ab-libris

- buy the book!

ab-libris ltd. PO Box 287, Durham, DH1 5YZ

Tel: +44 (0)709 209 1131 Fax: +44(0) 709 209 1131 Email: ablibris@ablibris.com

Ab-libris e-book example

An Introduction to Knowledge Engineering by Peter Smith

This document is a sample of the above book. This sample contains the cover, contents page and a single chapter (chapter2). The full version can be purchased on line from www.ablibris.com.

ab-libris computing classics

An Introduction to Knowledge Engineering

Peter Smith

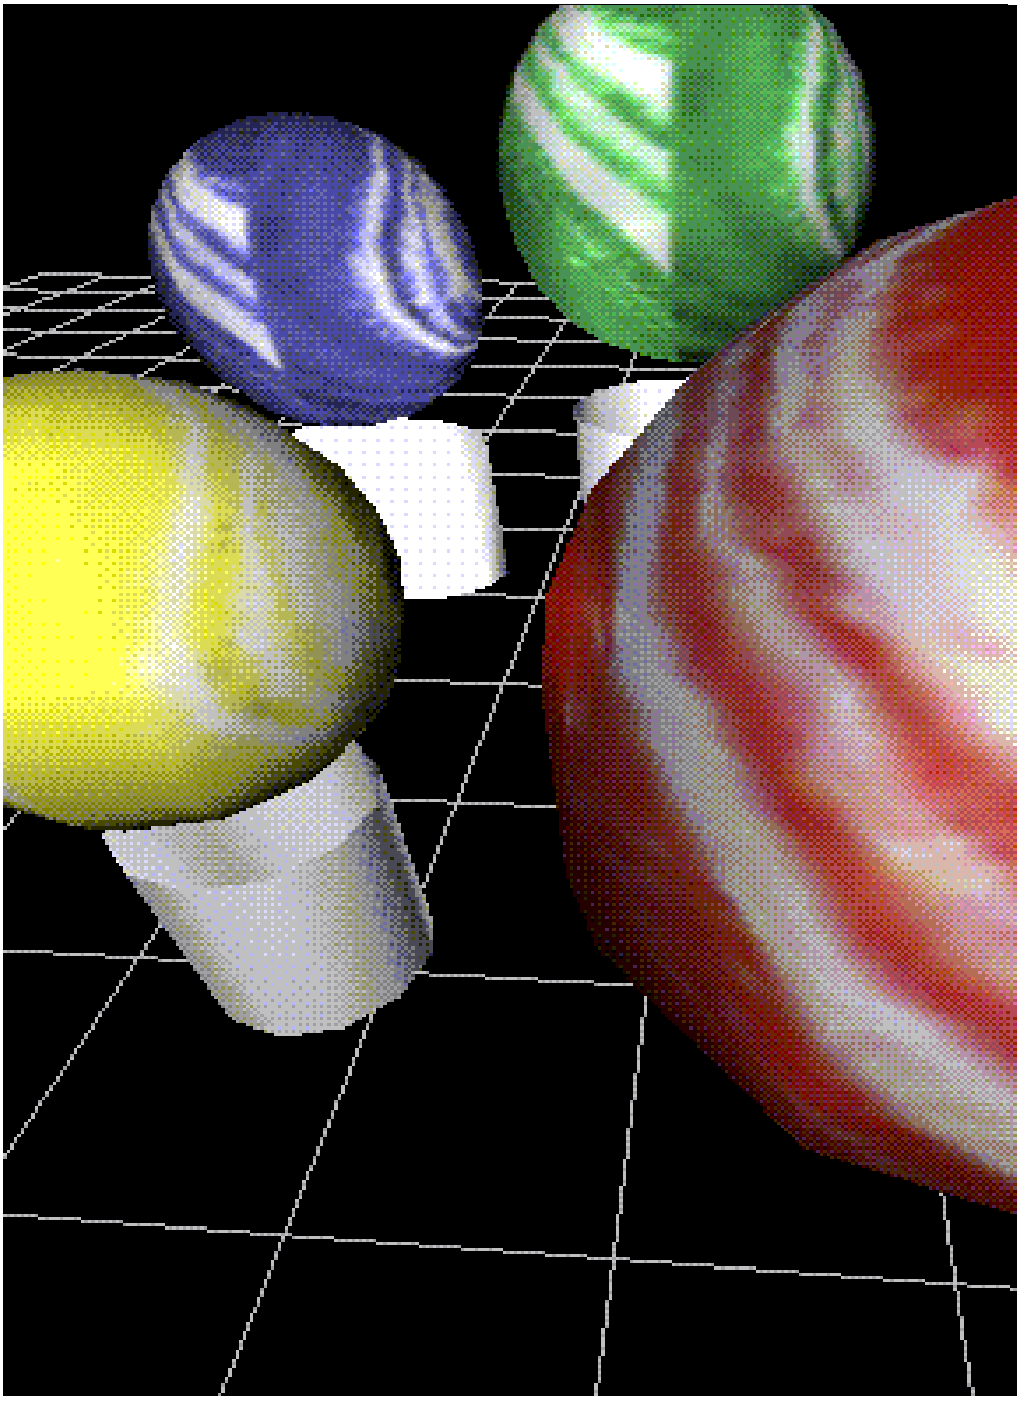

Original title:

An Introduction to Knowledge Engineering
copyright 1996 by International Thomson Computer Press

British Library Cataloguing in Publication Data

Smith, Peter
An Introduction to Knowledge Engineering
(Ab-libris computing classics series)
1. Knowledge Engineering

ISBN: 1-903561-01-9
Typeset in Durham by Ab-Libris Ltd.

Copyright notice

Contents

Acknowledgements from the Original Edition.

I would like to express my sincere thanks to the following:

Special thanks to Colin Hardy, Ali Jawad, Dave Leonard and C Lin for providing the basis for the case study material.

British Gas plc. for permission to use the forecasting case study.

Finn Gronskov and Prolog Development Center for material on ESTA.

Intelligent Environments for use of CRYSTAL material.

Paul Ross, Mark Lansbury, Dimitrios Markakis, Stephen Ng, Jason Kuo, Shi-Ming Huang, Xin Chen, Helen Edwards and Alan Gillies for their important contributions.

Dr Theodoulidis of UMIST for his invaluable comments and encouragement throughout the project; Dave Hatter for starting the ball rolling and Samantha Whittaker and Nikki Vaughan for seeing it through, Simon Kendal for providing useful comments and checking the draft.

And finally, but most importantly to Marie, Ashleigh, David and Laura for putting up with me while I sat in the corner for endless hours tapping away at a keyboard.

Peter Smith 1995

Preface

This book is about knowledge engineering; the development of knowledge based systems (KBS). KBS are complex pieces of software which attempt to encapsulate human knowledge within a computer system. They have much in common with traditional software products and suffer from the same difficulties and problems. In addition, however, the development of a KBS poses a new, and very different, set of challenges to the knowledge engineer.

It is the aim of this book to provide a practical introduction to the topic of knowledge engineering. This is achieved by covering all of the important aspects of the discipline in a practical manner.

