

## Expert Systems Technology: Just a Flash in the Pan?

R. LAYTON

*Department of Computer Science, University of the Witwatersrand, Johannesburg, South Africa*

Expert Systems Technology (EST) is becoming more and more widespread in both its range of applications and the public interest. It appears to be capable of performing tasks unmatched by other techniques. Within the mining industry a large body of unsolved problems now appear to have a solution. This paper examines the current state of EST both in the context of the history of Artificial Intelligence and current efforts at building mining expert systems. It provides recommendations for those wishing to get involved, and outlines the problems for those who believe that expert systems are capable of anything and everything. The paper also provides a detailed analysis of limitations of current EST that must direct research into more fundamental, commonsense areas than have been extensively covered before.

### Introduction

This paper deals with some issues concerning the effective utilisation of Expert Systems Technology (EST) within problems of information modelling in the mining industry. A short history of Artificial Intelligence (AI) is presented to give a general perspective on the temporal location of this very new science. This is followed by a brief introduction to some of the mining and geological projects that have been done using EST in the past or are currently under development. The next section explains the general process of building expert systems. This is followed by a section that identifies some of the things to watch out for that indicate that a problem is potentially viable for EST. The closing section covers what we still have to do in order to build expert systems that can simulate common-sense understanding of problems.

### A short history of AI and EST

AI can be said to have begun life at the Dartmouth conference of 1956.<sup>1</sup> It was the first time that researchers from around the world had gathered for an extended period to consider what John McCarthy described as 'Artificial Intelligence', which is loosely

defined as a machine or a computer program that behaves as if it was intelligent by human standards. Today many consider this term a misnomer, but no replacement term has yet gained the same popular acceptance, and the subject is changing so fast that it is unlikely any suitable descriptive term will be found to last more than a few years. So we are stuck with the term 'AI' to describe the general field of the modelling of intelligent human capabilities. Other terms in common usage such as Expert Systems, Knowledge Representation, Natural Language Understanding, Machine Vision, Robotics, and Cognitive Science, describe its branches.

At the Dartmouth conference the general feeling expressed by the participants was that all problems involving 'intelligence' appear to require a similar approach to symbolic processing. No distinction between problems classes had been identified and exploited. The games of checkers and chess, computer vision and natural language understanding were considered to be of the same magnitude, at least that is what the successes with the toy problems had indicated.

These toy demonstration problems were

implemented on the basis of a philosophy that holds that methods used for micro-worlds are directly applicable to macro-worlds. This proved to be an incorrect assumption.

In the 1960s a general problem-solving system was built by two of the Dartmouth attendees, Allan Newell and Herbert Simon.<sup>2</sup> It was called GPS - the General Problem Solver - and whereas it appeared to be general-purpose, it required much more effort to program than the solution was generally worth. Even coding the simple toy problems was a large exercise. This led to the need to distinguish different problem types and developed into what we now know as Expert Systems, in which the concept of 'domain knowledge' is emphasised.

The development of these large expert systems took us through the 1970s as we realised that games and other formalised systems are easy to implement, but real-world problems are much more difficult. We learnt not to underestimate the time and costs needed to build a large real-world system. The standard quotation was '10 man-years and \$2 million' to get a system up and running - including hardware and software. Whereas many 'knowledge engineers' (professional expert system builders) currently use small tools to solve demonstration problems, the above estimate is not exaggerated for the large complex problems we are now trying to solve.

The 1980s saw the drive towards the expert system 'shell' concept in which the basic inference procedures and representational structures were taken out of the existing large systems, such as PROSPECTOR (to be dealt with later), so that knowledge from other compatible domains could be installed at a fraction of the cost.

The computer industry never misses a good opportunity and since the start of the 1980s the amount of 'shells' on the market has grown exponentially. At the last count there were 500+ available to choose from, each with distinctive features such as novel approaches to uncertainty propagation (often quite unstatistical!), powerful user interface facilities, but few with a respectable set of successful and documented case studies.

This phenomenon has given the public the impression that AI has now come of age. The shells on the market are primarily based upon the general-purpose (and in general inadequate) representational scheme of *production rules*, and it is the belief of this author that we are now making the same mistakes in our claims of the generality of method as were discovered to be unfounded thirty years ago. Only now it is not the GPS system but the 'shells' that are showing us the error of our ways. It is not that all shells suffer from this problem, many are designed specifically with an open-ended architecture, but most restrict the user to a particular way of solving a problem.

Another development has been the specialised AI workstation, generally incorporating languages such as PROLOG, LISP and Smalltalk, powerful user interfaces based upon landscapes and windowing, powerful shell-building facilities, very fast computation, interfacing to other conventional languages and a whole host of utilities. These have proved popular among the more serious and committed users.

