Canada

EXPERT SYSTEMS

REPORT No. 1/85

OFFICE OF INDUSTRIAL INNOVATION

INNOVATION

Government of Canada / Gouvernement du Canada
Regional Industrial Expansion / Expansion industrielle régionale

EXPERT SYSTEMS

Alain Letendre

Office of Industrial Innovation
Department of Regional
Industrial Expansion

October 1985

TABLE OF CONTENTS

Page

FOREWORD

This technology situation report is the first of a series on emerging technologies to be published by the Technology Assessment Directorate of the Office of Industrial Innovation.

The purpose of this study is to demystify the technology of expert systems by informing the Canadian public of the opportunities, uses, limitations and challenges presented by the latest developments in this field.

It is aimed primarily at investors who may be attracted by the potentially phenomenal growth in this sector as well as Canadian software firms that may see in expert systems an opportunity to diversify into one of the most promising fields of artificial intelligence. We would also like to make potential users, including the government, aware of the possible benefits of the use of expert systems in products and processes. The study will, therefore, be widely distributed to these target groups.

In its effort to promote and diffuse technology advancements, the Technology Assessment Directorate emphasizes the development of networks of interested individuals, as well as the linkage between users and suppliers of these technologies.

Readers wishing to participate in those networks or comment on this report should feel free to contact the undersigned.

Opinions and views expressed in this report are those of the author and are not necessarily endorsed by the Department of Regional Industrial Expansion.

Alain Letendre
Office of Industrial Innovation
Department of Regional Industrial Expansion
Ottawa, Ontario
K1A 0H5

(613) 995-3741

SUMMARY

The main objective of this study is to present an overview of expert systems which will make Canadian industry more aware of the technology, as well as inform it of the opportunities, uses, limitations and challenges raised by the latest developments in expert systems.

Expert systems may be defined as software programs that can solve problems at a level of difficulty which would normally call for a human expert. Depending on the problem, the system may either replace or provide assistance to an expert. They are programs which attempt to computerize a number of intellectual functions and will also be used increasingly to give a degree of intelligence to various machines such as robots and laboratory instruments.

Unlike conventional software programs, which are based on algorithms, expert systems use a mechanism called an inference engine to search, manipulate and interpret the contents of a database containing knowledge relating to a given field. This knowledge includes facts, models and, most important, heuristics. These are analogous to the rules of thumb used by an expert in carrying out his work. An example of heuristics at work in the game of tic-tac-toe would be: "If you go first put your X in the centre."

Expert systems are of interest at two levels. First, the information industry is attracted by a world market for expert systems that could top U.S. $1.2 billion within five years. Second, users are stimulating demand because they can see that it would be possible to recover their investments very rapidly. Several good examples of potential benefits can be found on pages 4 to 6.

Expert systems are applicable in all situations where a large quantity of data must be analyzed or where a problem cannot be solved by means of an algorithm. The solution, however, must never depend on common sense.

Development costs will be recovered more rapidly if the application is in a capital-intensive field where experts are few and far between and where unused equipment costs are relatively high.

Here are a few examples of situations that are particularly suitable for expert systems: There is no local expert readily available and advice is needed in a hurry (e.g., a drill breaks during the drilling of an oil well); a fast-breaking situation which is too stressful for a human

being to deal with (e.g., an emergency in a nuclear power station); a number of specialists are about to leave a firm and it is difficult to find replacements for them.

Existing systems have been applied to a wide variety of functions ranging from data analysis and interpretation to the formulation of concepts and planning. The applications affect nearly all possible areas of endeavour: mine and oil exploration, medicine, engineering, chemistry and biotechnology, mathematics, marketing, agriculture, law, accounting and many other fields.

It is important to bear in mind that the technology has only recently reached the commercial market and that most systems are still at the experimental stage. Furthermore, we are only now beginning to use a number of these systems routinely. With rare exceptions, these systems are very costly because they are custom-designed to meet the specific needs of given clients.

Nevertheless, there are three clear trends that will lead to explosive growth in the number of systems in use. First, simple and inexpensive systems are currently being developed which can run on personal computers. While limited in scope, they may be quite useful. Second, we will soon see software packages for large expert systems that will be mass-marketed to offset development costs. Even large systems, therefore, will be within the reach of medium-sized businesses. Finally, for those who wish to develop their own systems, expert systems shells (or frameworks) are now being sold by the major consultants and producers in the field. These shells are like expert systems with an empty database. All that is required to obtain an operational prototype is the addition of the knowledge from a specific field.

Those tempted to undertake the development of a large expert system should realize that development costs may be very high, specialists are few and far between, the computers required are relatively costly, the programming languages used are little known outside expert circles, and the systems still suffer from limitations that may, in some cases, become very serious.

The fact remains that expert systems are among the first commercial applications of artificial intelligence and that the opportunities for applications are so great that firms like Northern Telecom, General Electric and GM have launched the development of such systems for internal applications. The rate of growth for the technology is phenomenal. The numbers of simple systems running on personal computers or imbedded in hardware like scientific instruments should increase dramatically.

To date, the United States has, without doubt, been the leader in the field, followed closely by Europe. Japan, however, has spent enormous sums of money to take the leadership in the race. This has led other governments to establish costly programs to support these new technologies.

In Canada, university activity in expert systems is at least as intense as in most European countries. In the private sector, many consulting firms and software manufacturers have begun work in the field. A growing number of firms have begun to develop their own expert systems, though the trend is not as marked as abroad. Within the federal government, a number of departments have shown interest in the technology, and this should significantly increase the demand for such products and services.

1.0 INTRODUCTION

1.1 Objectives

The main objective of this report is to give an overview of expert systems which will demystify the field and inform Canadian industry of the opportunities, uses, limitations and challenges raised by this technology. The study could indirectly facilitate the dissemination of expert systems and increase the interest of potential investors.

1.2 Links with Artificial Intelligence

Expert systems are one of the first commercial applications of artificial intelligence. Other markets for artificial intelligence are: natural language applications (human languages, not formal languages such as those used in data processing); computer-assisted instruction; visual recognition systems; robotics and integrated manufacturing systems; and speech recognition systems. (The boundaries between these various fields are not always clearly defined.)

Many analysts have claimed that the field of expert systems will be the artificial intelligence market with the highest growth rate over the next five years. What is certain is that it is one of the most talked-about technology.

The underlying principle for expert systems is the use of a large quantity of practical knowledge. They are also based on a new philosophy of artificial intelligence which has abandoned, at least for the moment, the idea of developing a general theory of intelligence founded on a single problem solving algorithm. The new philosophy assumes that intelligence is based on an enormous quantity of knowledge. This includes first principles, such as that an object falls to the ground if it is left suspended in mid-air; the most complex of scientific principles; as well as less concrete aspects, such as beliefs and expectations.

1.3 Definition

Feigenbaum, one of the pioneers in the field, defined expert systems as follows:

> An 'expert system' is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution.
>
> (Feigenbaum and McCorduck)

This definition clarifies the basic difference between expert systems and ordinary computer programs. The latter make use of algorithms, i.e., they tell the computer in detail what steps to follow to solve a problem. The program, therefore, always proceeds in the same invariable fashion.

Expert systems, however, use their knowledge in various different ways depending on the problem to be solved. A later section on the technology of expert systems will demonstrate in greater detail how they work.

The knowledge used in expert systems includes facts, models and heuristics.

Facts are the widely shared public knowledge, commonly agreed-upon among practitioners.

Heuristics are rules of good judgment, i.e., rules for plausible reasoning, rules for estimating or 'the art of good guessing'. These are the little-discussed and almost private rules which characterize the way a human expert works in a given field. Such a rule for the game of tic-tac-toe would be: "If you go first, place your X in the centre." Another such rule, which could be used to discover new laws, would be to place a priority on considering extreme cases.

Heuristics are an elegant way of getting around the problem of considering the multitudinous possible combinations in attempting to solve certain problems. The chess problem of king-and-knight against king-and-rook, for example, has two million possible configurations. This particular problem may be solved using only 30 rules.

2.0 SIGNIFICANCE OF EXPERT SYSTEMS

2.1 The Purpose of Expert Systems

Until recently, only the primary and secondary sectors were suitable for automation. The arrival of artificial intelligence was said to be the dawn of a revolution in which certain types of intellectual work could begin to be computerized (Letendre 1983).

Expert systems claim to be an answer to problems generally associated with knowledge, which is perishable, rare and difficult to apply.

Society's organizations present a problem. An employee, for example, becomes experienced in a position and then is promoted to a higher level: his expertise is lost from the first position. Another example is that of expertise which is currently the preserve of a small number of employees, such as tinsmiths, who are near retirement. This expertise will often be lost when the employee retires because no new staff have been trained as replacements.

Just as no individual can aspire to excel in a large variety of activities, expert systems currently available are also subject to such limitations.

Expert systems could be a solution to problems that cannot be dealt with using current data processing methods. They will certainly become a way to preserve and disseminate a vast range of expertise. The value of experts, therefore, will be amplified when their expertise is made available to all. **Expert systems will be the new Gutenberg revolution.**

In the shorter term, they will be used to manage information, advise professionals, act as electronic assistants, for example, by advising a surgeon on the best approach to take for a specific surgery or signalling the existence of contradictions between hypotheses and generally accepted knowledge.

By means of a few examples, we will demonstrate the usefulness of a few expert systems. These examples are all taken from Feigenbaum and McCorduck's book "The Fifth Generation - Artificial Intelligence and Japan's Computer Challenge to the World."

Knowledge engineers at the Stanford Research Institute International, working with scientists of the U.S.

Geological Survey, built PROSPECTOR, a consulting system in geology. The program has knowledge about geology and mineralogy in general; as well, it has been given specific knowledge about particular regions, such as the Mississippi River basin and the major mountain ranges of the United States. In 1982, the expert system was used by a company exploring for minerals in the state of Washington. The system successfully identified a molybdenum deposit, valued at more than $100 million. The deposit was previously known, but its extent was not.

A large American firm recently developed an expert system to diagnose failures in steam-driven electric power generating plants. The diagnostic of the system is based on chemical measurements taken in the effluent of the plant. A prototype of the expert system, based on an actual 1981 plant failure and a four-day shutdown that cost the company $1.2 million, succeeded in reasoning its way to a correct diagnosis of the difficulty in a matter of seconds.

The French oil company, Elf Aquitaine, contracts out oil well drilling services to oil companies. Elf would like to keep its own drilling experts at the well site all the time because errors can be extremely costly in both money and time. Occasionally wells that have cost $1 or $2 million to drill must be abandoned because of serious mistakes made in recovering from drilling problems. Elf Aquitaine contracted with Teknowledge Inc. with the help of an Elf drilling specialist to develop an expert system called the DRILLING ADVISOR. This system diagnoses a variety of drilling problems and offers recommendations for corrective action as well as recommendations for preventing further similar problems. One of the reasons the system is of such great interest is that waiting for an expert to arrive can cost more than $100,000 per day.

Hitachi is currently developing a system for the management of an integrated circuit production line. (Northern Telecom appears to be developing a similar system independently.) In this field, the percentage yield of good chips is critical to profitability. If the yield systematically begins to drop, the manufacturing experts may sometimes take a few days to diagnose the source of the problem and take remedial action. Even if the expert system brings only a marginal improvement, millions of dollars could be saved every year.

In biotechnology, even though powerful tools are available to researchers to determine the sequences of DNA bases, the process can be tedious, difficult, and error-prone. IntelliGenetics, a small company created by Dr. Feigenbaum, developed an expert system to assist biologists and geneticists in the analysis of sequences and in the interpretation of experimental results. The system, which is available commercially, not only increases the productivity of the biotechnology specialists, but is also more efficient than human experts.

Cognitive Systems of New Haven, Connecticut, developed COURTIER, a system for consulting on the selection of a stock portfolio. It will soon be installed in the entrance halls of the branches of a major Belgian bank. COURTIER will conduct an interview with the customers to obtain information about their financial position. Then, on the basis of the customers' current portfolios, the best available investment analysis techniques, the normal practices in the country and current market conditions, COURTIER will advise them about whether to buy or sell shares. The system can be interrupted at any point and the customers can express their own feelings about the recommendations or request general investment advice.

The best-known example is without doubt that of Digital Equipment Corporation (DEC), the mini-computer manufacturer, which customizes each computer to conform to specific customer needs. Because of the large number of modules that need to be put together for each machine, subject to large numbers of constraints and conditions, about 20 per cent of their systems were showing problems of one kind or another.

An expert system called EXCON (EXpert CONfigurer) was developed to advise DEC personnel about the configuration of the company's computers. The system has been used regularly since 1980 and it checks each order submitted by sales staff. The system first checks to see if the selected configuration is feasible and complete regarding types and lengths of cables, components, memory size and so on. It then recommends additions or corrections and prints out a report to help manufacturing staff assemble the computer correctly.

The system is said to have a 99 per cent success rate, and so saves the company from costly errors.

Because mistakes are often produced at the time of customer order, it is important to catch the mistake at this time rather than at the time of manufacture. Because it is also less expensive to isolate problems at this stage, the company has already begun the development of an expert system for its sales force, EXSEL (EXpert SELling assistant). EXSEL will help DEC sales staff to identify the configuration best suited to meet customer needs, and to estimate the delivery dates and requirements for preparing the facility.

For obvious reasons, it is impossible for outsiders to be certain how much EXCON is saving DEC. The estimate in the artificial intelligence field is U.S. $20 million per year. What is certain is that the system is profitable. Otherwise DEC would not have followed up its initial EXCON expert system with a complementary system.

2.2 Markets

Artificial intelligence is a technology which has just come out of the laboratories. As a result, market projections can be no more than rough estimates. Moreover, data available on current markets are often contradictory. This section is provided, therefore, only as a general guide.

According to DM Data Inc., the North American market for artificial intelligence products and services could reach U.S. $719 million by 1985 (Figure I). Much of the market will be the sale of computers optimized for the use of artificial intelligence languages.

If we exclude the sale of these computers and related computer languages, the North American market for artificial intelligence **applications** would total U.S. $358 million (Figure II). Visual recognition systems (especially robotics and automated inspection) account for 47 per cent of this market. Expert systems follow in second place, with 18 per cent of the market, equal to the natural language programs whose applications are really interfaces in querying data banks. Anyone who has ever attempted to use a software program like the popular DBase II will understand the attractiveness of a more user-friendly interface.

Many analysts agree in predicting that the market for artificial intelligence goods and services could grow dramatically, beginning in 1986.

Voice-activated terminals, robots capable of recognizing complex shapes in a controlled environment and electronic consultants are only a few examples of the range of new products that could reach the market within the next six to seven years.

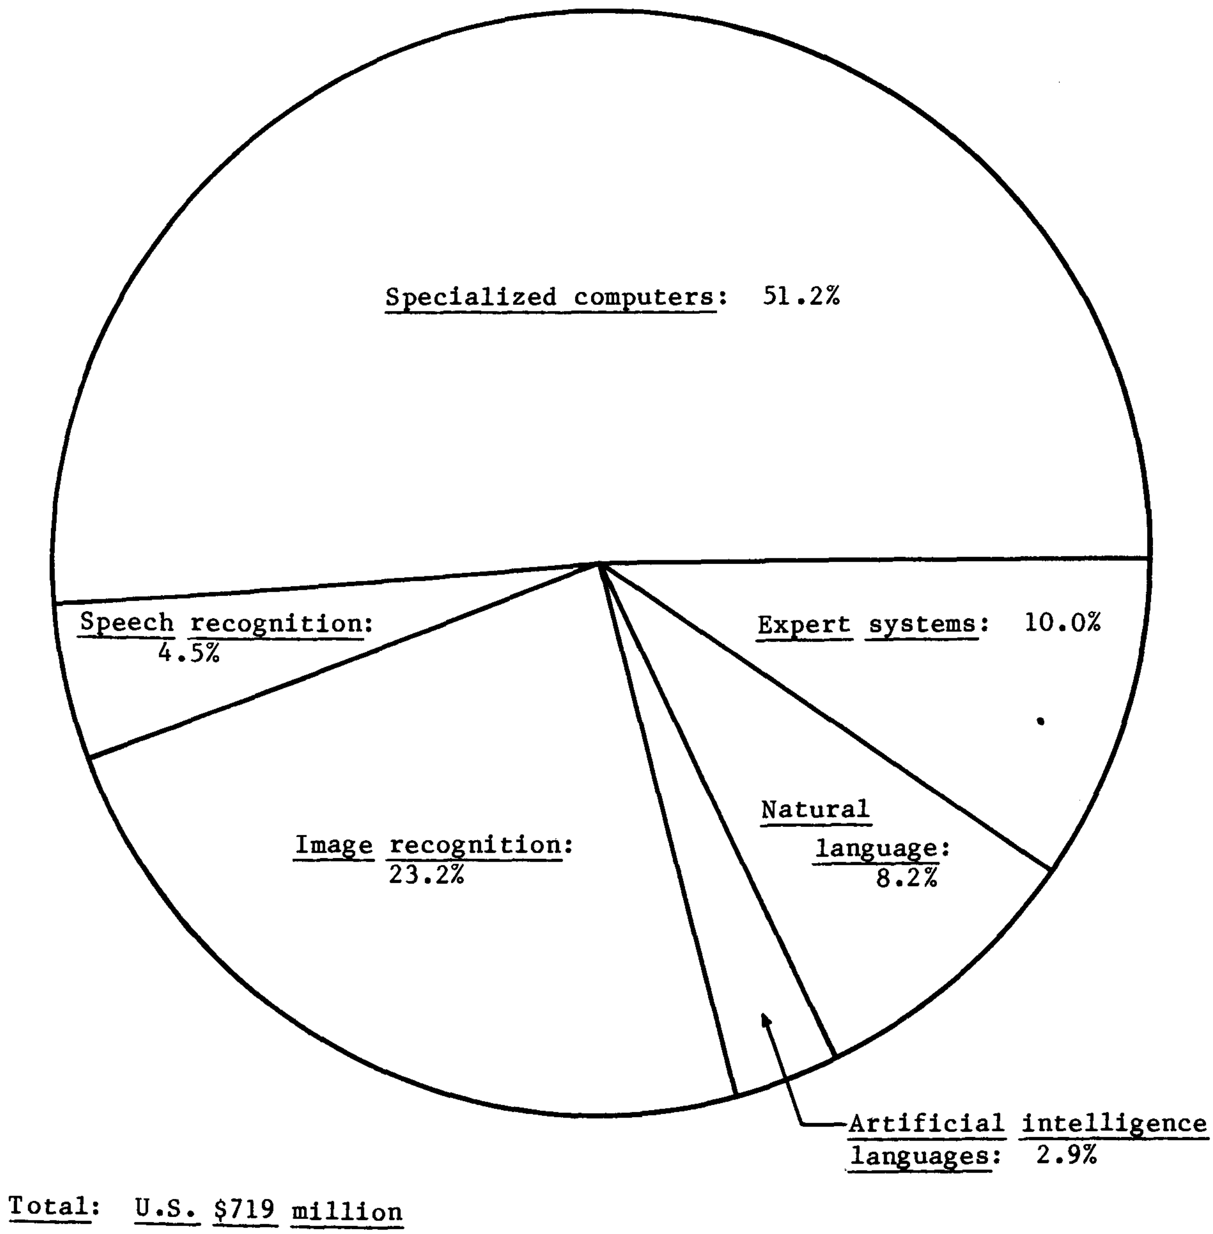

Figure I. Total artificial intelligence market by product for 1985.

Source: AI Trends '85.

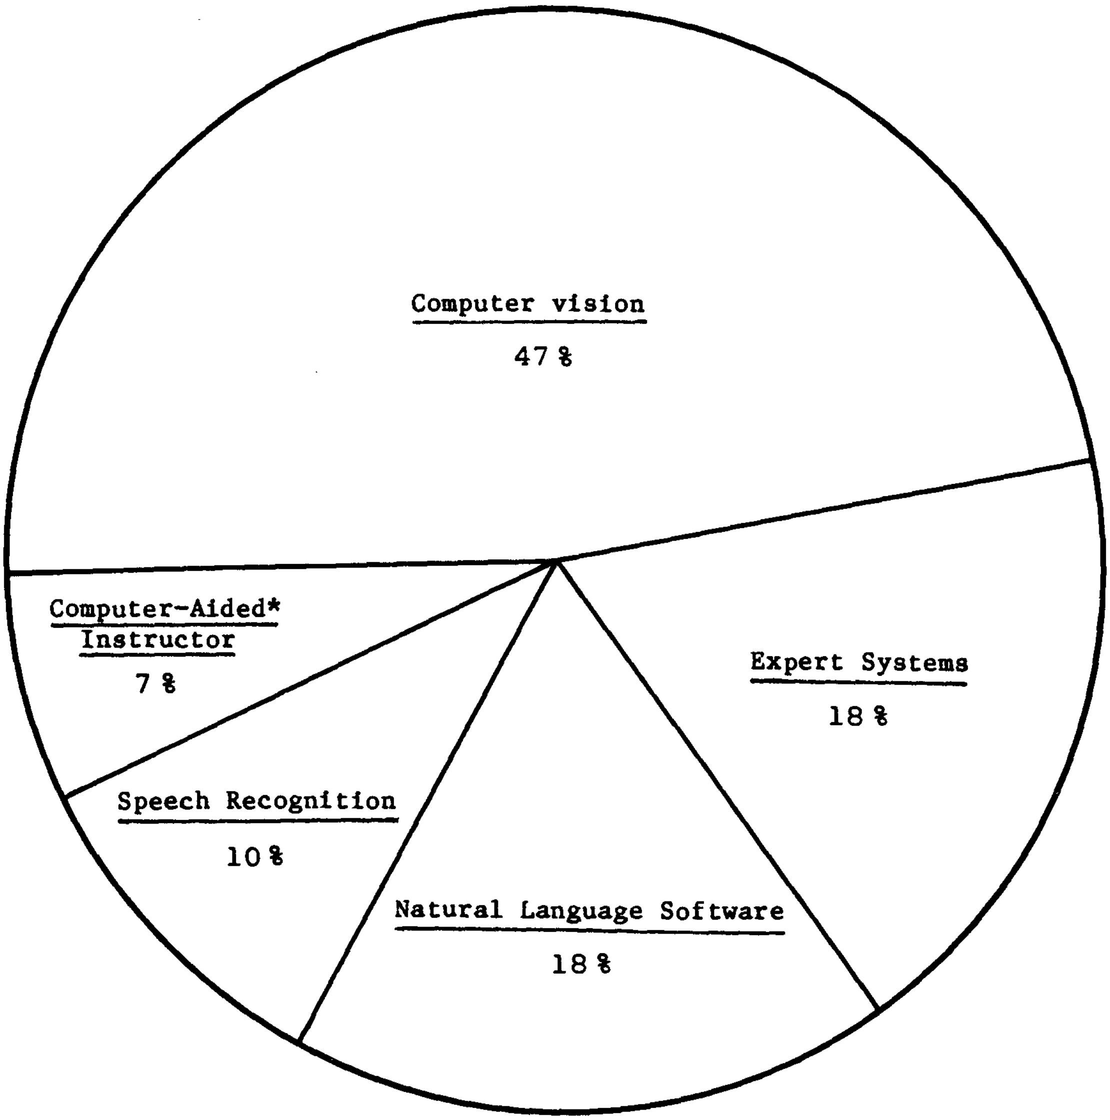

Total: U.S. $358 million

Figure II. Breakdown of sales for each market segment of artificial intelligence applications.