Chapter I provides an introduction to the subject and answers the question "What is knowledge engineering?". Chapter 2 focuses upon the subject of KBS, while Chapter 3 explains the life cycle which must be followed when developing a KBS. Chapter 4 provides detailed coverage of knowledge acquisition, which is often seen as the most important stage in the knowledge engineering process (and often described as a "bottleneck"). Chapter 5 covers knowledge representation, while Chapter 6 presents KBS implementation in a programming language and two expert system shells. Chapter 7 presents a series of practical lessons and considerations. Finally Chapter 8 provides coverage of the management issues which are important in knowledge engineering.

A series of four practical case studies are introduced in Chapter 2. These are then referred to throughout the book, in order to maintain a practical bias. The four case studies include a quality-control system, a forecasting system, a monitoring system and an integrated system (which integrates KBS technology with traditional information systems). These case studies are based upon real KBS development projects in which the author has been involved.

Chapter 2
An introduction to knowledge-based systems

OBJECTIVES

In this chapter you will learn:

- about the different forms of knowledge based systems (KBS);
- the history of KBS;
- about some of the applications of KBS;
- how to choose an application area for KBS treatment;
- the structure of a KBS;
- something about the case studies which we are going to use in the remainder of this book.

2.1 Introduction

2.1.1 Background

The earliest computers were extremely large awe-inspiring machines which were used by scientists to perform rapid calculations. In order to obtain an answer from these strange looking beasts the user would first have to decide on a method of solving a problem. He or she would then have to translate the method of solution into a set of instructions called a computer program. The program would invariably have to be written in a language that was unintelligible to the average human being and would consist of a series of strange commands which related to the way in which the computer's electronic components operated.

Since those early days computer designers have striven to make computers easier to use, to bridge the gap between the human and the machine. The languages in which the user has to write a computer program have been designed so that they resemble, as closely as possible, the English words and phrases and mathematical symbols which we use in everyday life. The computer hardware itself has seen dramatic reductions both in price and size, accompanied by similarly dramatic increases in its capacity to store and process data.

But still the computer is only a machine. We cannot (yet) talk to it in English as we would a human being. (There is, however, much progress in this area of research and systems are now appearing that can understand natural language). We cannot give it an everyday problem to solve unless it has already been programmed with an appropriate method of solution. We cannot expect it to learn from its mistakes and apply previous experience and common sense to the solution of problems. Or can we?

Is it even sensible to begin to try to emulate the human brain? The brain is an extremely complex machine. It contains about 10^{10} neurons. This is more cells than there are stars in the Milky Way. The brain deals with sight, hearing, taste and touch all at the same time. In addition it has the ability to store and retrieve vast quantities of information and to control our emotions, feelings and desires. Can a computer ever truly emulate such a complex structure?

Artificial intelligence (AI) is the name given to a field of research which is attempting, with varying degrees of success, to produce computer systems which can accomplish the above. In other words, the aim of AI research is to provide computers with similar intellectual abilities to those of a human being. This is obviously an awesome task, particularly as many of our own intellectual abilities and mental processes are not yet fully understood, even by experts in psychology.

One area of AI research which is quite advanced, however, is that of producing knowledge-based systems (KBS). Indeed, KBS development, or knowledge engineering, is one of the few areas of AI which has moved out of the research laboratory and into the real world. Knowledge engineering is beginning to fulfil a real potential in industrial and commercial applications.

As discussed in Chapter 1, a KBS is a computer system which embodies knowledge about a specific problem domain and can thus be used to apply this knowledge to solve problems from that problem domain. Many KBS are often termed expert systems, as they behave in a similar manner to a human expert when solving a problem. In fact, the terms KBS and expert system are often used interchangeably. This isn't quite right, as there are many KBS which do contain knowledge, but are not really expert systems. For instance an intelligent database system may be knowledge based, but it is not an expert system, because it has not been designed to apply expertise to solve problems. We will stick to using the term KBS throughout this book, to mean any system which has knowledge embedded within it.

2.1.2 Categories of KBS

Knowledge-based systems include:

- Expert systems - this is probably the most common category of KB we mentioned above the terms KBS and expert system are often used interchangeably. An expert system is any system which emulates the thought processes, decision-making processes and/or actions of a human expert (or group of experts). Most of the examples which we will meet in this book could be categorized as expert systems.
- Intelligent database systems - these are database systems which have added intelligence. This often takes the form of an intelligent front end which makes it easier to access the information in the database.
- Intelligent tutoring systems - these are educational systems which attempt to model a human tutor.
- Intelligent CASE tools - CASE (computer-aided software engineering) involves the use of computers to automate parts of the software development process. In recent years, some CASE tools have had intelligence (knowledge) added to them. Such tools try to replicate the knowledge of a software engineer and are the integration of CASE and KBS.
- Integrated or hybrid systems - these are systems which integrate KBS approaches and traditional information systems approaches. One of the case studies which you meet later in this chapter concerns an integrated system.

Waterman (1986) suggests the following categories of expert systems:

- Interpretation - systems which can be used for interpreting data in an expert manner.
- Prediction - systems which can be used to predict the outcome of a particular scenario.
- Diagnosis - systems which can diagnose the cause of a particular problem.
- Design - systems which can be used to design and configure objects.
- Planning - systems for planning and scheduling sequences of actions and events.

- Monitoring - systems which can monitor the state of a physical system by comparing observed data to expected values.
- Debugging - systems which can be used to prescribe cures for faults.
- Repair - systems which can implement repairs.
- Instruction-systems which can instruct users how to perform a particular action or group of actions.
- Control - systems which can control and govern overall system behaviour. The above categories covers most possible KBS. Some of the above categories of knowledge overlap and many systems can fit into more than one category.

A KBS will take the knowledge and experience of a human (or group of humans) and make it available 24 hours a day, every day of the year. The decisions made will be consistent and reliable, and it is possible for the system to be used from a number of different locations and by relatively inexperienced staff.

Many KBS can also learn from their mistakes and gain experience from their successes and failures, just as a human would. The system may also be able to explain the reasoning behind the way in which it has arrived at a particular conclusion.

You may ask 'Why do we need a computer to do such tasks for us? We have plenty of human experts to solve our problems for us.' But is that true? Consider the following properties of human experts:

- They are scarce.
- Their services are expensive.
- They are usually very much in demand and are therefore very busy.
- They are mortal.

Computer systems do not suffer from the above drawbacks.

Consider the following scenario:

Many human experts spend a lot of their time solving trivial problems, some of which are very similar (Figure 7). These are the day-to-day problems which they may meet all of the time, but cannot be ignored and still have to be dealt with. Because of their specialist training and expertise such experts are also in heavy demand not only to help solve day-to-day problems, but also to help with more difficult and interesting problems.

In many cases these very experts are the only people we can approach to help solve problems of a certain type. This may be because only they possess the specialist knowledge needed to solve the problems, or simply because they are the only people who are, by law, allowed to undertake such work. Without doubt these experts would be much happier if they could devote more of their time to solving more of the really interesting problems and have some help with the trivial day-to-day problems.

