KNOWLEDGE-BASED EXPERT SYSTEMS — THE STATE OF THE ART IN THE US

F. Hayes-Roth

Teknowledge Inc
Palo Alto, CA

SUMMARY

This paper aims to describe the current state of knowledge systems technology and its commercialisation in the US. First, knowledge systems are defined and placed in a historical context. The introduction is concluded with a preview of major ideas. The paper will assess the technological state of the art and will survey the current state of commercialisation. Finally, some anticipated future trends will be discussed.

INTRODUCTION

Knowledge-based expert systems, or 'knowledge systems' for short, have evolved over 15 years from laboratory curiosities of applied artificial intelligence into the targets of significant technological and commercial development effort [1]. These systems employ computers in ways that differ markedly from conventional data processing applications and they open up many new opportunities. Recently, many commerical and governmental organisations have committed themselves to exploiting this technology, attempting to advance it in dramatic ways and beginning to adapt their missions and activities to it. Some major events which have occurred recently in the US include:

1 The US DoD claimed that knowledge systems would become the front-line of the nation's defence in the 1990s and initiated its $500 million five year Strategic Computing Programme.

2 IBM licensed and sold its first AI software program (INTELLECT) and publicly endorsed the field as relevant and applicable.

3 IBM's principal competitors teamed up to form MCC and identified knowledge systems as an area of primary concern.

4 General Motors took an equity position in a knowledge engineering firm to hasten the introduction of knowledge systems into their business [2].

5 At the 1984 National Conference on Artificial Intelligence in Austin, Texas, IBM, DEC, Tektronix, Data General, Apollo and TI announced new AI products for sale.

Locating expert systems within the computing field

It is not possible to give a precise definition of 'expert systems', any more than it is possible to give a precise definition of an 'expert'. Rather than seeking elusive precision, this paper will adopt three common practical approaches: it will state what an expert system is, what it is not and how one is recognised.

What is an expert system?

An expert system is a knowledge-intensive program that solves problems that normally require human expertise. It performs many secondary functions as an expert does, such as asking relevant questions and explaining its reasoning. Some characterstics common to expert systems include the following:

- They can solve very difficult problems as well as or better than human experts
- They reason heuristically, using what experts consider to be effective rules of thumb and they interact with humans in appropriate ways, including natural language
- They manipulate and reason about symbolic descriptions
- They can function with data which contains errors, using uncertain judgemental rules
- They can contemplate multiple, competing hypotheses simultaneously
- They can explain why they are asking a question
- They can justify their conclusions.

What is not an expert system?

Compared to a human expert, today's expert system appears narrow, shallow and brittle. It does not possess the same breadth of knowledge or understanding of fundamental principles. It does not apparently think as a human does: perceiving significance, jumping to conclusions intuitively and examining a single issue from diverse perspectives. Rather, the expert system of today simulates an expert's thinking rather grossly. It reaches the same major decisions by elucidating many of the relevant criteria and making many of the same educated guesses that an expert would if forced to verbalise the thought process. Unlike a human, however, the expert system does not resort to reasoning from first principles, drawing analogies or relying on commonsense. In addition, today's expert systems do not learn from their experience.

In contrast to advanced data processing systems, today's expert systems seem specialised and unusual. Where conventional DP systems amass and process large volumes of data items algorithmically in order to automate time-consuming clerical functions, expert systems ordinarily address small tasks typically performed by professionals, in a few minutes or hours: interpreting, diagnosing, planning, scheduling and so forth. In order to accomplish these tasks, an expert system makes judicious use of data and reasons with it. In contrast to the algorithmic data processing approach, the expert system generally needs to search a large space of possibilities or will construct a solution dynamically.

How to recognise an expert system

Expert systems today generally serve to off-load some difficult task from a human professional. Find a computer that performs a function that previously required an expert. Locate the expert or team of experts who now maintain the program's knowledge base. Determine if the knowledge in the system is accessible: can you read it, ask for explanations and justifications that exploit it and can you modify it? Confirm that the system stores a substantial body of knowledge and that it reasons with that knowledge in flexible ways. Although any good programmer can implement a procedure integrating a few heuristics, incorporating hundreds or thousands of heuristics into a computer system requires knowledge engineering. Systems employing this magnitude of judgemental knowledge are expert systems.

