integrated data and knowledge base, Postgres will also provide object-oriented management and programming features.

We believe that by about 1991/92, research and development activities in the area of integrating expert systems with relational databases will result in commercially available products.

HARDWARE DEVELOPMENTS TO IMPROVE EXPERT SYSTEMS PERFORMANCE

The inferencing techniques used by expert systems are very inefficient when they are run on conventional data processing hardware. We predict that, within the next three years, two hardware developments will have a dramatic impact on the operational efficiency of expert systems.

The first is the emergence of reasonably priced parallel-processing architectures that can be used to process logic programs or rules in parallel. These hardware developments are very much at the heart of the Japanese fifth-generation project.

The architectures of ICL's DAP (distributed array processor) and the Inmos Transputer are forerunners of the kinds of inexpensive parallel-processing hardware that will be used for expert systems in the early 1990s. (A new company has been formed to exploit the DAP, and it has announced that the product will soon be able to operate in a Vax environment.) The trend can be seen from products that are already available. Gold Hill Computers has a concurrent version of its Common Lisp for use with Intel's IPSC Hypercube family of multiprocessors; Bolt Beranack & Newman has its Butterfly machine; and Thinking Machines Corporation provides its Connection Machine. All of these products are priced at a small fraction of the typical multimillion-dollar price of a Cray supercomputer.

The second hardware development that will improve the performance of expert systems is the ability to search data files in parallel. Early examples of such equipment include the Britton Lee database engine, Teradata's DBC, and ICL's CAFS (content addressable file store) disc controller.

IMPROVED KNOWLEDGE-ACQUISITION TECHNIQUES

One of the main bottlenecks in developing an expert system is the process of acquiring the knowledge that will be used to build a knowledge base. We now describe a research project and three other developments associated with knowledge engineering that we believe will help to ease the bottleneck within the next five years.

KADS — A RESEARCH PROJECT TO DEVELOP A KNOWLEDGE-ENGINEERING METHODOLOGY

At present, there is no generally accepted methodology for acquiring the knowledge to be stored in a knowledge base. However, there are several research groups investigating this area, including Esprit Project 1098 (sometimes called the KADS project). This five-year, 80-man-year project began in September 1985. The project staff come from STC, Scicon, and the Polytechnic of the South Bank in the United Kingdom, from Cap Sogeti in France, from the University of Amsterdam in the Netherlands, and from SCS GmbH in West Germany.

The aim of the project is to provide automated tools for the development of expert systems. By the middle of 1987, a methodology for knowledge acquisition had been developed, together with a prototype tool for automating the methodology. (KADS is actually the name of that prototype, but the name is used to refer the whole project.) The development of the prototype represents the end of the first half of the project. The second half is to provide tools that assist in the design process itself. Figure A2.3 shows the three stages of the knowledge-engineering process: problem analysis — for which the KADS prototype is the tool; design — which is currently being researched by the project team; and development — which already is supported by a growing number of commercially available tools.

The KADS project started with the problem-solving modelling concepts developed by the cognitive psychology department at the University of Amsterdam and combined these with the concepts of systems life-cycle process modelling and data modelling. The net result is a set of concepts describing a four-layer knowledge model, which are similar to the concepts of data modelling:

- Domain definition, which is equivalent to the entity model.
- Inference map, which is similar to the process-entity linkage model.
- Task description, which is equivalent to the process model.
- Problem-solving strategy description, which has no equivalent in today's data models.

We now describe in more detail the only element of knowledge modelling that has no parallel in data modelling — the problem-solving strategy description.

An expert system uses a knowledge base in two main ways to arrive at a conclusion — forward chaining and backward chaining. Forward chaining is used where there is a large body of evidence that is used by the expert system to arrive at a conclusion. For example, given all the facts relevant to an application for a business loan, an expert system could come up with an answer of 'yes', 'no', or 'no recommendation'. The inferencing technique used here would be that of forward chaining — working from the given facts to the correct answer.

By contrast, backward chaining works backwards from nominated answers (or 'goals', to use the correct technical term), and the system tests the given facts to see if any of the nominated answers fit. This method is most useful where there is a large number of possible answers, but not all the information is available to start with. Such situations typically occur in scientific research, where there are dozens of working hypotheses — for instance, whether a sample of seismic readings indicates or excludes the presence of certain minerals. The advantage of backward chaining is that many of the nominated answers can be rejected by the presence or absence of a few key facts.

In many problem-solving situations, it would be useful to have an expert system that could start with backward chaining to see if the answer can be obtained quickly, but that would, after some dead ends, switch to forward chaining with a subset of the known facts. The solution found in this way would then be tested against the complete facts by backward chaining, beginning the whole process again. Eventually the iterations through backward chaining and forward chaining would either converge to provide an answer to the problem or would indicate the need for further evidence (or facts). Such a process is termed a problem-solving strategy, and the KADS prototype is able to model this process.

