
Report 77-29
Stanford -- KSL

A Rule-Based Approach to the Generation of Advice and Explanations in Clinical Medicine.
Edward H. Shortliffe. Aug 1977

77-29

# A RULE-BASED APPROACH TO THE GENERATION OF ADVICE AND EXPLANATIONS IN CLINICAL MEDICINE

Edward H. Shortliffe
Department of Medicine
Massachusetts General Hospital
Boston, Massachusetts

The MYCIN system is a large computer program designed to help physicians select antibiotics for patients with septicemia or meningitis. Although the system's emphasis is on the selection of appropriate therapy for critically ill patients, it also necessarily assists with certain aspects of infectious disease diagnosis. This paper provides a brief overview of the program, describing in particular its scheme for the representation of clinical knowledge and the ways in which this representation facilitates both the generation of advice and the explanation of decisions. The encoding of knowledge in production rules, which are analyzed for advice generation by a goal-oriented rule interpreter, also permits a simple but powerful approach to natural language understanding. Despite its limitations, this approach provides an effective explanation capability without addressing many of the complex problems encountered in computational linguistics. The flow of information between a user and the MYCIN system is compared to that which occurs when a physician seeks the advice of a human infectious disease consultant.

## 1. INTRODUCTION

As medical knowledge has expanded in recent decades, it has become evident that the individual practitioner can no longer hope to acquire enough expertise to manage adequately the full range of clinical problems that will be encountered in his practice. The general practitioner has accordingly become rare, and today's primary care physicians are beginning to graduate from family practice residencies which recognize that "family doctoring" is a subspecialty in itself. Thus when a patient's problem clearly falls outside the area of the attending physician's expertise, consultations from experts in other subspecialties have become a well-accepted part of medical practice. Such consultations are acceptable to doctors in part because they maintain the primary physician's role as ultimate decision maker. The consultation generally involves a dialog between the two physicians, with the expert explaining the basis for his advice and the nonexpert seeking jus-

---

The MYCIN Project is located at Stanford University School of Medicine and is supported by BHSRE Grant No. HS01544. Much of the work described in this report was undertaken by other project members, notably A.C. Scott and W.J. Clancey, who have devoted much of their time to improvements in the general question-answerer, and R. Davis, who did most of the work on the reasoning status checker and on knowledge acquisition capabilities. Other researchers involved with the work have included J.S. Aikins, B.G. Buchanan, L.M. Fagan, and W.J. van Melle (computer scientists), V.L. Yu, S.G. Axline, and S.N. Cohen (physicians), and S.M. Wraith (clinical pharmacist). The research described has been undertaken on the SUMEX computing facility, also located at Stanford University, and supported by BRB/NIH Grant RR-00785.

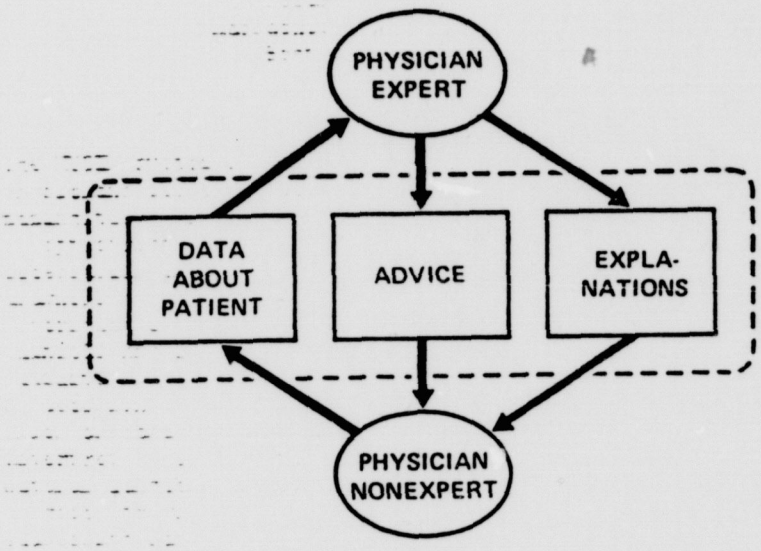


FIGURE 1 - Diagram summarizing the flow of information between physician and expert in the human consultation process.

tification of points he finds puzzling or questionable. A consultant who offered dogmatic advice he was unwilling to discuss or defend would find his opinions were seldom sought.