Now that we have the subject more or less defined, let us consider where it fits into the field of computing. Figure 1 places the area of expert systems in its historical context. The figure locates the beginnings of expert systems around 1980-1981, when efforts began to commercialise the technology. The first company formed exclusively to promote expert systems (in the field of genetic engineering) was IntelliGenetics. Shortly afterwards, Teknowledge was formed as the first knowledge engineering company. These companies were a spin-off from the Heuristic Programming Project at Stanford University, which had led the development of knowledge engineering during the 1970s.

Figure 1: Expert systems in the evolution of computing

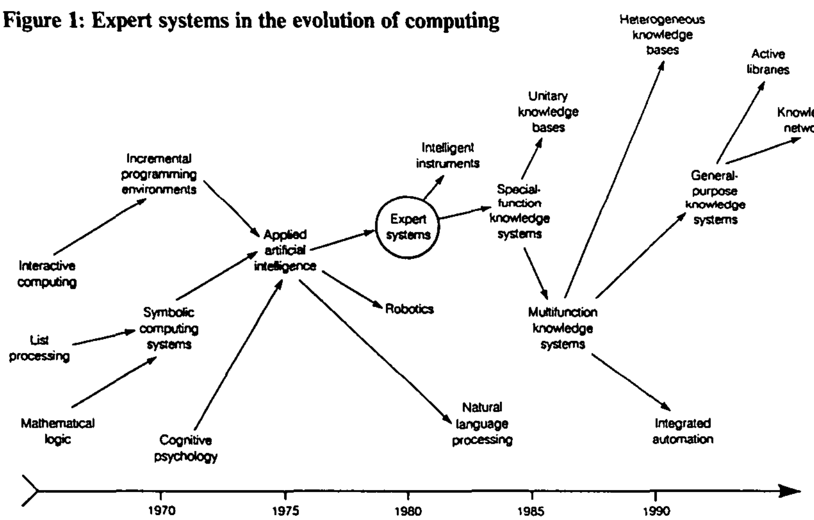

The field of applied artificial intelligence merged techniques from three areas: symbolic programming, cognitive psychology and work on incremental programming environments [3]. Out of a decade of work, three primary subfields emerged: expert systems, natural language and robotics (including vision, speech and locomotion). Presently, a few small commercial companies in the US occupy positions in each of these subfields.

To date, researchers in knowledge engineering have focused primarily on creating artificial experts for tough practical problems. The Stanford group were pioneers in this orientation, developing the first truly expert system DENDRAL, the first successful learning system META-DENDRAL, and a variety of applications in medicine. This group had a recipe for success with two key ingredients as follows:

1 Attack problems amenable to the techniques of applied AI.
2 Consider only important, difficult and high-value problems.

In more detail, here is a recipe followed by many US knowledge engineering efforts:

1 Seek problems that experts can solve via telephone communication.

2 Choose a problem that experts can solve in a 3 minutes to 3 hours timespan.

3 Choose a problem whose solution requires symbolic reasoning.

4 Prefer high-value problems.

5 Rule out problems where different experts disagree about solution correctness.

6 Rule out areas where you cannot solve initial problems with a limited subset of the expert's total knowledge.

7 Select an initial class of problems to solve that require only a subset of knowledge.

8 Identify some training problems and collect the expert's problem-solving protocol.

9 Build a knowledge base that contains explicit and declarative representations of the expert's concepts and heuristic reasoning rules.

10 Develop an initial expert systems to solve the training problems as the expert does.

11 Ask the expert to review the system's solutions and its lines of reasoning.

12 Augment the system to accommodate the expert's critique.

13 Apply the system to more training cases and augment its knowledge base incrementally.

14 Evaluate the system's performance on novel test cases.

Preview of the ideas

In surveying the state of the art and the state of commercialisation, a few key points emerge, as follows:

1 Knowledge systems transform book knowledge and private knowledge into an active inspectable form capable of performing high-value work.

2 Knowledge today takes many different forms, necessitating a variety of tools or instruments for knowledge engineering.

3 Knowledge systems must be integrated with data processing systems, but knowledge engineering work differs greatly from conventional software work.

4 A number of companies today make money from knowledge engineering products and services.

5 One cannot predict the direction of commercialisation.

6 Steady, rapid growth of both the technology and the commercial applications should continue through to the turn of the century.

The subsequent sections explain these points and provide additional details.

