EXPERT SYSTEMS IN THE 1980s

E A Feigenbaum

Stanford University
CA

ACKNOWLEDGEMENT

Material on the INTERNIST and SECS programs was adapted from the 1980 SUMEX-AIM Annual Report to the National Institutes of Health.

EXPERT SYSTEMS IN THE 1980s

INTRODUCTION

In the space allotted one can only briefly summarise what there is to say about expert systems – where we are and where we will go in the 1980s – and point the reader to references. Examples will be given of modern work on expert systems, but only in briefest description. Some of these (perhaps more than a fair share) will be drawn from the work of the author's own group, the Stanford Heuristic Programming Project.

ARTIFICIAL INTELLIGENCE, EXPERT SYSTEMS, AND KNOWLEDGE

Artificial Intelligence research is that part of Computer Science that investigates symbolic reasoning processes, and the representation of symbolic knowledge for use in inference. It views heuristic knowledge to be of equal importance with 'factual' knowledge, indeed to be the essence of what we call 'expertise'. In its 'expert systems' work, it seeks to capture the expertise of a field, and translate it into programs that will offer intelligent assistance to a practitioner in that field.

Expert systems

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. The knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the best practitioners in that field.

Knowledge

The knowledge of an expert system consists of facts and heuristics. The 'facts' constitute a body of information that is widely shared, publicly available, and generally agreed upon by experts in a field. The 'heuristics' are mostly private, little-discussed rules of good judgment (rules of plausible reasoning, rules of good guessing) that characterise expert-level decision making in the field. The performance level of an expert system is primarily a function of the size and quality of the knowledge base that it possesses.

THE EXPERT SYSTEMS OF THE 1970s

In our own fields of specialisation, we all understand how long it takes to acquire the knowledge and experience to perform well and that when we perform we are more 'knowledgeable' than 'smart'. DENDRAL (001), MACSYMA (002), and MYCIN (003) were ke programs in making knowledge the focus — setting the dominant paradigm for AI's efforts of the 1970s.

INTERNIST

The INTERNIST program is probably the most knowledge-intensive expert system. Developed by Pople and Myers (004), INTERNIST performs the task of differential diagnosis in the broad field of internal medicine. Its knowledge base encompasses almost 500 diseases and more than 3000 manifestations of disease.

Because INTERNIST is intended to serve a consulting role in medical diagnosis, it has been challenged with a wide variety of difficult clinical problems: cases published in the medical journals, cpc's and other interesting and unusual problems arising in the local teaching hospitals. In the great majority of these test cases, the problem-formation strategy of INTERNIST has proved to be effective in sorting out the pieces of the puzzle and coming to a correct diagnosis, involving in some cases as many as a dozen disease entities.

Although this consultative program is designed primarily to aid skilled internists in complicated medical problems, the program may have a spin-off as a diagnostic and triage aid to physicians' assistants, rural health clinics, military medicine, and space travel.

MOLGEN (Molecular Genetics)

MOLGEN's task (005,006) is to provide intelligent advice to a molecular geneticist on the planning of experiments involving the manipulation of DNA. The geneticist has various kinds of laboratory techniques available for changing DNA material (cuts, joins, insertions, deletions, etc); techniques for determining the biological consequences of the changes; various instruments for measuring effects; various chemical methods for inducing, facilitating, or inhibiting change; and many other tools. MOLGEN offers planning assistance in organising and sequencing such tools to accomplish an experimental goal.

There are actually two MOLGEN systems that design experiments in rather different ways. One is based on the observation that human scientists rarely plan experiments from scratch, but instead start with an abstract or 'skeletal' plan that contains the entire experimental design in outline form. The major design task is to instantiate the details of the plan by finding tools that will work best in the given problem environment.

The other approach used in MOLGEN to design experiments involves the interactions of steps in a plan as it is being formulated. A method called 'constraint posting' is used to make the interactions between steps explicit. Constraints are dynamically formulated and propagated during hierarchical planning and are used to coordinate the design of subparts of an overall plan.

MOLGEN has been tested in the planning of a large number of relatively routine experiments and one rather difficult experiment (rat-insulin gene cloning). Subpackages of

MOLGEN, particularly one for DNA sequence analysis, are currently widely used by molecular biologists. It is likely that extensive use of MOLGEN will be made by the new 'genetic engineering' industry.