Fig. 1 shows a schematic view of the consultation process. The physician nonexpert gives information about his patient to the expert in response to questions and, in return, receives advice and explanations. Thus there are actually three kinds of information flow between the physician and his consultant. This paper describes a computer program, termed MYCIN, which models the consultative process by attending to all three kinds of information. It is our conviction that programs which ignore the explanation pathway will fail to be accepted by physicians because they will see in such systems too severe a departure from the human consultation process (in which the primary physician is provided with sufficient information to allow him to decide whether to follow the offered advice).

MYCIN is a LISP program designed to serve as a clinical consultant on the subject of therapy selection for patients with infections. The program may be envisioned as interposed between the expert and nonexpert in much the way that the large box is positioned in Fig. 1. The difference is that the human expert can offer only general knowledge to the program, not patient-specific decisions. The program thus becomes the decision maker, using general medical knowledge from experts to assess a specific patient and to give advice plus explanations for its judgments.

Fig. 2 details the organization of MYCIN relative to the human consultation process depicted in Fig. 1. As before, the nonexpert offers data about his patient and in return receives both advice and, when desired, information via one of two internal explanation mechanisms (the "General Question-Answerer" or the "Reasoning Status Checker"). The basis for all decisions is domain-specific knowledge acquired from experts ("Static Knowledge"). A group of computer programs (the "Rule Interpreter") uses this knowledge, and data about the specific patient, to

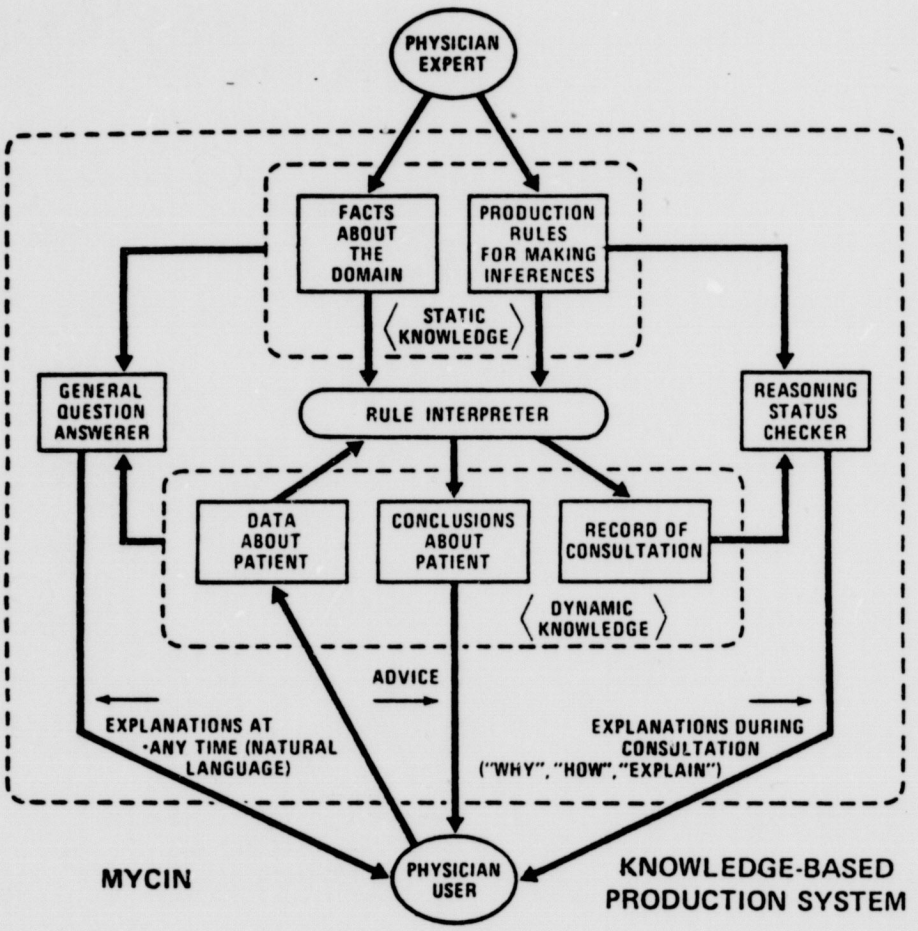


FIGURE 2 - Diagram summarizing the organization and flow of information within MYCIN. The correlation between this design and the human consultation process depicted in Fig. 1 is discussed in the text.

generate conclusions and, in turn, therapeutic advice. It simultaneously keeps a record of what has happened, and this record is available to the explanation routines if the physician asks for justification or clarification of some conclusion the program has reached. In the remainder of this paper some details of each of these system components will be presented. More extensive discussions of this material are also available [1-4].