TECHNOLOGY STATE OF THE ART

Problem-solving engines organise the activity of knowledge systems to solve problems. To understand these engines, we need to relate their implementation to their design and intended purpose. Today's knowledge systems aim to solve specific problems. A knowledge engineer analyses the problem to be solved and then adopts an overall approach generally consisting of several things, as follows:

1 A problem-solution paradigm, such as top-down refinement or multidirectional opportunistic search.

2 A general knowledge system architecture that makes specific choices regarding the key system design questions, including what kinds of knowledge to represent, in what formalism, for what kinds of inference and affording what kinds of flexibilities.

3 A specific problem-solving strategy that determines what knowledge to apply in what order.

Today's problem-solving engines provide specific devices for implementing the knowledge engineer's choices.

These engines provide a specific knowledge representation formalism and related interpreter, a high-level control architecture and executive, and an inference procedure and related inference engine. The knowledge representation formalism may include useful ways to describe conceptual taxonomies and conditional heuristic rules. A conceptual taxonomy includes relationships among types of object, class and individual, and determines how properties of one apply to another. For example, many problem-solving engines can accommodate and exploit these kinds of fact:

The Warsaw Pact includes the USSR, Poland, Hungary, ...

The USSR is a country.

Cuba is a client state of the USSR

All Warsaw Pact countries are client states of the USSR.

Every client state of any country will do whatever that country does regarding participation in the Olympics.

Heuristic rules represent judgemental knowledge. In fact, the last example above actually represents such a heuristic. Most current problem-solving engines provide a stylised if then formalism for representing such heuristic knowledge. Some examples of other kinds of heuristics that various problem-solving engines can exploit are as follows (the uncertain nature of all should be evident to the reader):

- If you see an oily sheen on the surface of the water, hypothesise that oil has been spilled
- If you hypothesise that oil has been spilled but cannot smell petroleum or feel slime, rule out that hypothesis
- If you have competing hypotheses, you can identify the correct one by trying to rule them all out
- If all hypotheses but one have been ruled out, accept that one
- If you need to schedule *n* events, schedule the most constrained one first
- If you need to determine how constrained an event is rate as highly constrained any event whose time and date are already fixed; rate as weakly constrained any event which is a weekly date with a friend etc.

Once the knowledge engineer adopts a particular problem-solving engine, the knowledge base for the application needs filling out. Knowledge bases generally consist of conceptual taxonomic relations and rules. From our experience, some generally valid heuristic rules about knowledge bases are as follows:

- An interesting demonstration of the technology requires only 50 rules
- A convincing demonstration of the power of a knowledge system requires about 250 rules
- A commercially practical system may require as few as 50 rules
- An expert level of competence in a narrow area requires about 500-1000 rules
- Expertise in a profession requires about 10 000 rules
- The limits of human expertise are at about 100 000 rules.

To help amass and maintain these rule bases, most research and commercial knowledge engineering tools provide various automated aids. These generally provide a range of support, from automatic name recognition on input and spelling correction to line-of-reasoning traces, knowledge base browsing and automate facilities for system testing and validation.

Forces affecting knowledge system technology in the US

Figure 2 shows the ways in which we can now see the technology expanding. In the upward direction, a variety of user needs are encouraging advances in knowledge systems technology. These needs determine the economic and functional niches the technology occupies. Specifically, knowledge systems address problems that arise from difficulties in retaining, transmitting and applying know-how. For example, knowledge systems can preserve fragrance-blending knowledge beyond the retirement of a perfume manufacturing expert or can disseminate to all world-wide equipment service facilities the diagnostic methods of a master mechanic who works in Dayton. In addition, they provide a means of employing know-how where it is needed, when it is needed, at great speed. This has many attractions for factory automation, process control, safety systems, military intelligence and weapons systems.