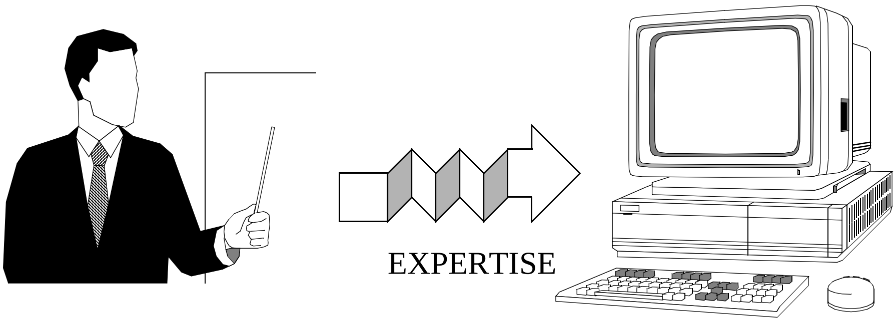

Figure 7. Many tasks undertaken by professionals could be done by a KBS.

Professional experts who might fall into the above category include:

- lawyers
- solicitors
- doctors
- architects

All of the above professionals are expensive to hire and extremely busy. Wouldn't it be useful, and indeed sensible, if at least some of their expertise could somehow be implanted within a computer system? A copy of this system could then be held at public places such as libraries and supermarkets and it could then be used as a cost-effective consultant on day-to-day problems. Of course for more difficult problems the system would need to advise the user to consult a real professional. Consider the scenario below:

John is a solicitor who deals largely with divorce cases. This is quite a lucrative business, particularly in current times where divorce is becoming more and more common. Many of the cases which he deals with are, of course, relatively straightforward and hence he finds them quite boring to deal with. He often boasts to his friends that he could deal with them 'in his sleep'. Yet he clearly can't turn cases away just because they are straightforward and uninteresting. These people need help and he may be the only one who can give them the help that they need. He would also be turning away lucrative business, and anyway he needs the money to help pay for his large mortgage and two cars.

What John really needs is a KBS which could screen for the trivial cases and help answer a lot of the straightforward questions for him. This KBS would have implanted within it some of the more basic knowledge relating to divorce law. It would also have a bank of case studies which it could draw on, thus bringing the benefit of past experience to bear on a case. Perhaps John could even get his assistant to help operate the system.

This would free up John's time for other things. He could then concentrate on more interesting cases, or take on even more cases and increase his salary. Either way, the benefits of such a system are quite obvious.

There are many other fields of enterprise where such a system would be of immense value. There are just as many other fields of human expertise where a computer system will never be able to match the capabilities of a human being. These include those activities which involve originality, intuition, creativity and artistic ability. For example, one could never imagine a computer being able to paint a great master, write a (good) play or compose music. That is how it should be. There need to be some things which only humans can do, or we enter into the dangerous scenario of the computers being able to really think for themselves and trying to replace us.

There are also legal implications of using a computer to give advice and solve human problems. These legal implications have yet to be really examined. For example, if a doctor uses a KBS to advise as to the treatment of a patient suffering from an infectious disease, and that patient dies, who is responsible? Is it the knowledge engineer who programmed the KBS? Is it the human expert(s) from whom the knowledge was extracted in the first place? Is it the company who sold the KBS? Is it the doctor who used the KBS as a basis for a diagnosis? Such questions as these remain unanswered for the time being. This means that great care must be taken when deciding upon an application area for a KBS.

The next sections of this chapter discuss some of the applications of KBS technology in more detail.

2.2 A history of KBS

2.2.1 Early systems

Early work on AI began in the 1950s and 1960s. This work was trying to make computers which could think as humans. However, much of the work did not produce any really practical results, and AI began to get itself something of a bad name.

Some of the earliest KBS were developed during the 1970s and earlier. These included one of the most well known and well documented systems, namely MYCIN.

MYCIN is a KBS which was designed at Stanford University in the USA to deal with problems in the treatment and diagnosis of infectious diseases. The system, therefore, comprises a large knowledge base which contains facts and rules about the forms and causes of infectious diseases and can thus be used to aid the clinician in diagnosis.

Another system which was developed around the same period as MYCIN is PROSPECTOR, which was designed by SRI International in association with the United States Geological Survey, to assist geologists during mineral exploration work. Both of these early expert systems contain large amounts of knowledge which have been drawn from well understood and researched problem domains.

Another famous early expert system was DENDRAL, which acted as a chemist's assistant in interpreting the data from mass spectrography. DENDRAL was developed by Ed Feigenbaum of Stanford University in the United States. Feigenbaum is seen by many as the father of expert systems.

In developing DENDRAL, Feigenbaum and his colleagues discovered that human chemists carried around enormous amounts of specialized knowledge in their heads, and it was impossible to do the job without having that specialized knowledge. The difficult task was to sit down with these chemists, watch them work and ask them questions about how they made decisions. Of course, they weren't always able to answer these questions as, although they knew how to do these things they found great difficulty expressing them in words. It was then necessary to figure out some way of representing the knowledge, these 'rules of thumb', in the form of a large KBS.

'In DENDRAL', Feigenbaum says (Feigenbaum and McCorduck, 1984), 'we looked not only for the relatively hard chemical knowledge about stability and mass spectral processes, but also for the relatively soft knowledge: how a particular scientist makes a particular kind of decision when he's not really sure, when there's a variety of evidence, a lot of ambiguities. How does he select?'

Each of the above systems, although successful as a research project, did not have any real success in business or industry. Probably the first real commercial application was XCON, which was developed by DEC (Digital Equipment Corporation) to configure their computers. Table 2 lists some early KBS.

Year	**System**	**Author**	**Task**	**Fate**
1956	No system	Newell, Simon and Shaw	Proved logic theorems	Laboratory prototype
1961	No system	Minsky and Slagle	Solved mathematical calculus problems	Laboratory prototype
1973	DENDRAL	Feigenbaum	Derived chemical structures from mass spectrograph	Algorithmic version sold well
1976	MYCIN	Shortliffe	Medical diagnosis of blood disorders	Given to hospitals, but not in every day use
1978	PROSPECTOR	Duda	Prospecting for mineral ores	Used a few times in commercial situation

