



Scientific DataLink

Report 79-26
Stanford -- KSL

Clinical Knowledge Engineering: The
MYCIN Project.
Edward H. Shortliffe,
Aug 1979

card 1 of 1

The MYCIN Project

E.H. Shortliffe

CLINICAL KNOWLEDGE ENGINEERING: THE MYCIN PROJECT

Edward H. Shortliffe, M.D., Ph.D.
 Heuristic Programming Project
 Departments of Medicine and Computer Science
 Stanford University, Stanford, California 94305

Abstract

The trend towards increased use of symbolic reasoning techniques for clinical decision programs stems from the dual goals of improving the performance and increasing the acceptance of such systems. This paper discusses one such program, the MYCIN System¹, and briefly describes additional areas of investigation that have grown out of the MYCIN experience. Some of the pressing basic research problems that remain are also identified.

1 Introduction

Since the early 1970's researchers at several institutions have begun to investigate potential clinical applications of symbolic reasoning techniques drawn from the branch of computer science known as artificial intelligence (AI) [25]. The term "artificial intelligence" is generally accepted to include those computer applications that involve symbolic inference rather than strictly numerical calculations. Such programs gain their power from qualitative, experiential judgments - codified in so-called "rules-of-thumb" or "heuristics", in contrast to numerical calculation programs whose power derives from the analytical equations used. The heuristics focus the attention of the reasoning program on parts of the problem that seem most critical and parts of the knowledge base that seem most relevant. They also guide the application of the domain knowledge to an individual case by deleting items from consideration as well as focusing on items. The result is that these programs pursue a line of reasoning as opposed to following a sequence of steps in a calculation.

It was perhaps inevitable that some researchers would consider applying AI methodologies to clinical medicine [20]. Major research areas in AI include knowledge representation, heuristic search, natural language understanding and generation, and models of thought processes - all topics pertinent to the problems that have been encountered by clinical decision making programs of the past [20]. Furthermore, AI researchers were beginning to look for applications to which they could apply some of the techniques they had developed in theoretical domains.

Among the medical programs using symbolic reasoning techniques are several systems that have been particularly novel and successful. Pople and Myers have developed a system called INTERNIST that assists with test selection for the diagnosis of all diseases in internal medicine [13]. This task has been remarkably successful to date, with the program correctly diagnosing a large percentage of complex cases selected from clinical pathologic conferences in the major medical journals. The program uses a hierarchic disease categorization, an ad hoc scoring system for quantifying symptom-disease relationships, plus some clever heuristics for focusing attention, discriminating between competing hypotheses, and diagnosing concurrent diseases [14].

¹Development of the MYCIN System was made possible by the SUMEX-AIM computing resource, which is supported by the Biotechnology Resources Program of the National Institutes of Health under Grant No. RR-00785. Dr. Shortliffe is the recipient of Research Career Development Award LM-00048 from the National Library of Medicine.

Reprinted by permission of the author.

Appeared in "Inference & Decision Making Processes in Medicine," Proceedings of the First International Workshop on Methodologies for Disease Control, August 25, 1979, Tokyo, Japan.

Weiss, Kulikowski, and Amarel (Rutgers University) and Safir (Mt. Sinai Hospital, New York City) have developed a model of reasoning regarding disease processes in the eye, specifically glaucoma [24]. In this specialized application area it has been possible to map relationships between observations, pathophysiologic states, and disease categories. The resulting causal associational network (termed CASNET) forms the basis for a reasoning program that gives advice regarding disease states in glaucoma patients and generates management recommendations.

For the AI researchers the question of how best to manage uncertainty in medical reasoning remains a central issue. The systems mentioned have developed ad hoc weighting systems and avoided formal statistical approaches. Others have turned to the work of statisticians and philosophers of science who have devised theories of approximate or inexact reasoning. For example, Wechsler [23] describes a program that is based upon Zaden's fuzzy set theory [25]. Our group has turned to confirmation theory for our model of clinical inexact reasoning [18].

