204

# Issues in Developing Expert Systems

Jack Minker

Department of Computer Science

University of Maryland

College Park, Maryland 20742

## 1. Introduction

The purpose of this note is to provide a brief overview of the field of expert systems, and to set forth some issues to be discussed in a panel session on the subject. The field of artificial intelligence has several objectives:

(1) The development of computational models of intelligent behaviour- both cognitive and perceptual.

(2) The engineering-oriented goal of developing programs that can solve problems normally thought to require human intelligence.

(3) The development of tools and techniques needed for the above two items.

The development of a system intended to meet the needs of users and is intended to provide expert advice falls into the second category.

The field of expert systems is relatively new. It extends back approximately twenty years, although it is relatively recent that such systems have been referred to as expert systems. Although there has, in the past few years, been a great deal of work on this subject, the actual accomplishments have, at best, been modest. In using the term modest, it is meant that with respect to having expert programs used by individuals in their daily work, there are relatively few such systems. In the following section we briefly note some of the expert systems that have been developed, and their status. In the last section we discuss several issues that must be addressed if expert systems are to become a reality. These issues are posed for discussion, and no positions are taken on them. The intent of the panel discussion is to explore the issues in great depth.

## 2. Background in Expert Systems

Early work in artificial intelligence was oriented towards providing general approaches to probelm solving. It was realized that some of the problems being attacked were, perhaps, more difficult than anticipated. This was particularly true with work in machine translation and natural language processing. Efforts to apply theorem proving techniques to arbitrary problems in diverse domains introduced combinatorial explosions. A move was therefore made towards specializing problems and building into application areas special knowledge focused on the domain of application. Systems that focus on specific problem domains, building in knowledge specific to that domain, have come to be called expert systems.

There have been several phases in the development of Expert Systems. This may be illustrated by efforts leading to one successful system, MACSYMA, whose function is to act as an expert in the area of formal integration of functions. It is perhaps of interest to note that throughout the development of MACSYMA the term "expert system" was never applied. The first stage was the demonstration that it was possible to perform symbolic integration on a computer. Jim Slagle's system, SAINT, was developed and was subsequently tested. It succeeded in passing an examination in integral calculus at MIT. Although it was successful, it was intended to be a research program, and not a finished product that could be used by scientists and engineers. Following SAINT, Joel Moses developed SIN, which was able to perform integration in a more powerful way than the Slagle program. It had more general rules built into it, and more explicit answers to problems whose integrals were already known. In the third stage, the system, MACSYMA, was developed to meet the day- to-day needs of scientists and engineers. Thus, after a long research and development stage, a final product was developed. Although MACSYMA is a successful system, in the sense that it is currently in use by scientists and engineers, many individuals fail to refer to it as an expert system. Neither MACSYMA, SAINT nor SIN are referred to as expert systems since the term was not in vogue when they were developed. It is clear, however, that all three systems would be referred to as expert systems were they developed today.

There are several stages in the engineering of an expert system:

Phase 1 - Research in which the feasibility of developing an expert system in a specific domain is established.

Phase 2 - Development of and experimentation with a prototype system.

Phase 3 - Field test the prototype system.

Phase 4 - Use of the expert system in the field.

In discussing the status of a particular "expert system", it is useful to distinguish its stage of development. There are four expert systems that are routinely in use: MACSYMA; DENDRAL(Feigenbaum et al. [1971]), an expert system that analyzes mass spectral patterns to determine the chemical structure of unknown compounds; R1 (McDermott [1981]), an expert system to determine computer layouts and configurations; and PUFF (Osborn et al.[1979]), an expert system that interprets pulmonary function tests.

A list of some representative expert systems and their domains of application appears in Table 1. A number of useful articles on expert systems appear in books (Michie [1979], Hayes-Roth et al.[1983], Webber et al. [1981], Szolovits et al.[1982]). Several comprehensive surveys have been written on expert systems (Duda et al.[1983], Buchanan [1982], Nau [1983]). See Reggia [1982] for a comprehensive list of references in expert systems oriented primarily towards medical applications.

| Expert System | Domain | Reference |
|---|---|---|
| AQ11 | Diagnosis of plant disease | Chilausky et al.[1976] |
| CASNET | Glaucoma assessment and therapy | Weiss et al.[1978] |
| DENDRAL | Mass spectroscopy interpretation | Feigenbaum et al.[1971] |
| Digitalis Advisor | Digitalis dosing advice | Gorry et al.[1978] |
| Dipmeter Advisor | Oil exploration | Davis et al.[1981] |
| El | Analyzing electrical circuits | Stallman et al.[1977] |
| Internist-I | Internal medicine diagnosis | Miller et al.[1982] |
| HASP & SIAP | Ocean Surveillence (signal processing) | Nii et al.[1982] |
| KMS | Medical consulting | Reggia [1980] |
| MACSYMA | Mathematical formula manipulation | Moses [1971] |
| MDX | Medical consulting | Chandrasekaran et al. [1979]. |
| Microprocessor EXPERT | Protein electrophoresis interpretation | Weiss et al[1981] |
| MOLGEN | Planning DNA experiments | Martin et al.[1977] |
| MYCIN | Antimicrobial therapy | Davis et al.[1977] |
| PROSPECTOR | Geological mineral exploration | Hart et al.[1978] |
| PUFF | Pulmonary function test interpretation | Osborn et al.[1979] |
| R1 | Computer layout and configuration | McDermott et al.[1981] |

TABLE 1 - Representative Expert System

Expert systems have been implemented using a variety of different approaches:

(1) Embedding control and inference in a program written in a language such as FORTRAN or PASCAL (Bleich [1972]).