EDWARD H. SHORTLIFFE

## 2. KNOWLEDGE REPRESENTATION

### 2.1 Static Knowledge

Static knowledge refers to all data that are constant in the program and unchanging from one consultation to the next.

#### 2.1.1 Facts About The Domain

Much of the knowledge MYCIN requires is simple statements of fact about the domain. These can generally be represented as attribute-object-value triples, or as predicate statements. For example:

| | | |
|---|---|---|
| | (GRAMSTAIN E.COLI GRAMNEG) | i.e., the gramstain of e.coli is gram negative |
| ::: | (STERILESITE BLOOD T) | i.e., the blood is normally sterile |
| ::: | (STERILESITE MOUTH NIL) | i.e., the mouth is not normally sterile |

#### 2.1.2 Production Rules

In addition to simple facts, MYCIN requires judgmental knowledge acquired from experts and available for use in analyzing a new patient. Judgmental knowledge in MYCIN is expressed as production rules [5] which define certain preconditions (the PREMISE) that allow a conclusion to be reached (the ACTION) with a specified degree of confidence (the "certainty factor" [6]). Although such rules are stored as LISP list structures, a series of routines is available for translating them into English. For example:

PREMISE: If the stain of the organism is gramneg, and
the morphology of the organism is rod, and
the aerobicity of the organism is anaerobic,
ACTION: Then there is suggestive evidence (.7) that the
identity of the organism is bacteroides.

Note that the purpose of this rule is determination of organism identity. Rules are classified and accessed in accordance with their purpose as described below.

### 2.2 Dynamic Knowledge

Dynamic knowledge refers to all data that are variable and change from one run of the program to the next.

#### 2.2.1 Data About The Patient - Acquired From The User

MYCIN asks questions of the user, driven by a reasoning algorithm described below. These questions generally ask the user to fill in the "value" in an attribute-object-value triple (e.g., "what is the patient's name?"), or to give the truth value of a predicate (e.g., "is the patient a compromised host?"). Thus these data may be represented, once acquired, in precisely the way that facts about the domain are represented in the static knowledge base (see 2.1.1).

#### 2.2.2 Data About The Patient - Generated By The Program

When the preconditions in the PREMISE of a rule are found to hold, MYCIN executes the ACTION portion of the rule and generates a new "fact" which can, once again, be represented as an attribute-object-value triple. As discussed in 2.1.2, conclusions may also have a confidence value associated with them, thereby requiring that the triple be expanded to a quadruple:

| | |
|---|---|
| (IDENTITY ORGANISM-1 BACTEROIDES .7) | i.e., the identity of ORGANISM-1 is bacteroides, with certainty factor of .7 |

Premises may be similarly expanded. Furthermore, by generalizing this scheme to include representation of data acquired from the user, the physician may be asked to express his confidence in the answer he gives when MYCIN asks a question.

#### 2.2.3 Maintenance Of A Record Of The Consultation

A history of the consulta-

tion is the third variety of dynamic knowledge. The details of representation need not be described here, but these data include records of which rules succeeded, which rules were tried but failed, how specific decisions were made, how information was used, and why questions were asked.

## 3. THE PRODUCTION SYSTEM

### 3.1 The Rule Interpreter

This series of routines analyzes rules in the static knowledge base, determines whether they apply to the patient under consideration, and if so draws the conclusions delineated in the ACTION portions of the rules. This process would quickly become unmanageable as system knowledge grew if there were not a mechanism for selecting only the most relevant rules for a given patient. This is accomplished by a goal-oriented approach described in detail elsewhere [1,7]. Briefly, as the rule interpreter examines the PREMISE of a rule, it notes whether the relevant data needed to determine the truth of each precondition are already known. If not, it digresses to examine those rules which make conclusions about the data needed by the first rule. The PREMISE conditions of those rules may, in turn, invoke additional rules, and in this way a reasoning network relevant to the first rule is formed. As described in 2.1.2, since rules are classified according to their purpose, it is easy to identify all rules which may aid in determining the truth of a specific precondition. The entire process is initiated by invoking a specific "Goal Rule" which defines MYCIN's task and is the only rule necessarily invoked for every consultation. When MYCIN can find no rules for determining the truth of a precondition, it asks the user for the relevant data. If the physician does not know the information either, the invoking rule is simply ignored.