Table 2: Some early knowledge based systems

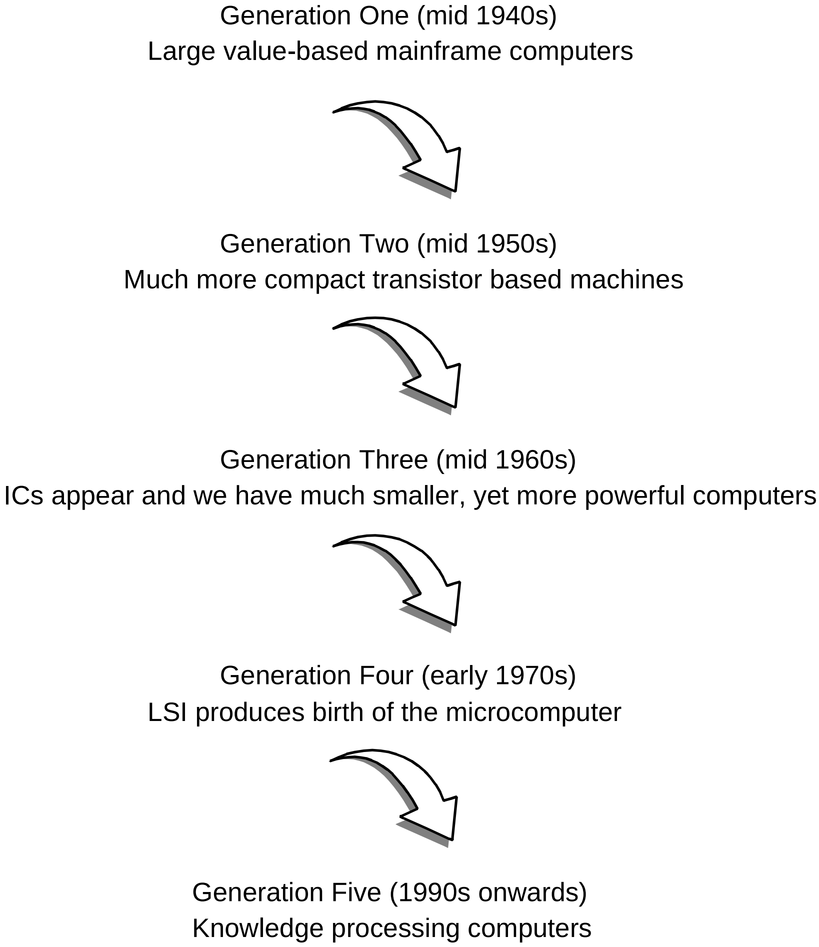

Figure 8. The five generations of computers.

2.2.2 The AI wilderness

In 1972 the Lighthill Report on Artificial Intelligence appeared (Lighthill, 1972). Sir James Lighthill, a distinguished physicist, was commissioned by the Science Research Council of Great Britain to evaluate the state of AI and recommend whether further funding should be given.

The Lighthill Report said that AI was too theoretical and that it had little potential for real use in the business world. This report virtually stopped the support of AI research in the UK in the 1970s. Similar situations arose in other parts of the world, making the road for people working in KBS a long and hard one.

However, AI did recover from this setback and much useful work has been undertaken in recent times. The remainder of this section gives you an idea of some of the most important developments in KBS over the years.

2.2.3 The Japanese fifth-generation project

It is often said that the history of computers can be divided into four generations, stretching from their inception in the mid-1940s to the present day (Figure 8). These four generations are each described below.

2.2.3.1 *Generation one* The first generation of computers emerged in the mid-1940s. These were extremely large, valve-based machines which were so vast that they filled an entire room and needed several people to operate and look after them.

2.2.3.2 *Generation two* In the mid-1950s advances in electronics resulted in the valve being largely superseded by the much smaller transistor. This enabled computer manufacturers to construct what was to become known as the second generation computer which was a much more compact, transistor-based machine.

2.2.3.3 *Generation three* Further advances in electronics during the mid-1960s enabled component manufacturers to combine entire circuits on a single device. These devices were known as integrated circuits (ICs). The use of ICs in the production of computers heralded the advent of the third generation computer, which was, once again, much more compact than its direct predecessor.

2.2.3.4 *Generation four* In the early 1970s another new age dawned with the mass production of much more powerful ICs. Each of these ICs contained a much larger amount of circuitry and a corresponding larger number of individual components. This new technology was known as large scale integration (LSI) and the resultant devices were termed 'chips'. The fourth generation of computers thus consisted of machines which were constructed entirely of such chips and were, once again, much more compact than previous models. It was also around this period that the first microcomputers began to appear. This was the birth of the personal computer (PC).

2.2.3.5 *So what is a fifth generation computer?* In 1981 the Japanese announced the Fifth Generation Computer Project, whose aim was to produce machines which would form the next generation of computer systems. However, unlike the previous four generations of computers, the fifth generation machine was not seen as being purely a result of advances in computer hardware. Rather, the vision of the fifth generation computer was that of a computer system which not only used all of the recent advances in both hardware and software technology, but also took a very different and new approach to the way a computer system was built and worked.

All of the previous generations of computers were essentially data-processing and/or number-crunching engines. In order to use such machines, the operator had to program the computer with an appropriate method of solution to the problem in hand. The philosophy behind the fifth generation computer was that the user would be able to communicate with the machine in natural language (for example English) and that the computer would contain sufficient knowledge to be able to understand and solve real problems. The fifth generation computer was thus seen as a highly sophisticated and powerful knowledge-processing and problem-solving machine.

Basically what the Japanese were attempting to produce was a system which resembled, as closely as possible, a human being in its intelligence and knowledge. Such a computer system was obviously a bold leap from the systems which were available at that time. Now, over a decade later, such systems are still not available. However, many advances have been made in KBS and 'thinking' machines, and the Fifth Generation Programme kick-started the rest of the world to think very seriously about KBS and AI.

2.2.4 ALVEY

Although the ALVEY programme has now reached completion, its impact on KBS (and information technology) research in the UK has been so great in recent years that no review of the area would be complete without mentioning it. The ALVEY programme for advanced information technology was launched in 1982 as a direct response to the Japanese Fifth Generation Programme. The programme encompassed many different aspects of information technology (IT), one of which was KBS. The budget for the programme was around 350 million pounds over a period of five years and it was based upon the concept of promoting joint collaborative research projects between universities and industry.

The ALVEY programme succeeded in starting off research in all areas of KBS/AI including work on vision systems and image processing (to produce computers which can 'see'), natural language processing (to produce computers which we can converse with in English), robotics and industrial KBS. These developments were supported by the funding of collaborative research projects workshops, conferences and technology transfer and development clubs.

This resulted in many exciting advances and helped to forge real link between industrial organizations and universities. Unfortunately, subsequent funding was not provided at a similarly high level and hence many research ideas and prototype systems were not able to be transformed into real products.

The KBS component of the ALVEY programme comprised four large projects known as 'demonstrator' projects as well as numerous small research projects. The four demonstrator projects were so named because their aim was to demonstrate how useful KBS technology can be to business and industry. The projects involved:

- the development of a KBS to aid in the formulation of chemical mixtures
- the use of KBS techniques in the production of mechanical systems (in. particular diesel engines and helicopter gearboxes);
- the development of a decision support system for air crews and a flight-deck simulator;
- the use of KBS techniques for stock control.

2.2.5 KBS research within Europe

The European Union (EU) has also set up a number of large research programmes in the information technology area. One of the most important of these programmes is the ESPRIT (European Strategic Programme of Research in Information Technology) programme.

In 1988, the British Minister for Industry and Consumer Affairs commented on the role of ESPRIT:

Britain is part of Europe and is working with other member states to secure the single market. This surely means it is right to put emphasis on European collaborations. European-scale effort is often needed when large investments will be required for the research itself or for the subsequent production and marketing of the product. ESPRIT provides the opportunity for both our companies and our universities to participate in research on a European scale and to develop the commercial relationship that will be needed. (Smith and Rada, 1994)

The ESPRIT programme has provided funding for a large number of multinational collaborative KBS projects involving universities and commercial companies since the programme began in the early 1980s. ESPRIT has thus been of enormous significance in creating a truly collaborative community of European researchers, developers and users (Smith and Rada, 1994). Through programmes such as ESPRIT, Europe has gained a lead in KBS development methodologies such as KADS (as discussed in Chapter 3).

Current ESPRIT projects are focusing more and more on producing practical computer software which can be used to benefit European industry. The EU is particularly keen to fund projects which will not only undertake useful research, but will also result in products that can be used in business and industry. This means that the KBS projects which are currently being funded by the ESPRIT programme are likely to result in more real, working KBS.