Source: AI Trends, March 1985.

Figure III shows that a high proportion of the current market is captured by sales of services (especially consulting). The 1993 market will consist primarily of software, with the share of the hardware remaining constant and the services sector undergoing a relative decline.

Figure IV estimates that the total world market for artificial intelligence products will be about U.S. $4.0 billion in 1990. Each new study usually increases the figures forecasted by earlier studies. For example, A.D. Little Inc. forecasts a world market of U.S. $11 billion for 1990 (i.e., 4.5 per cent of the data processing market) and employment totalling 100,000 person-years.

The market for expert systems will grow rapidly - 86 per cent per year over the next five years. Sales will reach $1.2 billion by 1990, i.e., slightly more than one-quarter of the total market for artificial intelligence goods and services. The magazine 'AI Trends', in its March 1985 issue, published an eloquent analysis of the industry. It established that the American market in 1985 would be U.S. $73.7 million (Figure V).

Only a few years ago, expert systems applications gravitated around the military and diagnostics sectors; it would now appear, however, that the application sectors undergoing the greatest expansion are finance and manufacturing (Figure VI).

To conclude, expert systems are currently generating as much interest in venture capital circles as biotechnology did in the late 1970s and early 1980s.

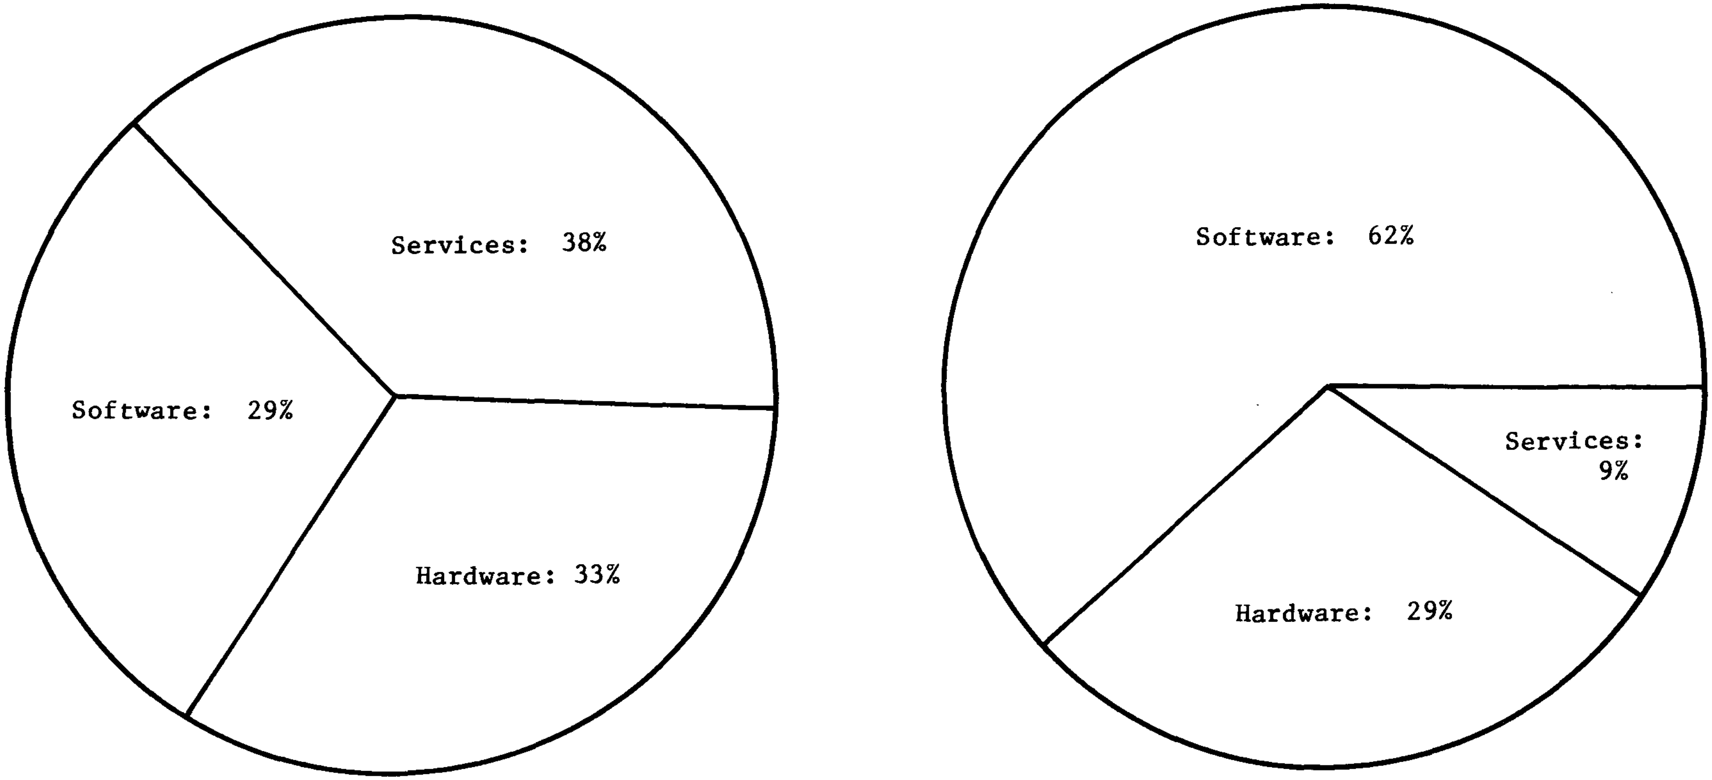

1983 1993

Figure III. Markets for artificial intelligence products and services.

Source: "Commercial Products Begin to Emerge from Decades of Research." Electronics, November 3, 1983, p. 128.

Segment	1983	1987	1990	Growth Rate (%)
Expert systems	16	800	1,243	86
Natural language	18	190	1,090	79
Computer-aided instruction	7	30	100	46
Visual recognition	30	230	860	60
Speech recognition	10	50	230	56
Hardware	25	200	500	54
Total market	106	1,500	4,023	68

Figure IV. World market for artificial intelligence ($ million).

Source: <u>Electronics Week</u>, 12 November 1984.

	1982	1983	1984	1985
AI Research Group			0.2	1.0
Advanced Information & Decision Syst.	1.8	2.7	3.9	6.0
Applied Expert Systems (APEX)		-	2.0	4.0
Artelligence Inc.			-	0.1
Brattle Research			0.6	1.0
Carnegie Group			2.1	7.0
California Intelligence			-	0.1
Computer*Thought		0.5	0.8	1.0
Digital Equipment Corp.			0.2	2.0
Expertelligence			-	0.1
Expert-Knowledge Systems Inc.			-	0.8
General Research		0.1	0.3	1.0
Iconics	1.0	1.0	2.0	4.0
Inference Corp.		0.2	1.0	2.0
Intellicorp	0.4	1.5	2.1	3.5
Intelligent Software Inc.			0.1	0.2
Jeffrey Perrone & Assoc.			0.1	0.2
Migent Software Inc.				-
Palladian			-	0.1
Perceptronics		2.0	4.0	10.0
Radian Corp.			-	0.3
Reasoning Systems Inc.				-
Smart Systems Technology	0.3	1.0	2.0	3.0
Software Architecture & Engr.	1.3	2.0	2.5	5.0
Syntelligence		0.1	1.5	2.0
Systems Control Technology (STC)	0.5	1.0	1.3	1.3
Teknowledge Inc.	1.5	2.2	3.0	5.0
Verac Inc.			1.5	3.0
Others	2.5	3.0	7.0	10.0
TOTAL	9.3	17.3	38.2	73.7

Figure V. American market for expert systems by company (in U.S. $ millions).

Source: AI Trends '85.

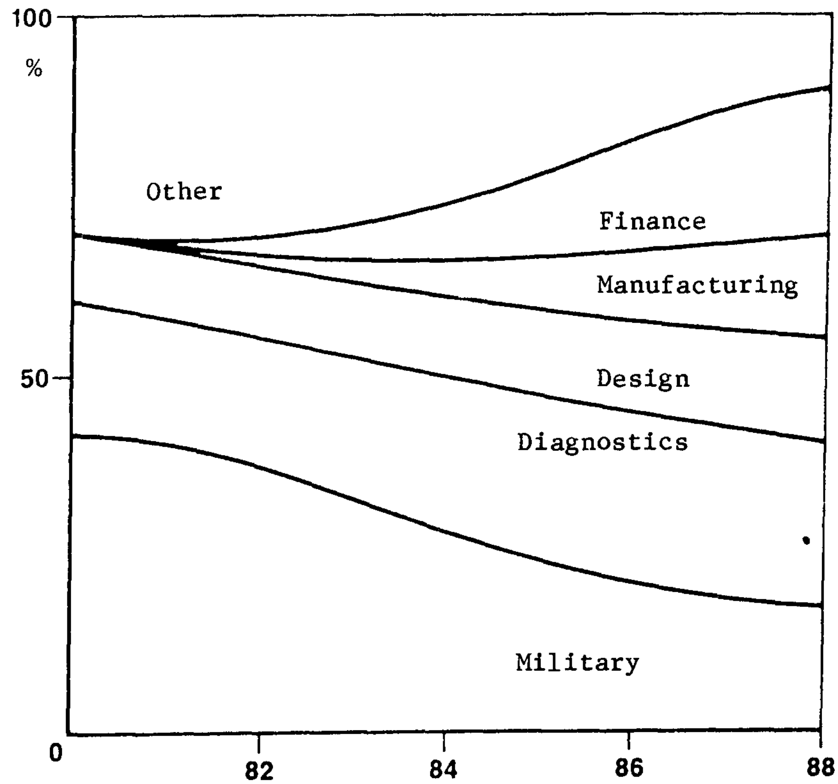

Figure VI. Trends in expert system applications.

Source: AI Trends '85.

3.0 APPLICATIONS AND USES

3.1 Basic Conditions

Experts must be available and able to articulate their knowledge and clearly state their approach. Their expertise must be based on their knowledge, judgment and experience. Ideally, the knowledge in question should not be a matter of dispute among specialists.

3.2 Types of Situations in which Expert Systems May Be Applied

Expert systems are applicable in all circumstances where a large number of data must be analysed, especially for problems suffering from combinatorial explosion as in the example of configuring computers at Digital Equipment Corporation.

This approach should also be used where traditional data processing techniques are inapplicable. This is the case, for example, for problems that cannot be described using algorithms.

The expertise in question must be relatively scarce and the system eventually produced would have to be used as frequently as possible. An example of an ideal situation would be the case of an application in a capital-intensive field where there were few experts and where idled equipment costs were relatively high. It goes without saying that a high return on investment is desirable.

Here are examples when an expert system could be used advantageously:

- No expert is present and advice is required rapidly (e.g., drilling platform);

- A situation is changing too rapidly and is too stressful for a human expert (e.g., nuclear power plant emergencies);

- A task requires continuous monitoring;

- A situation where it would be useful to combine the expertise of a number of specialists;

- Specialists are about to leave a firm and it is impossible to replace them;

- Experts could use assistance in improving their productivity, helping them to manage information and providing them with a tool to allow them to verify inconsistencies between various hypotheses;

- Few experts are available and new ones must be trained.

In a few years, when we will have learned to rapidly and easily update knowledge in expert systems, we will also be able to use them under conditions in which technology is changing rapidly.

3.3 Types of Problems that Can Be Solved

The success or failure of the development of an expert system will largely depend on the selection of the application, i.e., the problem that the system will try to deal with. Dr. Hayes-Roth, of Teknowledge, recommends choosing a problem that requires between three hours and three weeks for an expert to solve. This recommendation eliminates problems which are either too simple or too complex.

The problem should be clearly defined and given accurate limits. It may require a great deal of knowledge, but not an in-depth knowledge of the world around us. The knowledge should be input in the form of rules as much as possible and be described in a few hundred words.

To keep the number of rules reasonable, the problem selected should not require an analysis that is too intensive. The problem-solving method chosen will use symbolic reasoning rather than common sense, which, unfortunately, has not yet been simulated on computers. The task must therefore be cognitive, and this excludes, for example, the selection of a stock portfolio or tennis instruction.

The problem-solving method must be a structured one. An example would be programming the manufacturing operations of a plant or the development of an insurance policy.

It goes without saying that data must also be available on the problem. The existence of case studies is also essential if the system's operation is to be checked.

Although only a few years ago many people anticipated grandiose applications, current trends are for systems that are more modest and realistic, i.e., modest tasks which are well-defined, narrowly limited and uncritical for the organization.

The successes to date of limited intelligence expert systems in specific fields, however, have demonstrated the great relevance of the approach for industrial applications of artificial intelligence technologies.

3.4 Machine Applications

So far we have seen that expert systems are primarily used as electronic consultants or assistants to human experts.

Expert systems, as described in Letendre (1983), are also a new approach to artificial intelligence in which it is assumed that intelligent behaviour is based on a large quantity of knowledge. Expert systems, therefore, may also be an integral part of a many so-called intelligent systems.

For example, it is now possible to have a variety of instruments that can operate automatically and that are also capable of interpreting their own results. There is an electrocardiograph with such capabilities; another example is a system currently being developed to interpret remote sensing images.

The Japanese Ministry of Transport is currently undertaking the development of an intelligent ship capable of auto-piloting in ports. The system will have the knowledge of a professional pilot and also be able to diagnose a vast range of problems involved in operating the ship.

The HASP/SIAP system is a passive sonar surveillance system that interprets sounds in a noisy environment.

An interesting application is the SYNTHEX system (SYNTHesis EXpert System) developed by researchers at the Centre national d'études des télécommunications (CNET) in France. The system translates written texts into spoken words. Rather than the nasal sounding voice that science fiction addicts have learned to love, SYNTHEX can produce speaking styles which range from a normal reading voice to the intonation of an auctioneer or an encyclopedia salesman.

In the longer run, we will see super-robots, intelligent weapons, machines that will type out dictated texts and even independent and flexible factories.

3.5 Applications by Function

Expert systems are difficult to classify by function because they often carry out more than one specific task or function.

Figure VII was compiled on the basis of the writings of Gevarter, Miller and a number of other writers. It gives a partial list of the best-known expert systems classified by field for each main function. In the 1984 Inventory of Expert Systems, there is a detailed description of a large number of systems.

The figure shows that expert systems can diagnose, analyse and interpret data; carry out analysis, design and planning tasks; formulate concepts; interpret signals; monitor; advise the user of an instrument; improve the performance of computer-aided instruction systems; automate data acquisition; facilitate the construction of expert systems; act as an intelligent assistant or consultant; manage; program; and, to a lesser degree, learn by experience and interpret images.

Though not mentioned in Figure VII, expert systems can also control, explain, assign tasks, verify, process, signal, plan missions and do many other kinds of tasks.

Another classification of expert systems is the simpler Hayes-Roth, Waterman and Lenat classification shown in Figure VIII.

	FIELD	SYSTEM	INSTITUTION
Analysis	Electric circuits	EL	M.I.T.
	Mathematics	MACSYMA	M.I.T.
	Mechanical problems	MECHO	Edinburgh University
	Naval forces threat analysis	TECH	Rand/NOSC
	Evaluation of structural damage caused by	SPERIL	Purdue University
	earthquakes		Rutgers University
	Digital circuits	CRITTER	
Automatic programming		CHI	Kestrel Inst.
		PECOS	Stanford University
		LIBRA	Stanford University
		SAFE	VSC/ISI
		DEDALUS	SRI
		Programmer's Apprentice	M.I.T.
Computer-assisted instruction	Electronics	SOPHIE	Bolt Beranek and Newman
	Medical diagnosis	GUIDON	Stanford University
	Mathematics	EXCHECK	Stanford University
	Operation of power generating station	STEAMER	Bolt Beranek and Newman
	Diagnostics	BUGGY	Bolt Beranek and Newman
	Causes of rain	WHY	Bolt Beranek and Newman
	Game teaching	WEST	Bolt Beranek and Newman
	Game teaching	WUMPUS	M.I.T.
		SCHOLAR	Bolt Beranek and Newman
Concept formation	Mathematics	AM	CMU
Consultant (intelligent assistant)	Weapons distribution	BATTLE	U.S. Navy Centre for applied research in art. intelligence
	Medecine	Digital Therapy Advisor	M.I.T.
	Radiology	RAYDEX	Rutgers University
	Computer sales	XCEL	CMU
	Medical treatment	ONCONCIN	Stanford University
	Nuclear station	CSA Model-Based Nuclear Power Plan Consultant	GA Tech

Figure VII. Current existing systems classified by function.
Sourcer: Miller. The 1984 Inventory of Expert Systems.

	FIELD	SYSTEM	INSTITUTION
Data analysis and interpretation	Geology	DIPMETER ADVISOR	M.I.T./Schlumberger
	Chemistry	DENDRAL	Stanford University
	Chemistry	GAL	Stanford University
	Geology	PROSPECTOR	SRI
	Protein crystallography	CRYSALIS	Stanford University
	Determination of causal relations in medecine	RX	Stanford University
	Determination of causal relations in medecine	ABEL	M.I.T.
	Oil well data	ELAS	AMOCO
Design	Computer configuration	XCON	C.M.U.
	Automatic programming	PECOS	Yale
	Circuit synthesis	SYN	M.I.T.
	Chemical synthesis	SYNCHEM	State University of N.Y.
Diagnostics	Medicine	PIP	M.I.T.
	Medicine	CASNET	Rutgers University
	Medicine	Internist/Caduceux	Pittsburg University
	Medicine	MYCIN	Stanford University
	Medicine	PUFF	Stanford University
	Computer breakdown	DART	Stanford University/IBM
	Medicine	MDX	Ohio State University
	Computer breakdown	IDT	DEC
	Nuclear reactor accidents	REACTOR	E G & G Idaho
Expert system construction		ROSIE	Rand
		AGE	Stanford University
		HEARSAY III	USC/ICI
		EMYCIN	Stanford University
		OPS 5	CMU
		RAINBOW	IBM
	Medical diagnosis	KMS	MO University
	Medical consultation	EXPERT	Rutgers
	Diagnostics for electronic systems	ARBY	Yale/ITT
	Medical consultation using temporal data	MECS-AI	Tokyo University

Figure VII. (cont'd)

	FIELD	SYSTEM	INSTITUTION
Image understanding		VISIONS	U. of Mass/Amherst
		ACRONYM	Stanford University
Knowledge acquisition	Medical diagnosis	TEIRESIAS	Stanford University
	Medical consultation	EXPERT	Rutgers
	Geology	KAS	SRI
Learning by experience	Chemistry	METADENDRAL	Stanford University
	Heuristics	EURISKO	Stanford University
Management	Automated plant	IMS	CMU
	Project management	CALLISTO	DEC
	Petroleum data modelling	PHOENIX	Schlumberger-Doll Res.
Monitoring	Patient respiration	VM	Stanford University
Planning	Chemical synthesis	SECHS	U. of California Santa Cruz
	Robotics	NOAH	SRI
	Robotics	ABSTRIPS	SRI
	Interplanetary flight	DEVISER	JPL
	Movement Planning	OP-PLANNER	RAND
	Molecular genetics	MOLGEN	Stanford University
	Mission planning	KNOBS	MITRE
	Task assignment	ISIS-II	CMU
	Experimental Design in Modecular Genetics	SPEX	Stanford University
	Medical diagnosis	HODGKINS	M.I.T.
	Aircraft carrier operations	AIRPLAN	CMU
	Tactical targets	TATR	RAND

Figure VII. (cont'd)

	FIELD	SYSTEM	INSTITUTION
Signals interpretation	Speech understanding	HEARSAY II	CMU
	Speech understanding	HARPY	CMU
	Automatic acoustics	SU/X	Stanford University
	Ocean surveillance	HASP	System Controls Inc.
	Ship sensors	STAMMER-2	NOSC San Diego/SDC
	Evaluation of performance of left ventricle	ALVEN	University of Toronto
	Evaluation of military positions	ANALYST	MITRE
User advice	Structural analysis	SACON	Stanford University
	Computer programs		

Figure VII. (cont'd)

Category	Example of problem targeted
Interpretation	Infer description of situation on the basis of data obtained from sensors
Prediction	Infer probable consequences of specified situations
Diagnostics	Infer system malfunctions on the basis of observable data
Design	Configure objects within constraints
Planning	Develop a series of actions to be carried out
Monitoring	Compare observations to identify weak points
Repair	Specify solutions to remedy failures and develop a plan for solving the problem
Instruction	Diagnose and isolate errors and modify student behaviour
Control	Interpret, predict, repair and monitor system behaviour

Figure VIII. Categories of expert systems.
Source: Hayes-Roth, F.; Waterman, D.A.; and Lenat, D.B., Building Expert Systems, Addison-Wesley, Reading, Mass (1983).