Ventilator Management (VM)

A ventilator is a piece of medical equipment that assists a critically ill patient with his breathing. The task of the VM program (007) is to provide real-time advice to clinicians about patients undergoing mechanical ventilation. For VM, the ultimate therapeutic goal is to remove the patient from the ventilator, and VM gives advice on ventilator adjustments that expedite this goal. VM use is being explored at the intensive care unit of the Pacific Medical Center, which has an on-line computer-based patient-monitoring system that automatically obtains measurements of approximately 30 physiological parameters. VM is intended as an extension of this system that will do the following:

- Provide a summary of the patient status, easily understood by the clinician
- Recognise untoward events in the patient/machine system and provide suggestions for corrective action
- Give advice on adjustment of the mechanical ventilator based on an assessment of the patient status and therapeutic goals
- Detect possible measurement errors
- Maintain a set of patient-specific expectations and goals for future evaluation.

One interesting aspect of the VM program is that it works with multiple streams of data that are sent by the monitoring system over a period of time. Interpreting these items of data requires a time-based perspective of the patient's status, eg, a recommendation that a patient is ready for a change in therapy is made after VM analyses both current and recent past physiological statuses.

The experimental VM program at Stanford and the Pacific Medical Center is currently undergoing a second phase of research and development.

EMYCIN

EMYCIN (008) provides a framework for building consultation programs in various domains. It uses the domain-independent components of the MYCIN system, notably the production rule mechanism and backward-chaining control structure. Then for each particular consultation domain, the system builder supplies the rules and parameters of that domain to produce a functioning program.

EMYCIN is not, in itself, an expert system, but a software tool for builders of expert systems. It provides a useful environment with emphasis on speeding the acquisition and debugging of the knowledge of the new domain.

To date, EMYCIN has been successfully applied in the domains of pulmonary function (PUFF - see below), structural analysis (SACON (009)), and blood clotting disorders (CLOT). A skilled user of EMYCIN can, with the help of a domain expert, program a small consultation system very quickly, often in less than one week.

PUFF (pulmonary function disease diagnosis)

PUFF (010), a computer program for the interpretation of standard laboratory measures of pulmonary function, was implemented in EMYCIN. About 50 quantitative parameters are calculated from the measurement of lung volumes, flow rates, and diffusion capacity. PUFF uses these measurements, as well as the patient's history, the diagnosis of the referring physician, and its knowledge of three types of pulmonary disease to produce a report and diagnosis of the patient's condition. The PUFF system is now in routine use at the Pacific Medical Center in San Francisco. 85% of its reports are signed, unaltered, by the physician, after review.

SECS (Simulation and Evaluation of Chemical Synthesis)

The development of new drugs and the study of how drug structure is related to biological activity depends upon the chemist's ability to synthesise new molecules as well as his ability to modify existing structures, eg, incorporating isotopic labels or other substituents into biomolecular substrates. The SECS program (011) assists the synthetic chemist in designing stereospecific syntheses of biologically important molecules.

A spin-off of SECS, XENO, has been programmed to predict the plausible metabolites of a given xenobiotic in order that they may be analysed for possible carcinogenicity.

A META-SECS top-level plan generator has been outlined to reason, using synthetic principles, and conclude plans which will then be used to guide the existing SECS program in synthetic analysis. The first-order predicate calculus is being used to represent the synthetic strategies, and an inference processor is currently being designed.

The programs of the SECS project have been used substantially among chemists in drug and chemical industries and in university laboratories.

EXPERT SYSTEMS RESEARCH IN 1980

Here are two examples of recent expert systems research. Both problem areas – diagnosis of equipment failure, and VLSI design – will be important tasks for expert systems in the 1980s.

DART: diagnosis of computer system malfunction

The purpose of this new Stanford-IBM project is to explore the use of causal, structural, and teleological models of computer systems in diagnosing computer system faults. The major components of the research are the development of an adequate 'machine-definition' language for encoding such models and the identification of general diagnostic techniques. Each model includes a description of the system's 'anatomy' and 'physiology' and knowledge of how this structure realises the system's function in terms of the general principles of design. The practical goal of the research is an automated diagnostician for a contemporary computer system like the IBM 4331. Other intended by-products of the research include a program for assisting human field engineers, a training facility for field engineers based on the device models, and feedback on hardware design.

VLSI design aids project

The overall objective of the VLSI design aids project is the development of intelligent, high-performance computer tools to assist in the design of very large-scale integrated circuits. The technology of integrated circuits will soon permit between 10^5 and 10^7 gates to be in a single circuit; the problem of designing circuits of this size is considerable. The tools developed in this project will be integrated into the design automation system that is being developed by the Stanford Center for integrated systems. The initial focus of the project is on design problems at the layout level.