Thus the latest trend in AI has been what we knew all along. All problems are not equal. In fact some have been so difficult to represent that a new branch of AI is the area of commonsense reasoning, or *naive physics*. In this branch of AI, problems are considered that have a great deal of information associated with any decision.

Before we begin considering how to build these systems we must first understand a great deal more about the problems themselves and must start researching new conceptual systems. This view, expounded by Hayes,<sup>3</sup> has led to a great resurgence of fundamental research in AI over recent years. We are now primarily concerned with modelling things so obvious that you could pass them off as trivial and time-wasting. But to the computers in which we are attempting to instill thinking they are not obvious at all. In fact they are proving to be the most difficult problems so far encountered in AI.

It is the author's opinion that this is where a concentrated effort of AI in mining should be: on a discovery and a formalisation of a fundamental base of knowledge and not so much on short term projects that

if not analysed correctly in advance are likely to reach the same dead-end of much AI development of the past thirty years. It is not that the short-term problems are not important, many of them can be solved within existing EST methods. It is more that the potential benefits of these new methods will be so great that they will lead us more easily to make major advances in the knowledge of the discipline in which they are concerned - areas such as geology, mining, minerals processing and minerals marketing.

### **Prior EST projects and in mining-related applications**

Before considering an approach to ES development in mining applications, we firstly consider existing projects within the general domain of mining.

The first is PROSPECTOR, one of the most quoted of the early (late 1970s) large-scale expert systems that appeared to be demonstrating that AI was at last achieving its aims of embedding thinking in silicon. A much publicised success at 'discovering' an overlooked molybdenum deposit worth \$100M (the figures go up every time it is referenced) has now become a legend. Perhaps its importance stems from the fact that it remains one of the only major successes that have come to light in the past thirty years concerning the practical value of these techniques. And we know too well that no matter how promising the potential of a new method, if it does not achieve success but remains in the laboratories, it is (in the economic sense) a waste of time.

Work on PROSPECTOR was done at SRI International during the years 1974 to 1983. A good introduction to the work done is contained in the reports from the working group.<sup>4,5</sup> This system demonstrated that it is possible to build large systems, given sufficient experts, sufficient knowledge engineers and researchers, sufficient machine power, sufficient time and sufficient money. But even all of this has not achieved the economic success that matches the cost of development. It has however demonstrated some possible avenues for the encoding of difficult problems

Its major contributions were the usage of the inference net structure, connecting spaces and rules; the

method of Bayesian propagation applied to uncertain knowledge; and the importance of multiple representation schemes, in particular the merging of rules with semantic nets. Overall it remains one of the most popular areas of discussion in AI coursework and in 'AI applications in business' guides.<sup>6</sup>

The DIPMETER ADVISOR,<sup>7</sup> developed at the Schlumberger-Doll AI Laboratories, aids in the interpretation of the dipmeter logs.

The DRILLING ADVISOR<sup>8</sup> assists with borehole drilling problems.

The MUD expert system<sup>9</sup> assists the user with the diagnosis and remedying of problems of the 'mud' or drilling fluids, in oil drilling operations.

It has been noted recently by Waterman<sup>6</sup> that on his scale of implementation levels of expert systems (encompassing - demo prototype - research prototype - field prototype - production model - and, finally, the commercial system) all the geological expert systems mentioned were at the research prototype stage except MUD (field prototype) and PROSPECTOR (production prototype as at its termination). These projects have all been undertaken by teams with much money, manpower, machines and motivation but only lacking in suitable methods. This perhaps gives credence to the previously-expressed advice concerning the difficulty of actually obtaining usable systems and the critical need to re-examine our methods.

### **Domain analysis and the EST development cycle**

Over the past few years we have learnt more and more about the process of developing successful expert systems and methodologies about this process that can be formalised so as to be of use to others. Slowly, an Expert Systems Development Cycle is crystallising that is a correlate of the Software Design Cycle around which the area of Software Engineering is based. Now, in Knowledge Engineering we have our own method. This is called the method of Expert Systems Technology (EST).