3.6 Applications by Field of Activity

Expert systems may be applied to almost all fields of human endeavour. The following are examples of fields in which expert systems are currently being used or developed.

- mine and oil exploration, as well as other geological applications;
- the oil industry - for managing and diagnosing problems related to oil wells and drilling platforms;
- medicine - for medical diagnosis and assistance in treating various diseases;
- engineering - for the diagnosis of equipment failures and maintenance including applications for steam turbines, aircraft and diesel locomotive engines, computers and telecommunications networks;
- manufacturing - for production planning, scheduling, shipping and production problem diagnosis;
- the design of integrated circuits, structural analysis and identification of mechanical problems;
- financial services such as banking and insurance, including exchange rate forecasts, personal finance and investment consulting, and the development of insurance policies;
- chemistry and biotechnology - for planning experiments and interpreting results;
- mathematics - for solving complex equations and deriving new algorithms;
- sales - for helping staff to configure products in the most suitable way for customers;
- management - to facilitate the use of data bases;
- agriculture - to forecast the theoretical yields of crops under varying conditions and to diagnose diseases and recommend appropriate treatments;
- buildings - to provide automatic monitoring and to optimize energy use;
- the control of nuclear reactors:

- air traffic control;
- law and accounting - to automate the interpretation of taxation regulations, auditing and the drawing up of wills;
- programming - to improve the productivity of programmers by automating routine tasks;
- computer-aided instruction - to improve the capabilities of existing systems, for example, by leading the student to develop hypotheses or by explaining their weaknesses and by adapting to individual learning rates;
- military applications such as target classification, automatic surveillance, aircraft carrier management.

Figure IX is a summary compilation of a number of expert systems. Each example is accompanied by a brief description of the system.

3.7 Currently Used Systems

Although hundreds of expert systems are being developed, an overview of the literature shows that very few are routinely being used in the operations of the firms involved in this field. Most are still experimental systems. However, in view of the considerable activity in the field, the relative newness of the technology and the interest of major firms like Xerox, IBM, Schlumberger, Elf-Aquitaine and General Electric, it is more than likely that this situation will change drastically within two to three years.

The following list gives examples of systems being used routinely; they do not include shells of expert systems currently being marketed almost everywhere: MACSYMA, MOLGEN, PUFF, DENDRAL, XCON, SMP, PICON (Process Intelligent CONtrol), ADA TUTOR, MICROPROCESSOR EXPERT, STEAMER and DIPMETER ADVISOR.

3.8 Technology Diffusion

We have just seen that the technology of expert systems is still in its infancy. In the life curve of a technology, expert systems are still at the early stage of the curve - the product is entering the market.

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Agriculture	Counsellor	A system for diagnosing diseases which recommends treatments using ICI products and which also gives an economic analysis of potential yields.	ICI (United Kingdom)
Biotechnology	Molgen	Molecular genetics experiments planning.	Stanford University
	Microprocessor Expert	Interpretation of protein electrophoresis results.	
Chemistry	Chrysalis	Interprets the diffraction diagrams of protein molecules.	Stanford University
	Dendral	Generates plausible chemical structures representing organic molecules on the basis of data obtained by mass spectroscopy.	Stanford Unviersity
	Genoa	A descendant of Dendral which interprets chemical structures.	Molecular Design (California)
	Meta-Dendral	Discovers new rules for the behaviour of fragments in mass spectroscopy.	Stanford University
	SECS	Organic chemistry synthesis planning.	Princeton University
Communications	ACE (Automated Cable Expertise)	Uses data from a data base and monitors telecommuncations networks to determine repair and maintenance requirements.	Bell Laboratories

Figure IX. Expert system applications by field of specialty.

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Data processing industry	DART	Isolates and recommends repairs on IBM computers.	Stanford University and IBM
	DOC	Assists service personnel in identifying common failures in information systems.	Prime Computer
	PS1	Writes simple computer programs based on English descriptions of tasks to be carried out.	Kestrel Institute and Systems Control Corp.
	XSEL	Ensures that computer orders are correct before undertaking construction and assembly.	Digital Equipment
	XSITE	Helps in preparing sites for computer installation by determining size of room required, electrical requirements and air conditioning, among other things.	Digital Equipment
	XCON	Helps in the configuration of DEC's various computers.	Digital Equipment and Carnegie-Mellon University
Energy	________	Diagnoses and suggests corrective action for operating an electrical generating plant.	Batelle

Figure IX. (cont'd)

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Engineering	EL	Analyses electronic circuits.	MIT
	Mecho	Analyses mechanical problems.	Edinburgh University
	SACON (Structural Analysis Consultant)	Consults the users of a complex software package (Marc) for structural analysis. Marc normally requires a year of training before it can be used.	Stanford University
	____________	Assists in the design of the pipe layout in chemical process plants.	ICI Products (United Kingdom)
Finance and Accounting	____________	Forecasts exchange rates.	Security Pacific National Bank (Los Angeles)
	____________	Advises on personal finance matters.	Rand Corp. (for a large American bank)
Insurance	____________	Advises and assists insurance agents in developing better insurance policies for their clients.	Cognitive Systems (New Haven)
Law	LDS	Experimental system which models the decision-making process of lawyers and insurance adjusters in product liability cases.	Rand Corp.
	Taxman	An experimental system that makes use of implicit rules in taxation laws and which recommends a sequence of contractual agreements that a company could use to achieve its financial objectives.	Rutgers Universtiy

Figure IX. (cont'd)

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Management	KM-1	A management system which attempts to integrate database management systems and expert systems.	System Development Corp.
	RABBIT	A system which helps the user to query a database.	Xerox Palo Alto Research
	Human Edge Series: Sales Edge, Management Edge, Negotiation Edge and Communications Edge	(U.S.) $200 to $300 systems for the Apple, based on personality traits and information about oneself and others.	
Manufacturing	Callisto	Models, monitors and develops implementation plan, and manages large projects.	Carnegie-Mellon University
	IMACS	Helps improve management of workflow within a plant in conjunction with traditional systems.	Digital Equipment
		Coordinates shipping of products to consumers.	Digital Equipment
	ISIS	Develops task implementation plans.	Carnegie-Mellon University
		Diagnoses manufacturing problems on a production line for integrated circuits.	Hitachi
Mathematics	AM	Discovers mathematical concepts.	Stanford University
	LEX 2	Develops heuristics to solve problems in integral calculus on the basis of its experience in computing integrals.	Rutgers University
	MACSYMA	Decides which algorithm to apply in solving algebraic equations.	Symbolics (Cambridge, Mass.)
	SMP (Symbolic Manipulation Program)	Commercial system for solving mathematical equations.	Inference (California)

Figure IX. (cont'd)

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Medicine	CASNET	Diagnoses eye diseases.	Rutgers University
	Clini-Scan	Interprets results of blood tests.	__________
	Digitalis	Helps in prescribing treatment of digitalis to patients suffering from cardiac problems.	MIT
	Headmed	Prescribes anti-psychotic medication to mentally ill patients.	__________
	Internist/Caduceus	Diagnoses up to 500 different internal diseases.	Pittsburg U.
	Mycin	Diagnoses bacteriological infections of blood and meningitis, and recommends an antibiotic therapy.	Stanford University
	Oncocin	Monitors response of cancer patients to medication treatment and prescribes the most effective and least toxic chemotherapy.	Stanford University
	PIP	Diagnoses renal diseases.	MIT and Tufts Medical Centre
	PUFF	Diagnoses pulmonary dysfunctions and infections.	Pacific Medical Center and Stanford University
	RX	Statistical analysis assistance for chronic patients for whom data are available over a period of several years.	Stanford University
	VM (Ventilator Management)	Advises corrective action for patients under ventilator treatment.	Pacific Medical Center

Figure IX. (cont'd)

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Microelectronics	KBVLS1	An experimental system for developing integrated circuits based on the Meed-Conway method.	Stanford University/Xerox
	Palladio	Verifies integrated-circuit designs and suggests alternatives.	Fairchild/Stanford University
Military	Acronym	A collection of interrelated expert systems that specialize in image interpretation.	________
	Analyst	Combines information from various sensors.	Mitre Corp.
	HASP/SIAP	A passive surveillance system which interprets signals in a noisy environment. Infers position and type of vessel in the Pacific on the basis of data transmitted by a submerged network of acoustics sensors.	System Control
	SU/X	Identifies aircraft from their flight characteristics and tracks them on the basis of data obtained from sensors by developing hypotheses about their position and speed and extrapolating.	Stanford University
	________	Robot-guiding and navigation system.	Batelle
	________	Classifies targets for the U.S. Navy on the basis of images obtained from side looking radar.	Hughes Aircraft (Advanced Information and Decision Systems Division)

Figure IX. (cont'd)

FIELD	SYSTEM NAME	DESCRIPTION	INSTITUTION
Oil Industry	Dipmeter advisor	Analysis of data collected by sensors.	Schlumberger
	Drilling Advisor	Advises well-drilling superintendent of ways to solve problems resulting from stuck drilling bits.	Teknowledge (for Elf Aquitaine)
	SPILLS	A crisis management system to help locate and identify accidental chemical spills.	Rand Corp.
	______	Diagnoses malfunctions of a nitric acid reactor.	Batelle
Resources	Hydro	Consultation system for hydrological resource problem-solving.	SRI International
Training	______	Teaches how to operate a steam generator.	Bolt Beranek and Newman
	MDIS (Maintenance and Diagnostics Information System)	A knowledge acquisition system for various fields which helps new staff to diagnose and repair system breakdowns.	Boeing Aerospace
	Ada Tutor	Computer-aided instruction system.	Computer Thought (Texas)
	Guindon	An intelligent computer-aided instruction system which teaches by soliciting responses to technical questions and then providing the correct answers.	Stanford University
	Prospector	Calculates a plan for finding likely deposits; used to determine where to carry out exploration drilling.	SRI International for US Geological Survey
Transportation Industries	CATS-1	Expert system for diagnosing and maintaining electric diesel motors.	General Electrics

Figure IX. (Cont'd)

There is still a big gap between the experimental product and the type of expert system that could be marketed successfully. Finding the right application for expert systems is not a trivial problem.

One of the main obstacles to the dissemination of expert systems is their cost, at least for large systems. Almost without exception, all such systems are customized for a single client. As well, the development of expert systems may require a great deal of effort. Finally, the systems operate in sophisticated environments: the development tools are costly and most systems operate on specialized machines. Prices are falling rapidly as competition becomes more intense and as integrated circuits optimized for these computers are developed. Furthermore, cheap development tools and applications, which can be run on personal computers, are now reaching the market.

A typical price for a large business system, such as the PICON expert system, is between U.S. $150,000 and $200,000 (the price includes a Lisp machine).

The Lisp (LISt Processing) and Prolog (PROgramming in LOGic) languages, used to develop expert systems, are little known outside the field of artificial intelligence. There are so many dialects specific to a given make of machine that it becomes difficult to transfer expert systems from one machine to another.

Although there are differing opinions in the literature, it appears that people who were not involved in developing the systems sometimes find them difficult to use. This problem is common to any software package but expert systems are, if anything, easier to use. Also, natural language research should solve any problems of this kind.

Other obstacles to the diffusion of this technology are the slow speed of some systems, their unreliability in certain situations, their lack of versatility and problems related to user acceptance.

The question of whether potential users will use the systems or not is psychological. If they perceive that the system can help them solve problems or find better solutions, and the approach is not too expensive, the system will probably be accepted. An important condition is that potential users should be involved in developing the system. The problem involves organizational change about which there are many books in every faculty of administration.

In conclusion, some practical systems are now being developed by groups of end-users at relatively low cost. Although the gains are modest in comparison with large systems, they are, nevertheless, significant. Furthermore, the experience gained this way will be invaluable for subsequent and more ambitious projects.

4.0 TECHNOLOGY

4.1 Basic Technology

In data processing, it was realized very early that not all problems could be solved by algorithms and that problem-independent heuristics could not cope with combinatory explosion. Logic alone cannot handle uncertain facts and vague axioms.

It was rapidly understood that human experts, moreover, have an enormous amount of knowledge and can clarify a problem, suggest the correct approach, assess the reliability to information and determine whether the solution is reasonable and plausible.

Unlike conventional software, expert systems derive their power from heuristics, not algorithms, and like human experts, have a large knowledge base.

4.1.1 Structure of Expert Systems

The basic principle underlying expert systems is the separation of the knowledge or database from the control structure. This differs from conventional software, in which data are an integral part of the program. Classical software programs are, therefore, very difficult to modify.

An expert system consists of a knowledge base, which includes facts and heuristics; an inference engine that uses the knowledge base; and a working memory to keep track of the status of the problem and the system and to store in memory the information entered by the user. Figure X describes the basic structure of expert systems and Figure XI gives additional details. The structure of the next generation of expert systems is described in Figure XII.

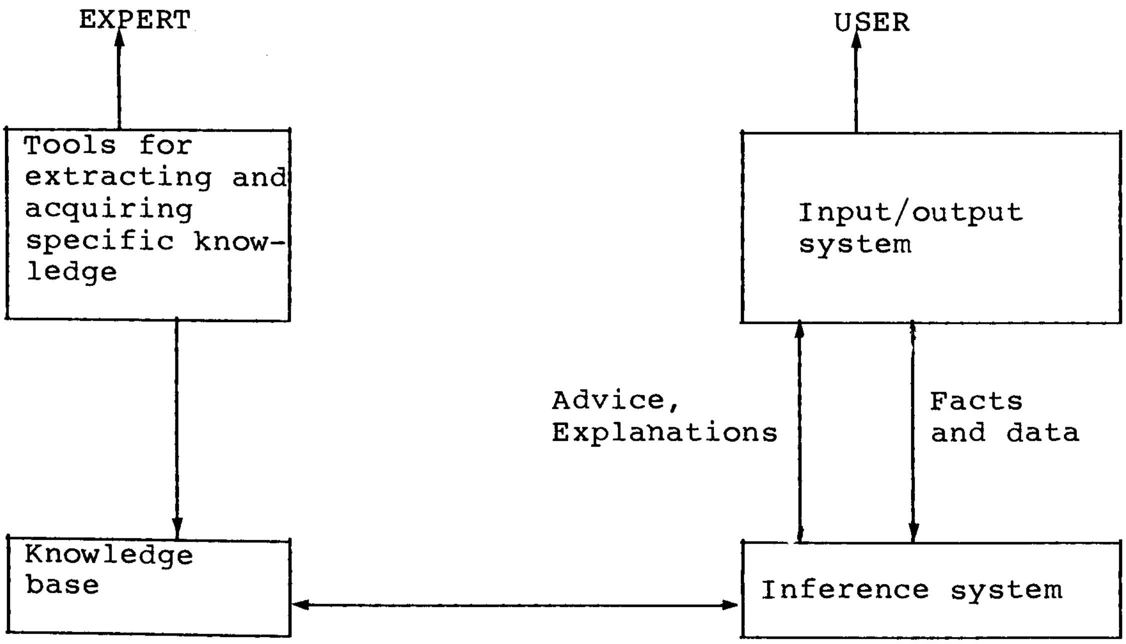

Figure X. Basic structure of an expert system.

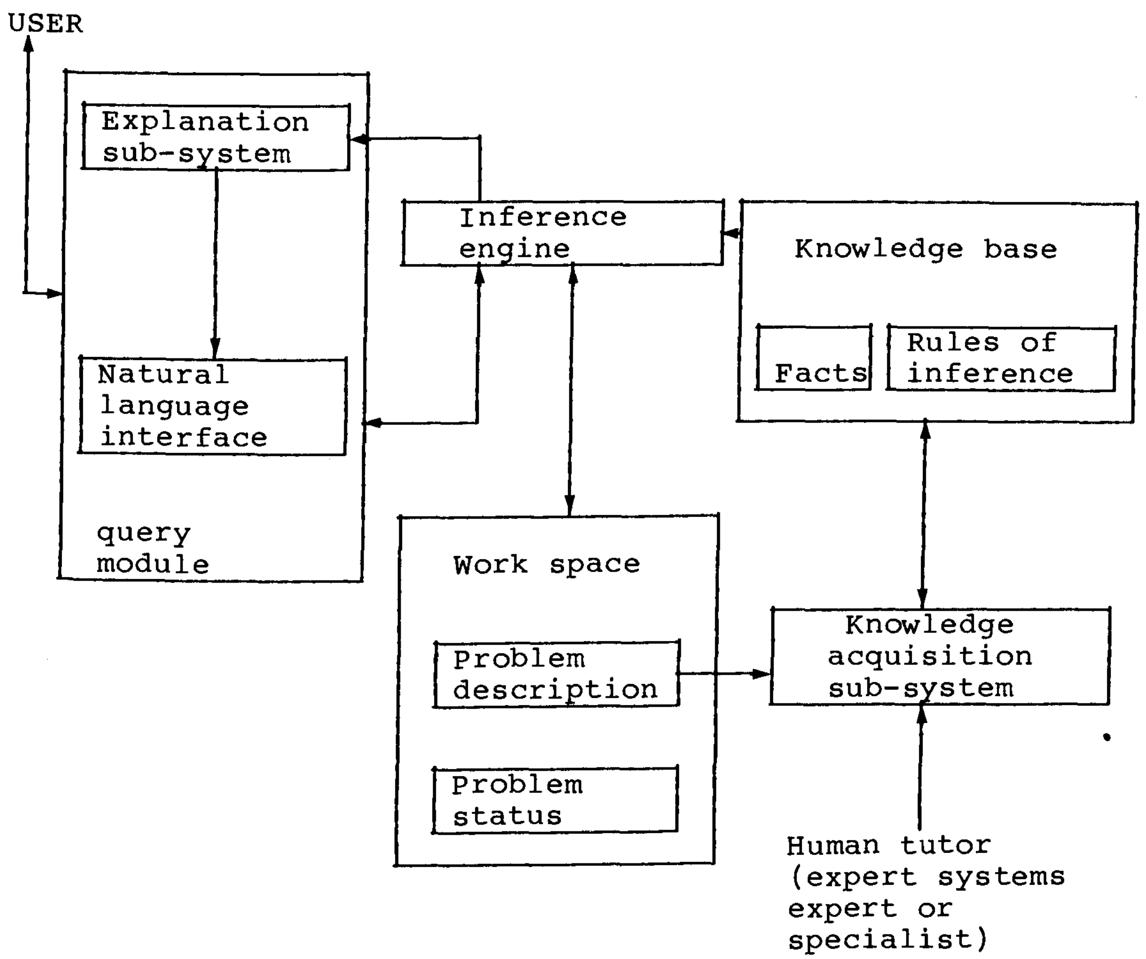

Figure XI. Detailed structure of an expert system. Source: P. Kinnucan. "Computers That Think Like Experts." <u>High Technology</u>, January 1984.