### 3.2 Maintenance Of Initiative In The Hands Of The Physician

As was discussed above, a physician is not likely to accept a system such as MYCIN if the program simply asks a series of questions and then presents a piece of dogmatic advice as it terminates execution. The production system has therefore been provided with a series of "interrupts" that allow the physician to digress with questions of his own or to demand justification for the line of questioning on which MYCIN has embarked during the consultation. Whenever the program asks a question, the user can temporarily refuse to answer and instead call on the explanation capabilities described in the next section.

## 4. EXPLANATIONS

### 4.1 The Reasoning Status Checker (RSC)

This component of the explanation system deals with most questions that arise during the consultation itself. Because the context of current reasoning about the patient is well-defined, the physician can be given a great deal of information on the basis of a few simple commands that do not require natural language processing. These commands are briefly described below; more extensive details of their implementation are also available [3,4,7]. As shown in Fig. 2, the reasoning status checker (RSC) uses only the knowledge base of rules and the current record of the consultation; the general question-answerer described in 4.2, on the other hand, has access to all static and dynamic knowledge.

4.1.1 The WHY Command Whenever MYCIN asks a question, the physician may prefer not to answer initially and instead to inquire about the reasoning underlying the questioning. Thus he may simply respond with the command WHY (i.e., "Why do you think the information you are requesting may be useful?"). Since all questions MYCIN asks are generated by rules, and since the rules are selected according to their purpose as discussed in 3.1, an English translation of the rule under consideration generally serves as an adequate response to the WHY query. The RSC

therefor responds by displaying the current rule. In addition, it places an identifying number before each of the preconditions in the PREMISE and indicates whether the condition is (a) already known to be true, or (b) still under investigation (note that one of the latter group of preconditions will have generated MYCIN's current question to the user). The physician can in turn inquire why the displayed rule was selected by asking WHY a second time, and the RSC will accordingly display the next rule in the reasoning network (see 3.1).

An experienced user will learn to seek higher level explanations by entering WHY followed by a modifying number. The number tells the RSC the relative size of the reasoning leap (roughly the number of chained rules) to compile into a single explanation. If the physician finds the leap is in fact too large, the command EXPLAIN provides a breakdown of the component reasoning steps.

4.1.2 The HOW Command As discussed in the preceding section, when MYCIN displays a rule in response to the WHY command, it labels each precondition in the PREMISE with a unique number. The physician may then respond to the displayed explanation by entering HOW followed by one of the identifying labels. If the reference condition is one that MYCIN has already concluded to be true, the RSC assumes that the physician is asking "HOW did you decide that the specified precondition is true?" and answers by citing the relevant rules that it used to make the decision. If, on the other hand, the cited condition has not yet been fully investigated, MYCIN assumes the physician is asking "HOW will you decide if the specified precondition is true?" and responds by citing the rules it intends to try, only some of which may actually succeed.

4.2 The General Question-Answerer (GQA)

The general question-answerer (GQA) is a more comprehensive explanation system which, at any time during or after the consultation session, has full access to all static and dynamic knowledge in MYCIN (Fig. 2). Since it cannot make simple assumptions based on context, as the RSC can do, and since it is important to keep the system useable by a novice, the GQA must accept and answer questions expressed in natural language. MYCIN's rule-based knowledge representation scheme, and some techniques borrowed from early work in computational linguistics [8-10], permit a straightforward but powerful approach to interpreting simple English questions without contending with several of the complex problems of natural language understanding. The details of this approach have been described elsewhere [4,11].

