



The Potential of Knowledge Based Systems in the Mineral Industry

**Brenda Romans
Intelligent Networked Computing Laboratory
HP Laboratories Bristol
HPL-93-97
November, 1993**

**knowledge based
systems**

The practical uptake of KBS (knowledge based systems) research by industry has been slow but it is beginning to increase as KBS techniques and technology become understood and accepted. KBS can improve the profitability of a company, but in order to do this the right technology has to be implemented in the right way for the organization concerned. This paper does not attempt to define a formula for the foolproof adoption of particular KBS technologies by the mineral industry. It introduces KBS and describes the problems that can be addressed by them, and the problems that should be considered during their development. It draws attention to the importance of expert knowledge in an organization and how that knowledge may be distributed across functional boundaries. The key to the profitable assimilation of KBS in an organization is to focus on acquiring knowledge and storing it in a manner that will facilitate its use by any section of that organization.

Internal Accession Date Only
Presented at APCOM 93, *Applications of Computers and Operations Research in the Mineral Industry*

© Copyright Hewlett-Packard Company 1993

1 Introduction

The uptake of KBS (knowledge-based systems) in industry has been dependent on the nature of the industry concerned. For example, there has been enormous interest in KBS and AI (artificial intelligence) in medicine even though there are ethical and legal barriers to implementation; medical researchers are prominent in the field of AI. The interest shown by the mineral industry has, in general, been similar to that shown by the manufacturing industries: a cautious investigation followed by spot implementations of the technology in disparate places.

The APCOM conferences can be viewed as a barometer indicating the importance placed on KBS technology within the mineral industry. APCOM papers are gathered under section headings that reflect either computing technologies or functions within the industry such as mineral exploration and mine design. The early KBS submissions were originally incorporated under their application headings but now there are also specific technology sections for expert systems and robotics. In a competitive industry, organisations are slow to publish the results of innovative technology that might help a competitor gain advantage. Consequently the reporting of KBS in use will not give a true picture of their current extent.

After briefly characterising KBS and putting them in perspective in the AI world, this paper describes the types of problem that can be addressed using KBS, with examples from the mineral industry. It goes on to consider the problems involved in implementing KBS and suggests methods for dealing with them. Section 6 discusses the importance of knowledge acquisition for knowledge based systems, and suggests how this activity alone can benefit an organisation.

Knowledge-based systems

In the study of AI, workers strive to emulate some part of human activity such as movement (robotics), vision, or reasoning. KBS have developed out of this work and they can be defined as "computer systems that emulate some defined part of the reasoning capability of an expert". For this reason KBS are also known as expert systems. Other types of system that deal with problem solving under the AI umbrella are neural nets, genetic algorithms, induction systems, and case-based reasoning. All four "learn" by example rather than by incorporating human knowledge and are not considered in this paper, (for more information see Carbonell 1990).

A KBS can be thought of as comprising three parts:

- A set of interfaces, to humans and machines, for acquiring and dispensing information.
- An inference engine; a mechanism for manipulating the knowledge in order to solve problems.
- A knowledge base containing the distilled wisdom of one or more human experts in a particular domain.

Walters (1988) and Waterman (1986) describe the fundamentals of KBS in detail. KBS research considers the design, development and implementation of these components and has a very practical bias.

2 Problems that can be addressed using KBS

KBS can approximate part of the reasoning capabilities of a human, and thus they will enable this expert behaviour to be made available when the expert cannot, but within very severe boundaries of operation. The most obvious problem that can be addressed using KBS is a shortage of skilled staff. The best state-of-the-art KBS technology will not replace experts, and it will not do so in the foreseeable future. It can however, be used to capture a fraction of their expertise in a particular domain and to make this available to less skilled (but more readily available) staff. This is of particular importance where the site requiring the expertise is remote (as are most mines) and skilled staff are difficult to recruit and keep - or where these skills are just plain scarce. It is also useful to note that the most valued experts in an organisation are often those approaching retirement age.

A different situation can arise when experts are available, but the geographical extent of the organisation is so large that they have to deal with problems at a distance. By incorporating a suitable subset of the expert knowledge into a KBS the people at the problem site can be guided through a troubleshooting procedure. With the information thus gathered it will be much easier for the expert to deal with the problem remotely.