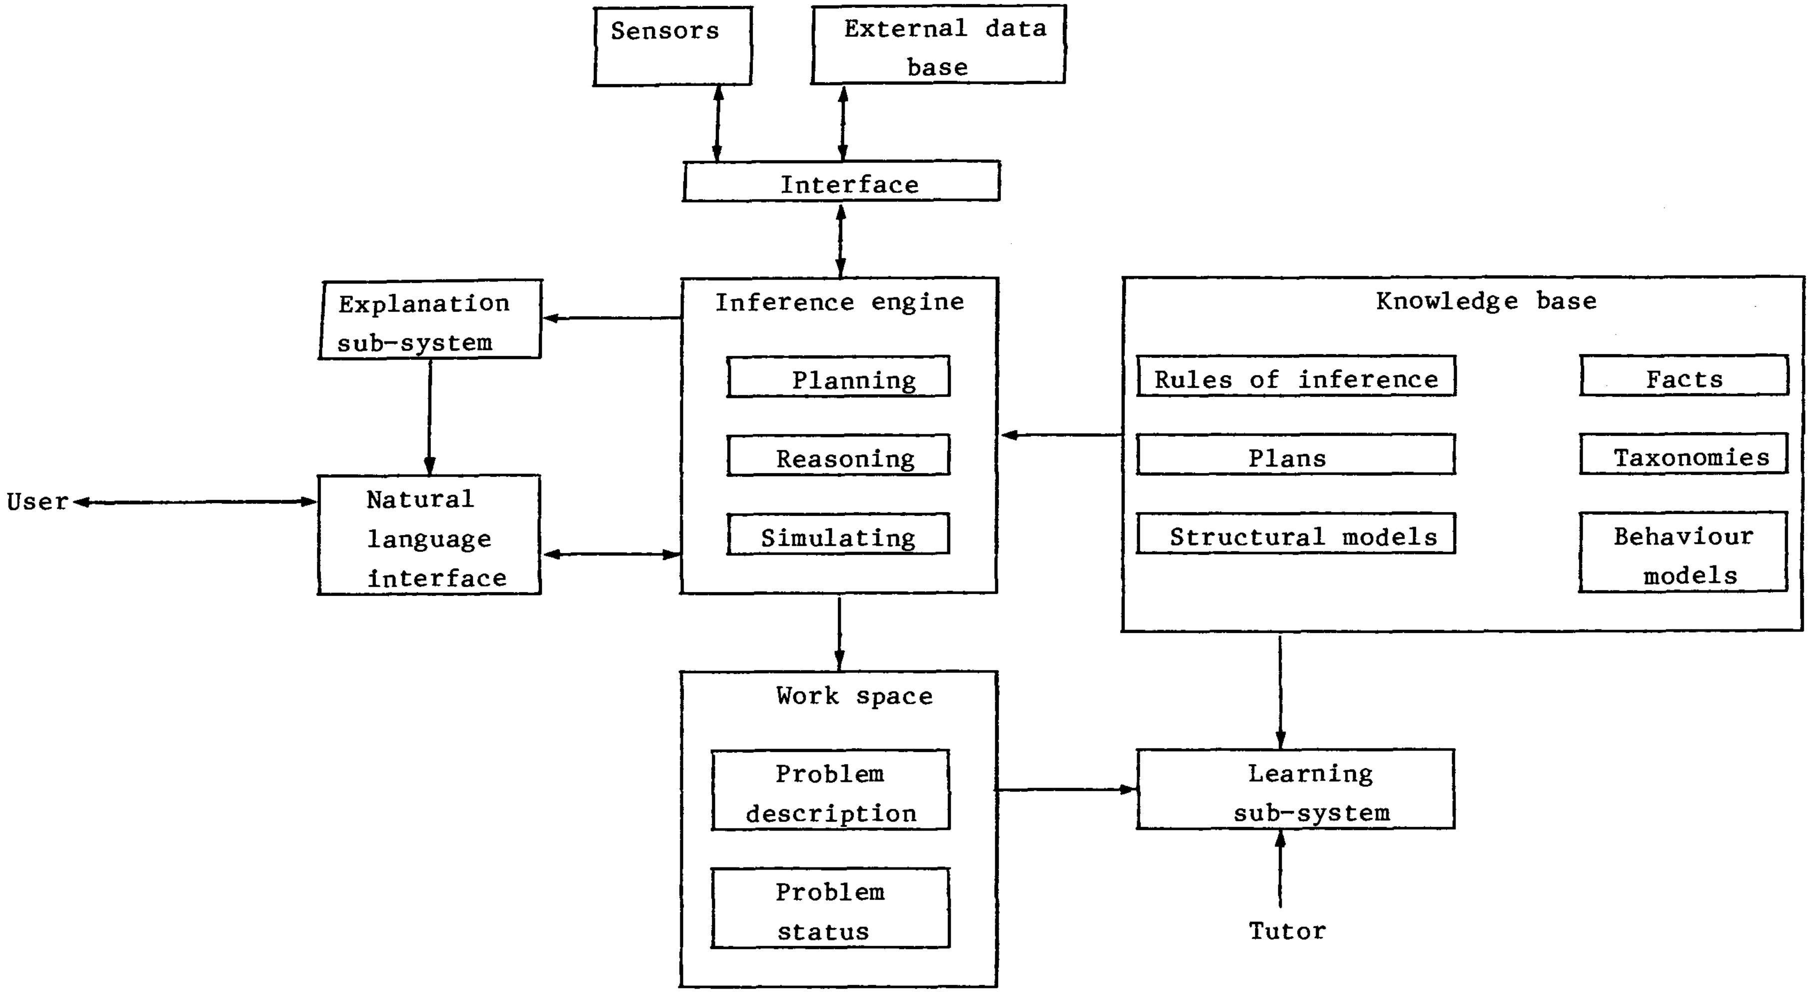

Figure XII. Future expert system.

Source: P. Kinnucan. "Computers That Think Like Experts." High Technology, January 1984.

Randall Davis, an expert in the field, stated the following principles for an expert systems architecture:

- the inference system (usually called inference engine) must be separated from the knowledge base to facilitate updating or modification of the knowledge base;

- the formalism used to represent knowledge must be as consistent as possible to keep the system as simple as possible;

- the inference engine must be kept as simple as possible to facilitate the explanations that the system will give of its 'reasoning' and to facilitate the knowledge acquisition process;

- redundancy must be used as much as possible to compensate for incomplete or imprecise knowledge.

4.1.1.1 The Knowledge Base

The knowledge base contains all public information about a given field. It includes terms and facts, including ways to classify and store knowledge, as well as standard procedures. Also included are the heuristics, which are personal rules of thumb used by experts, such as searching for important information in the first paragraphs of a journal article, or the analytic approach that involves searching for extreme cases. Heuristics are difficult to obtain from an expert because the expert is often not aware of the details of his own professional approach. The knowledge base may also include data about the logical structure of knowledge in a given field.

The following examples show the contents of the knowledge base in a number of systems.

The knowledge base of the DENDRAL system, which generates plausible structures for organic molecules on the basis of mass spectrometry, contains the following information: rules for deriving constraints on molecular structure on the basis of experimental data, a procedure for generating structures that meet the constraints, and rules for predicting the mass spectra of the structure generated.

Digital Equipment's XCON system for the configuration of the company's computers knows the hundreds of components that are used in the manufacture of the computer. It also knows the rules for determining when to undertake the next sub-task and the rules for properly carrying out each sub-task. The system includes approximately 1,200 rules; this number increases regularly.

MYCIN, the first expert system, diagnoses bacteriological infections and recommends therapy using appropriate antibiotics. Its knowledge base includes rules relating patient data to infection hypotheses, rules for combining certainty factors and treatment rules.

A final example is the AM system for discovering mathematical concepts. It contains elementary set theory concepts, heuristics for generating new concepts by combining elementary concepts and heuristic rules for determining the level of attractiveness of the ideas generated.

Figure XIII shows these examples in more detail.

As we have seen, expert systems are based on rules, though there are exceptions. The inference engine includes a rule interpreter which decides which rule to apply and then fires it.

Rule-based systems are by far the most common. Rules are normally of the 'situation-action' type (if...then). This is what is called 'production rules'. By combining them one to the other, we obtain a line of reasoning that takes the following form: if...and...and...or...then...if...then...

The following MYCIN rule is often quoted in its ordinary language form (as opposed to formal data processing languages):

> If (i) the infection is meningitis and (ii) the organisms were not detected by culture coloration and (iii) the type of infection could be viral and (iv) the patient is severely burned, then it is probable that Pseudomonas aeruginosa is one of the organisms responsible for the infection.
>
> (Duda and Shortliffe, 1983)

SYSTEM	APPROACH	KNOWLEDGE BASE	OVERALL DATA BASE
AM (Discovers mathematical concepts)	1. Begins with elementary set theory concepts 2. Searches for possible connections which may be generated on the basis of these concepts 3. Selects the most interesting of these and follows through on them	1. Elementary set theory concepts 2. Heuristics for generating new concepts by combining elementary concepts 3. Heuristics to determine a level of interest which will be used in rejecting bad ideas	1. Plausible concepts
DENDRAL (Generates plausible organic molecule structures on the basis of mass spectra)	1. Derives constraints based on data. 2. Generates possible structures 3. Makes mass spectra predictions 4. Compares predictions to the data	1. Rules for deriving constraints on molecular structures on the basis of experimental data 2. Procedure for generating structures which meet the constraints 3. Rules for predicting spectra on the basis of data	1. Mass spectrometry data 2. Constraints 3. Possible structures
MYCIN (Diagnoses bacterial infections and recommends antibiotic therapy)	1. Represents reasoning in the form of condition-conclusion rules and degrees of certainty 2. Begins with a diagnostic hypothesis and verifies whether it is supported by the evidence	1. Rules relating data to infection hypotheses 2. Rules for combining certainty factors 3. Treatment rules	1. Patient's medical history and tests 2. Current hypothesis 3. Status of system 4. Conclusions and rules used

Figure XIII. Example of contents of various knowledge bases.
Source: W.B. Gevarter. An Overview of Expert Systems. U.S. Dept. of Commerce, 1982.

SYSTEM	APPROACH	KNOWLEDGE BASE	OVERALL DATA BASE
MYCIN (continued)	3. Evaluates all hypotheses 4. Assigns treatment for all diagnoses that are strongly supported		
XCON (Configures computers on the basis of customer orders)	1. Breaks the problem down into sub-problems 2. Resolves each sub-problem, and moves on to the next	1. Properties of components 2. Rules for determining when to move on to next sub-problem 3. Rules for resolving each sub-problem	1. Customer orders 2. Current task 3. Status of system

Figure XIII. (cont'd)

Rules of this kind have a considerable attraction because the system becomes strongly modular. If we wish to add a rule, the program does not have to be completely rewritten; the latest rule can be added to the earlier rules. Human experts also use rules in their field of expertise.

Figure XIV is a simple example of a network of rules that makes it possible to determine the nature of an animal from a number of its characteristics. Thus, (i) if the animal has hair, (ii) if it has hooves, (iii) if it has long legs and (iv) if it has dark spots, then it is a giraffe.

The average complexity of current systems involves approximately 1,000 rules with a fairly high level of variance. DIPMETER, which measures the slope of geological formations on the basis of readings from four sensors, has only 90 rules; MYCIN has 500 rules; CATS-1 (for the diagnosis and repair of diesel engines) has 1,500 rules and XCON has more than 3,000.

A large number of rules is not necessarily synonymous with high performance. Although the number of rules generally increases with development, it may sometimes be best to keep the number down. DIPMETER, which had its knowledge base reduced from 150 to 90 rules, actually improved in performance.

Not all systems are rule-based. Some are based on **cause-effect relations** networks, such as CADUCEUS, PROSPECTOR, HARPY, MACSYMA, DIGITALIS. Figure XV describes a simple example of a procedural system. This approach may be used to develop complex recommendations.

Another type is based on **semantic networks**, which represent interrelations as in the example below taken from Barker's report (1983):

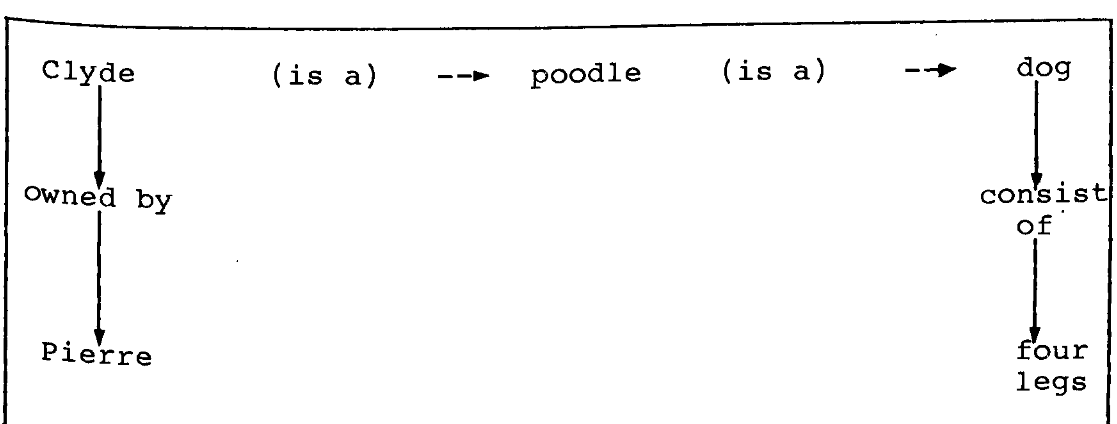

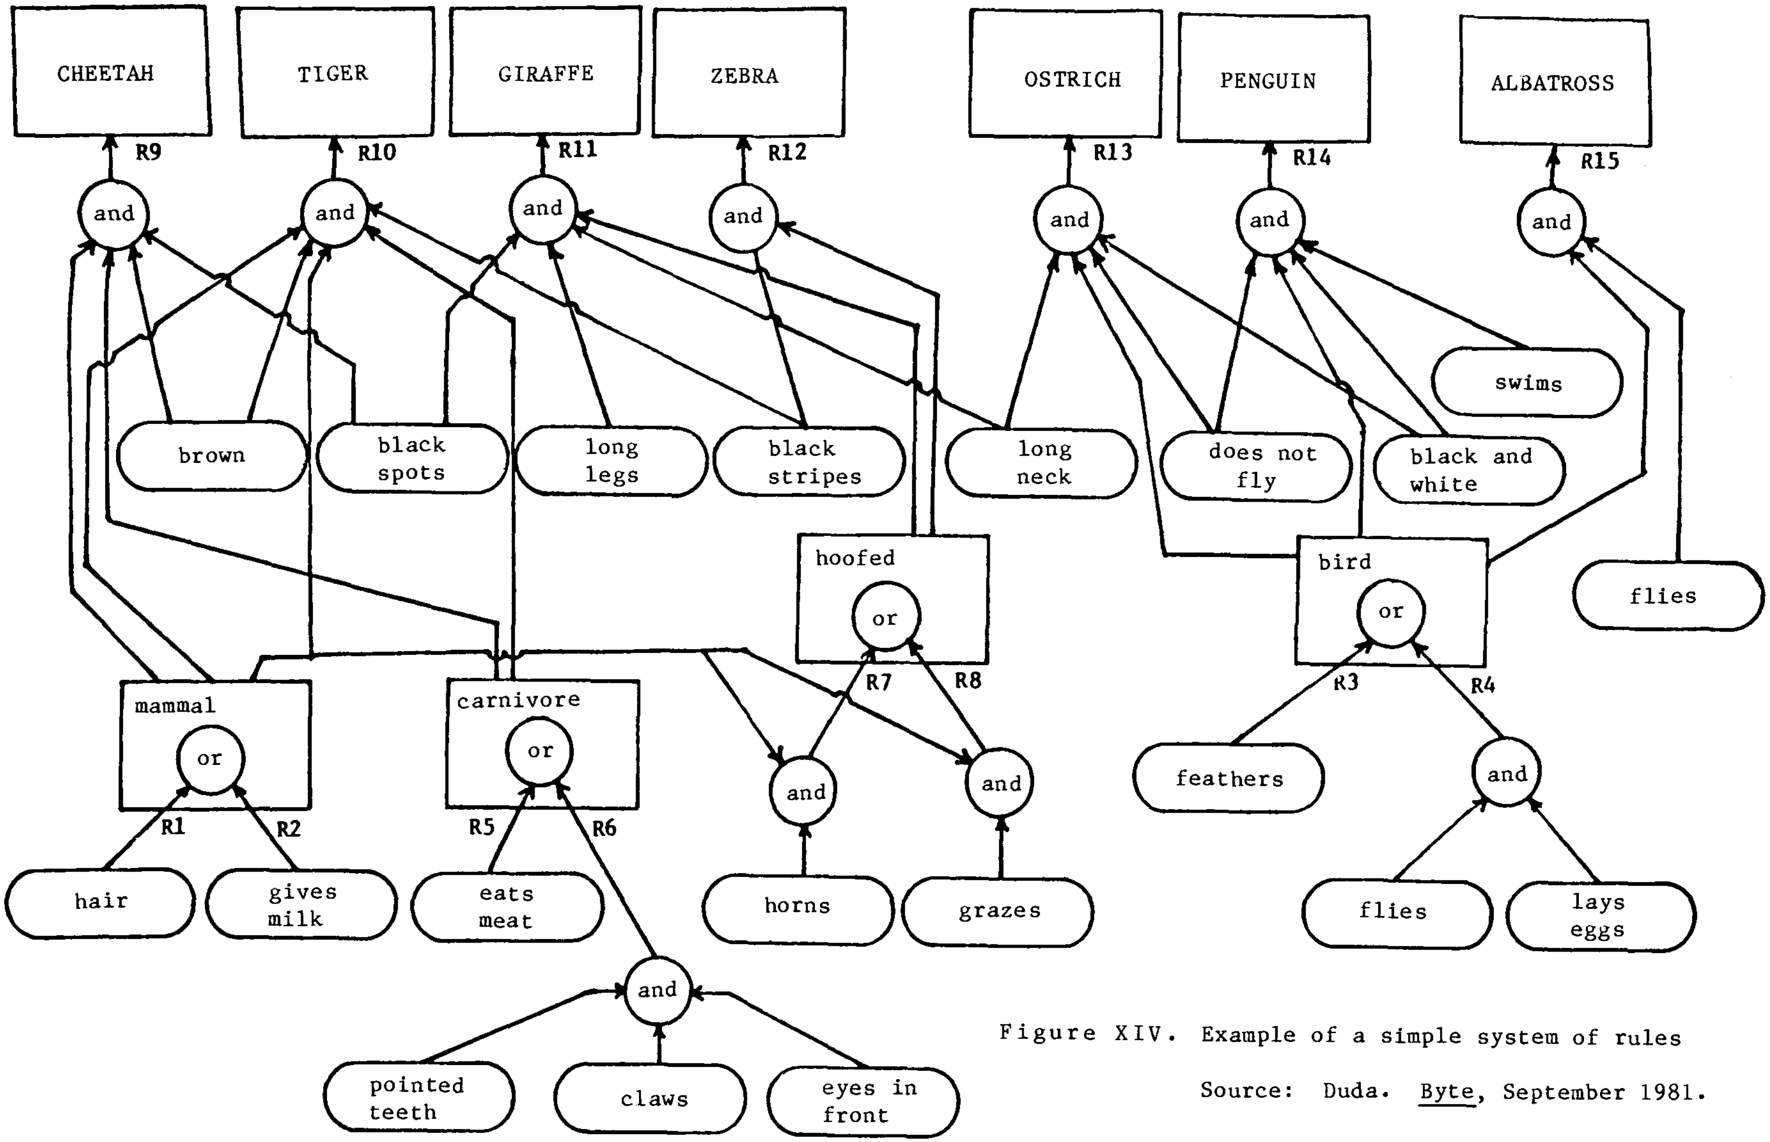

Figure XIV. Example of a simple system of rules

Source: Duda. Byte, September 1981.

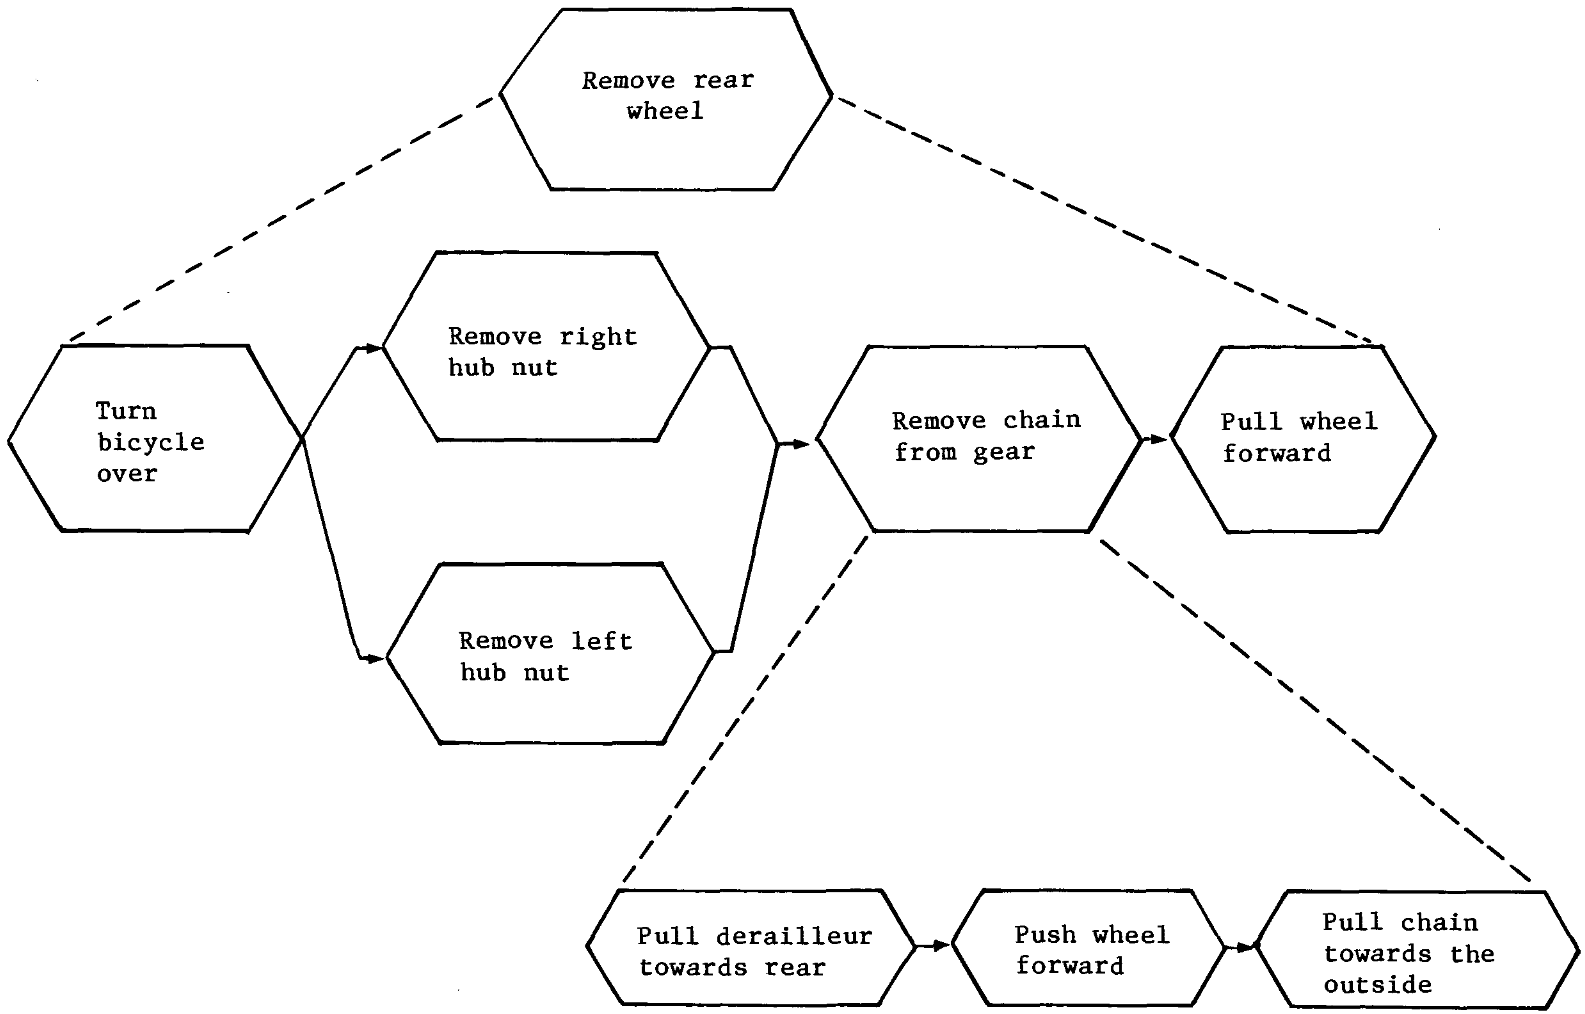

Figure XV. Example of a procedural network.
Source: A. Letendre. La prochaine génération de logiciels: l'intelligence artificielle. MOSST, 1983.