2 The MYCIN System

Our research group has used symbolic reasoning techniques to develop a consultation program known as the MYCIN System [19]. We have cited a variety of design considerations [16] which motivated the selection of AI methodologies, or what we have chosen to call "knowledge engineering" [9]. Our primary concern was that the program be useful to physicians; we therefore emphasized the selection of a problem domain in which physicians had been shown to err frequently, namely the selection of antibiotics for patients with infections. We have also identified human issues that we feel are crucial in making the system acceptable to physicians:

- (1) it should be able to explain its decisions in terms of a line of reasoning that a physician can understand;
- (2) it should be able to justify its performance by responding to questions expressed in simple English;
- (3) it should be able to "learn" new information rapidly by interacting directly with experts;
- (4) its knowledge should be easily modified so that perceived errors can be corrected rapidly before they recur in another case; and
- (5) the interaction should be engineered with the user in mind (in terms of prompts, answers, and information volunteered by the system as well as by the users).

All these design goals were based on the observation that previous computer decision aids had generally been poorly accepted by physicians, even when they were shown to perform well on the tasks for which they were designed. We felt that barriers to acceptance were largely conceptual and could be counteracted in large part if a system were perceived as a clinical tool rather than a dogmatic replacement for the primary physician's own reasoning.

Knowledge of infectious diseases is represented in MYCIN as production rules [4], each containing a "packet" of knowledge obtained from collaborating experts [19]. A production rule is simply a conditional statement which relates observations to associated inferences that may be drawn. For example, a MYCIN rule might state that "if a bacterium is a gram positive coccus growing in chains, then it is apt to be a streptococcus."

MYCIN's power is derived from such rules in a variety of ways:

- (1) it is the program that determines which rules to use and how they should be chained together to make decisions about a specific case;
- (2) the rules can be stored in a machine-readable format but translated into English for display to physicians;

(3) by removing, altering, or adding rules, the system's knowledge structures can be rapidly modified without explicitly restructuring the entire knowledge base; and

(4) the rules themselves can often form a coherent explanation of system reasoning if the relevant ones are translated into English and displayed in response to a user's question.

Associated with all rules and inferences are numerical weights reflecting the degree of certainty associated with them. These numbers, termed certainty factors, form the basis for the system's inexact reasoning [18]. They allow the judgmental knowledge of experts to be captured in rule form and then used in a consistent fashion.

The MYCIN System has been evaluated regarding its performance at therapy selection for patients with either septicemia [26] or meningitis [27]. The program performs comparably with experts in these two task domains, but as yet it has no rules regarding the other infectious disease problem areas. Further knowledge base development will therefore be required before MYCIN is made available for clinical use; hence questions regarding its acceptability to physicians cannot yet be assessed. However, the required implementation stages have been delineated [17], attention has been paid to all the design criteria mentioned above, and the program does have a powerful explanation capability [15].

3 Recent Related Research

Over the last several years, our project's experience with MYCIN has led to several additional experiments. All of these have attempted to define the range and limitations of production rules as a methodology for knowledge engineering.

3.1 Teiresias

Randall Davis has examined techniques for selecting problem solving strategies and for acquiring knowledge interactively from human experts. His program, named Teiresias [5], has capitalized on the encoding of so-called "meta-level" knowledge (i.e., knowledge about the knowledge in the domain and how it is represented in the program) which can in turn guide efficient selection of rules and effective interactions for knowledge acquisition. Although the MYCIN rule-base was used for much of his early work, the techniques interface with production rules in a generalized fashion which has permitted Teiresias to function effectively in a variety of additional domains which are non-medical.

3.2 EMYCIN

As MYCIN matured, it became evident that the program, once stripped of its infectious disease knowledge, could provide the essential elements of a general purpose expert consultation mechanism. William vanMelle has been developing this pure or "essential" MYCIN (called EMYCIN) [22], and there are now several examples of consultation systems that are being created by developing a new set of rules to be run by the EMYCIN program. These applications have ranged from a psychopharmacology therapy advisor [10] to a program which assists with structural analysis for aircraft design [2].

3.3 CENTAUR

Jan Aikins has experimented with the implementation of new knowledge structures, termed prototypes, which make more explicit the relationships among rules in a MYCIN-like knowledge

base [1]. Similar to "frames", the prototypes define high level concepts and hypotheses which are individually invoked and which in turn guide rule selection. Early experience with this melding of production rules and prototypes suggests that the addition of prototypes both generates consultations which appear more focussed and strengthens the kinds of explanations which can be generated. The program, termed CENTAUR, has been developed in the setting of PUFF [12] (a program for pulmonary function test interpretation which is itself a MYCIN derivative), but describes general techniques which may prove adaptable to any expert system in which knowledge is largely codified in production rules.