The KADS prototype runs on a Sun workstation with Unix, C, and Quintus Prolog. It has been used for at least two projects in STC and ICL (which is now part of STC). The first project was found to be too complex for KADS, and, in fact, would have been too complex for today's expert systems software. However, without KADS, the project would have proceeded to the design stage or development stage before this was discovered. The second project concerns support for operators in the multilayer circuit-board production plant. The knowledge model was completed during the first

Figure A2.3 Knowledge engineering and the KADS project

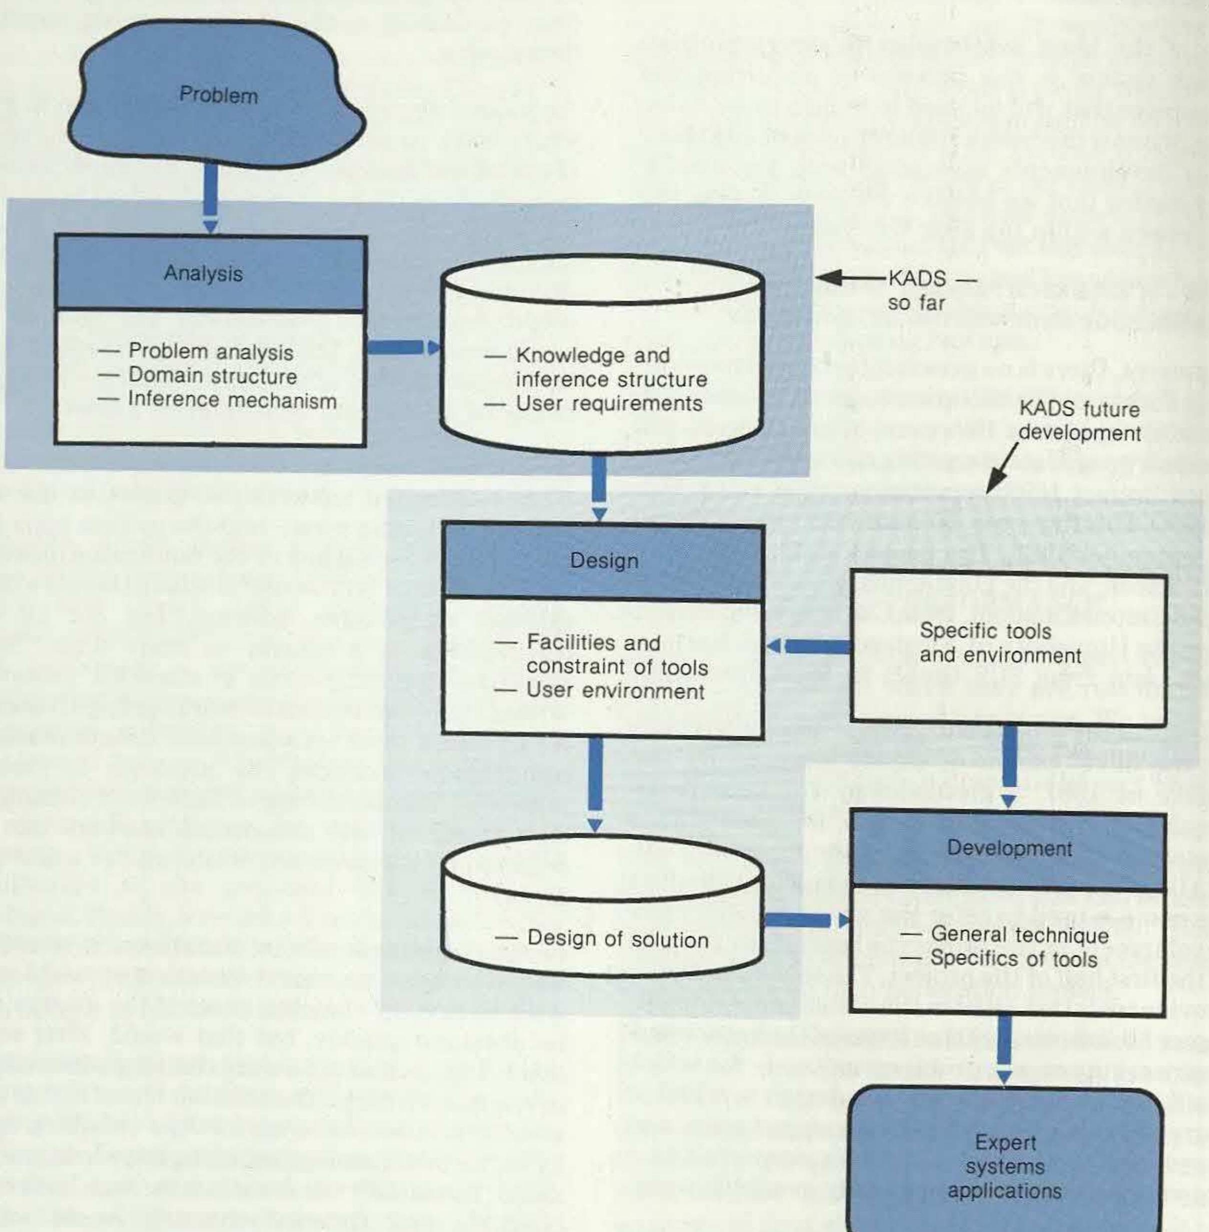

(Source: Hayward et al — STC/Esprit: KADS Project)

half of 1987, and the project has now proceeded to the design stage.

The advantage of KADS is that the knowledge model is independent of the way in which the expert system is designed and implemented. This independence is achieved because KADS formalises the understanding of knowledge by creating separate views of the facts (in the domain-definition layer), of the rules about using the facts (in the inference-map layer), of the business procedures (in the task-description layer) and of the control mechanisms (in the strategy-description layer). By doing this, KADS formalises the best practices of knowledge engineering, just as in the 1970s entity modelling formalised the best practices of database analysis and development.

Although it is too early to say whether KADS will lead to commercially viable products, or even whether it will gain universal acceptance as entity modelling has done, it is nevertheless a major step forward in the practical application of knowledge engineering.

BLACKBOARDING

One of the difficulties in acquiring the knowledge to load into a knowledge base is that it is sometimes difficult to get the experts to agree among themselves. A technique that is used to enable experts at least to be aware of each other's views is blackboarding. The blackboarding technique allows an electronic 'blackboard' to be shared by several experts who have different perspectives on a particular area of expertise. Blackboarding is an extension both of the computer-conferencing techniques of the early 1980s and of the more recent electronic-mail techniques. Experts communicate through the blackboard, which is structured and hence can be used to add meaning to the discourse. A few expert system products (ART for example) include blackboarding facilities.