The cycle begins with **identification** in which the goals of the system are defined, and superstructure is placed on the project in terms of money, people



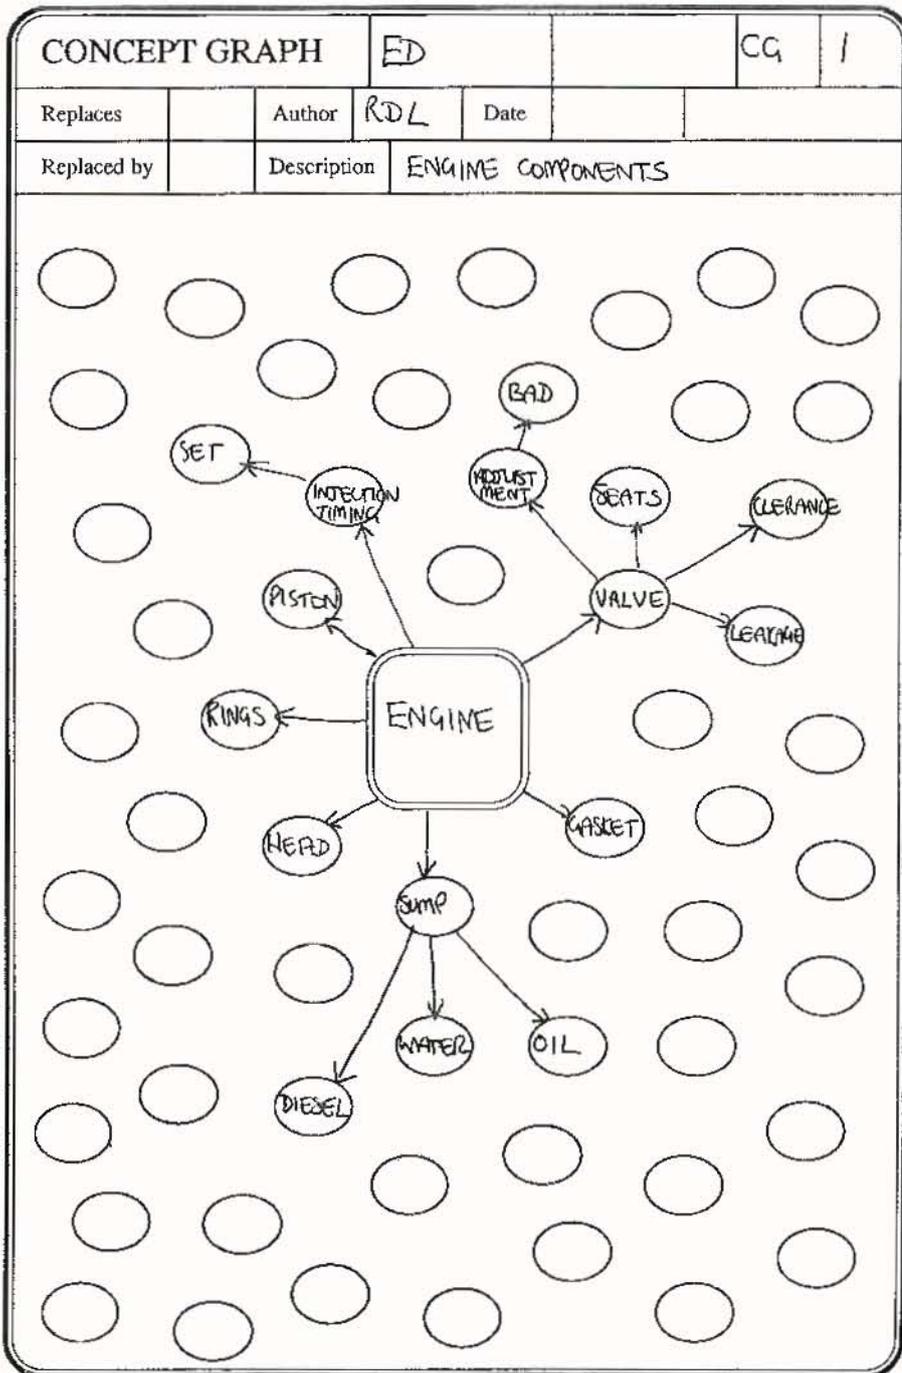


FIGURE 2. The concept graph

wrong; what do we do when we have found a solution or have been looking too long for one; and finally what to do if the problem is outside the area of expertise of the expert system.

A methodology in daily usage and under continual development by the author<sup>10</sup> assists knowledge engineers in this highly unquantifiable situation by using a collection of note-taking techniques. These techniques are referred to generally as *DANCE* (Domain ANALYSIS

Consultation Environment), and specifically as *domain dictionary*, *concept graphs* and *process forests*. The domain dictionary is used to record the terms and definitions used by the expert when referring to the domain of the problem. The concept graph is used to document relationships existing between these terms. The process forest indicates procedures, actions, sequences of events within the system being studied or problem-solving approaches taken by the experts in