A fourth type of knowledge representation is the **frame model**, which represents objects in terms of standardized properties. For example, birds have a habitat, anatomical features, colours, etc. The example below is taken from Barker's report (1983):

```
FRAME : Person

Name:     John Jones
Age:      if necessary (current date less date of birth)
          range:           (0 - 100) years
Father:   if necessary (husband of mother)
Mother:   Mary Jones
Date of birth:  6 January 1950
Telephone:  758-3741
            if changed:  update
```

A final type is the model based on **mathematical logic.** For example, "All men are mortal, Socrates is a man, therefore Socrates is mortal." A formal example would be:

```
IN (John, Paris)
IN (x, Paris) =   not FAR (x, Louvre)
-------------------------------------
not FAR (John, Louvre)
```

(Barker, 1983)

The above example reads as follows: John is in Paris, someone in Paris is not far from the Louvre, therefore John is not far from the Louvre.

Only two systems, however, have been developed using this approach: EXPERT EASE and TIMM. It should become more popular, however, with the arrival of the language Prolog (PROgramming in LOGic).

For more information about these various approaches, consult the special October 1983 issue of IEEE Computer.

One final important point is that **the approach adopted must be intimately related to the nature of the problem to be solved.**

4.1.1.2 The Inference Engine and How It works

The inference engine is a sequence of logical operations for manipulating and combining knowledge base rules. It is an interpreter of rules whose operations involve comparing chains of characters, joining them, separating them and substituting one chain for another. Here are a few general examples.

The DENDRAL system begins by using available data to derive initial constraints and then generate a series of possible structures. On the basis of these structures the system predicts corresponding mass spectra and compares them with initial data.

Digital Equipment's XCON system attacks problems by subdividing them into smaller sub-problems: correct faults in customer order, place components in computer cabinet, cable, etc. The system completes one sub-task and then moves on to the next.

The MYCIN system takes a different approach. It takes each diagnosis hypothesis and evaluates whether it is supported by the evidence. It searches for a treatment for each highly plausible diagnosis.

From these few examples, we can see that the system must always search the knowledge base. This leads quickly to the problem of combinatory explosion. One way of getting around the problem involves breaking the problem down into units that are more readily handled. Another approach is to generate all possible solutions and to eliminate the branches of the decision tree that do not meet a number of predetermined criteria.

The system, therefore, searches for a rule that satisfies the condition(s) and executes it. When the system executes a rule, the status of both the system and the database changes: a number of new rules are activated and others are deactivated.

The rule interpreter (the control structure) uses a precise strategy to find active rules and to apply them, while at the same time ensuring that the search space is searched as effectively as possible.

We mentioned previously that the various rules are placed end-to-end to create a line of reasoning. The main approaches are forward chaining, backward chaining, a combination of both and a number of other similar approaches.

In forward chaining, the system reviews the rules until it finds one whose antecedent (the 'if' of the 'if...then' rule) matches the available information. The system then activates this rule and deactivates others, applies the rule, updates the database and begins the process again until the goal is achieved or until no other rules are available. This approach is also called the data-driven strategy, because the data are the starting point for the process, and the goal to be achieved is not known.

The strategy has been applied successfully in data analysis, design, diagnosis and concept formation problems.

The problem with this approach, however, is that the system will attempt to derive all possible branches of the search space.

A related approach is the event-driven strategy. Each step is based on new data or on the results of previous steps: the system reacts to events. The approach is applicable to situations that change over time, primarily in real-time operations such as monitoring and control. It may also be applied to planning.

Backward chaining begins with a known conclusion and attempts to find a rule that leads to the conclusion. The system knows the 'then' part of the 'if...then' rule. It reviews the rules and tries to find one whose antecedents (or conditions) lead to the known conclusion. Following an initial pass, if the problem has not been solved, the 'if' becomes the new goal and the system searches for a rule that leads to it. The system is applied recursively until the problem is solved. If it cannot solve the problem, the system will ask the user for additional information.

This approach is used when the goal is known or in those cases where the system can formulate a hypothesis which it will later attempt to verify. This, of course, is why the approach is also called a goal-driven strategy. The goal or the conclusion is known and the system attempts to find the chain of rules that leads to this goal. A motor that will not start is one example. We know the effect and we are looking for the cause.

This strategy may be used in diagnostics problems (e.g., MYCIN) or for planning.

The danger of the approach is that if there are too many sub-goals to be met it will rapidly lead to combinatory explosion.

A combination of the above strategies may also be used. A data-driven strategy may be used to develop a number of hypotheses based on the data and then each hypothesis may be used as if it were a goal. When new facts are discovered in the backward-chaining research, the system checks whether new goals are suggested.

Another example would be the convergent relaxation technique, which uses both approaches and combines the solutions at some intermediary position. It is used in situations such as speech understanding, in which there is much uncertainty.

Although these rules are predetermined, it is not necessarily possible to predetermine the use to which the system will put them.

For example, the EURISKO system (which discovers new concepts) has been used to play the Traveller computer game. To the great surprise of everyone, the system fired on some of its own space ships to strengthen its fleet. Because the speed of the fleet is determined by the speed of the slowest ships, the system chose to keep only a fleet of small, rapid ships. Since no one had thought of this strategy, EURISKO won the tournament.

A new approach uses meta-knowledge about both the contents of the knowledge base and the various strategies that make it possible to choose a suitable approach. It is like an expert system within an expert system. The system, therefore, could examine only a few groups of rules depending on the context or change the order of the rules to examine them more efficiently.

The following example of meta-rule is taken from the EURISKO system:

> If all members of a set unexpectedly satisfy some rare property, then increase the interestingness rating of that set and the heuristics that led to its definition.
>
> (Lenat, 1984)

Another approach would be to treat the other rules as data. This could produce the following example of a meta-rule: "If the level of confidence in a rule falls below 20 per cent, then it must be discarded."

Figure XVI describes which strategies were used for a few systems. To date, it would appear that the forward chaining strategy is the most popular.

4.1.1.3 The Man-Machine Interface

The majority of these systems communicate with the user in a natural language (as opposed to a formal language like Fortran or Basic). This is clearly possible only in a very restricted domain and using a conversation style which is very structured.

Most systems communicate strictly by searching for certain key words or by developing sentences using predetermined patterns. If there is any ambiguity, the system will ask for explanations. For example, the system could ask for an estimate of certainty.

The most interesting feature of expert systems, however, is their capacity to explain how they reason. The approach is disconcertingly simple. The system prints out the series of rules it used in its line of reasoning. The explanation may or may not be in natural language. The approach is not satisfactory, however, when a large number of rules have been applied. Research is under way to produce summary explanations.

This is extremely important because even experts would not dare use the system without understanding the reason behind each recommendation or conclusion. Moreover, this facilitates the search for defects in the system. It also helps with later modifications.

An example of a dialogue between the system and a user is given on the page after Figure XVI.

4.1.2 Computers

Although expert systems can operate on conventional minicomputers and some simple systems can function on personal computers, special computers have been developed for artificial intelligence applications.

Such machines have been optimized for working with the Lisp language. Some are currently being developed to work with Prolog.

System	Function	Field	Strategy			
			Forward chaining	Backward chaining	Combined	Event-driven
Mycin	Diagnosis	Medicine		X		
Dendral	Interpretation of data	Chemistry	X			
EL	Analysis	Electronic circuits	X			
Guidon	Training	Medicine				X
KAS	Knowledge acquisition	Geology	X			
Meta-Dendral	Learning	Chemistry	X			
AM	Concept formation	Mathematics	X			
VM	Monitoring	Medicine				X
GAI	Interpretation of data	Chemistry	X			
EXCON	Design	Computers	X			
ABSTRIPS	Planning	Robotics		X		
NOAH	Planning	Robotics		X		
Molgen	Design	Genetics			X	
SYN	Design	Electronic circuits	X			
HEARSAY-II	Interpretation of signals	Speech recognition			X	
Harpy	Interpretation of signals	Speech recognition	X			

Figure XVI. Strategies used by the inference engine of selected expert systems.

Source : Miller, K. The 1984 Inventory of Expert Systems.

Figure XVII.

```
Please enter information about the patient.

        Name       Age        Sex       Race
1) **  Pt219      42 YEARS   FEMALE    CAUCASIAN

2) Are there any positive cultures obtained from Pt219?
** NO
3) Are there any pending cultures of Pt219?
** YES

        site       date       stain       method of
                   collected  examined?   collection
                                          [urine,sputum only]
4) **  CSF        3-Jun-75   YES         N/A

        stain of smear      morph
        [type NONE if no organism was seen on the smear]
5) **  NONE
                                    .
                                    .
                          {more questions}
                                    .
                                    .

13) Has Pt219 recently had symptoms of persistent headache or other
    abnormal neurologic symptoms (dizziness, lethargy, etc.)?
** YES
14) Has Pt219 recently had objective evidence of abnormal neurologic
    signs (nuchal rigidity, coma, seizures, etc.) documented by
    physician observation or examination?
** YES


The CSF cultures will be considered to be associated with MENINGITIS.




34) Is the meningitis a hospital-acquired infection?
** NO
35) Is Pt219's clinical history consistent with EPIGLOTTITIS?
** NO
36) Is Pt219's clinical history consistent with OTITIS-MEDIA?
** NO
37) Has the patient ever been to an area that is endemic for
    coccidiomycoses?
** UNKNOWN
38) Is Pt219 a burn patient?
** WHY

[i.e. WHY is it important to determine whether PATIENT-219 has been
seriously burned?]

[4.0] This will aid in determining the organisms (other than those
      seen on cultures or smears) which might be causing the infection.
  It has already been established that
     [4.1] the infection which requires therapy is meningitis, and
     [4.2] organisms were not seen on the stain of the pending csf
           culture (CULTURE-1), and
     [4.3] the type of the infection may be bacterial
  Therefore, if
     [4.4] Pt-219 has been seriously burned
  then
     there is suggestive evidence that Pseudomonas-aeruginosa is one
     of the organisms (other than those seen on cultures or smears)
     which might be causing the infection
  Author: Yu
  [RULE578]

   [back to question 38...]
   ** HOW 4.3
```

Figure XVII. (cont'd)

```
[i.e. HOW was it established that the type of the infection may be
bacterial?]

The following rules concluded about: the likelihood that the type
of the infection may be bacterial

               Rule          Yes       No

   [4.5]      RULE500       (.4)
   [4.6]      RULE501                 (.17)
   [4.7]      RULE502       (.1)
   [4.8]      RULE524       (.42)
   [4.9]      RULE504       (.28)

Thus, it has been established that the type of the infection may be
bacterial.
```

My therapy recommendations will be designed to treat for organisms that either are very likely or, although less likely, would have a significant effect on therapy selection if they were present. It is important to cover for the following probable infection(s) and associated organism(s):

```
INFECTION-1 is MENINGITIS

  + <ITEM-1>  MYCOBACTERIUM-TB [from clinical evidence only]
  +    ...    COCCIOIOIOES     [from clinical evidence only]
  +    ...    CRYPTOCOCCUS     [from clinical evidence only]

   43) Ooes Pt219 have a clinically significant allergic reaction to any
       antimicrobial agent?
   ** NO
              [Considering ETHAMBUTAL-ANO-INH-ANO-RIFAMPIN
               for use against INFECTION-1...]
   44) Is Pt219 pregnant or breast feeding?
   ** NO
   45) Oo you have reason to suspect that Pt219 may have impaired renal
       function?
   ** NO
   46) Pt219's weight in kilograms (or <number> POUNOS):
   ** 51.3

[REC-1] My preferred therapy recommendation is as follows:
        Give the following in combination:
        1) ETHAMBUTAL
        Oose: 1.28g (13.0 100mg-tablets) q24h PO for 60 days
              [calculated on basis of 25 mg/kg]
              then 770 mg (7.5 100mg-tablets) q24h PO
              [calculated on basis of 15 mg/kg]
           [The dose should be modified in renal failure.
           I can adjust the dose only if a creatinine clearance
           or a serum creatinine is obtained.]
        Comments: Periodic vision screening tests are recommended
                  for optic neuritis.

        2) INH
        Oose: 513 mg (5.0 100mg-tablets) q24h PO
              [calculated on basis of 10 mg/kg]

        3) RIFAMPIN
        Oose: 600 mg PO q24h
        Comments: Administer dose on an empty stomach
```

Figure XVII. (cont'd)

Although I suspect that fungus may be a possible pathogen, in the absence of definitive evidence, I will not recommend antifungal therapy at this time. Please obtain samples for fungal, TB, and viral cultures, cytology, VDRL (blood and CSF), coccidioides complement-fixation (blood and CSF), cryptococcal antigen (blood and CSF), viral titers (blood and CSF). An infectious disease consult may be advisable.

Example of session with the Mycin system.

The user answers the questions asked by Mycin on the lines beginning with two asterisks (**). On line 38, Mycin explains its reasoning.

(Reproduction of an article by Duda and Shortliffe, "Expert Systems Research." Science, 15 April 1983, pp. 261 - 268.)

Such computers are expensive. They average between U.S. $30,000 and $70,000, although prices are beginning to drop rapidly. They all have a large core memory, powerful disk drives, high resolution graphics and often have utilities such as text and graphics editors. They all use a mouse connected to a windowing system to facilitate browsing through the system. It is possible, therefore, to examine the contents of a network node and its relationship with other nodes.

Figure XVIII is a compilation of Lisp computers available on the market. Machines produced by Apollo Computer, Data General, Digital Equipment, Hewlett-Packard and Sperry are soon to be added to this list.

4.2 Assessment of the Technology

4.2.1 The Ideal System

Before assessing existing expert systems and those in the development stage, we should begin by establishing a comparison standard consisting of a description of what the ideal expert should be. Human experts ideally:

- solve problems in their field;
- explain the reasons for their choices;
- learn from experience;
- restructure their knowledge on the basis of experience;
- break established rules where required;
- determine whether or not information is relevant.

The performance of human experts also gradually declines as their limits are approached.

Manufacturer	Model	Price (U.S. dollars)
Lisp Machine Inc. (Culver City, Calif.)	LISP Machine Lambda Machine Picon Machine	$ 80,000 N/A $150,000
Bolt, Beranek & Newman Cambridge, Mass.	Jerico	N/A
Symbolics Inc. Woodland Hills, Calif.	LM-2 3600	$ 75,000 $ 60,000
Three Rivers Computers Inc. Pittsburg, Penn.	Perq	$ 25,000
Texas Instruments Austin, Texas	Explorer	$ 52,000
Xerox Electro-optical Systems, El Segundo, Calif.	1100 1108 1132	$ 45,000 $ 27,000 $180,000

Figure XVIII. Lisp computers available on the market.

4.2.2 Advantages and Uses of Existing Systems

The advantages of expert systems are clear: they work without getting tired, they are immortal, they can be improved and reproduced where required, and they are easy to improve and maintain when compared with conventional computer programs.

Other advantages are:

- permanent availability of human expertise;
- greater accessibility to this expertise;
- the possibility for the human expert to readily get a second opinion;
- the exploration of alternatives through the use of 'what if' questions in a given field;
- the possibility of solving problems requiring a great deal of knowledge in situations where data are incomplete and where no human experts are available;
- **the opportunity to solve problems that cannot be solved as well by humans or by conventional data processing techniques, such as cases where a large number of solutions must be explored.**

Once again, the XCON system is an example of the usefulness of expert systems.

At Digital Equipment Corporation, configuring a computer may involve a thousand different steps and require from 25 to 125 items of information about each of the 7,000 components that may be used in manufacturing the computer. This is an enormous quantity of information to master. Most of all, no human being could evaluate all the possible interrelations. An option, for example, may indirectly involve the addition of a new bus and a new back plane. To date, XCON has filled approximately 150,000 orders.

4.2.3 Performance

The performance of expert systems is often surprising - if we restrict ourselves to a number of well-chosen examples. In general, however, few systems surpass the performance of human experts.

The DIPMETER system (for analysing oil exploration data) gives a correct interpretation in its field of expertise 90 per cent of the time. The XCON system, is accurate 95 to 98 per cent of the time, probably because of all the experience that has been acquired using this system.

The PROSPECTOR system has used geological, geophysical and geochemical data obtained from preliminary exploration to identify and successfully locate a molybdenum deposit in the Mount Tolman region in the state of Washington. The system did not, however, identify the full extent of the deposit.

The CADUCEUS system has been used informally in Pittsburg. The system rapidly identified the disease of a man, whose symptoms had puzzled doctors, as systemic amyloidosis. This is a rare and sometimes fatal disease in which protein concentrates in the vital organs and destroys them. A liver biopsy later confirmed the diagnosis.

4.2.4 Limitations

Despite the positive things said thus far, expert systems have serious limitations, mainly because they are still in the early stages of development. For example, we do not fully understand why some designs work well.

4.2.4.1 Knowledge Representation Limitations

Current applications of expert systems are limited by available representation methods, which means that applications are often limited to fields where 'if...then' rules may be applied.

Some types of knowledge are difficult to represent on a computer, for example, processes, beliefs, functions, shapes and dynamic systems. Nevertheless, elegant ways have been found to represent certainty and time.

One of the first representational problems attacked was the representation of uncertainty. Early work focused on the use of probabilities and certainty values (indicating a degree of confidence in facts or conclusions) established by the expert building the system. Each MYCIN rule, for example, is accompanied by a degree of certainty. A more sophisticated and more experimental method is fuzzy logic, which is a continuum of values between true and false.

The concept of time has not been studied as much. One approach is the situational calculus, invented by McCarthy and Hayes, which uses sequences of actions. This approach makes use of the concept of 'situations' which change only when enough actions have taken place or when data indicate a new situation. Each situation determines what changes and what remains fixed in a frame describing a stereotyped situation. The VM system uses this approach to monitor the breathing of patients.

Rules are too restrictive as a form of knowledge representation because they eliminate subtleties and hide causal relationships that are only implicit in this form of representation. The other approaches are too restrictive and inflexible as well.

The future will **perhaps** belong to the combined use of rules, to deal with routine problems; and logic, to deal with more complex problems.

Current systems present a problem because they are applicable only to limited fields. PUFF, for example, is restricted to pulmonary functions. The CADUCEUS system would be more useful if it were provided with anatomical knowledge or if it were provided with information about the progress over time of various diseases.

To ensure that the final product is consistent, only one expert is normally consulted in developing a system. Rules often have to cooperate with one another; for example, one rule may become active only after another has been applied. It is, therefore, difficult to build multidisciplinary systems. One approach could be to use sub-systems, each of which would work on one aspect of the problem, and combine the results at the end. The sub-systems would exchange messages in a common work space called a blackboard.

One of the main limitations of existing systems is that they have only surface knowledge, i.e., heuristic knowledge, based on experience and on the recognition of key points. They do not have deep knowledge, i.e., knowledge of basic principles, structures, the functions and behaviour of objects, and causal relations - they do not have any knowledge about the world. None of the expert systems used in medicine, for example, knows what a human being is. The PUFF system, which specializes in pulmonary diseases, would be of no use to a surgeon because it doesn't know that the lungs are located inside the human body.

Surface knowledge is like the knowledge of a mechanic who can adjust a carburetor without knowing why what he does makes it work better. Deep knowledge is like the knowledge of an engineer who understands all aspects of how a carburetor works and who can predict the outcome of his manipulations and explain why things occur the way they do.