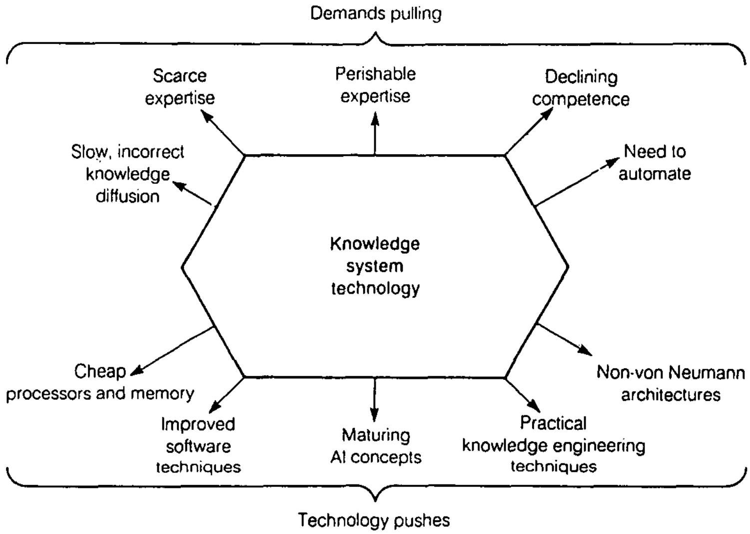

Figure 2: Forces of pull and push, that are expanding the technology

In the downward direction, the field will expand in response to technology drives. These include improved forms of conventional hardware and software for symbolic computing, less expensive memory and processors, standardised and increasingly reliable knowledge engineering techniques and possibly novel non-von Neumann architectures (such as those that the DARPA and Fifth Generation programmes now pursue). Throughout the world today, many groups are active in these areas. However, the author does not anticipate any near-term breakthroughs that will cause a discontinuity in the technology's development. Rather, steady improvements in these directions should reduce the cost, expand the capability and increase the reliability of the technology. This will make an already practical technology much more so.

Assessing technological capacities

To assess the state of the technology, we need better measures than those currently in existence. Figure 3 estimates the levels and rates of change of some key technology measures.

Item	1984 level	1984-1985 change
Knowledge system prototypes under development	70	+50%
Knowledge systems being deployed	15	+100%
Knowledge systems being maintained	10	+200%
Knowledge engineering departments established	15	+150%
Senior knowledge engineers	40	+50%
Knowledge engineers	150	+50%
Knowledge engineer trainees	300	+100%
Applied AI graduate students	250	+20%
Undergraduate students in AI	2000	+50%
LISP or PROLOG programmers	2000	+50%
LISP or PROLOG installations	400	+100%

Figure 3: Estimated measures of current US technology capacities

To assess the rate of progress, we need a method of measuring the productivity of a knowledge engineering team. How can you measure this? Although there are many differences between the systems that have been built thus far, a reasonable first-order approximation treats them all as collections of heuristic rules. Typically, a knowledge engineer and an expert collaborate to create the rules that enter a system. Using the range of 500–1000 rules as the typical size of an interesting application, the author has estimated the rate of system development over numerous applications.

Figure 4 displays development times in terms of engineering person-hours per rule. The graph shows approximately a doubling of productivity every two years over the last 15 years. These improvements can be attributed to two factors:

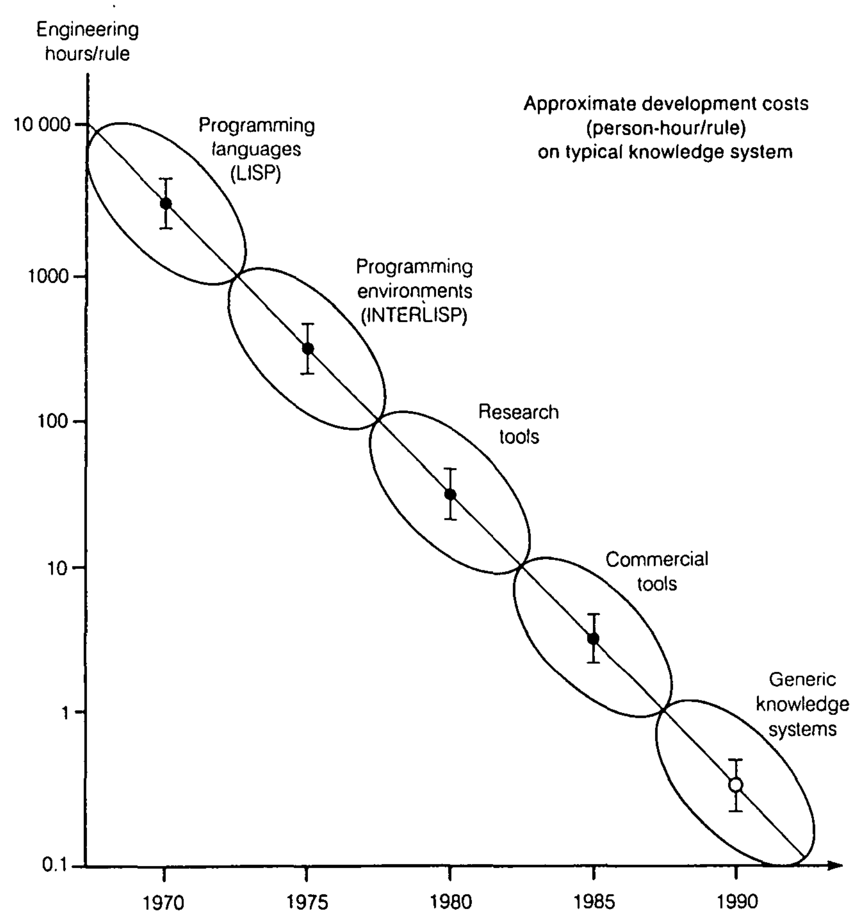

Figure 4: Technology efficiency in knowledge engineering doubles every two years

1 Technological advances in the best available methods for carrying out knowledge engineering.
2 Typical improvements due to learning within a given method.

The figure identifies five successive periods in the technology for knowledge system construction: programming languages, such as LISP: programming environments, such as INTERLISP: research tools, such as EMYCIN: commercial tools such as S.1; and anticipated generic knowledge systems which incorporate a user's own knowledge into a prefabricated heuristic problem-solving package (eg a personal financial planner). With the recent emergence of commercial knowledge engineering tools, many useful systems will reach deployable staus in a year or less.

Weaknesses in current technology

While many applications today seem straightforward, many others present difficulties that stretch the current technology. The barriers most frequently encountered include:

- A need for flexible and general natural language understanding, which may arise when users need to exercise initiative in directing the activities of a knowledge system
- A need to incorporate knowledge that is hard to represent, which often arises in situations which require reasoning with spatial or temporal problems
- A need to combine and unify the knowledge of multiple experts when no prior standardisation has occurred
- A need to apply broad bodies of knowledge quickly, as may arise in real-time command-and-control problems.

Applications today generally work around these technological shortcomings or necessitate some corresponding advance in the state of the art.

STATUS OF US COMMERCIALISATION

Knowledge systems, and the tools for creating them, have begun to enter the commercial world. The first systems to find regular use grew out of long-term academic research. DENDRAL, which determined molecular structure from mass spectroscopic and NMR data, developed over a period of 15 years at Stanford [4]. Numerous pharmaceutical companies access it on the national SUMEX-AIM medical computing network. Stanford recently granted an exclusive commercial licence on DENDRAL. Similarly, MACSYMA envolved over many years of MIT research [5]. Eventually it became a super-human expert in its area of speciality, the evaluation of mathematical expressions. It too has been commercialised (see Figure 5).

Figure 5: Knowledge system technology representative products

Area	Name	Vendor	Functions
End-user applications	Questware	Dynaquest	Personal computer HW/SW selection
	Logician	Daisy	Electronic design
	TK!Solver	Software Arts	Equation solver
	SMP/Macsyma	Interference/ Symbolics	Mathematical simplification and problem-solving
Commercial applications tools	S.1	Teknowledge	Industrial diagnostic and structured selection problems
	M.1™	Teknowledge	Microcomputer tool for small expert system applications
	OPS'	DEC	VAX AI programming language used in-house
Research and experimentation	KEE	IntelliCorp	Extended programming environemnt, frames
	LOOPS	Xerox	Extended programming environment; objects
	ART	Inference	Varied representations and inference techniques
	ROSIE™ •	Rand Corp	Legible, intelligible symbolic programming language
	KL-TWO	BBN/ISI	Knowledge representation schemes
Supporting systems	COMMONLISP	DEC/LMI/ Symbolics	New attempt to standardise LISP
	INTERLISP	Xerox/ISI	Mature programming environments
	FRANZLISP	Berkeley/ Tektronix	UNIX-based LISP
	LISP Machines	LMI/Xerox/ Symbolics/ Tektronix	Integrated graphics, personal workstation
	PROLOG	Silogic/Quintus/LMI	Formal semantics logic-based programming language
	LISP/VM	IBM	First IBM 370 LISP product