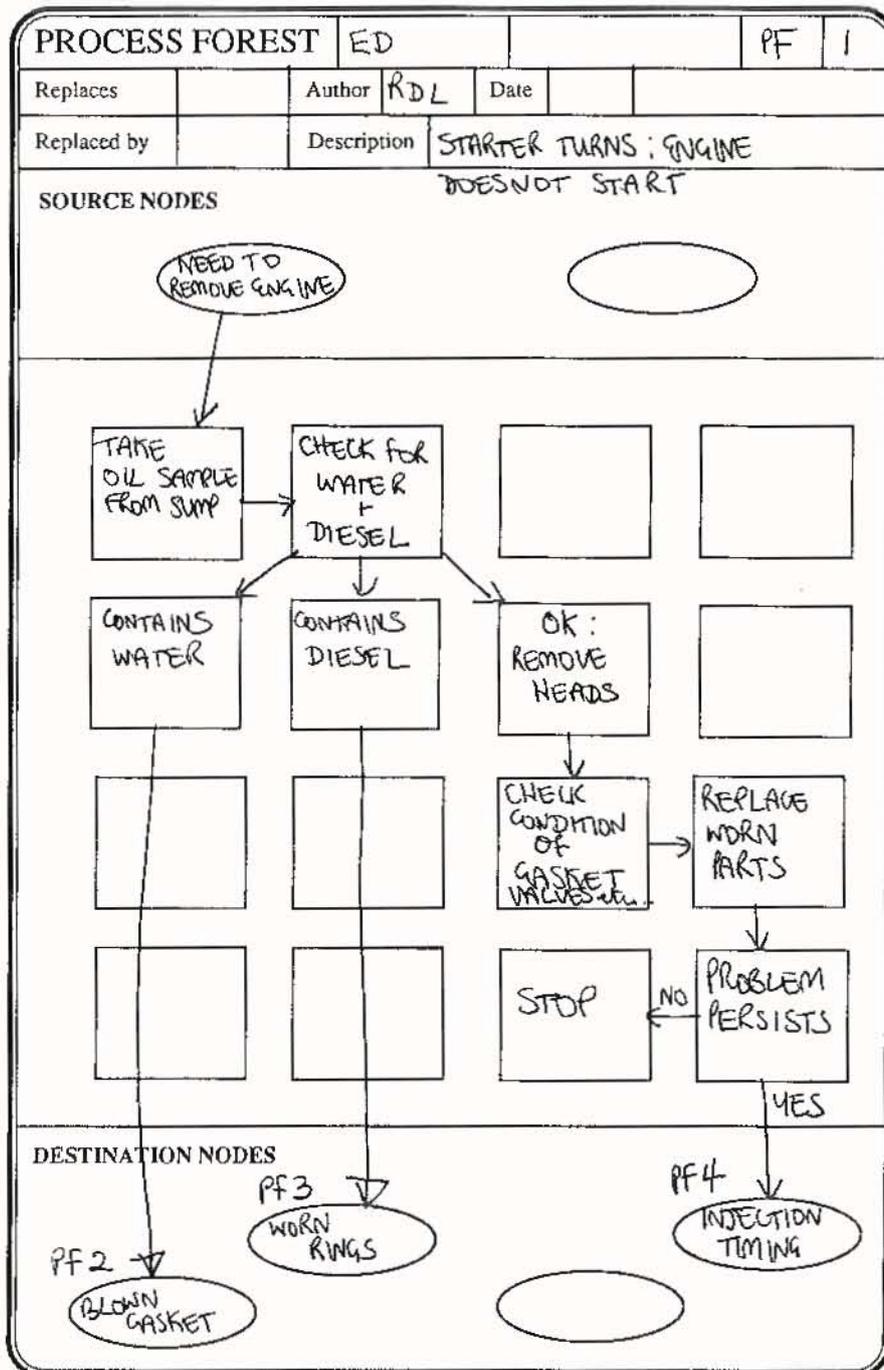


FIGURE 3. The process forest

solving familiar problems. Examples of these are given in Figures 1, 2 and 3.

The next phase, **formalisation**, is where the principles of problem solving are being quantified on paper - to examine the methods used by the expert. They might be pure logical reasoning, or involve a combination of these with well-proven quantitative methods, and even good guesswork. In general, the more the expert relies on guesswork and the less on formalised methods, the less

likely that conventional methods will be effective. There might be extra actions of a data-gathering nature that need to be performed. These can range from the very expensive (such as drilling more boreholes in gold exploration) to the very cheap (re-running a computer program for a database retrieval). Process forests<sup>10</sup> (Figure 3) allow the knowledge engineer quickly and easily to gather details of the processes used by the expert in such a way that maximum usage is made of the

expert's time. It has been the author's experience that the expert's time is the critical parameter in defining the problem initially.