When people cannot process information because it is arriving at too fast a rate for comprehension a KBS can be used to interpret some of that information. Without KBS, merely reducing the information (for instance by simple averaging) may lose some important details that would be observed by noting each individual reading. A KBS can be designed to read the information and provide details of any behaviour recognised to be of interest to the operator. This function is most useful when running a process that a) gives frequent updates on the state of the process parameters or b) bombards the operator with un-filtered alarms. The expert input to the KBS would be knowledge of the process, and the implications of the parameter values and alarms.

In other situations the information is too complex for the human operator to be able to deal with it in a timely manner. What is "timely" depends on the situation: finding a fault in a device whose inactivity is causing a smelter to be off-line is somewhat more urgent than diagnosing a problem in the engine of a backup haul truck. The information may be overwhelming because of complex interactions or because of the sheer volume involved. This type of problem can occur in many situations, (process control is an obvious candidate) and a diagnosis system can soon incorporate more information than an expert working from memory can be expected to handle. Another area where complexity and information processing can suitably be handled by KBS technology is in scheduling, where the complexity soon leaves the operator behind.

The task of teaching new engineers, or teaching existing staff to use new equipment has been investigated to some extent by most people developing KBS. If there is a body of knowledge that is to be used to give advice to operations staff then it follows that that knowledge could be used for training purposes. In fact people can "grow out" of an expert system, with frequent use they learn from the

system and find that they use it less and less while maintaining an improved performance.

By using KBS to alleviate these problems organisations can make more efficient use of human resources. By making the human operators more effective noticeable improvements in profitability can be made by:

- Reducing the down-time of critical components of the operation.
- Predicting impending problems.
- Optimising the management of physical resources.

It may be factors like these and not the more immediately related human angles that provide the impetus for an organisation to consider implementing KBS.

3 Successful KBS in the mineral industry

Fault diagnosis is a classic area for the research into, and deployment of, KBS. More research has been carried out in this area than any other, and as a consequence the practical benefits became available quite early.

The SHEARER expert system described by Perkin et al. (1986) diagnoses faults in a shearer, a machine for cutting coal from a seam on a longwall face. This system was deployed in British Coal mines to enable shearer faults to be located quickly to reduce shearer down time, thus reducing breaks in coal production. At a potential production rate of 300 tonnes per hour and at £30 per tonne there was real incentive to determine the cause of hold-ups quickly. The system was constructed to provide diagnosis advice to operators, and enabled them to perform diagnoses outside their range of experience. This was possible because the information contained was taken from a range of experts including operational engineers from the coal industry, and engineers from the machine manufacturer. The knowledge stored did not represent a deep model of the machine but a set of causal relationships, in this case faults related to symptoms. A deep model would contain detailed information on the structures and processes defining the machine and would be more difficult to develop and to work with.

An Exxon chemical plant in Scotland has approached the problem of health monitoring with visible success (Milne 1993). They have installed a system to monitor a gas turbine and give early warning of problems. As an additional advantage the information from the system enabled the engineering staff to identify the cause of a problem that had been present for years. As a result they were able to increase production and estimate that this alone is worth £1.5 million per year.(Intelligent Applications Ltd.1993)

The LINKman system has demonstrated in a different way how successful KBS can be in reducing costs. LINKman was originally developed to help operators control cement kilns, a highly skilled and difficult job requiring constant surveillance. As with SHEARER they took the best information available and built a system to help the operator to make the most effective use of the resources (fuel, feed and refractory) as well as to improve the quality of the product. This was so successful that the system has been developed further, more information has been added to the knowledge base, and it is being marketed commercially.

Rather than offering advice to the operator the LINKman system is able to act as an auto-pilot and can run the process alone for a large percentage of the time. This frees the operator for other tasks and helps reduce inconsistency problems that can arise with shift changes. Haspel et. al. (1993) report that the system has shown benefits in operation at nine sites around the world. Differences between the plants mean that different benefits are realised but overall they are reporting typical ranges of:

- 2.5% to 5.0% increase in production
- 2.5% to 5.0% decrease in fuel consumption
- 25% to 100% increase in refractory life

with best achieved performances in excess of these, and with quantifiable benefits in other areas.

4 Difficulties in developing KBS

While a knowledge-based system has the potential to solve certain problems and thereby increase productivity and profit, this cannot be achieved without some effort. In fact, considerable work can be involved in implementing KBS.

When developing a KBS there are methodologies such as KADS (Hickman 1989) that can help, but the types of problems that commonly arise in developing conventional computer systems also occur. Significant obstacles involve performing testing and maintenance. It is important that the system is designed and tested to confirm the following:

- That the information contained in the system is valid.
- That the system contains no coding errors.
- That the system performs the intended function.