One example of such a project is CIM.REFLEX. The CIM.REFLEX project (ESPRIT 2 Project 6304) used KBS in the production planning process in order to enable manufacturing companies to respond in a more flexible and dynamic manner. The project aim was to specify, develop, implement and evaluate a KBS for production planning which also addresses the tasks of product configuration and costing.

CIM.REFLEX had six partners. Two of these were universities, two were manufacturing companies and two were software houses. The project team covered three EU countries: Denmark, France and the UK.

CIM.REFLEX addressed the need for manufacturing companies to react quickly and with sufficient flexibility to customer demands. A core issue for the project was to meet practical industrial requirements as well as applying the latest research from KBS technology.

The complete CIM.REFLEX system gives knowledge-based decision support at the time a sale is made, using data which reflect the current situation in the factory. A combination of an expert-system programming language (Prolog) and the DECIBAC expert system shell has been used to develop the CIM.REFLEX system.

In order to do this the project team spent many months working with a large number of staff in manufacturing companies. Many knowledge-acquisition sessions were held to capture the knowledge of these staff. This included staff from the shop floor, from the planning department and from the accounts department. The system was designed to enable sales staff to use this knowledge when they make a sale.

Other programmes such as AIM (Applied Informatics in Medicine) have also provided funding for KBS work. For instance, the project GALEN has built a large KBS for medicine.

2.3 KBS applications

2.3.1 KBS in the UK

In order to identify current practices in KBS in the UK, the Department of Trade and Industry commissioned a survey of companies using KBS (DTI, 1992). The survey, which involved telephone interviews with 199 companies, uncovered many interesting findings about KBS users. The main findings were:

- KBS and expert systems are not just used by a few, very large, companies. Rather, KBS applications were found in both large and small companies in many different market sectors. In fact, 10% of the companies who were found to be using KBS were small companies with less than 100 employees. This indicates that KBS technology has the potential to help businesses of every shape and size, and that we are going to see more smaller companies using them in the future.

- The diversity of KBS applications is very large. This is because KBS technology is so versatile, and can be used in almost any field where human knowledge exists. The most popular application areas in the UK were found to be manufacturing and financial services, which together made up 62% of the sample.
- Applications currently in operation span a wide range of business areas. Manufacturing, production, marketing and customer service, and accounting/financial management are the most common application areas, but there are many other applications in areas such as training and regulation compliance.
- Systems performing diagnostic, advisory and assessment task types are most popular. Diagnostic systems are used to provide a diagnosis as to the cause or source of a particular problem, for instance why a particular part for a car is not being manufactured correctly. Advisory systems give advice as to how to perform a particular task, for instance how to set up a piece of complex machinery to make a new product. Assessment systems might be used for assessing the status of a particular manufacturing process or, in financial application areas, for assessing the risk involved in insuring someone or something.
- KBS can make knowledge widely accessible and can be used at many levels within an organization. Around one third of the companies who were using KBS said that their systems were being used by more than 50 people.
- KBS development does not have to take too long or be too complicated. Twenty nine per cent of the systems identified by the survey had been developed in less than six months.
- Personal computers are most frequently used for the development and operation of KBS. Many surveys have identified a trend towards using PC-based packages known as expert system shells for the development of small (and even some large) scale systems. We will cover the use of shells in Chapter 6.

Many of the companies interviewed were planning future work to enhance their KBS, and almost two thirds had new systems in development or planned. Thirteen per cent were even planning to sell their KBS as a commercial product. This means that, in the coming years, more and more KBS will be made available on the market.

2.3.2 Current applications

In one survey by Fahnrich of the use of KBS in the USA, the UK and Germany, the USA was credited with about two thirds of the systems, the UK had a little over one sixth, and Germany had a little less than one sixth (Figure 2.5). Particularly interesting was the difference in what the KBS were being used for.

The USA and UK concentrated on data processing, manufacturing and electrical engineering. The majority of the German systems were for mechanical engineering and plant construction. In a classification of Japanese expert systems by industrial field the top areas were all in heavy industry

KBS are now moving into the mainstream of computing. That is, more and more KBS applications are being integrated with traditional IT systems such as personnel systems, banking systems and so on. This means that KBS technology is becoming more widely used and that the skills of knowledge engineering are becoming more and more important and valued in the worldwide market place. It is these skills which we will aim to develop in you, the reader, throughout this book.

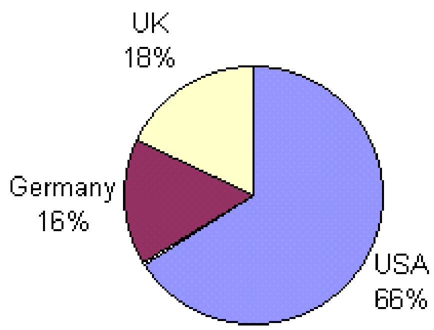

Figure 9. Proportion of KBS applications in the USA,UK and Germany

2.3.2.1 *Social applications*

In recent years there has been a large amount of research into the social implications of KBS; that is, systems which can be used to improve both the quality and quantity of advice and expertise available to the person on the street.

Systems include:

- a system which was designed to advise expectant mothers as to their maternity rights, including knowledge about eligibility for maternity benefit, maternity pay etc.
- a system, currently available on a PC, which gives advice on an employee’s rights with respect to dismissal, redundancy pay etc.

- a system which is designed to give home owners advice on planning law and answer questions such as “is my plan for an extension likely to be approved.

2.3.2.2 *Financial applications* The area of financial applications is currently one of the most popular applications for the use of KBS. KBS are in current use in many large financial institutions such as insurance companies, banks and finance houses.

The type of systems which are in common use include:

- KBS which aid bank managers when they are deciding whether or not to grant a loan to a particular customer;
- systems which give advice as to whether to grant a mortgage or not;
- systems which advise insurance companies as to the risk factor involved in insuring a particular individual or item;
- systems which are used by credit card companies to help them decide whether or not to issue an individual with a credit card;
- systems which have been devised to recognize and guard against computer fraud. This is becoming an increasingly large area. As computer fraud may often be difficult to detect, the knowledge which is encapsulated within such a KBS can be quite complex.

A survey by Kingston (1991) reported over 100 KBS applications in this sector in the UK and over 400 worldwide. Much of the work in this area remains secret. Companies are, without a doubt, developing (and using) some very sophisticated KBS “behind closed doors”. They are not likely to want to tell their competitors about pieces of computer software which are earning them lots of money and giving them a market advantage.

2.3.2.3 *Industry, manufacturing and the military* Many UK industrial and manufacturing companies have now introduced KBS into their daily operations. Applications in this area include:

- KBS which are capable of diagnosing various industrial faults, such as faults in aircraft, gas turbines and helicopters;
- systems which are designed to minimize plant downtime and avoid shut- downs, by identifying any potential problems as rapidly as possible — this is very important as, in large factories, a lot of money can be lost if things go wrong and the plant has to be shut down for any length of time;
- systems which are used to design and make small mechanical parts (such systems may form part of a much wider plant automation system which may also use robots) — this moves us some way towards

a goal of total factory automation, where much of the day-to-day work within a factory is undertaken by computers, machines and robots;

- military applications of KBS such as the identification of targets and potential threats to security — the applications of KBS in defence are wide and varied; it is often difficult to find out any details of such systems as the defence authorities around the world are not likely to go around giving away their secrets.