These two previous phases are collectively referred to by the author as **domain analysis** for the reason that in practice they are difficult to separate, and that their essential task is in defining the domain in detail. Once sufficient analysis has been performed, we are in a position to quantify our first decisions regarding the approach to the development of the expert system. These include its feasibility, its duration and cost, the tools needed, where the scope must be reduced, how to ensure short-term success and how to measure competence.

The final two phases are **implementation** and **testing**, but these are not considered within the scope of the present paper.

### **Potential areas of success for EST within mining applications**

Many readers may find the comments expressed here have been unnecessarily negative about the state of affairs in EST at present, and that some successes are achievable in the short term.

This sentiment has its case, and as the work of Miller<sup>11</sup> shows, there is scope for success given that the problem area is chosen correctly, that the domain is artificially constrained, and that we choose tools that can ensure short-term success, as opposed to tools that we are told are right for the job but that we are not accustomed to using.

So what exactly is the nature of the problems to which we can address our EST methods in the general area of the minerals industry? A number of examples are provided, which are categorised firstly into short-term projects that can be set up using existing skills and tools, and secondly longer-term projects, that will require extensive research but for which there should be associated benefits.

#### **Short-term projects**

Two particular project classes are suggested as being suitable for short-term expert system implementation. These are the area of geostatistical methods and of machine diagnosis procedures.

*Geostatistics* is regarded as both a science and an art.

The science indicates the mathematical methods that are proven to produce the best estimates and to increase our certainty in these estimates. The art is involved in how we actually apply the methods, in which regard it shares common ground with the elusive art of expert systems development itself.

The science of geostatistics is quantitative. Pick up any good book in the area and you will be presented with equations, algorithms and even complete programs. The art is qualitative and hence is not amenable to quantitative methods. An example of this point is the common disagreement that many geostatisticians have with the concept of automatic variogram model fitting. But yet owing to the complex interface required with geological knowledge, and the fact that geological knowledge has a wide span from the highly general, concerning the methods of deposition of minerals, to the highly specific, that apply exclusively to a local area in a single mine, this appears to be a large problem. In the *variogram model critic*<sup>11</sup> non-geostatistically expert users are required to create variogram models. It is seen here that much of the complexity in developing systems that can reason effectively about the modelling of variograms is reduced by considering only the aspect of the criticism of existing models, and avoiding the complexity of generating completely new models. The complex cases are singled out effectively for expert advice. The justification for using EST is that it catches the majority of cases in which the model is correct, or so close to correctness that the natural robustness of the geostatistical technique will carry these minor deviations. This in turn saves expert time. Only the difficult and erroneous cases are submitted to expert analysis, and it is a part of the job of the expert system to detect this. This is currently the subject of a further report.<sup>12</sup>

*Machine fault diagnosis* involves the finding of faults within specific classes of machines, and of proposing or actually performing remedial action to correct the fault.

Many different types of machines are used in the mining environment, and the knowledge about them concerns such things as: their specification; the

manufacturer's suggested approach to fault diagnosis; their intended usage; their age; their actual usage history; their service history; their breakdown history; known weak points associated with the machine type or manufacturer; the number and geographic location of the experts; the availability of spares; the number and geographic location of the machines; the cost caused directly by machine failure or inefficiency (e.g. having to restart processes); and the loss accountable for by downtime (e.g. production loss).

The job of the expert system, taking these factors (and others) into consideration, is to assist in the determination of the actual cause(s) of the observed symptoms by isolating the potential causes of failure or inefficiency and consequently recommending that certain actions be taken. This might include the recommendation the expert himself must be on hand as the diagnosis is beyond the scope of the expert system, or alternatively the solution is beyond the capabilities of the operator. The types of knowledge that will need to be encoded will centre around the general cause-effect relationships that exist within the structure of the machines and which of this general set of relationships the expert would consider, and in what order, in analysing a given set of symptoms on a given machine.

To get the system productive as soon as possible the restrictions that would be suggested in the domain are: to concentrate only on a particular class of machines; to analyse only a particular class of problem (thus the system will have to recognise anything that falls outside of this class); and only to diagnose the cause of the problem and not initially to recommend remedial paths of action.

### **Longer-term projects**

The extension of any short-term project by increasing its domain is a possibility for a longer-term EST project. As we know from the difficulty of the extension of the toy world problems to the real world, this extension to larger domains will give rise to problems that were not noticed in advance. But in the process of having reached a successful implementation in a small domain we will naturally be enthusiastic about the extension and probably oblivious at the same

time about any of these *problems of scale*.