With the last five years, many companies have applied knowledge engineering to their own internal problems. DEC reports a $10 million annual savings from one expert system for configuring VAX orders [6]; Schlumberger claims that its knowledge-based Dipmeter advisor will increase their revenues from interpretation services significantly; Elf-Aquitaine expects to field a drilling advisor system that they say should reduce total drilling costs by more than one per cent, and numerous other companies have undertaken confidential and proprietary applications

Very recently, three new types of commercial application have emerged. The first of these consists of tools for the construction of knowledge systems. The second consists of specialised hardware and systems software for general AI programming. The last new product category includes the commercial expert system, a problem-specific artificial advisor. These have arisen in a variety of areas including CAD, data interpretation, computer selection and mathematics problem-solving.

Figure 6 illustrates the current and projected state of commercialisation. The figure comprises a kernel hexagon, representing the major ingredients of knowledge system technology, and three rings of increasing commercialisation. The first ring shows the types of commercial products now available. These consist of supporting hardware, the software, knowledge engineering tools and equipment, and knowledge applications. The next ring portrays the primary areas of new product focus anticipated in the next three years. In the outermost ring, the figure indicates the key mid-term commercial targets which companies will reach by 1990.

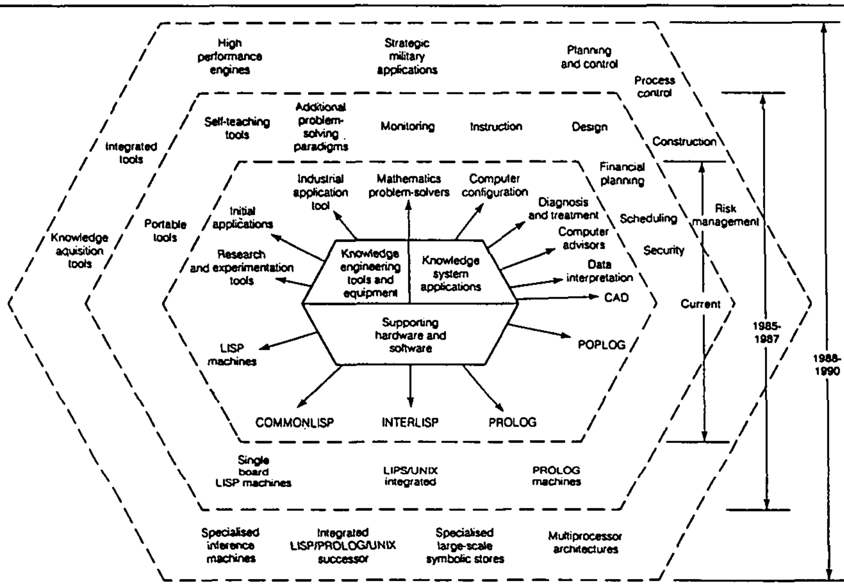

Figure 6: The commercialisation of knowledge systems technology

Figure 5 lists a representative set of products now available. The table categorises these in four ways: end-user applications, commercial application tools, research and experimentation tools, and supporting systems.

The supporting systems consist primarily of LISP and PROLOG and their variations. To date, the vast majority of applications work has employed LISP. Over the almost 20 years that AI and knowledge engineering professionals have worked in LISP, they have developed elaborate programming environments. These programming environments offer many productivity aids to programmers. However, they require considerable computing resources. This has motivated the development of specialised, personal LISP workstations. In Europe and Japan, PROLOG captures as much interest as LISP does in the US. Although it does not afford a mature programming environment yet, PROLOG offers a simple and powerful knowledge programming system based in logic. Vendors of supporting systems emphasise performance, price and compatability with existing and planned computing systems and standards.

In summary, the current commercial situation reflects a wide diversity of products and a relatively immature market. A number of companies now have profitable lines of business based in knowledge system technology. Many independent forecasts of market growth suggest that demand should increase rapidly for at least a decade. Although this paper has sketched the current state of commercial activity and suggested what seem to be the most likely near-term growth paths, no one can predict accurately how a new industry will develop. Because the technology addresses a fundamental problem — retaining, distributing and applying know-how electronically — it should continue to prosper. On the other hand, we should expect this technology to combine with other emerging technologies in surprising ways. The final section discusses some of these possibilities.