A consequence of this surface knowledge is that expert systems do not know when to break the rules and have no way of knowing what is relevant and what is not. They do not know the spirit of the rules, nor do they know what can and cannot be done. Expert systems have no idea of their own limitations.

This could have serious consequences. The MYCIN system will never prescribe tetracycline to a child under eight years old because it stains the teeth of young children. There are, nevertheless, occasions when aesthetic considerations should be set aside, e.g., to save a life.

The following are other knowledge representation problems:

- Some problems require more than one type of representation and cannot be manipulated readily by a single strategy or a single inference engine. Examples include problems related to syntax, phonetics and spectral frequency in speech recognition.

- Others require a number of rules which is impossible to manage. A system that provides advice on repairing breakdowns, for example, must have information about all possible types of breakdowns instead of simply being provided with a set of rules for the detection and correction of problems.

4.2.4.2 Knowledge Acquisition Limitations

The acquisition of knowledge is the collection and coding of information for the original system and later updates. This involves the addition, elimination or modification of rules. It requires evaluating a piece of information and its interaction with other rules, as well as selecting an initial form of representation and an initial structure for describing the chosen application.

Extracting knowledge from experts is often a laborious process because they are usually not aware of the approach they take to solve the problem. It is also usually difficult for experts to express themselves in a way that is compatible with the formalism of knowledge representation methods.

This stage in building an expert system is usually carried out by a specialist called a 'knowledge engineer'. As one would expect, there must be close contact between the expert and the knowledge engineer. One of the problems will be to ensure continuity in their joint work, because building an expert system may take from one to two years.

Tools may be used to facilitate the process. The simplest are editors, used to enter and modify rules. In a later section on the building of expert systems, we will examine a number of commercially available tools such as system shells that often simply require filling in the blank spaces.

Others are intelligent interfaces for interviewing the expert, or expert systems, specialized in the field of 'knowledge engineering'.

The TEIRESIAS system, for example, assists an expert in the construction or modification of a system based on MYCIN-like architecture. The system uses a model of the MYCIN knowledge base to determine whether new information can be incorporated with previously entered data. The system uses the model to make suggestions to the expert being interviewed.

IBM is currently developing a system shell which will support frames and inference rules. It will also include mechanisms for interrogating the user and to explain the reasoning process used by the system. Moreover, the system will include a mechanism for extracting data from other systems such as sensors or databases. This should considerably reduce the number of questions the user will have to answer. For example, a system specialized in the diagnosis of computer breakdowns could obtain data directly from the computer being examined.

Because data acquisition is a bottleneck in the building of expert systems, a mechanism must be developed in the long run to allow such systems to learn by themselves. For example, the system could deduce rules on the basis of examples or directly read books on the subject.

Few systems, however, even experimental ones, have the capacity to learn by themselves. The META-DENDRAL system is an expert system that models the process of developing a theory for generating general rules of fragmentation in mass spectrometry. These rules take the same form as those used by DENDRAL. The method used by META-DENDRAL involves generating, verifying and fine-tuning a set of rules obtained from spectra. The rules obtained in this way are usually excellent ones.

Another approach is the one used by EURISKO, which derives new heuristics associated with new concepts as these concepts are discovered. The system 'reasons' by analogy. In the example of designing a space fleet mentioned earlier, EURISKO designed a symmetrical fleet because this was a concept it had learned while designing integrated circuits.

Two other examples are INDUCE/PLANT and LEX. The first infers rules for diagnosing plant diseases on the basis of symptom data; the second learns new heuristics for selecting rules to apply as it solves problems in integral calculus.

The validation of encoded information is another problem, because there is no methodology to ensure a priori that coded data will be intelligible and usable.

4.2.4.3 Inference Engine Limitations

New methods must be developed to manage large databases and allow them to be searched in a reasonable amount of time. Performance tends to fall off rapidly as the database grows in size. If the system contains more than several thousand rules, it becomes imperative to use meta-rules to reduce the effective work space. The emergence of new parallel architectures for computers will help eliminate this problem. Another approach will make use of new database management methodologies.

Inference engine concepts are still limited and without much variety. We do not yet have systems that operate by analogy or extrapolation, or that can modify instructions or perform any other behaviour that a human expert would use to cope with a given problem.

Inference engines have predetermined and, therefore, inflexible mechanisms. They could be more efficient were they able to choose from among several strategies, depending upon the problem to be solved.

Expert systems have serious difficulties with unfamiliar situations or situations not specifically anticipated. A system cannot, for example, deal with two diseases present at the same time even though it can solve the problem for each disease individually. Similarly, a system for providing advice on automobile repairs would very likely be useless in diagnosing that a potato had been stuffed into the tail pipe. We are, therefore, far from having created systems that can adapt to new situations.

The performance of expert systems is usually fragile and unpredictable when the systems are pushed to the limits of their knowledge. The performance does not fall off gradually, as we might desire, because the systems cannot know when they are dealing with facts outside their area of expertise. A system will never ask for help when it is faced with a new situation.

Finally, existing systems have no way of verifying and evaluating the plausibility of their own conclusions. This sometimes leads to truly naive recommendations.

4.2.4.4 User Interface Limitations

Communications with the user are often inflexible and structured; explanations are limited to a list of the rules which have been applied. It would be desirable to tailor explanations to the user, making the systems much more convincing. This is currently impossible, however, because the systems have no model of the user and its needs.

Interfaces with the user are often so poor that it is necessary to call on an expert to use them. Expert systems can sometimes be very difficult to use without special training. For example, when the XCON system was transferred from the laboratory to the plant, the success rate dropped from 90 per cent to 60 per cent because users were unfamiliar with the system and used incorrect data. These problems were different from those identified previously. It is true, however, that all software requires a training period if it is to be used correctly.

4.2.4.5 Maintenance Problems

Expert systems now in use cannot automatically restructure their knowledge base, which makes rapid updates difficult. (There are, however, experimental systems that can do this.)

System maintenance is usually carried out by an applications expert or a 'knowledge engineer'.

4.2.4.6 Other

Other limitations are:

- the nature of knowledge causes difficulties because it is sometimes incomplete or inaccurate, poorly specified and almost always acquired gradually;
- knowledge is still limited by the memory and sequential nature of current computer architecture;
- building an expert system may be a laborious process;
- the accuracy of results must often be analysed by an expert.

4.2.4.7 Conclusion

The intelligence of existing expert systems is well below that of human beings. Expert systems might be compared to idiot savants who may indeed be true geniuses in restricted fields, yet completely deficient in every other intellectual respect. Expert systems only imitate expertise and have no common sense, for this would require too much basic knowledge.

4.2.5 R & D Required

Buchanan (1981), who has been cited by a number of writers, identified the following R & D needs:

- improved knowledge acquisition systems, as well as methods for reducing knowledge acquisition costs and selecting appropriate problems;
- learning by example;
- better explanation sub-systems and interfaces that are more natural and easier to use;
- improved development tools;
- better architectures and improved inference procedures;
- effective techniques for consulting several experts and for acquiring knowledge from various sources;
- better approaches for modelling time;
- methods for developing better hypotheses and for obtaining an improved knowledge of the world (e.g., knowing that a chicken in an oven cannot fly);
- the ability to make use of physical and biological causal models and to use them in connection with other knowledge;
- general planning methods;
- methods for searching very large databases;
- forms of reasoning by analogy;
- techniques for using formal deduction methods;

- approaches for parallel processing;
- new knowledge representation methods such as multiple models (e.g., electronics black box and circuits in electronics).

Many other fields have been mentioned in this report. In addition, other authors have identified more specific research fields:

- interactive aids for validating, examining and facilitating the identification of problems in very large databases;
- the automatic discovery of new heuristics;
- methods for inferring new inference rules on the basis of past experience;
- the representation of naive physics (e.g., an unsupported object will fall);
- new forms of reasoning (e.g., by default, by analogy or by synthesis).

4.2.6 The Future

4.2.6.1 Technology

Systems of the future will have more than 10,000 rules by the 1990s. By the end of the following decade they will have a billion. Because such a large number of rules will require fast machines, they will function on computers operating in a parallel architecture.

Automated methods will have been found to translate various documents directly into the knowledge base. An electronic circuit, for example, could be transferred automatically to a system specialized in diagnosing malfunctions of electronics circuits.

We will begin to see the first non-rule-based systems before 1990. In the early 1990s we will begin to see systems with sophisticated natural language interfaces and the technology will begin to proliferate. By the end of the 1990s we should be able to make use of systems capable of self-learning.

Figure XIX, taken from the special issue of IEEE Spectrum on fifth generation computers, gives a good summary of the evolution of this technology.

4.2.6.2 Applications

Guesswork is used at this stage to determine future applications of such expert systems. Medicine has been one of the most popular applications field for expert systems up to now. However, very few systems are routinely used in our hospitals. While such systems are not reassuring to the patient when used by a non-specialist, they become an exceptional tool in the hands of a specialist.

Applications for expert systems are virtually unlimited. They could facilitate decision-making in air traffic control, help in the synthesis of new medicines, coordinate plant operations, reduce downtime in the distribution of tasks, monitor equipment, optimize production plans, or adjust computer-aided instruction systems to the level of each individual student.

In the home, expert systems could provide advice on furnace repair, nutrition, gardening or even tax returns. It will also be possible to build a system that reads newspapers and summarizes those events likely to interest each reader.

Field	Current Systems	Next Generation
Knowledge base	10^3 to 10^4 rules	10^6 to 10^8 rules
Knowledge acquisition	1 rule per hour	10 rules or more per hour Automatic acquisition
Interface	Structured language Limited graphics	Natural language Speech understanding Artificial vision

Figure XIX. Growth of technology.

Source: IEEE Spectrum, November 1983

Gevarter (1983) and Miller (1984) identified the following opportunities for expert systems in various fields:

- civil engineering
  design, planning, scheduling and control
- equipment
  design, monitoring, control, diagnostics, maintenance, repair and instruction
- command and controls
  analysis of intelligence, planning, targeting, threats assessment and communications
- defence
  target identification, automatic control, electronic warfare
- occupations
  consulting, medical instruction and diagnosis, analysis in law, accounting, management, finance and engineering
- education
  instruction, evaluation, diagnostics, concept formation and experience-based knowledge development
- image processing
  interpretation of photographs, cartography
- software
  instruction, specifications, design, production, verification and maintenance
- consumers
  intelligent games, finance and investment, purchasing, shopping, query of databases
- intelligent assistant
  assist users of data processing systems
- office automation
  intelligent systems
- control processes
  automation of plants
- exploration
  space, prospecting.

4.2.6.3 Social Impact

We are on the threshold of a new Gutenberg era. Expert systems have the potential to disseminate expertise wherever it is required. Rural communities, for example, would benefit from urban medical expertise if rural doctors had access to the various expert systems specializing in the major fields of medicine. In the future, a physician might be sued for malpractice for not having consulted an expert system.

Such systems will have to be accessible by means of a simple telephone call to an expert system data bank. Already the SUMEX-AIM network (Stanford University Medical Experimental computer for Artificial Intelligence in Medicine) allows outside terminals to have access to the computers of Stanford University and Rutgers University (N.J.).

Thus, society's knowledge will be more readily accessible to everyone. Before retiring, experts will want to immortalize their knowledge in a reproducible expert system rather than in a book.

Expert systems might, however, create a dependency or discourage initiative if overused. The technology could inhibit technological or scientific progress by providing sterotyped solutions to problems that have previously been perceived in a different perspective.

5.0 DEVELOPING AN EXPERT SYSTEM

5.1 The Process

The process for building an expert system is as follows: a problem is selected and defined, a nucleus of sub-tasks is constructed, the required knowledge is coded, the system is evaluated and attempts are made to gradually improve the systems. These steps are carried out with the help of an expert systems specialist.

The specialist (generally called a knowledge engineer) interprets, integrates answers to his questions, draws analogies, searches for exceptions and raises a number of design problems.

An expert must be found and persuaded to cooperate in the venture. Meanwhile, the knowledge engineer must read available literature to have an overview of the field and become familiar with the technical jargon. Only then can the problem be delimited. The key point here is to properly understand the structure of the problem to be able to choose a form of representation, a knowledge base structure and an appropriate control structure.

Knowledge, as was defined earlier, consists of facts, heuristics, rules of common sense, models and any other sort of knowledge that may help to solve problems. This includes the knowledge of the system builder and the user, as well as any general problem-solving techniques.

The expert is then asked to describe how he solves the problem. This can be laborious because the expert is often not aware of how he arrives at his conclusions. The knowledge engineer's role is to help the expert discover that information. Here the specialist asks the expert if specific facts are indeed correct or if certain approaches can be used. He must also verify whether a procedure is generally accepted or if it is a peculiarity of the expert. The process of extracting expertise from a human expert may indeed be difficult and time-consuming.

The next stage involves coding and programming the system. In technical terms the set of rules must be written. This information becomes the 'database' for problem-solving; it is stored in such a way that it can be accessed by means of natural language queries.

The system is run and errors are corrected by repeating the whole process. The expert will sometimes say that this is not how he would go about solving a given problem and attempts are then made to refine the rules. Also, the problem may be redefined.

There are many ways of checking the system's operation. The most common is to take a case cited in a specialized journal and to compare the system's solution to the solution given in the article. With MYCIN, the evaluators could not tell whether the diagnosis came from a consultant, an internist, a medical student or a computer.

Once the prototype is built, its knowledge base is gradually and continually increased by adding rules for handling exceptions. When it came out of the laboratory, XCON had only 200 rules. It now has more than 3,000, partly because of the extension of the system to other Digital Equipment products.

5.2 Tools to Facilitate System-Building

Luckily, many commercial tools are available to help in building expert systems. The best known are **expert system shells** (or frameworks). In simple terms, these are complete expert systems whose knowledge base is empty. All one needs to do is to add the necessary information for specific applications.

Such products generally include one or more techniques for representing knowledge and may apply different reasoning processes depending on the fields of application. They often include a variety of interfaces to facilitate interaction with the expert, including editors and mechanisms to translate the knowledge expressed in English into a form that is understandable to the machine.

These shells sometimes make it possible to build a very simple initial prototype in only a few days. Figure XX gives a few examples of shells. Not all are available commercially and this is an incomplete compilation.

These tools are not always easy to understand or to use. Some were not intended to be used commercially and, therefore, not supported by the developer. As well, they are also often based on out-of-date technology.

SYSTEM	INSTITUTION	COST	COMMENTS
a) shells			
ART (Automated Reasoning Tool)	Intference Corp.		Inference procedures
Emycin	Stanford University		Version of MYCIN without a knowledge base
ES/P Advisor	Expert System International	$1,300	
Expert	Rutgers University		
EXPERT-EASE	Expert Software International (Scotland)	$695	For IBM-PC
GEN-X	General Electric		
KES (Knowledge Engineering System)	Software Architecture and Engineering	($16,000 for VAX) ($4,000 for PC)	May be used on VAX or IBM-PC
LES (Lockheed Expert System)	Lockheed	N.A.	Internal use only
MI	Teknowledge	$12,500	Runs on IBM-PC
Personal Consultant	Texas Instruments	$30,000	This shell supports up to 400 rules
Rule Master	Radian		Runs on VAX or IBM-PC
SI	Teknowledge (Palo Alto)	$50,000	Runs on VAX
SAGE	SPL (United Kingdom)		
SAVOIR	ICI (U.K.)		
______	Epitek (Sweden)		
______	Infologies (Sweden)		
TIMM	General Research (Santa Barbara)		

Figure XX. Tools for building an expert system

SYSTEM	INSTITUTION	COST	COMMENTS
b) Languages			
FRL	MIT		Knowledge representation language
HEARSAY III	U. of Southern California		General system based on HEARSAY II. Includes a mechanism for determining context as well as a blackboard
KEE (Knowledge Engineering Environment)	Intelli Corp. (Palo Alto)	$60,000	A language that supports frames
KL-ONE	Bolt, Besanek and Newman		Knowledge representation language
KRL	Stanford University		Knowledge representation language
LOOPS	Xerox	Can be copied for a nominal fee.	Object-oriented language supporting rules
MRS	Stanford Univeristy	N.A.	Knowledge representation language
OPS5	Carnegie-Mellon University Digital Equipment	$5,000	LISP-based language to facilitate use of rules
UNITS	Stanford University		Knowledge representation language and interactive knowledge acquisition system. Supports frames and rules
ROSIE	Rand Corp.		Rule-based language. Translate from English to INTERLISP

Figure XX. (Cont'd)

SYSTEM	INSTITUTION	COST	COMMENTS
SNEPS	SUNY (Buffalo)		Knowledge representation language
c) Other			
AGE (Attempt To Generalize)	Stanford University		Expert system to guide the user in building a system
KAS	Stanford Research Institute		Supervises interaction with expert to build or improve network-based system such as PROSPECTOR
Krypton	Fairchild Corp.	N.A.	Combines a shell for describing objects with a declarative knowledge representation formalism
RLL			Expert system to help in building another expert system
TEIRESIAS	Stanford University		Expert system to facilitate transfer of human knowledge to the system

Figure XX. (Cont'd)

On the other hand, they may make it possible to save a lot of time in building an expert system. Kinnucan (1984) reported that it took only four months to develop a prototype for diagnosing breakdowns in telecommunications networks.

The price of these tools varies considerably, depending on what is needed. The SAGE skeleton made by Great Britain's SPL costs CDN $4,000 for a version that runs on the Apple and $20,000 for a mainframe version. Texas Instruments has just launched a system that can support up to 400 rules for U.S. $3,000. A Scottish company is selling a small system for the IBM PC for U.S. $695 and Level 5 Research (Florida) just announced the INSIGHT shell which supports up to 400 rules and costs only U.S. $95.

These prices are dropping rapidly. The Scottish system that is now sold for U.S. $695 was initially sold for $2,000. The price dropped in only a few weeks.

The trend will probably be towards the marketing of skeletons specialized in well-defined areas such as data surveillance, repair, diagnostics and any other fields for which common approaches may be found.

The buyer must be careful because a high price does not necessarily mean a high quality. Some systems are expensive and weak while others are cheap and powerful. Prices sometimes include training and a right to make many copies. Other times, it includes only the system itself.

In addition to these shells, a number of accessories have been developed, including knowledge acquisition systems. There are expert systems available to evaluate whether a given problem can be dealt with by the use of expert systems. Figure XX gives a few examples.

A number of knowledge representation languages are available to facilitate the writing of rules (or any other form of knowledge representation). (See Figure XX for a few examples.)

These tools are complemented by the standard artificial intelligence languages.

First, there is LISP (LISt Processing). The main variants are INTERLISP (XEROX, Bolt, Beranek and Newman, SRI, University of Southern California, Stanford University), FRANZLISP (Berkeley and Carnegie-Mellon universities), MACLISP (MIT and Stanford University), T (Yale), NIL (MIT), LOGLISP (Rochester) and COMLISP (COMmon LISP, Carnegie-Mellon University and University of Utah).

In Europe and Japan, it appears that PROLOG (Logic Programming) is the most popular of artificial intelligence languages. In Canada, the Prologica Company sells a version of Prolog for $33,500. Another language is SMALLTALK.

Although the most common languages such as PASCAL and FORTRAN may be used, they are not as efficient.

The choice of a language must depend on three factors: familiarity of staff with a given language, availability and the nature of the problem to be solved.

5.3 Level of Effort Required

It is difficult to estimate what level of effort will be required to develop an expert system because each case is different. In 1965, 38 person-years were needed to develop DENDRAL. Nevertheless, the experience gained in developing DENDRAL has made it possible to considerably reduce the time required.

In 1982, it was estimated that an average of five person-years were required to develop an average system. Today, depending on the sources consulted, it is estimated that it takes between seven person-months and five person-years. A prototype, complete enough to prove the project's feasibility, can usually be developed within three to six person-months.

There is no use in attempting to accelerate the process by adding person-years because coordination and communication problems then appear. If the team is too small, on the other hand, it takes too long to code the system.

The optimum size for a development team is somewhere between two and five persons. Normally, there is an expert in a selected field, a specialist in expert systems and a specialist in natural language interfaces on the same team.

A **prototype** sophisticated enough to be used in the field will often cost approximately $100,000 for a large system. Before it can be used in day-to-day operations a complex system is likely to cost $1 million or more. It is estimated that Digital Equipment Corporation spent $10 million for XCON. It all depends on the complexity of the system and the level of performance desired.

Computers used to develop these sytems are very expensive. If required, accessories such as modems and printers are added, the total equipment budget will vary from $75,000 to $100,000 per researcher.

Including equipment, software and staff, a large prototype could easily cost from $1 to $3 million. A small system could be developed for less than $90,000.

5.4 Foreign Consultants

Figure XXIII is a list of firms offering products and services related to expert systems which are often cited in the literature. In view of the relatively large number of Canadian firms working in this field, they will be mentioned only in the section of the report devoted to Canada.