A medium-sized potential application is the *interpretation of geophysical instrument readings* as required, for example, in downhole logging of coal boreholes. This process is not as complex as the geophysical procedures used in the oil industry, both in terms of the number of tests actually done and in terms of the difficulty of getting the actual results, owing to the depth of the hole and the environment of offshore drilling rigs. The standard solution is to use a geophysical company to do the entire job, including the interpretation. These companies should make the data tapes available for their clients' own databases. By analysing the logs to determine the qualitative cut-off points for each of the readings we can provide a basis for the simple analysis of the geological structures, and in particular the location, thickness and the quality of the coal. Most of these can be determined by the comparison of the readings and can be set up as a set of rules that first classify the rock type, secondly classify the coal found into grade categories, and thirdly estimate the grade parameters for the coal.

A more complex problem in this domain that could also be considered is the tuning of the rules to fit new properties. There are general rules that are applicable. There are also refinements that can be made by comparing geological analysis of the core to the geophysical readings obtained. This enters the realm of machine learning as the goal is to determine the suitable rules. Since machine learning is very much in its infancy it is better to consider how an interactive system could be built to assist the geologist or geophysicist in the process of rule modification.

In some cases expert systems are not considered as being capable of functioning independently in a domain but can serve the purpose of decision support, or the 'right hand man'. They can assist with weighing up facts in a complex domain, or in providing a model with which to identify and control the reasons why one reaches decisions in highly qualitative problem areas.

One basic example, the go/no go decision in terms of actually starting a new mining venture, assuming that the initial exploration has been done and a basic understanding gained of the quantity and quality of the

ore. In gold exploration one ends up with a small set of borehole data; a statistical analysis performed on this data; a geological analysis of the core samples; the projections of future price and demand for the ore; and the possible cost of setting up and running the mine and the processing plant. EST methods can be used where quantitative methods fall short: in embedding the reasoning behind the decisions, and the assumptions upon which the decisions are based, hence allowing these assumptions themselves to be manipulated for the purposes of discovering potentially sensitive areas.

### **Potential areas of difficulty for EST within mining applications**

The examples given previously indicate some areas in which the author feels gains could be made within the realm of existing EST tools. The actual size and complexity of the project and hence the selection of these EST tools (both hardware and software) requires analysis of the particular case in the manner as indicated previously. Of great importance is that we begin small and learn to grow. To start too large might end in expensive failure that destroys management confidence in the entire potential of EST. To start too small allows accusations of dealing exclusively with toy problems.

This section now exemplifies some of the problems within the mining environment that are too difficult to examine at present and should be considered as research goals, as opposed to the problems mentioned previously which helped identify problems of potentially successful implementation with existing technology and methods.

The proposed research will examine those areas too complex to model at present, the results of which should eventually provide us with a new generation of EST tools and the knowhow to use them.

It is widely hoped that the next generation of expert systems tools will have the ability to reason using a great deal of commonsense. Much research on AI is proceeding towards this goal. The area of research is referred to in the AI literature as *commonsense reasoning*, *naive physics*, or *qualitative physics*.

In the first generation expert systems the inference procedures were regarded as common ground and could

simply be standardised so that the same inference engine could be applied to other problems.

The movement to the second generation will be to add to this general-purpose inference engine a very large base of commonsense abilities as givens. One of the major problems with current expert systems tools is the necessity continually to state the obvious. With this restriction removed we could concentrate more on the problem domain itself.

In examining some of the areas in which *commonsense* is applicable to the problems existing within the mining industry, a set of categories is examined under which specific points are made. This given set of categories also contains many highly overlapping areas, and any one aspect of geology or mining can be seen within many of these commonsense categories. It is the author's thesis<sup>13</sup> that this overlap, or redundancy, is perhaps the most essential quality of commonsense understanding, and this natural redundancy must be accommodated effectively to secure a good, working commonsense model.

#### **Time**

Firstly, let us consider the ore itself. Rocks and minerals with which we are interested have more history than anything that we previously considered in the entire realm of knowledge of the physical world, outside of the universal truths (which are apparently timeless) and the subject areas of astronomy and cosmology. Thus the ability to reason effectively about histories in particular and time in general must be at the top of our list of commonsense priorities. We refer to the creation of an orebody in terms of geological time. We refer to the analysis of this same orebody in terms of years and months. R. Buckminster Fuller<sup>14</sup> considers the effect of including the time taken to create the ore in his economic equation of the value of the ore, and deduces from his model that owning a gallon of petroleum should be sufficient to make us millionaires. Temporal reasoning is of paramount importance in our analysis and exploitation of orebodies.