(2) statistical pattern classification techniques as the basis of making conclusions. For example, Bayseian (Ben-Bassat [1980]), and linear discriminant function (Faught et al. [1979]), have been proposed.

(3) Developing cognitive models of diagnostic reasoning Reggia [1981].

(4) Production rule based systems (Davis et al. [1977]).

3. ISSUES

There are a wide range of issues that have to be addressed before expert systems can reach full maturity. These range from the philosophical to the moral to research issues and to user acceptance. It is not intended that all of the items need be addressed before such systems can become a reality, but that a number of these issues must be developed before the field can reach maturity.

1. Philosophical Issues
   a. What is meant by knowledge and how does one differentiate between data and knowledge?
   b. Will it ever be possible to capture all knowledge in a domain of real interest?
   c. Can one deal with systems in which there are significant gaps in knowledge, and how can one assess the effectiveness of such systems?
   d. How does one differentiate an expert system from an application program?
   e. Can an expert system exhibit intelligence in the same sense as attributed to humans?
2. Moral and Sociologic Issues
   a. Are there classes of expert systems that should never be attempted: they are morally repugnant ?
   b. What are the legal problems? Who is responsible for adverse reactions when a medical expert system incorrectly diagnoses a patient?
   c. What are the potential social consequences of expert systems and

are they for the good, or wil they lead to major social problems?

d. Who should build expert systems: domain experts, computer scientists, or both?

3. Research Issues

Knowledge Acquisition and Representation

a. How does one identify and encode knowledge?

b. What characteristics should a knowledge representation formalism have?

c. How does one express temporal knowledge and physiological mechanisms involved in the evolution of disease processes?

d. How does one represent exceptions to situations?

e. How does one explain the basis for the decision criteria and/or rules used in a knowledge based system?

f. Should knowledge be augmented by using causal and mechanistic links that represent functional behavior?

g. How does one obtain large, reliable data/knowledge bases?

Inference and Uncertainty

a. How does one deal with vagueness and ignorance? Are fuzzy logic (Zadeh [1978]) and statistical theories of evidence (Shafer [1976]) useful?

b. In what ways is logical inference useful?

c. Will indefinite data(i.e., data of the form p V q) be needed for expert systems? What are the implications with respect to the development of such a system or answers obtained during its use?

d. How can logical inference handle exceptions?

e. How is reasoning performed in the presence of ignorance and how can a reasoning system recognize the limits of its knowledge?

f. What is "common sense" knowledge, and how can it be embedded in expert systems?

Control

a. How is search controlled in an expert system?

b. What is needed to permit the user to exercise control and to understand what the expert system is doing?

c. Do current languages allow for control needed to find solutions to problems in an efficient manner?

Explanation

a. Explanation in terms of goals and its knowledge base is very useful. However, experts who provided a set of rules are likely to give explanations in terms of physiological mechanisms or disease processes. How can a system accomplish this?

b. How does one provide explanations to different classes of users? That is, how does one maintain models of users and provide explanations to the various users according to the implied intent of the user?

c. How can the user be aware of the significance of questions asked by an expert system? (e.g. the expert system may ask if a spinal tap has been performed, and the user should be able to understand why the question is being asked, as well as the fact that this test is potentially dangerous).

4. System Assessment and User Acceptance

a. How does one certify an expert system?

b. How does one assess system performance, particularly where the "correct" solution to the problems may not always be known? (e.g. medical diagnosis)

b. How does one obtain large, reliable databases?

c. How does one scale up a system from small experimental systems? What are the problems?

d. How does one develop user friendly systems?

e. How does one develop systems that can be transferred from the experimental laboratory to a remote user site?

f. When will cost-effective systems be developed?

g. How can user resistance to change be overcome?

h. How will new knowledge and changes be made at the user sites?

We have set forth some of the issues associated with developing expert systems. In the course of the panel discussion we will consider these issues.

## 5. REFERENCES

1. M. Ben-Bassat, et al. "Pattern-Based Interactive Diagnosis of Multiple Disorders-The MEDAS System", IEEE Trans. Pat. Anal. Machine Intelligence, 2, 1980, 148-160.

2. H. Bleich, "Computer-Based Consultation", AM. J. Med., 53, 1972, 285-291.

3. B. G. Buchanan, in Machine Intelligence, Vol 10, J. E. Hayes, D. Michie and Y.-H. Pao Eds. Wiley, New York, 1982, pp. 269-300.