During our focus-group discussions we heard about a novel extension of the blackboarding technique. Several financial institutions in the United Kingdom are considering clubbing together to develop an investment dealing-room blackboarding system (see Figure A2.4). What is interesting about this particular initiative is that instead of human experts with different perspectives, the inputs will come from other systems (such as Reuters and London's International Stock Exchange Topic financial information system). At the same time, previously developed specialised 'advisory' systems (some of which may be expert systems) will obtain information from the blackboard and contribute to it. The advisory systems contributing to the blackboard will include systems that recommend buy or sell decisions, portfolio management systems, foreign-exchange dealing systems, and so on. The user will access the composite information (as well as the component details) from the blackboard, thereby obtaining the best possible advice.

The users of this system will themselves be experts in a specific field — some in equities (stocks and shares), some in government bonds, and others in foreign-exchange or money markets. The blackboard will enable them to pool their expertise in the belief that the resulting dealing-room knowledge will be greater than the sum of the individuals' expertise.

FUZZY LOGIC

Fuzzy logic is concerned with converting qualitative values (high, low, very low, and so on) to likely numeric values. Thus, fuzzy logic could be used to convert an input that said 'high level of inflation' to a specific percentage increase. The conversion is carried out by using a function curve that specifies the likelihood of a specific rate of inflation being regarded as high (see Figure A2.5). Even though several expert system products incorporate fuzzy logic, the technique has not been put to much use. One possible reason for the lack of use of fuzzy logic may be that qualitative values are converted manually before they are entered into an expert system. At our UK focus-group discussion we heard about another possible reason.

In the work that Touche Ross has carried out on expert systems, it has tried to apply fuzzy logic where appropriate. In its experience, the single curve used to convert qualitative to quantitative values is not very useful in practice because it models a static situation. In the example given in Figure A2.5, a single curve would be sufficient if the rate of inflation never changed. In practice, however, people's perception of whether inflation is high or low is influenced by whether the previous level of inflation has been very high or very low. In other words, people's perceptions are influenced by their immediate past experiences. If inflation has been high, people will only accept that it is low when it has dropped substantially. Similarly, if inflation has been low, it will only be generally recognised as being high once it has increased substantially.