As a KBS holds and dispenses human knowledge in the real and changing world it should be capable of change so that it continues to reflect the real world and can incorporate the results of human learning experiences. This means that the system should be built in such a way that it is easy to maintain.

As described earlier a KBS can be thought of as having three components, a set of interfaces, an inference engine and a knowledge base. To focus on other common potential difficulties the next three sections consider each area in turn and briefly describe some of the decisions that must be made when developing each KBS component.

Interfaces

There are two principal ways of getting information into the system. It can be entered by an operator or it can be taken from sensors or another computer source such as a database. Accessing other machine sources may not be trivial but it will be governed by the protocols that exist between machines. The design of the user interface is more involved. The common problems that occur during the development of any software interface apply, but KBS offer a few extra pitfalls. KBS are commonly used interactively and are often intended for use by more than one category of user. Unfortunately a single common interface is not always able to

provide the facilities required by each of these groups. To compound the problem a KBS can also (and rightly) be seen as a training tool because of the information that it contains. While users can become more knowledgeable by using the KBS it is unusual for the same interface to be suitable for training novices.

In summary, because the rich information content of a KBS promises information for all, there is a temptation to force one interface to do all jobs. To gain user acceptance (without which the system is useless) the interfaces must be very carefully designed. Pauley (1992) offers practical advice on the design of the human interfaces and methods for gaining acceptance from the users.

Inference engine

There are many different ways of implementing the inference engine and there have been many books and papers written on the subject. The choice depends on such factors as:

- The type of knowledge involved.
- The task to be performed.
- The amount of knowledge needed.
- The complexity of the reasoning.
- The ability and experience of the system implementers.
- The ability and experience of the proposed system maintainers.
- Whether interoperability with existing systems is required.
- Whether the policy of the organisation is to develop a one-off system or to experiment for the future.

It is obvious that such decisions are very specific to the task and to the organisation involved and therefore particular recommendations cannot be sensibly made here. It is useful to note however that organisations will often allow adventurous individuals to experiment with KBS as a sideline, using whatever KBS tools, techniques and methodologies come to hand. This is not to be discouraged but a follow-up behaviour of treating these small and singular experiments as though they were informed and thorough is unfortunate and common. A key fact behind the SHEARER, LINKman and Exxon successes is that behind them were well set up teams of people to investigate the technology and select the most appropriate for the job.

Knowledge base

The task of creating the knowledge base is split into two distinct parts: firstly obtaining the knowledge, and secondly putting it together in a coherent form that is compatible with the inference engine. The task of constructing the knowledge base is like that of constructing the inference engine and the above comments apply. In acquiring the knowledge however there are common difficulties that can be avoided with forethought. Some of the major points to consider in collecting expertise are:

- The expert must be truly expert.

- Using multiple experts is desirable from the point of view of information coverage but a difficult thing to manage.
- Knowledge from other (e.g. textbook or case study) sources can be used but has to be carefully integrated with the expert information. This will take up expert time.
- Knowledge elicitation takes a lot of expert and knowledge engineer time.
- Testing the knowledge requires the participation of the expert and can take as much time as the initial knowledge elicitation.
- There are many techniques that can be used to elicit knowledge and their selection depends on such things as the task, the nature of the information and the personality of the expert. Suitable methods have to be selected for each case.
- Care must be taken to represent the knowledge in a workable format.

The major point to be made here is that when deciding to build KBS it is important to ensure that enough expert time is available. This is often difficult to arrange as experts are in high demand.

5 How to take advantage of KBS. . .

It may appear from looking at the list of problems that have to be considered in the development of KBS, that they are difficult to develop. They are certainly not easy but, as demonstrated by the examples given (and there are other notable successes both in the mineral and other industries), there are large benefits to be realised if the right approach is taken. Some guidelines to finding that approach are given below.

Choice of project

Firstly and most importantly: *do not use KBS where they are not appropriate or not necessary*. KBS can be used where human knowledge is required to solve the problem. As an example, in some cases of smelter operation, engineers have spent years attempting to model the system mathematically without a lot of success. There are usually several operators who are very good at running the process and their knowledge can be used to advantage in a KBS. In contrast, in ore reserve estimation geostatistics is used successfully to model the deposit. It would be unusual for a mine manager to base his operations on the advice of a venerable old miner saying "dig here". However, a KBS could be used to provide an interface to a geostatistics package to enable a geostatistician to work more efficiently.

This leads to another point: do not be afraid to integrate a KBS into existing systems - it is merely another source of information.