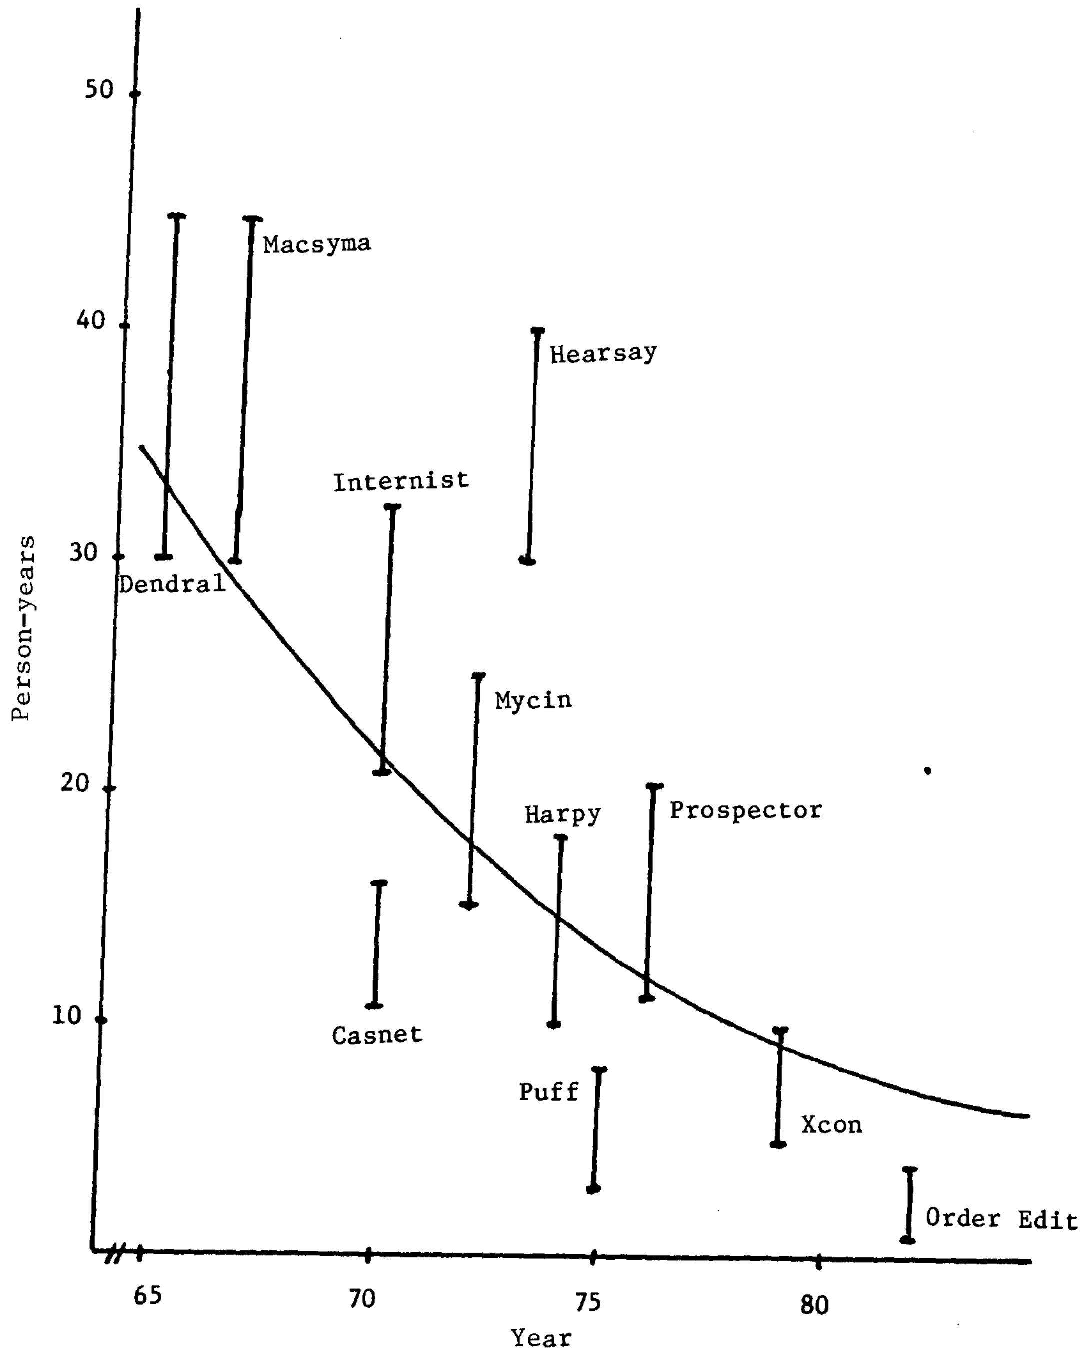

Figure XXI. Time required to complete a system.

Source: Sperry Corp.

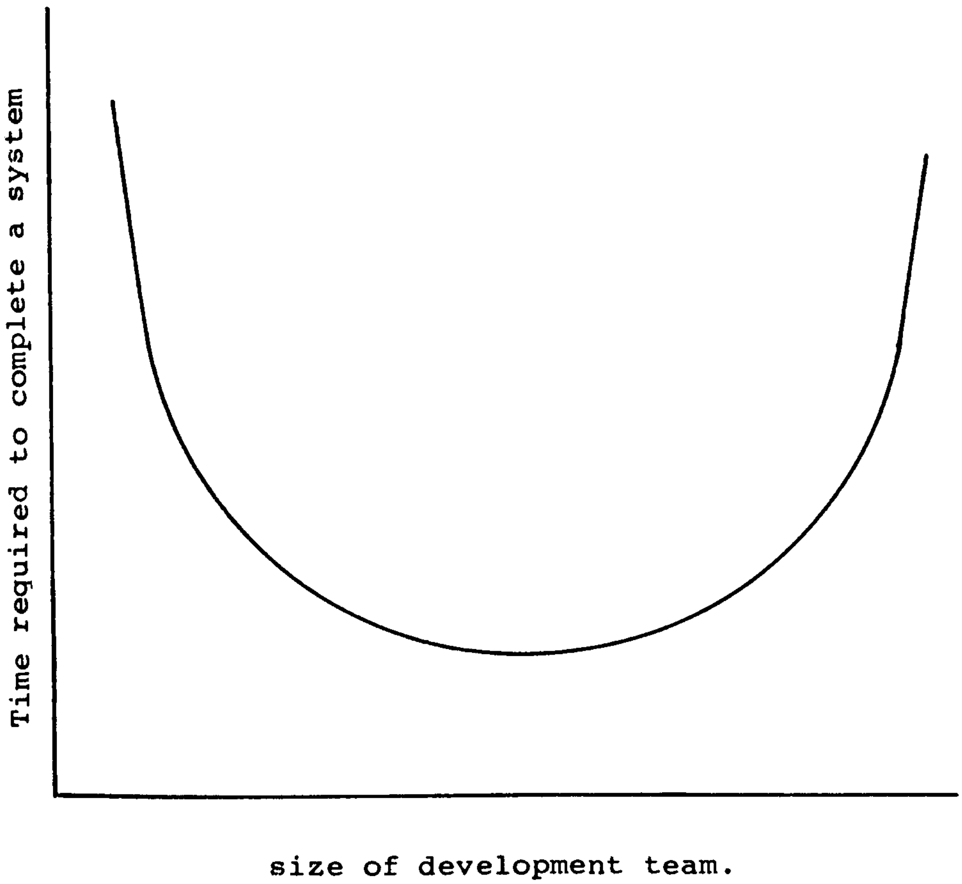

Figure XXII. Time required to complete a system by size of development team.

Advanced Information and Decision Systems (Mountain View, Calif.)	(custom systems)
Applied Expert Systems (Cambridge, Mass.)	(custom systems)
Bolt, Beranek and Newman	
Brattle Research (Cambridge, Mass.)	(custom systems)
Carnegie Group Inc. (Pittsburgh, Penn.)	(shells)
Cognitive Systems (New Haven, Conn.)	(financial applications)
Digital Equipment	(shells)
Eantech (Palo Alto, Calif.)	
EG&G (Idaho)	(nuclear applications)
Epitek (Sweden)	(shells)
Expert System International (U.K.)	
Fairchild Advanced Research Laboratory	(consortium with French firm)
Framentec (See Teknowledge)	(consortium with French firm)
General Electric	
General Research (Santa Barbara, Calif.)	
Helix (U.K.)	
Hewlett-Packard	
Human Edge Software (Calif.)	(systems based on psychology principles)

Figure XXIII. Firms offering products or services in the field of expert systems.

IBM	(data processing operations)
ICI (U.K.)	
Inference Corp. (Los Angeles, Calif.)	(shells)
Infologics (Sweden)	(shells)
IntelliCorp (Palo Alto, Calif.)	(shells)
IntelliGenetics (Palo Alto, Calif.)	(custom systems and shells)
Intelligent Terminal (Edinburgh, Scotland)	
Lockheed	
MITRE Corp.	(air traffic control applications)
Molecular Design Ltd. (Hayward, Calif.)	
Radian	
Rand Corp.	(shells)
Schlumberger-Doll Research	(oil exploration applications)
Silogics (Los Angeles, Calif.)	(custom systems)
Smart Systems Technology (McLean, Virginia)	(expert training and custom systems)
Software Architecture and Engineering (Arlington, Virginia)	(shells)
SPL (Abingdon, U.K.)	

Figure XXIII. (cont'd)

Stanford Research Institute	
Syntelligence (Menlo Park, Calif.)	(banking applications)
Taylor Instruments, Division of Combustion Engineering (Rochester, N.Y.)	(plant controllers)
Teknowledge (Palo Alto, Calif.)	(custom systems) (shells) (expert training)
Texas Instruments	(shells)
Verac (San Diego, Calif.)	

Figure XXIII. (cont'd)

6.0 ACTIVITIES ABROAD

6.1 United States

6.1.1 Government

Most government initiatives are in the military sector.

The main sponsor of artificial intelligence research in general, and of expert systems in particular, is without doubt DARPA (Defence Advanced Research Project Agency). DARPA will contribute $600 million over a five-year period under the Strategic Computing and Survivability program.

The purpose of this program is to develop super-intelligent military computers on the basis of microelectronics research, using new computer architecture and artificial intelligence. The main research fields are **expert systems**, speech recognition, vision systems and natural language understanding.

The first application will be an autonomous system capable of moving by itself and carrying out independent action. DARPA anticipates that this will be a driverless vehicle used for reconnaissance and munitions handling. Using visual sensors, it will be capable of travelling at speeds of up to 60 kilometres an hour over distances of up to 50 kilometres. The sensors will cooperate with an expert system to avoid obstacles, identify landmarks and reproduce a map of the terrain.

The next priority will be an assistant for pilots to free them from tasks they must currently carry out themselves. Through the use of expert systems, speech recognition and graphics technology, the system will manage flight and arms systems during the course of battle.

The third priority is a combat management system which will be able to deal with the enormous complexity of modern military systems. The system will predict events that are likely to occur, suggest actions to be taken, do detailed planning of any such action and resolve possible conflicts that might arise between different goals. An air traffic control system is already being developed for the USS Carl Vinson aircraft carrier.

The U.S. Air Force is particularly interested in applications for autonomous systems, robotics, electronics consultants, CCCI (command, control, communications and intelligence), as well as simulation. Research is under way at the following centres: Air Force Office of Scientific Research, the Rome Air Development Center, the Avionics Laboratory, the Flight Dynamics Laboratory, the Human Resources Laboratory, the Air Force Institute of Technology and the Air Force Management Research Centre.

The U.S. Navy's interests cover just about every applications imaginable. Research is carried out primarily at the following centres: the Office of Naval Research, the Navy Center for Applied Research in Artificial Intelligence, the Naval Ocean Systems Center and the Naval Underwater Systems Center. Work is also being done at the Naval Surface Weapons Center, the Naval Air Development Center, the Naval Weapons Center and the Naval Training Equipment Center.

The only known U.S. Army centre for research in this area is the U.S. Army Engineering Topographic Laboratories at Fort Belvoir in Virginia.

Another Department of Defense program, STARS (Software Technology for Adaptable Reliable Systems), will modernize military data processing. Knowledge-based systems is one of the fields covered by the program. The others are reusable modules, high-level languages, applications generators, specialized computer architectures and standard methodologies.

Other sources of funding include the National Institute of Health, the National Science Foundation, the National Library of Medicine, the U.S. Geological Survey and NASA.

6.1.2 University

The main universities working in the field of expert systems are: Stanford, Rutgers, Carnegie-Mellon and the Massachusetts Institute of Technology. The others are the California Institute of Technology, Harvey Mudd College, the Illinois Institute of Technology, North Carolina State University, Rennselaer Polytechnic Institute, Ohio State University, Syracuse, Berkeley, Minnesota, New Hampshire, Cornell, Duke, Pittsburgh, Illinois and Pennsylvania.

6.1.3 Private Sector Work

In 1983, 13 firms established a cooperative research centre called the Microelectronics and Computer Technology Corporation in Austin, Texas. The chief executive officer of this corporation is Admiral Inman. Each firm is committed to involvement in at least one project for at least three years. The founding companies are: Advanced Micro Devices, Allied, Control Data, Digital Equipment, Harris, Honeywell, Martin-Marietta, Mostek, Motorola, NCR, National Semiconductor, RCA and Sperry. The operating budget for the first year was U.S. $50 million.

Research concentrates on four major topics: microelectronics, software, CAD/CAM and advanced computer architectures. The latter includes the following areas: artificial intelligence and expert systems, man-machine interfaces including speech and image recognition, database management and parallel processing.

Another initiative was launched by the Stanford University Center for Integrated Systems, a university-industry-government cooperative. This centre is attempting to develop expert systems that can assist in the design, checking and debugging of very large-scale integrated circuits (VLSI). Other research themes are concentrated on the chips themselves. The participating companies are: Digital Equipment, Fairchild, GE, General Telephone and Electronics, Gould/AMI, Hewlett-Packard, Honeywell, Intel, IBM, ITT, Monsanto, Motorola, Northrop, Phillips/Signetics, Tektronic, Texas Instruments, TRW, United Technologies and Xerox.

The following are only a few of the companies that are doing research in the field: Control Data, Digital Equipment, Eaton, General Electric, Harris, Hewlett-Packard, Honeywell, IBM, Martin-Marietta Aerospace, National Cash Register, RCA, Sperry, Westinghouse, Amoco, Bolt Beranek and Newman, Bell Labs, Computer Thought (Texas), GM, Hughes, Machine Intelligence (Sunnyvale Calif.), TRW and all the other companies mentioned in this report.

6.2 Japan

The major Japanese artificial intelligence program is known as the Fifth Generation Computer Program launched in April 1982 with the establishment of the Institute for New Generation Computer Technology (ICOT).

The project is sponsored by MITI (Ministry of International Trade and Industry). It will spend approximately CDN $500 million over the 10-year project. Industry will match this amount.

The goal of this program is to develop a computer that resembles human beings in several ways (e.g., inference, association and learning). It will lead to the development of a microcomputer that acts as an expert system to the user. The user of such a machine will have access to mainframe computers and to external knowledge bases through the use of telecommunications networks.

What will a fifth generation computer be able to do? It will be able to see, put forward different ways of solving specialized problems and detect and correct a number of problems by itself. It will be able to program itself on the basis of both stated objectives and a summary description of the process to be used in achieving these objectives.

Many people think that the creation of such a computer is impossible, but even a partial achievement of some of the objectives of the program would be a giant leap forward.

The first three years of the program were entirely funded by MITI at a cost of U.S. $45 million. The first stage involved reviewing, evaluating and restructuring existing knowledge in the field. Prototypes for a number of sub-systems were developed.

The next three years will be devoted to completing the sub-systems and to the development of basic algorithms and computer architectures.

The final stage of the project will integrate all of the above and develop software applications for the first fifth generation computer prototype.

Research efforts for the development of the equipment will follow two paths: an inference engine operating in parallel rather than in sequence as do all current computers, and a knowledge-based machine, i.e., a machine optimized for managing knowledge bases of approximately 100 billions bytes.

Software research will concentrate on a problem-solving system, an inference system, and a knowledge management system, as well as intelligent man-machine interfaces and automatic programming.

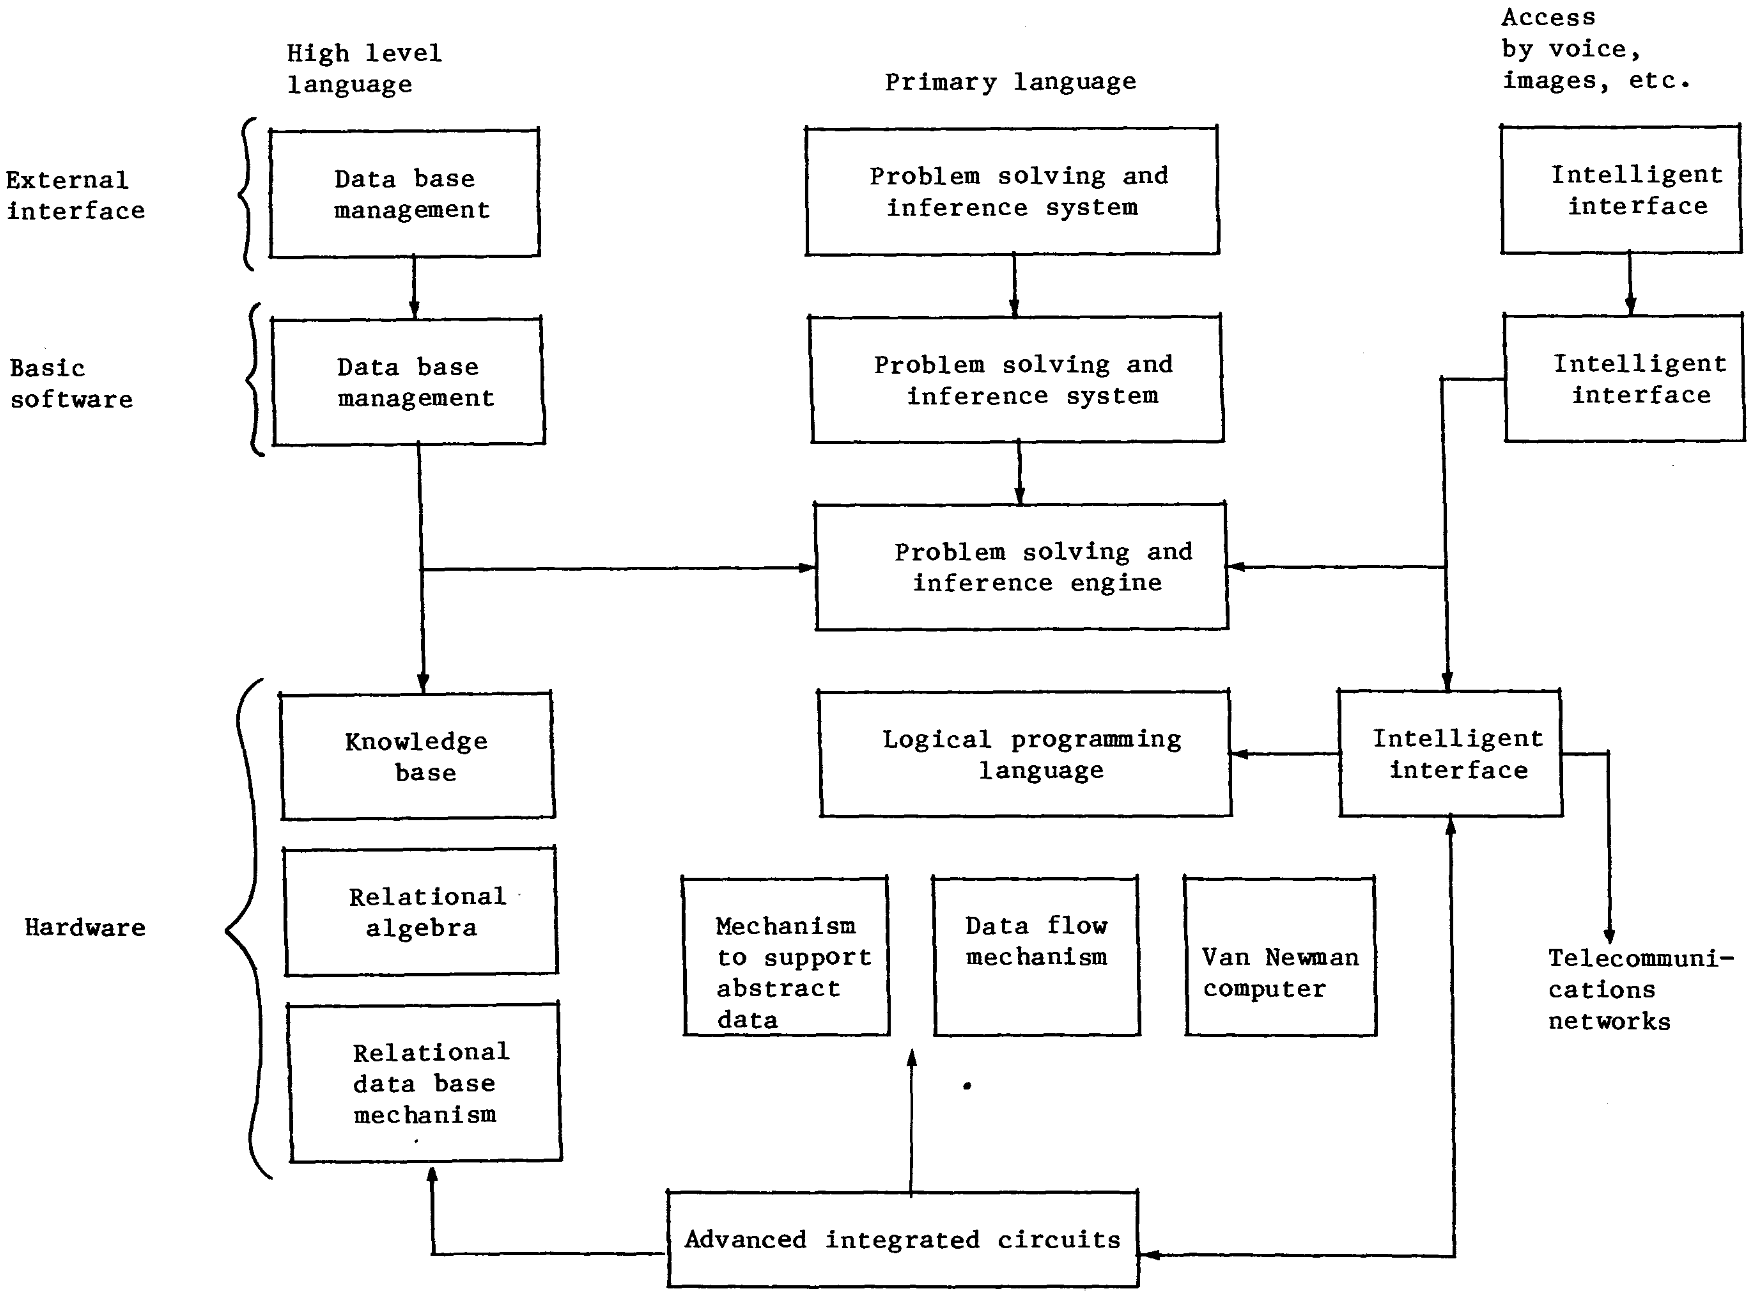

Figure XXIV. Fifth generation computer architecture. Source : Technocrat, January 1982.