Touche Ross believes that the fuzzy logic function curve should be extended to contain two elements (see Figure A2.6), one for when the trend in the

Figure A2.4 Example of an investment dealing room blackboarding system

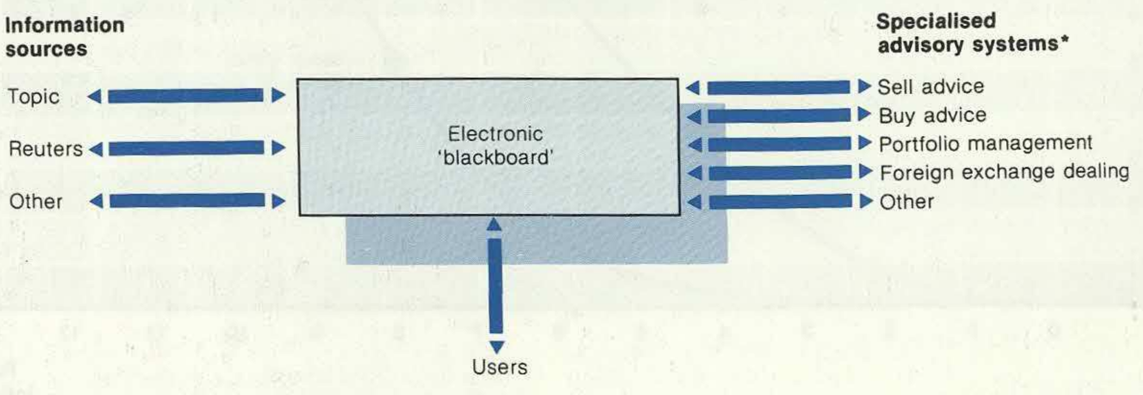

* The advisory systems may be conventional systems or expert systems

Figure A2.5 Fuzzy logic function curve

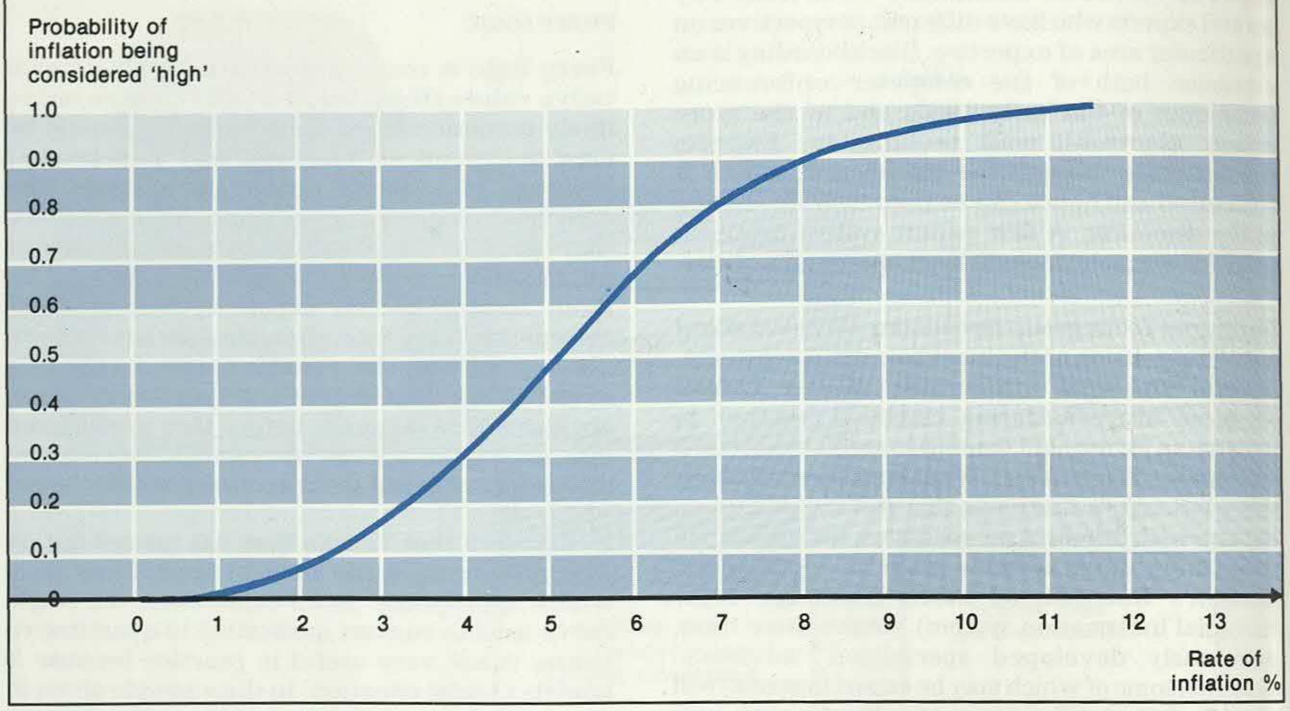

Figure A2.6 Enhanced fuzzy logic function curve

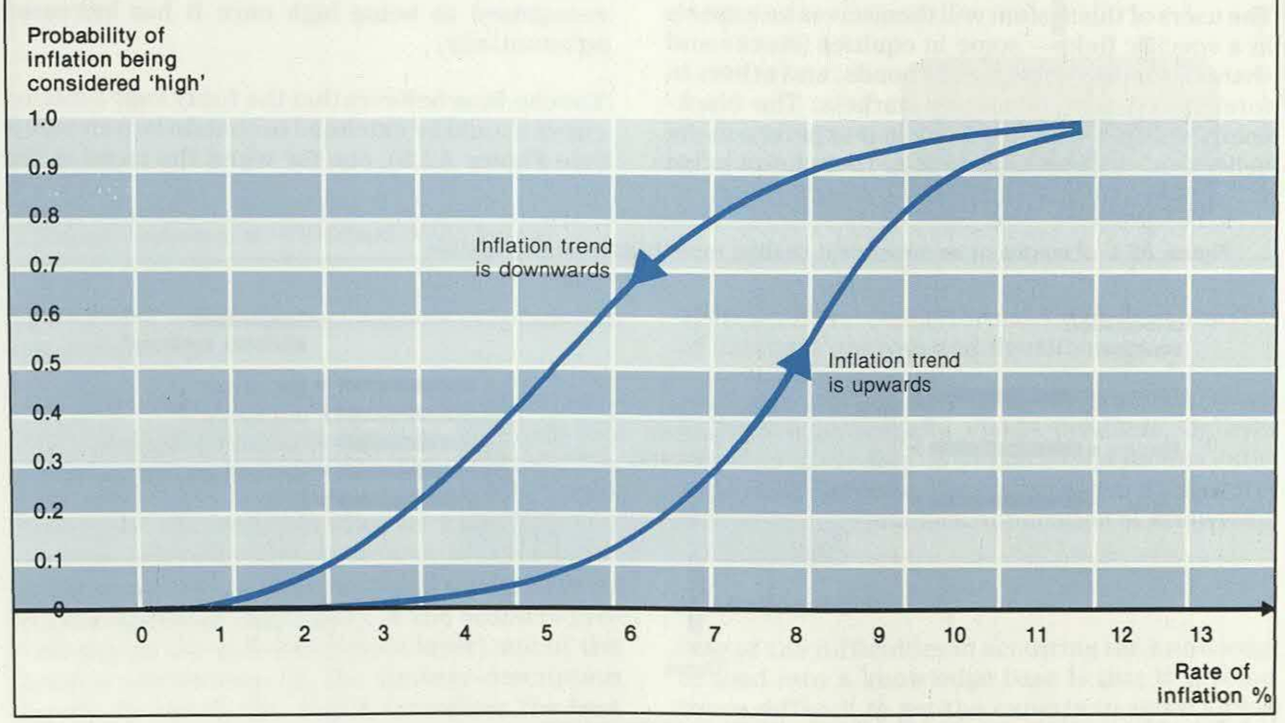

attribute being modelled is downwards, and one for when it is upwards. Thus, the fuzzy logic calculations will need to take account of the historic values (of inflation, for example) and whether the trend is up or down.

We believe that extending fuzzy logic to take account of these factors will allow the technique to model human expertise more realistically and enable it to be used for a wider range of problems.

ANALOGOUS REASONING

In 1985, researchers at Microelectronics and Computer Technology Corporation (MCC) in Austin, Texas, began a ten-year project (called CYC) aimed at overcoming knowledge-acquisition bottlenecks. CYC will make use of analogous reasoning, which is used when people try to apply knowledge and experience acquired in one domain to a totally different domain by drawing analogies. For example, expertise in military strategy and tactics might be applied to the treatment of diseases by drawing analogies between the two domains. The line of reasoning might be:

- Medical treatment is analogous to warfare against disease.
- Bacteria can be thought of as the enemy.
- An infection is an invasion by bacteria.
- Thus, containing a disease is similar to containing the enemy.
- Drugs are the weapon to contain the enemy.
- Resistance to drugs can be thought of as countermeasures by the enemy.

The analogy could be extended further to include concepts like infiltration and subversion, or even to applying the ancient military maxim of making friends with the enemy's enemy.

Thus, if there is an established knowledge base for one domain, it should be possible to apply it in an analogous way to another domain. If the CYC project is successful, the technique of analogous reasoning could be applied to exploring, teaching, problem-solving, and to new theory formulation — which would represent a major breakthrough in knowledge engineering.

Glossary of expert system terms

Algorithm
A clearly defined step-by-step procedure in computation, such as that used for finding square roots.

Artificial intelligence
The branch of computer science concerned with the development of machines that can reason, understand human language and speech, recognise the physical world around them through vision systems, move around the world, and solve difficult problems — machines that can, in other words, mimic the things that make humans appear intelligent.