During 1989, the UK Department of Trade and Industry undertook a large survey of the use of KBS in manufacturing (DII, 1989). The survey reported that, although the UK was somewhat behind developments in technology in the USA in the field of manufacturing, there were currently more live applications (54) than found in, for example, Japan (where 32 such applications were identified).

2.4 The components of a knowledge-based system

This section describes the major components of a typical knowledge-based system (Figure 10):

- knowledge base
- inference engine
- •user interface
- explanation facilities
- learning facilities

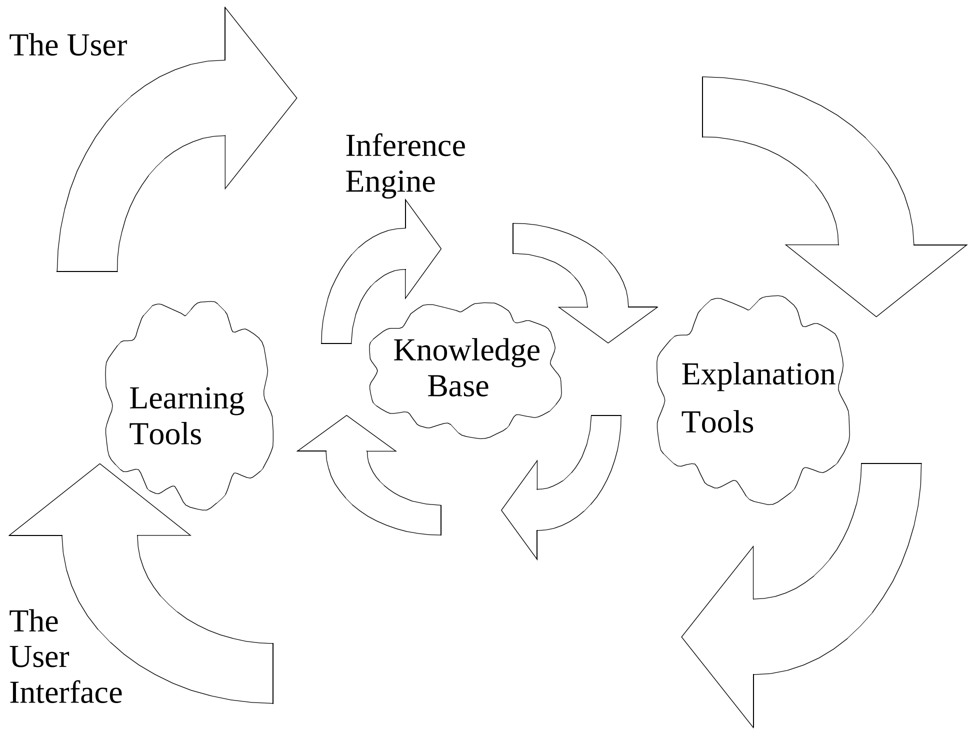

Figure 10. Diagram of a typical KBS.

2.4.1 The knowledge base

The knowledge base is obviously the heart of any KBS. It comprises human knowledge represented in some form. There are many ways which have been devised to represent knowledge within a computer and these are covered in detail in a later chapter.

One of the most common ways to represent knowledge is in the form of facts and rules. These facts and rules contain knowledge about the particular problem area from which the system draws its expertise (often termed the problem domain).

A fact is a clear, concise statement which expresses something which is true within the particular problem domain. For example, the following statements are all facts:

```
John is a man.
Susan lives in a house.
Tony drives a car.
Mary has brown hair.
```

A rule is generally of the form:

If *statement* then *statement2*.

where *statement1* and *statement2* are both expressions which may or may not be true. The rule is simply stating that if *statement1* is true then this implies that *statement2* is also true.

Examples of rules are:

```
If a person marries twice and he/she has not been divorced
then he/she is a bigamist.

If a person has $1,000,000 then he/she is a dollar million-
aire.
```

A KBS which is designed to aid medical practitioners in the diagnosis of abdominal pain would have a knowledge base which contained facts and rules about the likely causes of such pain and the particular symptoms related to each of these causes. Such rules might take the form:

```
If the patient has pain to the lower right of the abdomen
and the patient has vomited
then the patient may have appendicitis.
```

These facts and rules will usually be represented within the knowledge base in a form which the computer can recognize and deal with. This might involve the use of an expert system shell or an expert system language such as Prolog. Figure 11 shows some rules which are written in the form required for an expert system shell. The rules are taken from one of the case studies which is used in the remainder of this book.

2.4.2 The inference engine

In order to make use of the expertise which is embodied in the knowledge base, the system must also possess an element which can scan facts and rules, and provide answers to the queries given to it by the user. This element is known as the inference engine.

This is the thinking part of the KBS. The inference engine has the ability to look through the knowledge base and apply the rules to the solution of a particular problem. It is, therefore, the driving force of the KBS. It acts in a similar manner to the human brain in that it uses the knowledge from the knowledge base to provide an answer to a particular question (or a solution to a particular problem).

2.4.3 The user interface

The user interface is the means by which the user communicates with the computer system. Ideally this interface should be as English-like as possible to make it easy to use for inexperienced users. That is, an ideal system would

```
Find-Employee-Rule:If first priority group
                  and skills consideration
         and recommended moves and preferred work areas
Then display search result

First-Priority-Group-Rule:If staff type = internal
and staff resource centre=vacancy resource centre
and staff location=vacancy location
and staff grade=vacancy grades required
and staffavailability = yes
Then staff selected as first priority group staff

Staff-Type-Rule:ifstaff-type$="CS"
Then staff type = internal

Resource-Centre-Rules:Ifres_cent$=vres-cent$
Then staff resource centre=vacancy resource centre

Location-Rule:If site-code$--vsite-code$
Then staff location=vacancy location

Grades-Value-Convert-Rule:If convert grade descriptions into
values
Then grade evaluation

Grade-Evaluation-Rulel:If reql-value>=req2_value
and curr_grade_value<=req1_ value and
curr_grade_value>=req2_value
Then staff grade=vacancy grades required

Grade-Evaluation-Rule2:If reql-value<=req2_value
and curr_grade_value<=req2_value and curr_grade_value<=reql-
value
Then staff grade=vacancy grades required
```

Figure 11. Sample rules.

allow the user to type (or speak) his or her questions to the system in English. The system would then recognize the meaning of the questions, and use its inference engine to apply the rules in the knowledge base to deduce an answer. This answer would then be communicated back to the user in simple English.

Such an ideal system is, however, not likely to be available in the immediate future due to the difficulties which arise in trying to program a computer to recognize, and understand the meaning of, even the most simple English sentences and phrases.

One way in which AI is used in user-interface design is in the construction of adaptive user interfaces. As the name suggests, an adaptive user interface is one which adapts to the different needs of different groups of users. For instance a medical KBS might present a different series of screens and questions to an experienced consultant from those which it presents to a trainee doctor. In order to construct an adaptive user interface the computer must have knowledge about the different users of the system. It can then present each group of users with an interface which has been specifically designed to match their needs. A really clever adaptive interface will also be able to learn about users from the way in which they interact with the system. The interface can then change dynamically (i.e. over time) as it becomes more familiar with (i.e. gains more knowledge about) those people who are using it to communicate with the computer system.