6.3 Europe

The European Community Commission, in an attempt to reconstruct its technical base and reverse the deterioration in its electronics trade balance, launched its famous ESPRIT program (European Strategic Programme for R & D in Information Technology) in February of 1984.

The first stage of the program will cover five years and will cost approximately CDN $1.5 billion, half of which will be financed by the commission. Decisions concerning the second stage will take place after a 1988 review. Approximately 2,000 person-years will be involved each year - an enormous research effort.

All projects submitted must involve at least two companies from two different countries and concern competitive R & D. It is hoped that this approach will lead to the achievement of a genuine spirit of international cooperation. The following 12 firms are already involved in the program: General Electric (England), International Computers (England), Plessey (England), General Electric (France), Bull (France), Thompson-Brandt (France), AES-Telefunken (West Germany), Nixdorf (West Germany), Siemens (West Germany), Philips (Holland), Olivetti (Italy) and Societa Torinese Esercizi Telefonici (Italy).

The specific objectives of ESPRIT are:

- to improve man-machine interfaces in order to make computers easier to use;

- to make machines more intelligent by making it unnecessary for the user to worry about details;

- to integrate all communications networks;

- to improve system reliability and security.

These objectives can only be achieved by making enormous progress in both software and hardware.

It was on this basis that five major research topics were identified. Three of these topics involve technology required for progress: advanced microelectronics (e.g., geometries below the micron level), software technology (e.g., modularity and automation of the process) and advanced information processing. The other two topics are specific technology applications: office automation and computer-aided manufacturing.

The topic that most interests us is advanced information processing. It covers four major areas: knowledge engineering; pattern recognition and man-machine interfaces; advanced computer architectures; and computer-aided design methodologies.

The objectives of the knowledge engineering thrust are to acquire:

> the tools and technologies which will be needed for the practice of Knowledge Engineering in order to realize commercially and socially acceptable Knowledge Based Systems applications, of which Expert Systems, Decision Support and Computer-Aided Instruction are examples.
>
> (Steering committee decision)

The research will include: knowledge representation and inference techniques; knowledge acquisition and learning techniques; knowledge handling; dialogue and natural language; implementation languages; interpreters used in advanced knowledge representation formalisms; cognitive systems and how to measure them; and advanced systems applications.

Other research topics also deal with expert systems. Under the heading of software technology there is the following project: the development of software systems based using expert systems.

6.3.1 <u>Great Britain</u>

In the field of artificial intelligence in general, the main government contribution is the Alvey program which involves expenditures of £350 million over five years. Academic institutions will receive £50 million for research from the Ministry of Education and Science. The Ministry of Defence and the defence industry will each contribute half of up to £80 million for research. The Ministry of Trade and Industry and the non-military industry will each contribute £110 million towards research. The government's contribution will, therefore, reach £200 million.

The five research priorities are: (i) very large scale integrated circuits, (ii) software engineering, (iii) man-machine interfaces and (iv) knowledge-based intelligent systems. Approximately £40 million will be devoted to the last topic.

An introductory kit was prepared by the National Computing Centre to increase public awareness in Great Britain of this technology. Costing only a few hundred dollars, it includes a simple expert system shell.

Finally, firms will be able to send their personnel for training to government centres.

There is a special association devoted to expert systems, the British Society of Expert Computer Systems, chaired by the well-known Dr. Alex D. Agapayeff.

Among the many clubs that have been established is one which includes approximately 20 companies as members. It is supported by the Alvey Directorate, which has ordered a special expert system for quality control in an ethyl-oxalate plant.

The following universities are carrying out research in the field: Edinburgh, Essex, Sussex, Cambridge, South Bank Polytechnic and the Open University. Other institutions include the Imperial Cancer Research Foundation, the Ministry of Defence and the Ministry of Trade and Industry.

In this report we have already named a number of firms involved in research. Others include Marconi, Knowledge-Based Systems, Logica VTS and Plessey.

Several videocassettes on the subject have been prepared by the British and are available from the Open University.

6.3.2 West Germany

In 1984, the West German government approved DM 3 billion for the Department of Research and Technology to support information technologies over a five-year period. The measures adopted included tax incentives, venture capital, government purchases and measures to improve cooperation between government, university and industry laboratories.

The research areas chosen were the following:

- integrated circuits (DM 600 million);

- industrial automation (DM 530 million);

- micro-peripherals (DM 320 million);
- wide band digital networks and photonics (DM 260 million);
- new components (DM 200 million);
- knowledge processing and pattern recognition (DM 200 million);
- computer-aided design software (DM 160 million);
- new computer architectures (DM 160 million);
- the establishment of a research network (DM 100 million);
- basic research (DM 100 million);
- CAD applied to integrated circuits (DM 90 million);
- key components (DM 90 million);
- integrated optics (DM 90 million);
- high resolution television (DM 60 million).

The various research institutes involved are the National Centre for Research in Mathematics and Data Processing in Berlinghoven near Bonn, the Fraunhofer institutes in Stuttgart and Aachen and the universities of Hamburg, Kaiserslauten and Stuttgart.

The following companies are conducting research into expert systems applications for processes and production: Siemens, Rank Xerox, Digital Equipment, Triumph Adler and Nixdorf.

6.3.3 France

Work is being carried out at CNET (Centre national d'études des télécommunications), CERT (Centre d'études et de recherches techniques), the Commissariat de l'énergie atomique, the CERFIAT, the Institut de programmation de Paris, the Institut de mathématiques appliquées de Grenoble, the Laboratoire de recherche en informatique d'Orsay, the CRIN (Centre de recherche en informatique de Nancy), the INRIA (Institut national de la recherche en informatique et en automatique), the LAAS (Laboratoire d'automatique et d'analyse de système) in Toulouse, the Laboratoire de recherche en informatique, the Université de Paris Sud and the Université de Marseilles.

The following companies are known to be active in the field: General Electric (France), Cap Sogeti, Compagnie d'informatique militaire spatiale et aéronautique, Elf-Aquitaine, Thomson CSF, CII-Honeywell-Bull, Seri Renault, Matra and Schlumberger.

The major association is the Association d'intelligence artificielle et de simulation.

6.3.4 Holland

Expert systems have just begun to generate interest in Holland. The main centres of activity are the Free University of Amsterdam, the University of Delft, Philips, Shell and a number of small software firms.

6.3.5 Other

Research in Italy is fairly scattered with perhaps 50 researchers for the whole field of artificial intelligence. The main centres appear to be the Centre for Telecommunications Research (ELSAG) and the Honeywell subsidiary.

Sweden has a small but dynamic team working in the field at the universities of Linkoping, Uppsala and Stockholm. At least one company, Norsk Data, is known to be working in the field.

There is one centre of activity in Belgium - the Université libre de Bruxelles.

Finally, the Academy of Sciences in Moscow has launched its own fifth generation computer program for a five-year period at a cost of approximately U.S. $100 million.

7.0 CANADIAN ACTIVITIES

Interest in artificial intelligence and expert systems is relatively new in Canada.

In 1982, Alain Letendre, of the Ministry of State for Science and Technology, circulated a preliminary version of the report "The Next Generation of Software: Artificial Intelligence". The next year the Science Council of Canada commissioned a few studies in the field. The council has since organized two seminars on the subject. The same year, the Department of Communications, jointly with the Secretary of State, commissioned an important study by the Cognos Company of Ottawa, a shortened version of which will be published this year.

There was a great deal of new activity in 1984. In September, there was a conference on fifth generation computers and expert systems in Calgary; in October, there was a seminar on fifth generation computers at the University of Toronto; and in November, a seminar on Artificial Intelligence and Expert Systems was organized by the Centre de cours intensif of the Ecole polytechnique de Montréal.

Since then, there has been a deluge of seminars, conferences, courses and demonstrations.

7.1 The Federal Government

At the Department of National Defence attempts are being made to apply expert systems to sonar systems to combat fatigue and loss of attention. The Defence Research Establishment in Halifax is trying to apply the technology to the interpretation of acoustic signals.

The National Research Council has begun to develop an expert system for the injection moulding of plastics. A system is currently being designed for applications in robot welding, as well as one for the diagnosis of cardiac defects. Research has been undertaken to develop an expert system capable of assisting the development of design strategies. It could be used in engineering, architecture and industrial design.

The Communications Research Centre of the Department of Communications will soon begin to develop an expert system for diagnosing brain damage in young children.

The Natural Sciences and Engineering Research Council of Canada, through its strategic grants program, supports university research in artificial intelligence. In 1984-85, nearly $500,000 were devoted to expert systems under the theme 'Telecommunications and Data Processing'.

The Department of Transport recently began a general review of possible applications of expert systems in the transport sector. The department is also exploring the possibility of developing systems for the following applications: evaluation of the stability of truck haulers, air traffic control, equipment maintenance and inventory control.

Many other departments have identified applications for expert systems. The Department of Energy, Mines and Resources, for example, has shown interest in applying expert systems in the following fields: mineral prospecting, paleontology, seismic interpretation, remote sensing and interpretation of drilling data.

7.2 The Provinces

In Alberta, the Alberta Research Council is developing an expert system to forecast storms. Another project is under way to develop robotics applications.

7.3 University Research

Simon Fraser University is active in almost all areas of artificial intelligence. The key project at the university is the Automated Academic Advisor, a long-term project which attempts to integrate the concept of a dynamic database, knowledge representation techniques and natural language understanding to produce an operational expert system. The University of British Columbia is integrating artificial vision and knowledge representation for remote sensing applications in forestry.

The University of Alberta, in cooperation with the Alberta Research Council, has developed an expert system to forecast hailstorms. Another research project involves studying the optimum order of adding knowledge to an expert system.

The University of Saskatchewan is interested in computer-aided instruction applications. The SCENT project, for example, calls for an expert system which has knowledge about the types of errors normally made by students.

At the University of Toronto, research is directed by Dr. Tsotsos. Its fields of interest include medical applications and the development of a new generation of expert systems capable of reasoning by analogy and manipulating changing information. This group has just received a $1.5 million grant from the Natural Sciences and Engineering Research Council of Canada. The team has also developed the ALVEN system for analysing radiographs of the left ventricle. Another system, CAA, detects heart irregularities from electrocardiograms.

The University of Western Ontario specializes in very small expert systems.

At the University of Ottawa, a group under the direction of Dr. Skuce is working on the development of an expert system capable of answering questions on regulations, specifications and instructions. The system will later be applied as a tax advisor. In the Department of Electrical Engineering, work is being done on a system that can analyse remote sensing photographs of forests.

The INCOGNITO group (INformatique GOGNITive) at the University of Montreal, intends to work in expert systems applied to automatic programming and natural language.

At Dalhousie University, Dr. M. Horacek is interested in applications of expert systems to medical diagnosis; the Technical University of Nova Scotia is doing work on the use of logic programming in expert systems.

7.4 Industry

The Canadian expert system industry, as in other countries, has only recently been created. It consists primarily of very small companies. Given the small population of Canada compared with its enormous neighbour to the south, the Canadian industry appears to be at least as healthy as the American. (Figure XXV).

Except for Northern Telecom, which has undertaken the development of a system for monitoring and correcting an integrated circuit assembly line, few Canadian companies are involved in the development of expert systems for internal use. Cognos Incorporated, nevertheless, has identified a number of firms which are interested in using expert systems (Figure XXVI). We may assume therefore, that the level of diffusion of this technology is small in Canada compared with U.S.

7.5 Other Private Sector Activities

The Canadian Institute for Advanced Research (a private non-profit corporation) established research into artificial intelligence as one of its priorities through the Artificial Intelligence, Robotics and Society program. It assigned $10 million to this program between 1983 to 1988. Its research topics include artificial vision (University of British Columbia), robotics' sensors (McGill University) and expert systems (University of Toronto). For further details contact the president, Dr. Fraser Mustard, at (613) 963-1380.

7.6 Associations

The only association that specializes in artificial intelligence is the Canadian Society for Computational Studies of Intelligence, 243 College St., 5th Floor, Toronto, Ontario, M5T 2Y1.

It was founded by approximately 30 researchers in 1973 (which makes it the oldest national artificial intelligence association in the world). It currently has 300 members. It is now a special interest group within the Canadian Information Processing Society (CIPS).

Alan Campbell P.O. Box 2542 Smithers British Columbia V05 2N0	Shell
Atletic Inc. Nova Scotia (902) 424-6580	Expert system to facilitate the interpretation of geophysical data in the oil industry
Applied AI Systems P.O. Box 13550 Kanata, Ontario (613) 592-0084	Custom expert systems Shells (APES) Training Natural language Micro-PROLOG Golden Common Lisp Consulting
Canadian Artificial Intelligence Products Corporation (CAIPC) 106 Colonnade, Suite 220 Nepean, Ontario K2E 7P4 (613) 727-0082	Custom expert systems Natural language Computer translation Consulting
Cognicom 20 Richmond Street Suite 425 Toronto, Ontario M5C 2R9 (416) 366-4857	Consulting
Control Data Mississauga, Ontario	A machine specialized in knowledge processing is currently being developed

Figure XXV. List of main Canadian companies involved in expert systems.

Dogwood AI Research 1826 West 1st Avenue Vancouver, British Columbia V6J 1G5	Expert system in forestry Expert system for avalanche control
Expert Systems Corp. Vancouver, British Columbia	Signal processing system
ForceTen Enterprises Inc. Nova Scotia (902) 453-0040	Expert system applied to computer-aided instruction
Gomi AI Systems Kanata, Ontario	Medical applications
Intralogic Kitchener, Ontario	Prolog
LISP Canada Inc. 5252 de Maisonneuve Boul. W. Montreal, Quebec H4A 3S5 (514) 487-7063	LISP Machines LM-PROLOG ZETALISP-PLUS
Logicware Inc. 1000 Finch Ave. W. Suite 600 Downsview, Ontario M3J 2V5 (416) 665-0022	MPROLOG
MacDonald Dettwiler and Associates Airport Executive Park 3751 Shell Rd. Richmond, British Columbia V6X 2Z9 (604) 278-3411	Expert system for image processing

Figure XXV. (cont'd)

N.W. Artificial Intelligence Vancouver, British Columbia	-------------------------
Northern Telecom 33 City Centre Dr. Mississauga, Ontario L5B 3A2	Expert system being developed for monitoring and correcting an integrated circuit production line
Texas Instruments	Explorer Computer Skeleton (Personal Consultant)
Unitek Technologies 115, 10751 Shellbridge Way Richmond, British Columbia V6X 2W8 (604) 276-2429	Accounting and tourism applications
Xerox Canada Inc. 703 Don Mills Road Don Mills, Ontario M3C 1S2 (416) 429-6750	Xerox 1108 Interlisp-D

Figure XXV. (cont'd)

ORGANIZATION	INTEREST
Alcan	Process control Intelligent robots
Bank of Montreal	Financial advisor
Crack Resources	Seismic data analysis
Imperial Oil	Refinery programming Optimization of distribution network Seismic analysis
Noranda	Design of open pit mines
Ontario Hydro	Process control
Spar Aerospace	Intelligent space arm

Figure XXVI. Examples of potential applications identified by a number of Canadian companies.

Source: Cognos Inc.

BIBLIOGRAPHY

Barker, G.R. — Trends in Artificial Intelligence Research and Applications, Report submitted to the OECD Committee on data processing and communications, October 1983, 39 p.

Campbell, N.A.; Hollister, V.F.; Duda, R.O.; and Hart, P.E. — "Recognition of A Hidden Mineral Deposit by An Artificial Intelligence Program." Science, 3 September 1982, pp. 927-928.

D'Abrosio, B.D. — "Expert Systems: Myth or Reality?" Byte, January 1985, pp. 275-282.

Duda, R.O. and Gaschnig, J.G. — "Expert Systems Research." Science 15 April 1983, pp. 238-280.

Evanczuk, S. and Manuel, T. — "Practical Systems Use Natural Languages and Store Human Expertise." Electronics, 1 December 1983, pp. 139-145.

Fawcette, J. — "Pentagon Push for Software Development Tools." High Technology, October 1983, pp. 78-79.

Feigenbaum, E.A. and McCorduk, P. — The Fifth Generation - Artificial Intelligence and Japan's Computer Challenge to the World. Addison-Wesley, 1983.

Fox, J. — "Computers Learn The Bedside Manners." New Scientist, 29 July 1982, pp. 311-313.

Gellaire, H. — "Les systèmes experts : une percée de l'intelligence artificielle." La Recherche, May 1982, pp. 644-645.

Gevarter, W.B. — An Overview of Expert Systems. U.S. Department of Commerce, May 1982, 67 p.

Gevarter, W.B. — "Expert Systems: Limited but Powerful." IEEE Spectrum, August 1983, pp. 39-45.

Gevarter, W.B. — Expert Systems. Proceedings of the Army Conference on Application of Artificial Intelligence to Battlefield Information Management, April 1983, pp. 193-206.

Joseph, E. "A Knowledge-Based Era." Computer World, 1 November 1982, pp. 101-108.

Kinnucan, P. "Computers That Think Like Experts." High Technology, January 1984, pp. 30-42.

Lenat, D.B. "Computer Software for Intelligent Systems." Scientific American, September 1984, pp. 204-213.

Letendre, A. La prochaine génération de logiciels : l'intelligence artificielle. Ministry of State for Science and Technology, 1983, 139 p.

Lineback, J. "Lisp Machine Provides a Shell for Industrial AI Applications in One of First Expert Systems to Go to Work." Electronics Week, 27 August 1984, pp. 33-36.

Martins, G.R. "The Overselling of Expert Systems." Datamation, November 1984.

McCulla, G. and Cercone, N. Directions for Canadian Artificial Intelligence. Canadian Society for Computational Studies of Intelligence, 1984, 81 p.

Miller, R.K. "The 1984 Inventory of Expert Systems." Technical Insights, 1984, 183 p.

Myers, E. "Machines That Lisp." Datamation, September 1981, p p. 105-108.

Sharkin, J.N. "Expert Systems: The Practical Face of Artificial Intelligence." Technology Review, November-December 1983, pp. 72-78.

Sheil, B. "Family of Personal Lisp Machines Speeds AI Program Development." Electronics, 3 November 1983, pp. 153-156.

Sweitzer, D.L.S. Artificial Intelligence - The State-Of-The-Art. Prepared for the Science Council of Canada, February 1982, 70 p.

Tate, P. "The Blossoming of European AI." Datamation, 1 November 1984, pp. 85-88.

Wallich, P. "Progress Is Reported in Expert Systems and Software to Build Such Systems and in Object-Oriented Programing." IEEE Spectrum, January 1984, pp. 47-49.

Webster, R. and Miner, L. "Expert System - Programming Problem Solving." Technology, January-February 1982, pp. 63-73.

Zadeh, L.A. "Making Computers Think Like People." IEEE Spectrum, August 1984, pp. 26-32.

________ "AI Trends '85." DM Data, 1985, 136 p.

________ "Artificial Intelligence - The Second Computer Age Begins." Business Week, 8 March 1982, pp. 66-75.

________ "Experts Split over Future of Artificial Intelligence." New Scientist, 18 February 1982, p. 440.

________ "Fifth Generation Computer Project." Technocrat, January 1982, pp. 29-47.

________ Identifying Research Areas in The Computer Industry to 1995. Hoboken: Noyes Publications, 1984, 154 p.

________ "Lisp Machine Provides a Shell for Industrial AI Applications in One of the First Expert Systems to Go to Work." Electronics Week, 27 August 1984, pp. 34-36.

________ "Special Issue on Knowledge Representation." IEEE Computers, October 1983.

________ Strategic Computing. DARPA, October 1983, 69 p.

________ The AI Business - The Commercial Uses of Artificial Intelligence. MIT Press, 1984.

________ "The Emerging Artificial Intelligence (AI) Industry. DM Data, August 1983.

________ "TI's AI Machine Enters Market Already A Best Seller." Electronics Week, 8 October 1984, pp. 19-20.

________________ "Tomorrow's Computers." IEEE Spectrum, November 1983, 120 p.

________________ "When Bugs Cause the Computer to Crash, Prime's Prolog-Based Doc Tells What's Up." Electronics, 23 February 1984, p. 52.

________________ "Why Can't A Computer Be More Like A Man." The Economist, 9 January 1982, pp. 77-81.

________________ "Why Computers Can't Outthink The Experts." Fortune, 20 August 1984, pp 105-118.