Backward chaining or goal-driven reasoning (see also Forward chaining)
Reasoning from a conclusion or goal back through a set of rules to see whether the rules and facts available do lead to that conclusion.

Blackboarding
A system for allowing two or more experts to cooperate by posting messages to each other on an electronic 'blackboard'. Used in complex realtime expert systems.

Certainty factors
A method using probability parameters within rules for handling inexact information. The probability represents the expert's level of confidence that the rule is correct.

Cognitive science
A general term covering all those branches of science concerned with the study of understanding, including artificial intelligence, psychology, and linguistics.

Cognitive psychology
The study of how the human mind understands.

Combinatorial explosion
The extremely large number of ways that a small number of events can be selected and/or sequenced. For example, four different items can be sequenced in 24 ways; ten items in 3,628,800 ways; 20 items in 2.4 trillion ways.

Forward chaining or data-driven reasoning (see also Backward chaining)
Reasoning forward from facts and rules to see what conclusions they lead to.

Frames
A construct in a knowledge base that can accommodate a range of information about an object. Frames hold knowledge in structures rather than as rules. Some frame-based systems divide knowledge into 'classes' and 'subclasses' — for example ships, steamships, sailing ships, ketches, sloops, and so on. Properties of classes can be 'inherited' by subclasses, which reduces the data-input requirement when building a knowledge base.

Fuzzy logic
A method for handling inexact information in the form of non-numeric (value) judgements by quantifying approximate probabilities such as 'quite likely' or 'very high'.

Heuristics
Rules of thumb. The rules of expertise, good practice, and knowledge of the field.

Induction
Reasoning from specific instances to a general rule, deriving a rule from examples.

Inference
The process of reaching conclusions through applying logic.

Inference engine
The computer program in an expert system that works out the logical consequences of the rules and, sometimes, controls the whole operation of the system.

Inference mechanism
The strategy that is used by the inference engine of an expert system to deduce conclusions from the knowledge base.

Inheritance
The process by which properties of classes are passed down to subclasses — used in frame-based expert systems.

Knowledge acquisition
The process of assembling and structuring the knowledge of a domain from one or more experts.

Knowledge base
The information held in data files in an expert system that constitutes its domain expertise.

Knowledge-based systems
Computer systems that consist of large amounts of knowledge rather than algorithms; another term for expert systems.

Knowledge engineering
The process of building a knowledge-based system in cooperation with a human expert; corresponds to analysis, design, and programming in conventional computing.

Knowledge representation
The method used for storing knowledge in the knowledge base of an expert system. It can be in the form of rules, frames, semantic networks, or other representations.

Lisp
A computer programming language commonly used for expert systems work. It is designed for 'list processing' — the manipulation of text held in structures called 'lists'. Lisp is the preferred expert systems language in the United States.

Logic programming
Programming an expert system by expressing facts, relationships, and rules in logic statements.

Natural-language processing
Ways for computers to handle human language — for example, accepting instructions in ordinary English rather than in a programming language.

Object
A package of information in a knowledge base, generally corresponding to some real-world concept or entity, whose attributes and relationships with other objects can be extended and manipulated as part of the inference process.

Object-oriented programming
A technique in which entities in the real world are represented as independent 'objects' that send messages to each other. In a radar system, for example, the aeroplane is one object and the clouds are another, and the radar signals are the messages.

Paradigm
An example or pattern of a way of thinking or of generating knowledge. Usually used as a term for the reasoning mechanism of the inference engine.

Predicate calculus
A form of logic — a formal language of symbol structures used for symbol manipulation; relevant to symbolic processing. Predicate calculus is the basis of logic programming and relational database theory.

Production rule
The formal name for the type of rule commonly used in expert systems, of the form *if* (some condition) *then* (some conclusion).

Prolog
A logic programming language. Prolog is a high-level language capable of manipulating symbols and symbol structures, while providing extended facilities for expressing knowledge and using knowledge in a reasoning process. Prolog is the preferred expert systems language in Europe and Japan.

Rule
A statement about a deduction that can be made from a given item of information, possibly embodying heuristic knowledge and typically in the form of a production rule.

Rule-based systems
Computer systems in which knowledge is encoded as rules rather than as algorithms or frames.

Rule induction
Deriving rules automatically through the use of induction.

Search space
The area in which an expert system will seek to find a solution to the problems it is set. The larger the search space, the more likely it is to find a solution, but the longer it may take.

Semantic network
A diagnostic structure used to denote the relationships between objects. Similar to the structure of entity modelling.

Shell
A generalised expert systems application providing a structure or framework for a designer to build a knowledge base. The shell also provides the inference mechanism, together with its predetermined control strategy so that a usable expert system can be developed without requiring expert systems expertise or application-development experience.