4. B. Chandrasekaran et al., "An Approach to Medical Diagnosis Based on Conceptual Structures," Proc. Sixth Int'l Joint Conf. Artificial Intelligence, 1979, pp. 134-142.

5. R. Chilausky, B. Jacobsen, and R.S. Michalski, "An Application of Variable- Valued Logic to Inductive Learning of Plant Disease Diagnostic Rules", Proceedings Sixth Annual International Symp. Multiple-Valued Logic, 1976.

6. R. Davis et al., "The Dipmeter Advisor: Interpretation of Geological Signals," Proc. Seventh Int'l Joint Conf. Artificial Intelligence, Aug. 1981.

7. R. Davis, B. Buchanan, and E. Shortliffe, "Production Rules as a Representation for a Knowledge-Based Consultation Program," Artificial Intelligence, Vol. 8, No. 1, 1977, pp. 15-45.

8. R. O. Duda and E. H. Shortlffe, "Expert Systems Research", Science, Vol. 220, No. 4594, 15 Apr 1983, pp. 261-268.

9. E. Faught et al., "Cerebral complications of Angiography for Transient Ischemia and Stroke-Prediction of Risk", Neurology, 29, 1979, 4-15.

10. 10. E. Feigenbaum, G. Buchanan, and J. Lederburg, "Generality and Problem Solving: A Case Study Using the DENDRAL Program," Machine Intelligence 6, D. Meltzer and D. Michie, eds., Edinburgh University Press, 1971, pp. 165- 190.

11. G. Gorry et al., "Capturing Clinical Expertize-A Computer Program that considers Clinical responses to Digitalis", Amer. J. of med., 64,1978,452-460.

12. P.E. Hart, R.O. Duda, and M.T. Einaudi, A Computer-Based Consultation System for Mineral Exploration, Technical Report, SRI International, Menlo Park, California, 1978.

13. F. Hayes-Roth, D. Waterman, and D. Lenet, Eds. Building Expert Systems, Addison-Wesley, New York, 1983.

14. N. Martin et al., "Knowledge-Base Management for Experiment Planning in Molecular Genetics," Proc. Fifth Int'l Joint Conf. Artificial Intelligence, 1977, pp. 882-887.

15. J. McDermott and B. Steele, "Extending a Knowledge Based System to Deal with Ad Hoc Constraints," Proc. Seventh Int'l Joint Conf. Artificial Intelligence, 1981, pp. 824-828.

16. D. Michie, Ed, Expert Systems in the Microelectronic Age, Edinburgh Univ. Press, Edinburgh, 1975.

17. R. Miller et al., New England Journal of Medicine, 307,468(1982).

18. J. Moses, "Symbolic Integration: The Stormy Decade," Comm. ACM, Vol. 14, No. 8, 1971, pp. 548-560.

19. D. S. Nau, "Expert Computer Systems", Computer, 16, 63(1983), pp. 63-85.

20. H. Nii et al., AI Magazine 3,23(1982).

21. J. Osborn et al., "Managing the Data from Respiratory Measurements," Medical Instrumentation, Vol. 13, No.6, Nov. 1979.

22. H.E. Pople, "The Formation of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Synthetic Reasoning," Proc. Fifth Int'l Joint Conf. Artificial Intelligence, 1977, pp. 1030-1037.

23. J. A. Reggia, "Computer-Assisted Medical Decision Making", Applications of Computers in Medicine, M. Schwartz, Ed, IEEE, 1982.

24. J. A. Reggia and B. Perricone, "Knowledge-Based Decision Support Systems - Development Through High- Level Languages", Proc. Twentieth Ann.

Technical Symposium of Washington DC, ACM, June 1981, pp. 75-81.

25. J. Reggia et al., "Towards an Intelligent Textbook of Neurology," Proc. Fourth Annual Symp. Computer Applications in Medical Care, 1980, pp. 942-947

26. G. Shafer, A Mathematical theory of Evidence, Princeton Univ. Press. Princeton, New Jersy, 1976.

27. R.M. Stallman and G.J. Sussman, "Forward Reasoning and Dependency-Directed Backtracking in a System for Computer-Aided Circuit Analysis," Vol. 9, Artificial Intelligence, 1977, pp. 135-196.

28. P. Szolovits, Ed, Artificial Intelligence in Medicine, West View, Boulder, Colorado, 1982.

29. S.M. Weiss et al., "A Model-Based Method for Computer-Aided Medical Decision- Making," Artificial Intelligence, Vol. 11, No. 2, 1978, pp. 145-172.

30. S. Weiss and C. Kulikowski, "EXPERT - A System for Developing Consultation Models," Proc. Sixth IJCAI, 1979, pp. 25-40.

31. S. M. Weiss et al., In Proc. of the Seventh Int'l. Joint Conf. on AI, Univ. of British Columbia, Vancuvar, Aug 1981, pp. 835.

32. L. Zadeh, "Fuzzy Sets", Systems 1,3(1978).