Planning a project

Starting in a small way by buying an inexpensive KBS shell or toolkit is a practical possibility. This is a low cost way of experimenting with KBS technology, but there is a risk that negative results will be seen as failure if the exploratory nature of the work is not made explicit. A common mistake is to employ a student to perform some exploratory KBS work, and to have this disappear when the student leaves.

If the experiments are well documented it is possible to build on this knowledge and subsequently tackle more ambitious tasks.

If there is a problem that is amenable to solution by KBS and one which promises a large payback, then it is worthwhile seeking professional advice. There are AI professionals who have a proven track record in industrial consultancy, and companies providing KBS solutions with proprietary software. Some mining companies have AI or KBS sections that can perform this function in-house. KBS professionals can be recruited, or existing computer staff can be trained to bring them up to strength in chosen subjects.

Testing and maintenance

The development, testing and maintenance of KBS has a lot in common with that of conventional computer systems. It follows then, that taking a similar well-structured approach will yield the same benefits and reduce the risk associated with the project.

Testing KBS is a problem, and one that is receiving a lot of attention. There are several strategies for ensuring that your system is as error free as possible and performs as intended:

- Look to conventional computing in addition to KBS research for basic testing methods. (Myers 1979) (Hetzel 1988)
- Prepare for testing from the beginning of the project.
- Allow plenty of expert time for testing. If the system is to reproduce some of the behaviour of the expert then it is the expert who will be able to confirm that it is working correctly.
- Use the KBS as an advisory system rather than closing the loop on a process. Allow the user to see what is going on.
- Give the user as much information as possible. This will help with user acceptance and education as well as giving information that could help detect faults in the KBS reasoning.

Plan for maintenance and development from the start of the project. Human knowledge grows and a KBS must expand too if it is not to fall into disuse. It must be determined who is to carry out the maintenance, how it is going to be accomplished and how the system will be re-tested.

6 . . . without actually building one

Knowledge elicitation (KE) for KBS is the process of obtaining the knowledge necessary to perform a chosen task. It is often seen as the bottleneck in KBS creation as it can take some time to tease all the information from the expert(s). It is tempting to extract the knowledge in a format that will suit a particular tool rather than the task and the expert. When this happens the knowledge becomes obscured and difficult to use without that tool. It is better to treat the knowledge base as a resource in its own right. A knowledge base is a structured, coherent body of knowledge that can be used by any person requiring that information. A number

of different KBS tools can then be used to deliver this knowledge, or it may be disseminated without the need for additional processing.

Expert knowledge can be used without building a KBS. A long-realised benefit of KE is that the expert's performance improves as a result of the acute focus necessary during the process. The act of gathering information can also be beneficial because it causes the organisation look for and assess its expert resources. This can be illustrated by considering the use and availability of diagnosis knowledge to find faults in a device being manufactured. This knowledge concerns fault diagnosis at various stages in the design, manufacture and use of the device. In a factory, as in a mine or mill, the organisation is split into different departments that have separate and well defined tasks. This partitions the knowledge into sections that reflect the structure and processes of the organisation but not the content of the knowledge. Table 1 lists the type of diagnostic knowledge available in the different departments of a factory manufacturing a device such as a gravimeter.

Table 1. diagnostic knowledge available in a factory

Department	Function	Knowledge
Repair shop	Diagnose and fix faults	Common device faults and fixes.
Service	Diagnose faults and fix devices on the customer's site.	Common device faults and easy fixes. How individual customers use or mis-use the device. The device environment.
Manufacturing	Manufacture and test the device for faults introduced in manufacture.	Manufacture induced faults, including those induced by testing regimes.
Marketing	Understand the use to which the customer puts the device.	The ways in which the customer uses the device.
Development	Design the device and future modifications.	A deep knowledge of the device at a basic component level.

Communication between departments may be good but is typically conducted at a high level to discuss such items as shipping dates and customer requirements. Diagnosing faults in the device is a problem that affects all sections of the factory. As departments have well-defined responsibilities, each has enough information to deal fairly effectively with its own tasks. Problems occur when the expertise in one department is not sufficient to perform its function adequately. For instance, the repair shop may be unable to deal with novel failure modes introduced by a new process in manufacturing. A closer look at the knowledge available throughout the site shows that all the knowledge listed in the table could be brought to bear on the device diagnosis issue. Each department has information on device fault finding but it is not always explicit and is not often exchanged.