Consider the complexities facing the micropaleontologist who is comparing evidence of the fossils found in certain rocks with the known geological occurrence time-range of these fossils. He will

endeavour to determine bounds on the age of the rocks. In turn, the environment in which the rocks exist can be assessed for the possible presence of the target ore.

In examining a coal prospect we examine the prehistory, the depositional environment, and the weather conditions at the time of formation (among other things). We can infer that a sandstone band in the middle of a coal seam could have been caused by flooding of the swamp at a certain time. The 'geological record' of a locality is indicated by examining the core samples of a borehole. Here time is condensed into a very small space. A more impressive display of time is outlined when highway development cuts grooves through hills, such as that on I-70 on the road from Denver to Dillon through the Dakota Hogback which has been set up as a geological viewing point.

Geological interpretation is complex, but without it we are not in a position to account for the observed differences in the statistical analysis of the individual laboratory samples. The more we can account for grade and thickness differences geologically, the more correct our model is likely to be for predictive usage.

Temporal reasoning in mining, of the hour, day and year scales, occurs in the analysis of plans of exploitation. In turn this is related to economic time and the ebb and flow of the current of supply and demand, and the subsequent effect on the selling prices of the ore and the purchase prices of the required raw materials, machines and manpower.

#### **Space**

Another overriding quality of ore bodies from the view of both geological science and mining technology is the fact that they consume space.

From the geological view the extent of this space is a fundamental result of geological analysis and the subsequent economic potential analysis of the field. Spatial reasoning about the geomorphological environment of the deposition and the deformation in terms of the earth processes is a natural part of the practice of geology.

From the mining perspective, spatial reasoning relates to the shape and orientation of the orebody, and to the choice of equipment that can operate within these parameters. It also concerns the design of the main

material- and people-moving arteries of the mine.

#### **Gravity, density and mass**

Gravity appears to be the most fundamental force of nature. It certainly has proved the most elusive for theoretical physicists. Its effect is felt in virtually every action involving the movement or relationship of physical objects and substances on this (and probably every other) planet. We know that owing to density differences, water flows down through rock and oil flows up. We know about 'fluid gravitational containment vessels' but prefer to use the more common term 'tanks'. They prevent the fluid from falling to the local source of the force, the centre of the earth, until a fluid of a higher density or a rock of zero permeability is reached. In the case of tanks, cups, etc. the containing force is gravity.

Gravity is also the cause of our extreme concern with *up* and *down*, especially in connection with the movement of large objects up (such as a few million tons of coal), and spanners falling down things such as mine shafts. We are naturally more concerned with our movements one kilometre up and down than our taking a restful stroll one kilometre to the shop and back. And gravity alone is the responsible party.

In order to build expert systems that begin to understand a bit about the decisions we make regarding the movement of objects and substances, and about how they occur in nature, we will have to instill the same form of gravitational reasoning that we have come accustomed to since birth.

#### **Earth processes**

As any good earth scientist knows, there are a wide variety of earth processes that we can identify and with which we can explain the movement, deposition and transformation of materials that end up as the economically interesting ores that now concern us. For example :

*Wind*, that moves light objects regularly, and heavy objects rarely, that is subject to the forces of pressure difference, and that in turn causes waves.

*Tide*, that is governed by the celestial objects of the moon and the sun, and the force of gravity, acting on the sea, which itself is held in place by gravitational containment.

*Fire*, that transforms matter very quickly to a useless state of ash.

*Heat*, that is the responsible force in creating coal from peat, and a few other things as well.

*Earthquakes* and other forms of violent and massive physical distortion, a result of pressure caused when very big objects collide very slowly.

*Rivers*, the flow of water downwards by the force of gravity, avoiding obstacles it cannot shift, pushing along those it can, eventually reaching the sea either directly as a mouth, or diffusedly as a delta.

#### **Pressure and flow**

A physical state of a substance such that an adjacent substance with a different pressure will activate a *flow*. The wind is just the direct observation of a pressure difference between pockets of the air.

#### **Waves and cycles**

Seasons, and other constant period waves or cycles, with very high predictive content. Waves in the sea that cause, in conjunction with the tides, the movements of sand and pebbles.

#### **Chemistry**

The complexity of the organic chemical processes that created plants that gave rise to peats that become coal is essential to an understanding of petrographic interpretation that some say gives the maceral constitution of a section of coal the closest thing possible to a fingerprint. Deductive reasoning utilising these laboratory calculated relative components can form the basis for a thorough description not only of the expected coal quality but also the reasons that justify these expectations.

Inorganic chemical processes contribute to the metamorphic processes, always aided and abetted by pressure, heat, gravity, and the flows that caused the interacting substances to be at the same place at the same time.