Symbol
A string of characters that represent an entity, attribute, or relationship.

Symbolic processor
A computer specially designed to handle symbols (words and possibly pictures), rather than numbers.

System builder, or system toolkit, or system environment
A set of software tools for building an expert system, much more elaborate than a shell. The best known are ART, KEE, and Knowledge Craft.

Bibliography

General

Expert Systems in Europe. Frost & Sullivan, February 1985.

Goodall A. The Guide to Expert Systems. Learned Information, 1985.

Hayes-Roth F, Waterman D A, Lenat D B. Building Expert Systems. Addison Wesley, 1983.

Hewett J, Timms S, d'Aumale G. Commercial Expert Systems in Europe. Ovum, 1986.

Smart G, Langeland-Knudsen J. The CRI Directory of Expert Systems. Learned Information, 1986.

A Strategic Guide to Implementing an Expert System. system dynamics limited, November 1986.

Identifying opportunities

Canning McNurlin B. A three-year strategy for expert systems. EDP Analyzer, Volume 25, Number 3, March 1987.

Stow R, Lunn S, Slatter P. How to identify business applications of expert systems. Second International Expert Systems Conference, London, October 1986.

Expert systems, fourth-generation languages, and systems development

Experts turn to the 4GL. Network, June 1987.

Wolfe A. Software. Electronics, 16 October 1986.

Expert systems and relational databases

Advanced model for integration of DB and KB management systems. Inria Information, Number 13, April 1986.

Kerschberg L. Expert database systems. BCS Computer Bulletin, Volume 3, Part 2, June 1987.

Moad J. Building a bridge to expert systems. Datamation, 1 January 1987.

Nijssen G M. Knowledge Engineering, Conceptual Schemas, SQL and Expert Systems: a Unifying Point of View. University of Queensland, Australia, December 1986.

Stonebraker M, Rowe L A (editors). The Postgres Papers. Memo Number UCB/ERL M86/85. University of California, Berkeley, 5 November 1986.

Intelligent decision-support systems

Luconi F C, Malone T W, Scott Morton S M. Expert systems: the next challenge for managers. Sloan Management Review, Volume 4, Summer 1986.

Remus W E, Kotteman J. Toward intelligent decision support systems. An artificially intelligent statistician. MIS Quarterly, Volume 10, Number 4, December 1986.

Knowledge engineering

Hayward S. A Structured development methodology for expert systems. Online KBS 1986 Conference, June 1986.

Hirsch P et al. Interfaces for knowledge-based builders' control knowledge and application-specific procedures (ESE/VM). IBM Journal of Research and Development, Volume 30, Number 1, January 1986.

Lenat D B, Prakash M, Shephard M. CYC. Using Common Sense Knowledge to Overcome Brittleness and Knowledge Acquisition Bottlenecks. MCC Tech Report Number AI 055 85, 15 July 1985.

Future developments

Broughton P, Thomson C M, Leunig S R, Prior S. Designing systems software for parallel declarative systems. ICL Technical Journal, Volume 5, Number 3, May 1987.

Devlin K. Situation semantics. BT Network, Spring 1986.

Fisher E M. Building AI behind closed doors. Datamation, Volume 32, Number 15, 1 August 1986.

Nairn G. Out of the lab, on to the desktop. Infomatics, Volume 8, Number 4, April 1987.

Peltu M. Playing the AI game. Datamation, Volume 32, Number 23, December 1986.

Verity J W. The Lisp race heats up. Datamation, Volume 32, Number 15, 1 August 1986.

Weiszmann C, Messenheimer S. Parallel symbolic computing extends AI opportunities. Computer World, 22 December 1986.

Surveys and statistics

Expert systems are set to double annually. Computer Weekly, 16 October 1986.

Results of Survey on Trends in Expert Systems in Japan. JIPDEC, October 1986.

Squaring up to expert systems. Business Computing & Communications, March 1987.

Tehan P. Long road ahead for ai. Computing, 28 May 1987.

BUTLER COX FOUNDATION

Butler Cox

Butler Cox is an independent management consultancy and research organisation, specialising in the application of information technology within commerce, government and industry. The company offers a wide range of services both to suppliers and users of this technology. The Butler Cox Foundation is a service operated by Butler Cox on behalf of subscribing members.

Objectives of the Foundation

The Butler Cox Foundation sets out to study on behalf of subscribing members the opportunities and possible threats arising from developments in the field of information systems.

The Foundation not only provides access to an extensive and coherent programme of continuous research, it also provides an opportunity for widespread exchange of experience and views between its members.