If a suitable format and method were developed the knowledge could be recorded so that:

- The knowledge elicitation is straightforward.
- The required knowledge can be extracted from the knowledge base easily.
- The knowledge base can be constructed incrementally, as experts and new knowledge becomes available.

The information could be made available on paper or on a computer system. Using a KBS can add value by offering explanations, tuition or by intelligently searching the information to save time in diagnosis. Different KBS could be used for different users or functions in each department. Both formal and informal techniques for acquiring expert information have been investigated in KBS work, and many practical recommendations have been made e.g. (Burton et. al. 1987) (Kidd 1987) (Firlej and Hellens 1991).

The same situation can be seen in the mineral industry where there are pools of knowledge that are complementary but not exchanged. This may be the case because of a lack of communication, or because the knowledge is structured in a different way. For instance; knowledge of the fluctuation in composition and reactivity of a mined material may have an impact on:

- The degradation of any material or machinery that it comes into contact with.
- Ventilation management.
- Recovery plans (e.g. reaction times, chemical quantities, crushing regime)
- The value of the recovered material.

The knowledge held in each department could materially affect the operations in the others if a common knowledge base allowed it to be shared.

7 Conclusions

As research into the development of knowledge-based systems continues, the practical benefits are being understood and used in industry. It has been shown in three examples from the mineral industry that using KBS can lead to substantial savings. Operational improvements can be made by implementing KBS wherever human knowledge is required but is in short supply. They can profitably be used to provide expert help to technicians or to augment the performance of experts. Developing successful KBS however is not trivial. There are many methods, techniques and tools available and these need to be understood and matched to the problem to be solved.

The knowledge incorporated in a KBS is valuable in its own right. Using suitable techniques to capture and represent the knowledge can improve the performance of the expert and accumulate a body of knowledge that can be transferred to other people using a variety of means, not necessarily KBS. The functional divisions within an organisation take into account the responsibilities of those divisions but not the distribution of expertise within them. By looking at the demand for that knowledge within the organisation it is possible to identify other areas that could

benefit. With a knowledge base in place the task of developing KBS for different applications is then much more straightforward.

Acknowledgements

The author would like to thank Graeme Pauley of RTZ for many valuable discussions on the subjects presented.

References

- Burton, A. M., Shadbolt, N. R., Hedgecock, A. P., Rugg, G., A Formal Evaluation of Knowledge Elicitation Techniques for Expert Systems: Domain 1, in Research and Development in Expert Systems IV. ed. Moralee, Cambridge University Press, (1987)
- Carbonell, J. G., ed. Machine Learning: Paradigms and Methods, Elsevier, (1990)
- Firlej, M., and Hellens, D., Knowledge Elicitation - A Practical Handbook. Prentice Hall, (1991)
- Haspel, D., Henderson, W., A New Generation of Process Optimization Systems For the Cement Industry, Artificial Intelligence in the Minerals Sector Colloquium April 1993, eds. Denby and Schofield, IMM and Nottingham University, preprints, (1993)
- Hetzel, B., The Complete Guide to Software Testing. 2nd ed., QED Information Sciences, (1988)
- Hickman, F., Killin, J., Land, L., et. al. Analysis for Knowledge-based Systems: A Practical Introduction to the KADS Methodology. Ellis Horwood, Chichester, (1989)
- Intelligent Applications Ltd. Expert System Provides Automatic Diagnosis of Gas Turbine Malfunctions, Expert Systems Applications 9 no.1 (1993) 3
- Kidd, A., (ed.) Knowledge Acquisition for Expert Systems - A Practical Handbook. Plenum Press, (1987)
- Milne, R. W., Getting the Payback from Expert Systems, Artificial Intelligence in the Minerals Sector Colloquium April 1993, IMM and Nottingham University, eds. Denby and Schofield, IMM and Nottingham University, preprints, (1993)
- Myers, G. J., The Art of Software Testing. Wiley, (1979)
- Pauley, G. S., Successful Implementation of a Knowledge-based, Smelter Operations Scheduling System at the Palabora Mine in South Africa, IDASCO 92, IEE conference publication no. 352, (1992)
- Perkin, R. M. G., Pitt, M., and Price, A. E. Expert systems: Practical Applications in a Traditional Industry, in Research and Development in Expert Systems III, Proceedings of Expert Systems '86, ed. Bramer, Cambridge University Press, (1986) 231-241
- Walters, J.R., Nielsen, N. R., Crafting Knowledge-based Systems. Wiley, (1988)
- Waterman, D. A., A Guide to Expert Systems. Addison-Wesley, (1986)