3.4 VM

Larry Fagan's Ventilator Manager (VM) [8] deals with issues of real-time processing using production rule techniques. Clinical decision making programs have traditionally been very weak in dealing with the time relationships of disease. Such relationships can be complex and difficult to represent, but may be crucial to adequate data interpretation and patient management. In VM the task is to monitor post-surgical patients in the intensive care unit and to make recommendations regarding how best to wean them from respirators. Time-related issues include the generation of expectations and the need to measure and reinterpret clinical parameters on a recurring basis.

3.5 GUIDON

The knowledge codified in systems such as MYCIN should ideally be available for teaching as well as for advice generation. William Clancey has developed the GUIDON System [3] to demonstrate the ways in which knowledge stored in production rules can be utilized for intelligent computer-aided instruction. His work has required the consideration of a number of new issues, e.g., the need for a dynamic model of the student being taught, and the "hidden" knowledge presumed by many production rules which must somehow be made explicit if adequate student instruction is to take place.

4 Discussion

Whereas the computations used by other paradigms (e.g., Bayes' theorem, decision theory, pattern recognition, clinical algorithms) mostly involve straightforward application of well-developed computing techniques, artificial intelligence methods such as production rules are largely experimental. New approaches to knowledge representation, language understanding, heuristic search, and the other symbolic reasoning problems I have mentioned are still needed. Thus AI programs tend to be developed in research environments where short term practical results are unlikely to be found. However, out of this research are emerging techniques for coping with many of the problems encountered by the other approaches that have been applied in medical decision making domains. For example, AI researchers have developed promising methods for handling concurrent diseases [14], [24], assessing the time course of disease [8], acquiring adequate structured knowledge from experts [6], and implementing new techniques for inexact reasoning [18].

Szolovits and Pauker have recently reviewed some applications of AI to medicine and have attempted to weigh the successes of this young field against the very real problems that lie ahead [21]. They identify several deficiencies of current systems. For example, termination criteria are still poorly understood. Although INTERNIST can diagnose simultaneous diseases, it also pursues all abnormal findings to completion, even though a clinician often ignores minor abnormalities if the rest of a patient's clinical status is well understood. In addition, although some of these programs now cleverly mimic the reasoning styles observed in experts [7], [11], it is less clear how to keep the systems from abandoning one hypothesis and turning to another one as soon as new information suggests another possibility. Programs that operate this way appear to digress from one topic to another — a characteristic that decidedly alienates a user regardless of the validity of the final diagnosis or advice.

Still largely untapped is the power of an AI program to understand its own knowledge base, i.e., the structure and content of the reasoning mechanisms as well as of the medical facts. In effect, AI programs have the ability to "know what they know", the best working example of which is perhaps Teiresias [5]. Because such programs can reason about their own knowledge, they have the power to encode knowledge about strategies, e.g., when to use and when to ignore specific items of medical knowledge and which leads to follow up on. Such meta-level knowledge offers a new dimension to the design of intelligent assistant programs which we predict will be exploited in medical decision making systems of the future.

References

1. Aikins, J.S. "Prototypes and production rules: an approach to knowledge representation for hypothesis formation." Proc. of 6th Intl. Joint Conf. on Artif. Intelligence, Tokyo, Japan, August 1979.
2. Bennett, J.S. and Englemore, R.S. "SACON: A knowledge-based consultant for structural analysis." Proc. of 6th Intl. Joint Conf. on Artif. Intelligence, Tokyo, Japan, August 1979.
3. Clancey, W.J. "Tutoring rules for guiding a case method dialogue." Int. J. Man-Machine Studies, 11,25-49 (1979).
4. Davis, R. and King, J. "An overview of production systems." In Machine Representation of Knowledge (E.W. Elcock and D. Michie, eds.), New York: Wiley, 1976.
5. Davis, R. Applications of Meta-Level Knowledge to the Construction, Maintenance, and Use of Large Knowledge Bases. Heuristic Programming Project, Stanford University. Memo HPP-76-7, July 1976.
6. Davis, R. "Interactive transfer of expertise: acquisition of new inference rules." Proc. of 5th Intl. Joint Conf. on Artif. Intelligence, Cambridge, Mass., 1977.
7. Elstein, A.S., Shulman, L.S., and Sprafka, S.A. Medical Problem Solving: An Analysis of Clinical Reasoning. Cambridge, Mass.: Harvard Univ. Press, 1978.
8. Fagan, L.M., Shortliffe, E.H., and Buchanan, B.G. "Computer-based medical decision making: From MYCIN to VM." To appear in Automedica, 1979.
9. Feigenbaum, E.A. "The Art of Artificial Intelligence: Themes and case studies of knowledge engineering." AFIPS Conference Proc., NCC 1978. Vol. 47. Montvale, N.J.: AFIPS Press, 1978, p.227.
10. Heiser, J.F. and Brooks, R.E. "Design considerations for a clinical psychopharmacology advisor." Proc. 2nd Conf. on Computer Applications in Medical Care, pp. 278-285, Washington, D.C., 1978.
11. Kassirer, J.P. and Gorry, G.A. "Clinical problem solving: a behavioral analysis." Annals. Int. Med. 89, pp. 245-255 (1978).
12. Kunz, J.C., Fallat, R.J., McClung, D.H., Osborn, J.J., Votteri, B.A., Wil, H.P., Aikins, J.S., Fagan, L.M., and Feigenbaum, E.A. "A physiological rule-based system for interpreting pulmonary function test results." Memo HPP-78-19, Stanford Heuristic Programming Project, November 1978.