Membership of the Foundation

The majority of organisations participating in the Butler Cox Foundation are large organisations seeking to exploit to the full the most recent developments in information systems technology. An important minority of the membership is formed by suppliers of the technology. The membership is international, with participants from Australia, Belgium, France, Italy, the Netherlands, Sweden, Switzerland, the United Kingdom and elsewhere.

The Foundation research programme

The research programme is planned jointly by Butler Cox and by the member organisations. Half of the research topics are selected by Butler Cox and half by preferences expressed by the membership. Each year a shortlist of topics is circulated for consideration by the members. Member organisations rank the topics according to their own requirements and as a result of this process, members' preferences are determined.

Before each research project starts there is a further opportunity for members to influence the direction of the research. A detailed description of the project defining its scope and the issues to be addressed is sent to all members for comment.

The report series

The Foundation publishes six reports each year. The reports are intended to be read primarily by senior and middle managers who are concerned with the planning of information systems. They are, however, written in a style that makes them suitable to be read both by line managers and functional managers. The reports concentrate on defining key management issues and on offering advice and guidance on how and when to address those issues.

Selected reports

5 The Convergence of Technologies
8 Project Management
11 Improving Systems' Productivity
15 Management Services and the Microprocessor
17 Electronic Mail
18 Distributed Processing: Management Issues
19 Office Systems Strategy
20 The Interface Between People and Equipment
21 Corporate Communications Networks
22 Applications Packages
23 Communicating Terminals
24 Investment in Systems
25 System Development Methods
26 Trends in Voice Communication Systems
27 Developments in Videotex
28 User Experience with Data Networks
29 Implementing Office Systems
30 End-User Computing
31 A Director's Guide to Information Technology
32 Data Management
33 Managing Operational Computer Services
34 Strategic Systems Planning
35 Multifunction Equipment
36 Cost-effective Systems Development and Maintenance
37 Expert Systems
38 Selecting Local Network Facilities
39 Trends in Information Technology
40 Presenting Information to Managers
41 Managing the Human Aspects of Change
42 Value Added Network Services
43 Managing the Microcomputer in Business
44 Office Systems: Applications and Organisational Impact
45 Building Quality Systems
46 Network Architectures for Interconnecting Systems
47 The Effective Use of System Building Tools
48 Measuring the Performance of the Information Systems Function
49 Developing and Implementing a Systems Strategy
50 Unlocking the Corporate Data Resource
51 Threats to Computer Systems
52 Organising the Systems Department
53 Using Information Technology to Improve Decision Making
54 Integrated Networks
55 Planning the Corporate Data Centre
56 The Impact of Information Technology on Corporate Organisation Structure
57 Using System Development Methods
58 Senior Management IT Education
59 Electronic Data Interchange

Forthcoming reports

Competitive-Edge Applications: Myth or Reality?
Communications Infrastructure for Buildings
The Future of the Personal Workstation

Availability of reports

Members of the Butler Cox Foundation receive three copies of each report upon publication; additional copies and copies of earlier reports may be purchased by members from Butler Cox.

Butler Cox & Partners Limited
Butler Cox House, 12 Bloomsbury Square,
London WC1A 2LL, England
☎ (01) 831 0101, Telex 8813717 BUTCOX G
Fax (01) 831 6250

Belgium and the Netherlands
Butler Cox BV
Burg Hogguerstraat 791
1064 EB Amsterdam
☎ (020) 139955, Fax (020) 131157

France
Butler Cox SARL
Tour Akzo, 164 Rue Ambroise Croizat,
93204 St Denis-Cedex 1, France
☎ (161) 48.20.61.64, Fax (161) 48.20.72.58

Germany (FR)
Butler Cox (Deutschland) Ltd.
Richard-Wagner-Str. 13
8000 München 2
☎ (089) 5 23 40 01, Fax (089) 5 23 35 15

United States of America
Butler Cox Inc.
150 East 58th Street, New York, NY 10155, USA
☎ (212) 486 1760

Australia
Mr J Cooper
Consultants (Computer and Financial) plc Australia
Level 5, 303 Pitt Street, Sydney 2000, Australia
☎ (02) 2870400, Fax (02) 2870450

Ireland
SD Consulting
72 Merrion Square, Dublin 2, Ireland
☎ (01) 766088/762501, Telex 31077 EI,
Fax (01) 767945

Italy
SISDO
20123 Milano, Via Caradosso 7, Italy
☎ (02) 498 4651, Telex 350309, Fax (02) 481 8842

The Nordic Region
Statskonsult AB
Stora Varvsgatan 1, 21120 Malmo, Sweden
☎ (040) 1030 40, Telex 12754 SINTABS

Spain
Associated Management Consultants Spain SA
Rosalía de Castro, 84-2°D, 20835 Madrid, Spain
☎ (91) 723 0995