The user interface is, in many ways, the most important part of any software system. It can make or break the system. A good user interface will be easy to use and even encourage users to use the system. A poor user interface will mean that the users will not enjoy using the system and will find it difficult to use. This may mean that the system will be rejected by its potential user community and will, in many instances, not get used at all.

The subject of user-interface design is a very large field of study, and much research has been undertaken to try and improve the way in which humans can communicate with computers. This has involved a lot of experimental work in which human subjects have been observed using different forms of interface.

2.4.4 Explanation and learning facilities

Any really useful KBS should be able to explain the answers that it gives to its users. That is, the system should be able to explain exactly why it has given a particular answer, and what knowledge it has used to reach a particular conclusion. It should also be able to explain why it has been unable to give an answer.

This is very important if the KBS is going to be really accepted by the users. That is, users often want to know why a particular answer has been given. If the KBS cannot reassure these users by letting them know how it arrived at a particular answer or conclusion they are unlikely to really trust it. This may mean that it will not get used as much as it should and it may even not get used at all. A good KBS should also be able to learn from its experiences, just as a human does. Human knowledge does not stand still. Our knowledge changes over time as we learn from our past experiences and from our mistakes. A really useful KBS should be able to do the same. That is, the KBS should develop with time as it learns from its own experiences and mistakes. This means that the knowledge base will grow as the system is used. This is how it should be. After all, your own knowledge is growing and changing all the time. For instance, when you have finished reading this book, you should possess more knowledge than you did before you picked it up.

2.5 Which problem areas are best suited to KBS treatment?

2.5.1 When to use a KBS

As we have learnt in this chapter, KBS are gradually finding their way into many areas of modem life. So should we expect to see a KBS appearing in every office, shop and factory in the near future? Perhaps not. Only certain types of application are suited to KBS implementation. But how do you determine whether or not a particular problem area is suitable for the 'KBS' treatment? It is the aim of this section to answer this question.

The following rules give an indication of the sort of criteria which must be satisfied in order for a particular application area to benefit from KBS treatment:

- The field under study should be able to be reduced into a series of rules rather than mathematical formulae or equations. In particular, a KBS may not be applicable if the problem involves a large number of complex calculations.
- The field under study should be well understood so that well-defined knowledge can be formulated and represented in computer form.
- The field under study should not encompass problems which take too short (i.e. less than half an hour) or too long (longer than, say, one week) a time to solve.
- There should be general agreement among recognized experts in the field. It is no use if all the experts have different ideas or theories - in such a case whose knowledge would you computerize?
- The knowledge within the problem domain should be sufficiently large to warrant the development of a KBS of, say, more than about a dozen rules - for less it is probably more efficient to solve the problem manually.
- There should be one or more 'tame' experts who are agreeable to their involvement in the project.

Notwithstanding the above rules, however, any application that requires access to specialist knowledge is a potential area for the introduction of KBS technology.

2.5.2 Benefits

But what exactly are the benefits that can be gained by the introduction of a KBS?

There are many benefits that can accrue from the introduction of KBS technology within a particular application area. Some of the more obvious benefits include:

- A KBS can make knowledge and expertise much more accessible than would otherwise be possible.
- Using a KBS can be much cheaper than hiring the services of a real human expert.
- A KBS can be used to preserve knowledge which would otherwise be lost over time.
- Computers are not prone to human error in the application of their knowledge.
- KBS can be used to facilitate communication between humans themselves and hence improve their own knowledge.

2.6 The practical case studies

Throughout this book we will use four practical case studies to illustrate the major concepts of knowledge engineering. We will start here by introducing you to each of the case studies. The case studies have been chosen to illustrate the development of four different, but commonly found, KBS applications.

As you progress through the book, you will meet the case studies again and again. Each time, however, we will move a little further towards the development of the final KBS. By the time you have finished reading the book, we will have covered portions of the actual computer programs which were used to implement the KBS.

The case studies are:

1. a quality control system, which illustrates how KBS technology can be used to give workers advice on how to recognize (and rectify) faults in the products they are making and thus exercise good quality practices (Hardy, 1993);

2. a forecasting system, which shows how a KBS can be used to predict future trends (Leonard, 1992);

3. a monitoring system, which demonstrates how a KBS can be used to check on the status of an important piece of equipment (Jawad, 199 1);

4.an integrated system, which shows how KBS technology can be integrated with conventional computer software technology (Lin, 1993).

Each of the systems are taken from genuine practical experiences. They are based on KBS which were developed for the companies concerned.

2.6.1 The quality-control system

The company concerned in this case study is part of an international group of companies associated with the manufacture and distribution of vending and food services packaging. The particular problem which we will look at in this case study is the manufacturer of plastic cups, the sort of cups which are used in vending machines and which we drink our coffee or hot chocolate out of Each and every stage of product manufacture is handled by the company, from the raw materials through to packaging and distribution. One of its central concerns is the quality of the final product. This is monitored manually both by the operators and the Quality Assurance Department. Since the output of the machines can exceed 80,000 cups per hour, it is very important to keep the product within its specification, and to avoid having to stop the machines to put things right. This is because once a machine starts making cups incorrectly, it will continue to do so. This will soon result in a large number of spoilt cups and a lot of wasted production time.

Currently, the operators are sufficiently well trained to be able to deal with most of the common problems which can arise. However, when problems are unfamiliar, or do not respond straight away to adjustments to the machines, expert diagnosis is obtained either from more experienced operators or from production/development personnel. As is typical of such situations, having to bring someone in takes valuable time, costs money and takes the expert away from whatever he or she was involved in.

The task in this case study was, therefore, to develop a knowledge-based diagnosis and fault-finding system that could be used by the operators to assist in their operations. In particular, the aim of the KBS was to find any faults in the production of cups.

Whilst visually the problems affecting the cups are not large, even with very small problems the product is still flawed, and this can lead to problems and customer dissatisfaction. It is also hard to determine the cause of the problem, since it can happen at any point in the production process. There are nine typical problems that can occur with a cup, and the KBS was developed to recognize each of these and then give the operator appropriate advice.

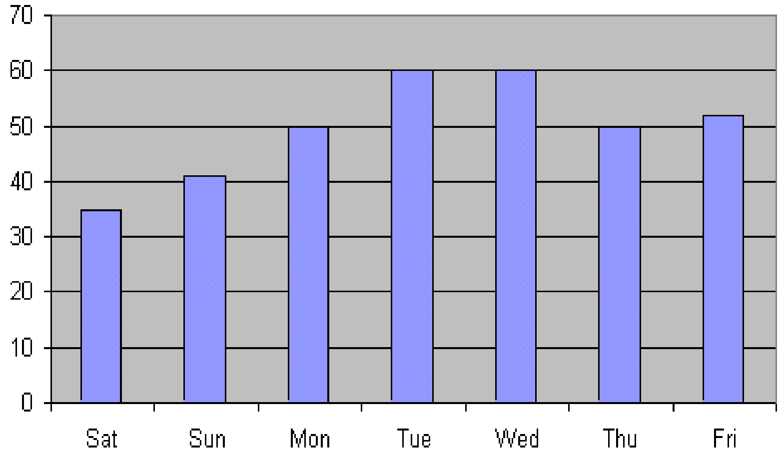

Figure 12. A KBS designed to estimate gas demand.