The project is currently developing a Heuristic Layout Program (HLP), which is a technology-independent knowledge-based program to automatically generate a geometric (mask-level) layout of an integrated circuit, given an abstract description of the circuit (eg, a hierarchical symbolic layout of the circuit). The program is intended to serve both as a layout tool and as a test bed for experimentation with various layout heuristics. HLP follows the hierarchical planning-with-constraints scheme developed as part of the MOLGEN project.

EXPERT SYSTEMS: WHERE ARE WE GOING IN THE 1980s?

Predictions are risky, and today's prophet is tomorrow's fool. Nevertheless, here are some of the author's views on the course of events in the 1980s.

New fields of application

Home entertainment and advice-giving: AI scientists have long used games as a vehicle for exploring new concepts. A game generally has a constrained knowledge base and is highly structured. There exists not only a considerable amount of competence within the AI community for constructing intelligent game programs, but also a vast consumer market for home entertainment of this sort. In the author's view, the money-making potential of this market will make microcomputer-plus-TV-based home entertainment the dominant market for expert systems. The concept of home entertainment is broader than games and includes consultation and advice about a broad range of subjects of interest to the consumer (eg, financial advice, and garden and plant care). Knowledge bases for these specialties will be assembled by experts in much the same way that the 'how-to' books traditionally have been.

'Intelligent agents' for computer-based systems

As we and our users are (jointly and painfully) aware, the use of complex computer software is a very knowledge-intensive activity. The manuals are too complex, filled with a mass of detail, easily forgotten; they are often poorly written and contain merely the 'facts' about the system, not the 'lore' necessary to facilitate use (the 'lore' is known to the 'system-hackers'). In the 1980s, expert systems that are 'expert' about the use of particular complex software systems will radically improve the 'quality of life' for a user attempting to apply the software. Such an expert system will have an extensive knowledge base provided by the system developers, and augmented by the user community itself, that can be employed to interpret and fulfil user requests. The manuals of the past will become active and inferential in the service of user needs. The market for such systems is as broad as the market for all software!

Signal understanding applications

Of the many applications to the fields of science and engineering that will be achieved, one class will be outstanding in its impact – signal understanding. The AI methodology that underlies expert systems work contains the most powerful techniques known for realising signal-to-symbol transformations, the essense of signal understanding (as opposed to signal processing). For situations of low signal-to-noise ratios, it makes little sense to use enormous amounts of computation for statistical calculations to tease the little signal from the noise, when most of the understanding can be readily inferred from the symbolic knowledge surrounding the situation.

Scientific developments of critical importance ('breakthroughs')

There are many important problems of knowledge representation, utilisation, and acquisition that must be solved, but the acquisition problem is the most critical 'bottleneck' problem. An expert system is 'expert' only because it is knowledge-intensive. This problem is now becoming sharply focused, and is beginning to receive the attention it deserves. It is clear that acquisition itself must be a knowledge-driven activity. In the 1980s there will be 'breakthroughs' in two styles of acquisition as follows:

1 The acquisition of domain-specific knowledge directly from recorded data (ie, essentially automatic).

2 Interactive transfer of expertise, in which a program guides and helps the expert in explicating and formalising his/her knowledge.

Similarly, the technologist/builder of the expert system needs tools for facilitating the construction of the system. Some like EMYCIN and AGE (from Stanford), EXPERT (from Rutgers), and ROSIE (from the Rand Corporation) are now available; and the 1980s will see a proliferation of such software packages representing, in effect, 'high-level languages' with which to write expert systems. The major inference methods of AI will be packaged up in this way, and will be coupled to powerful representation structures and formalisms.

Hardware

The microelectronic revolution of the 1980s will put the average expert system into a relatively inexpensive and small hardware package. This is an unremarkable prediction based on the progress of VLSI toward million-gate and ten-million-gate chips. The knowledge bases will be held in disk, bubble, or semiconductor memories of large sizes, accessed as needed. We will see a spectrum of expert systems, from those of modest scope of knowledge that may be packaged in small, highly portable machines costing less than a thousand dollars, to powerful aids for professionals that will reside in workstations costing a few tens of thousands of dollars. Very large knowledge bases will be stored in central repositories, accessed by communication network as needed.