#### **Liquids**

The essence of flow and containment. The difference between putting 1 kg of rock and 1 kg of water down on the grass in your garden is that the rock is still there after the process of 'putting' is complete. We should not have to spell this out in detail to an expert system

that deals with liquids and porous objects. It should be as obvious as we see it. Hayes<sup>3</sup> has studied liquids as a specific case of commonsense knowledge.

#### **Materials**

The substance out of which objects are constructed. The properties of the materials include elasticity, combustability, breakability, hardness, fluidity, liquidity, porosity, permeability, ... Since all physical objects are composed or either uniform or composite material, and that materials interact by chemical, by merging, or diffusive (such as dropping a glass of water on the floor) processes, an understanding of the materials and their properties is essential to our ability to reason about them.

#### **Other**

Some other primitives of commonsense that cannot be dealt with here are: structure, systems, processes and containment.

#### **Summary**

To summarise this section it is sufficient to note that of all the above cases of the so-called common sense areas of knowledge, it is quite frightening to realise how little we actually know about how we reason using them. It appears to be automatic, in-built, but highly effective. It is currently the greatest challenge in AI to discover this method.

### **Conclusions and Recommendations**

The sections of this paper have indicated where we show in EST at present, in terms of applicability to mining problems, and where research should be directed in order to get further.

The basic conclusions are firstly that some work can be done at present with existing tools, so long as we restrict our problems, start small and grow slowly. Secondly, problems that require substantial commonsense reasoning should be avoided or redefined to allow user guidance. Thirdly, there is a great need for fundamental research in order to tackle the larger problems requiring commonsense reasoning in the domain of mining expert systems.

The longer we look at the **flash in the pan** that **expert systems technology** is currently providing us with, the

more attractive it looks and the more money we are willing to commit to this venture. Is it what we want or are we after something bigger? If so, let us recognise this quickly, and not expect it to be what it is not.

Let us then be cautious about the nature of this 'flash' and question if it possesses the worth that we expect of it. As in the gold rush days, many saw flashes in the pan and invested their life savings for nought. We have not learnt much from the past if we are now doing it again!

### References

1. CHARNIAK, E. and McDERMOTT, J. *Introduction to Artificial Intelligence*. Reading, Mass. Addison-Wesley, 1985.
2. FEIGENBAUM, E. and BARR, A. *The Handbook of Artificial Intelligence*. London, Pitman, 1981.
3. HAYES, P.J. The second naive physics manifesto. In *Readings in Knowledge Representation*, R.Brachman and H.Levesque. Kauffmann Publishers, 1985.
4. DUDA, R.O., GASCHNIG, J., HART, P.E., KONOLIGE, K., REBOH, R., BARRET, P., and SLOCUM, J. Development of the PROSPECTOR consultation system for mineral exploration. Final Report, SRI Projects 5821 and 6415, SRI International, Inc., Menlo Park, Calif., 1978.
5. DUDA, R.O., HART, P.E., NILSSON, N.J., REBOH, R., SLOCUM, J., and SUTHERLAND, G.L. Development of a computer-based consultant for mineral exploration. Annual Report, SRI Projects 5821 and 6415, SRI International, Inc., Menlo Park, Calif., 1977.
6. WATERMAN, D.A. *A Guide to Expert Systems*. Reading, Mass., Addison-Wesley, 1986.
7. DAVIS, R., *et al.* The dipmeter advisor: Interpretation of geologic signals. *Proceedings of the International Joint Conf. on AI 1981* pp. 846-849.
8. HAYES-ROTH, F. The industrialisation of knowledge engineering. In *Artificial Intelligence for Business*. W. Reitman. Ablex, Norwood, NJ, 1984.
9. KAHN, G., NOWLAN, S., and McDERMOTT, J. A foundation for knowledge acquisition. *Proc. IEEE Workshop on Principles of Knowledge-Based Systems*. IEE Computer Society, IEE Computer Society Press, Silver Spring, MD, 1984.
10. LAYTON, R.D. Formalising the interaction between the expert and the knowledge engineer. Technical Report. Computer Science Department. University of the Witwatersrand. 1987.
11. MILLER, S. *et al.* The implementation of a gold ore reserve system. *APCOM 87*, Volume 3: *Geostatistics*, Johannesburg, SAIMM, 1987.
12. LAYTON, R.D. and MILLER, S. Quantitative and qualitative reasoning in a variogram model critic. Technical Report. Computer Science Department, University of the Witwatersrand. In preparation.
13. Layton, R.D. Exploiting redundancy in knowledge representation. Phd thesis. In progress, University of the Witwatersrand.
14. FULLER, R. BUCKMINSTER *Critical Path*. London, Hutchinson, 1983. p. xxxv.