13. Pople, H.E., Myers, J.D. and Miller, R.A. "DIALOG: A model of diagnostic logic for internal medicine." Proc. 4th Intl. Joint. Conf. on Artif. Intelligence., MIT, Cambridge, Mass., 1975.
14. Pople, H. "The formation of composite hypotheses in diagnostic problem solving: an exercise in synthetic reasoning." Proc. of 5th Intl Joint Conf on Artif. Intelligence. Cambridge, Mass, 1977. pp. 1030-1037.
15. Scott, A.C., Clancey, W., Davis, R., and Shortliffe, E.H. "Explanation capabilities of knowledge-based production systems." Amer. J. Computational Linguistics, Microfiche 62, 1977.
16. Shortliffe, E.H., Axline, S.G., Buchanan, B.G., and Cohen, S.N. "Design considerations for a program to provide consultations in clinical therapeutics." Proc. 13th San Diego Biomedical Symposium, 311-319, San Diego, Calif., February 1974.
17. Shortliffe, E.H. and Davis, R. "Some considerations for the implementation of knowledge-based expert systems." SIGART Newsletter, No. 55, 9-12, December 1975.
18. Shortliffe, E.H., and Buchanan, B.G. "A model of inexact reasoning in medicine." Math. Biosci. 23, pp. 351-379 (1975).
19. Shortliffe, E.H. Computer-Based Medical Consultations: MYCIN. New York: Elsevier/North Holland, 1976.
20. Shortliffe, E.H., Buchanan, B.G., and Feigenbaum, E.A. "Knowledge engineering for medical decision making: a review of computer-based clinical decision aids." Proceedings of the IEEE, September 1979.
21. Szolovits, P. and Pauker, S.G. "Categorical and probabilistic reasoning in medical diagnosis." Artificial Intelligence 11, pp. 115-144 (1978).
22. vanMelle, W. "A domain-independent production-rule system for consultation programs." Proc. of 5th Intl. Joint Conf. on Artif. Intelligence, Tokyo, Japan, August 1979.
23. Wechsler, H. "A fuzzy approach to medical diagnosis." Int. J. Biomed. Comp. 7, pp. 191-203 (1976).
24. Weiss, S.M., Kulikowski, C.A., Amarel, S. and Safir, A. "A model-based method for computer-aided medical decision-making." Artificial Intelligence 11, pp. 145-172 (1978).
25. Winston, P.H. Artificial Intelligence. Reading, Mass.: Addison-Wesley, 1977.
26. Yu, V.L., Buchanan, B.G., Shortliffe, E.H., et al. "An evaluation of the performance of a computer-based consultant." Comput. Prog. Biomed. 9, pp. 95-102 (1979).
27. Yu, V.L., Fagan, L.M., Wraith, S.M., et al. "Computer-based consultation in antimicrobial

selection - a comparative evaluation by experts." J. Amer. Med. Assoc. Vol. 242, 1979
(in press).

28. Zaden, L.A. "Fuzzy sets." Information and Control 8, pp. 338-353 (1965).

**Copyright © 1985 by KSL and
Comtex Scientific Corporation**

FILMED FROM BEST AVAILABLE COPY