2.6.2 The forecasting system

Forecasting is the practice of estimating, as closely as possible, future trends. This KBS aimed to improve and enhance the forecasting techniques used in British Gas, and in particular within one region. Its aim was to improve the accuracy of estimation of gas usage. This project aimed to investigate the use of a KBS to improve the accuracy of the short-term forecasting that is carried out daily within each region of British Gas (see Figure 12).

British Gas buy their supplies of natural gas from over forty privately owned gas fields situated around the coast of Great Britain, and transport it through its network of steel pipes to the gas regions and then to the customers.

The amount of gas needed daily must be ordered from the suppliers twelve hours in advance, and as the gas day traditionally runs from 6 a.m. one day to 6 a.m. the next, this means the gas wells are notified by 6 p.m. The amount of gas needed is calculated by British Gas Headquarters, primarily by summing the estimates of the amount of gas that each region expects to sell the next day.

To allow Headquarters time to firm up their estimate they require each region's estimate by 4 p.m. Each region is therefore required to forecast their estimation of the volume of gas for the following day 14 hours in advance of that day starting. Not only have they to estimate this far in advance but they are required to forecast as accurately as they can, as they can lose money if they get it wrong. It was, therefore, vital to develop a KBS which could accurately forecast the gas demand required by the region for the following day.

This is not easy as gas demand can depend upon several factors:

- seasonal factors (we use more gas in the winter than in the summer);
- the weather on a particular day;
- other factors, such as, for example, on the afternoon of the Cup Final most people will be in the house with the central heating on.

2.6.3 The monitoring system

This case study concerns a KBS which was designed to give advice on the initial setting and subsequent adjustment of a mechanical ventilator used in a hospital intensive care unit.

The ventilator is used to control the breathing of seriously ill patients who have often been rushed into hospital. Previous manual systems relied heavily on the intuition of the doctor as to how to set the ventilator, so as to stabilize the patient in as short a period of time as possible.

A ventilator is a mechanical device which ventilates a patient's lungs or in other words 'does the breathing for the patient'. All modem ventilators achieve this by periodically forcing gas into the patient's lungs.

In normal clinical practice, the initial settings of the ventilator are decided by the doctor on the basis of a number of known or estimated physiological factors such as weight, height, age, lung condition, heart condition or disease.

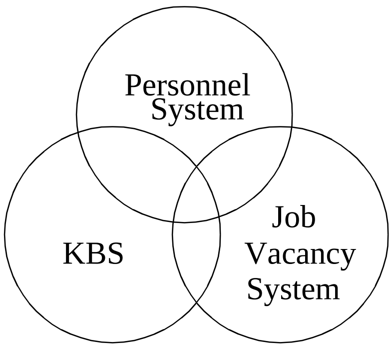

Figure 13. The integrated HRM system.

Ventilator adjustments by the doctor tend to be made according to rules that are not hard and fast, and there always exists an element of 'inspired guesswork' in the process. Measurements of blood gases, patient condition and ventilator adjustments must be made as frequently as necessary to keep the patient stable.

The goal of the KBS described in this case study is to assist the doctor in the intensive care unit by performing the following:

- predicting the initial setting of the mechanical ventilator;
- suggesting adjustments to the settings of the ventilator;
- summarizing the patient's status.

2.6.4 The integrated system

This case study concerns an application where a KBS in integrated with a conventional software system. This situation is becoming more and more common as KBS are accepted in the everyday business world. The system in this case study is a human-resource management system which was developed for a large organization. A human-resource management (HRM) system is a software system which is designed to make the best use of the people in the organization. In this case the main aim of the system was to identify which staff were best suited for particular projects and to identify the training needs of those staff.

The department employs approximately 4000 staff, and is sub-divided into a number of areas, each responsible for specific services. Each of these areas has a resource manager who is responsible for a number of projects. The duty of the resource manager is to fit a suitably qualified person to a specific job within each of the projects which they are responsible for.

Any vacancy not filled within any one area is then made available across the whole department. This involves staff being transferred between areas. If no suitable staff are available internally, the resource manager must turn to the use of external consultants. Because of the expense involved, these are usually only employed on high-priority, short-term projects.

The main task to be done by the KBS is to match staff with suitable projects. The KBS must, then, hold knowledge relating to staff skills, location, availability and other personal factors. The new system was required to use the existing personnel database and job vacancy database (Figure 13).

2.7 Summary

This chapter has given a feel for the field of knowledge-based systems. The history of KBS has been covered along with common application areas, both past and present. We have also introduced the case studies which are to be used as the basis for much of the following text.

2.8 Exercises

1.Using the categories defined by Waterman (section 2.1.2), place the following KBS into the most appropriate categories.

- A KBS which checks the condition of turbines in a power plant and informs the operators when there are possible problems. The operators must then make the decision as to whether the turbine needs shutting down and repairing.
- A KBS which asks a car buyer the attributes of the car which he/she is looking for. The system will then try and match a car to the customer's exact requirements.
- A KBS which is used by a shop-floor manager in a factory to generate work schedules.
- A KBS which is used by a personnel manager to produce a shortlist of applicants for a job.

2.Categorize the four case studies described in sections 2.6.1-4 according to Waterman's categories.

3.Following this question is an extract from the instructions for a new portable colour television. Study the extract and identify:

- five facts
- five rules

Write these facts and rules in English.

Your new colour television

Safety

1. **Moisture.** Avoid positioning the set where extreme dust or moisture is likely. The unit should not be exposed to dripping or moisture such as rain, dew, condensation, etc.

2. **Vibration**. Place the unit on a stable, level surface, avoiding places subject to strong vibration.

3. **Magnetism**. Do not position the set near strong magnetic fields such as loudspeaker enclosures as they can adversely affect the picture quality.

4. **Accidents**. Keep such things as flower vases, coffee cups, etc., which may get split, away from your television. If liquids should be spilled into the unit serious damage may result. If you spill any liquid into the machine, disconnect the mains plug from the wall socket or remove the car battery lead and consult qualified service personnel before attempting to use the equipment again.

5. **Cleaning**. Do not spray cleaner or wax directly onto the unit.

*Power Indicato*r

Plug the mains lead into the wall socket before switching on the power. When the power is switched on a power indicator will glow constantly during operation and stand-by modes. If the power indicator doesn't glow check the fuse in the plug and confirm power is available at the wall socket.

The Aerial

Your television is fitted with a built-in telescopic aerial. The aerial should be extended and rotated for best reception.

When using the unit in a caravan or car, or in areas of poor reception, an alternative aerial may be required. The type of aerial needed depends on the position of your television, its distance from the transmitter, local interference, etc. If you don't already have a suitable television aerial and your are uncertain of what to buy, you should seek the advice of a local professional aerial contractor.

If an alternative aerial is used, make sure that it is correctly installed. Take the aerial plug and insert in into the 75 ohm aerial socket which can be found on the back of the television.

Remote Control

The remote control handset has an operating range of approximately 7 metres. For best effect it should be pointed directly at the remote sensor on the front of the television. However, operation is possible up to an angle of 30'.

The remote control may not work if:

- he batteries are low.
- The batteries are inserted incorrectly.
- The path from the Handset to the television is obstructed (by furniture, etc.)