The knowledge engineering industry

As the pressure builds to apply Artificial Intelligence to a variety of expert system projects, a new industry (now in its infancy) will emerge. In a fashion analogous to that of the aerospace engineering industry, this new 'knowledge engineering' industry will transfer the developments of the research laboratories to the useful expert systems that industry, science, medicine, business, and the military will be demanding. High technology companies have already begun to build these groups internally (eg, Schlumberger Inc; Texas Instruments; and IBM). In the US, new small firms have been incorporated, and the best will have a dramatic growth in the 1980s. Expert systems methodology is at the place today that operations research was in the mid-fifties, as it moved from its post-World War II decade of development to its maturity of intense application.

Where are the knowledge engineers?

Limiting the pace of development of this industry will be the shortage of people – the new knowledge engineers. Everywhere, the shortage of highly-trained computer scientists and technologists is being felt. In the US, this is currently called 'the crisis in experimental computer science'. Few universities are training graduate students in expert systems methodology. Because of the heavy non-academic demand, few of the new PhDs are located in universities, so there is currently no 'amplification factor' in the production of trained people. Failure to solve this educational problem will make all predictions about expert systems in the 1980s seem to have been wildly optimistic. Significantly, only in the US has there been a vigorous discussion of this problem, leading one to the seemingly inescapable conclusion that the US (moving to solve the problem) will dominate this new industry throughout the 1980s.

THE FUTURE NEED AND PROMISE

There is a certain inevitability to expert systems research and application. The cost of computers will fall drastically during the coming two decades. As it does, many more of the practitioners of the world's professions will be persuaded to turn to economical automatic information processing for assistance in managing the increasing complexity of their daily tasks. They will find, in most of computer science, help only for those problems that have a mathematical or statistical core, or are of a routine data-processing nature. However, such problems will be rare, except in engineering and physical science. In medicine, management, and the military – indeed in most of the world's work – the daily tasks are those requiring symbolic reasoning with detailed professional knowledge. The computers that will act as 'intelligent assistants' for these professionals must be endowed with such reasoning capabilities and knowledge.

EPILOGUE: A NEW FRONTIER OF KNOWLEDGE

Beyond the 1980s, the larger impact on science, technology, and society will be the exposure and refinement of the hitherto private heuristic knowledge of the experts of the various fields of practice studied. The ethic of science that calls for the public exposure and criticism of knowledge has traditionally been flawed for want of a methodology to evoke and give form to the heuristic knowledge of scientists. Expert systems methodology is beginning to fill that need. Heuristic knowledge can be elicited, studied, criticised by peers, and taught to students. Perhaps it is less important that computer programs can be organised to use this knowledge than that the knowledge itself can be organised for the use of the human practitioners of today and tomorrow.

REFERENCES

001 BUCHANAN B G and FEIGENBAUM E A
DENDRAL and Meta-DENDRAL: their applications dimensions
Artificial Intelligence vol 11 (1978)

002 MOSES J
A MACSYMA primer
Mathlab Memo 2 Laboratory for Computer Science Massachusetts Institute of Technology (Oct 1975)

003 SHORTLIFFE E H
Computer-based medical consultations: MYCIN
American Elsevier New York (1976)

004 POPLE H E
The formation of composite hypotheses in diagnostic problem solving: an exercise in synthetic reasoning
Proc 5th Intl Joint Conf on *Artificial Intelligence* Boston MA (Aug 1977)

005 FRIEDLAND P
Knowledge-based hierarchical planning in molecular genetics
Doctoral dissertation Computer Science Department Stanford Univ (Sep 1979)

006 STEFIK M
Orthogonal planning with constraints, a program that plans synthesis experiments in molecular genetics
Doctoral dissertation Computer Science Department Stanford Univ (Sep 1979)

007 FAGAN L M, SHORTLIFFE E H and BUCHANAN B G
Computer-based medical decision making: from MYCIN to VM
Automedica vol 3 no 2 (1980)

008 VANMELLE W
A domain-independent production rule system for consultation programs
IJCAI-6 (1979)

009 BENNETT J S and ENGELMORE R S
SACON: a knowledge-based consultant for structural analysis
Proc IJCAI-6 pp 47-49 (1979)

010 OSBORN J, FAGAN L, FALLAT R, McCLUNG D and MITCHELL R
Managing the data from respiratory measurements
Medical Instrumentation vol 13 no 6 (Nov 1979)

011 ANDOSE *et al*
1980 SUMEX-AIM Annual Rep to the National Institutes of Health (1980)

012 BARR A and FEIGENBAUM E A (editors)
The handbook of Artificial Intelligence
William Kaufmann Inc Los Altos CA (To be published 1981)

013 NII H P and AIELLO N
AGE: a knowledge-based program for building knowledge-based programs
Proc IJCAI-6 pp 645-655 (1979)