FUTURE TRENDS

For nearly 500 years, books have provided the primary means of retaining knowledge and transmiting it to humans. To achieve excellence in a profession, a human has needed to study, interpret and memorise these books, then apprentice and train with someone who can clarify and illustrate the book's principles, then practice for years and learn practical rules from experience. The development of printing made an enormous impact on human culture by providing a means of distributing representations of human expertise to large numbers of potential practitioners. However, because it could not explain or apply its knowledge directly, the passive book left much work for the reader.

As technology progressed and economies advanced, knowledge transfer became a bottleneck of cultural development. In highly advanced fields such as medicine and electronics, knowledge creation outpaces knowledge diffusion and utilisation. In information processing fields such as military intelligence and earth resources, data acquisition rates outstrip analysis and interpretation capacities. In highly capitalised fields such as automotive and electronics manufacturing, global competition based on price and quality has highlighted the need to integrate and coordinate knowledge about all phases of product development. This challenge is exacerbated by a very significant acceleration of new technologies and a rapid shortening of product lifetimes. In all of these areas, the same point emerges: the computer has created both the need and the opportunity to enhance the distribution of knowledge. Knowledge systems address that need.

What will the future bring? Three basic trends are certain, from US perspectives at least. First, this technology will spawn many new and speculative products. Some of these will be incredibly successful. The big success may be an expert system for advising about wardrobe colours, one about job searching or one about improving your personal finances. Or perhaps the big success will be an inexpert, broad knowledge system: all about the Bible, a computerised pen pal or 'everything you always wanted to know about ...'. One cannot predict which consumer-oriented knowledge systems will succeed, but we can predict that most will fail commercially.

The second trend will see knowledge engineering continue to penetrate industrial and commercial organisation in a broad way. Unlike many information processing applications, knowledge system applications often create substantial near-term payback potential. Most large organisations will find many high-value applications of this technology. For example, the typical expert system may raise the average level of performance on a task by a factor of 2-10 over current practices. If the organisation processes thousands of transactions of this sort annually and the benefit of such performance improvement exceeds £100 per transaction, the company may realise an annual payback of more than a million dollars for each application. Since these numbers appear quite typical, more and more people will perceive the attractiveness of this technology in the years to come.

The third trend is toward closer integration of knowledge systems with data processing. Today, most computer-readable data elements reside in conventional databases. Many knowledge engineering applications need to access and analyse this data. In addition, many knowledge systems formulate plans for controlling other electronic systems. Both types of application will draw the two technologies closer together.

What industrial and commercial application areas will the technology emphasise? One can highlight several areas that arise repeatedly in studies of likely ways to exploit the technology. These include: expert systems for automotive and equipment repair; heuristic control systems for military functions and industrial automation; knowledge-aided design systems; knowledge-based planning aids; and automated interpretation systems for sensors and instruments. However, the technology will flow into all sectors of the economy, in an unpredictable fashion, because it provides an alternative means of supplying to industry an essential input — knowledge. In the long run, industrial use of knowledge systems technology will mature and the technology will become a standard input factor in most industries.

To summarise, experience to date in the US indicates that knowledge systems technology will find several paths to commercial markets. Before the highest-value applications of the technology become apparent, many more products and uses will need to be developed and tested. Somewhat ironically, the rate of success of knowledge engineering now depends on the efficiency with which the field can discover opportunities. That the technology will find several high-value targets is not, however, in doubt.

REFERENCES

[1] Hayes-Roth F, Waterman D A and Lenat D B. Building expert systems Addison-Wesley Publishing Co Inc, Reading MA (1983)

[2] Holusha J 'GM'S Silicon Valley stake' The New York Times Business Page (3 May 1984)

[3] Barr A and Feigenbaum E A 'The handbook of artificial intelligence' vols I and II William Kaufman Menlo Park C A (1981/1982)

[4] Lindsay R D, Buchanan B G, Feigenbaum E A and Lederberg J 'Applications of artificial intelligence for organic chemistry: the DENDRAL project' McGraw-Hill Book Co New York (1980)

[5] Martin W A and Fateman R J 'The MACSYMA system' Proc 2nd Symposium on 'Symbolic and algebraic manipulation' pp 59-75 (1971)

[6] McDermott J 'R1: an expert in the computer systems domain' Proc 1st Annual National Conf on 'Artificial intelligence' pp 269-271 (1980)