4.2.1 Questions About Static Knowledge The ability to retrieve information from the static knowledge base gives the GQA a tutorial capability. Since the static knowledge is acquired from experts, the GQA can essentially act as an intermediary between an expert and a physician seeking general information about the infectious disease field. The user might ask simple questions of fact (eg., "Which culture sites are normally considered sterile?") or questions regarding judgments stored in rules. Questions of the second variety are called "rule-retrieval" questions because they may be answered simply by identifying and displaying English versions of relevant rules from the knowledge base. Retrieval may be keyed to the rule PREMISE (eg., "How do you use the gram stain of an organism?"), the ACTION (eg., "When do you decide an organism might be a streptococcus?"), or to both the PREMISE and ACTION (eg., "Do you ever use the morphology of an organism to determine its identity?"). Furthermore, a question may deal with a specific rule (eg., "What is rule037?"). Note that none of these questions refers to a specific consultation and thus requires no access to the dynamic knowledge base (Fig. 2).

4.2.2 Questions About Dynamic Knowledge Although the RSC permits inquiries about the dynamic knowledge base, its scope is limited by the context of the current question being asked by MYCIN (see 4.1). If the physician wishes to ask more general questions regarding the status of MYCIN's reasoning, or if he wishes

to review the program's decisions after the consultation is complete and MYCIN is no longer questioning him, the GQA gives him free access to all information about the specific consultation. Once again, the user might ask simple questions of fact (eg., "From what site was CULTURE-2 obtained?") or questions regarding the basis for MYCIN's judgments. The second variety is again a "rule-retrieval" question, but is keyed to the consultation record in dynamic data rather than to the knowledge base of rules in static data. Thus questions may again reference the PREMISE (eg., "How did you use the gram stain of ORGANISM-1?"), the ACTION (eg., "What makes you think ORGANISM-2 might be a streptococcus?"), or both (eg., "Did you use the morphology of ORGANISM-1 to determine its identity?"). Note that these questions parallel the examples in 4.2.1 but that they are consultation-specific and thus request the retrieval not of all relevant rules, but only those that were actually used successfully in the specified context. Finally, one may again wish to ask about a specific rule (eg., "Did you use rule037 when considering ORGANISM-1?").

## 5. KNOWLEDGE ACQUSITION

The only component of Fig. 2 not yet discussed is the crucial step of acquiring domain-specific knowledge from experts and coding it for storage in the static knowledge base. When MYCIN was first being developed, such knowledge was acquired by extensive meetings during which infectious disease experts and computer scientists discussed specific patients and attempted to analyze and extract the individual facts and rules they were utilizing. Recently extensive work has been devoted to the problem of automating the knowledge acquisition process in sessions involving clinical experts interacting with MYCIN directly. This problem has been thoroughly explored in another publication [3] and will not be discussed further here.

## 6. CONCLUSIONS

A rule-based expert system is described which uses artificial intelligence techniques, and a model of the interaction between physicians and human consultants, to attempt to satisfy the demands of a user community that is often reluctant to experiment with computer technology. An ability to explain decisions, and thus to respond to simple questions expressed in natural language, is emphasized. The representation of expert knowledge in production rules facilitates greatly the generation of explanations without requiring solutions to several of the complex problems encountered in computational linguistics.

## References

[1] Shortliffe, E.H. Computer-Based Medical Consultations: MYCIN. Elsevier/North Holland, New York, 1976.

[2] Davis, R., Buchanan, B.G., and Shortliffe, E.H. Production rules as an approach to knowledge-based consultation systems. Artificial Intelligence 8, 15-45 (1977).

[3] Davis, R. Applications Of Meta Level Knowledge To The Construction, Maintenance, And Use Of Large Knowledge Bases. Doctoral dissertation, Stanford University, June 1976. Also available as AIM-283, Stanford Artificial Intelligence Laboratory, Stanford University, July 1976.

[4] Scott, A.C., Clancey, W.J., Davis, R., and Shortliffe, E.H. Explanation capabilities of knowledge-based production systems. Amer. J. Computational

Linguistics, Microfiche 62, 1977. Also available as TR HPP-77-1, Heuristic Programming Project, Stanford University, March 1977.

[5] Davis, R., and King, J. An overview of production systems. To appear in Machine Representations of Knowledge (Machine Intelligence 8, eds. E.W. Elcock and D. Michie), John Wylie, New York, 1976. Also available as AIM-271, Stanford Artificial Intelligence Laboratory, Stanford University, October 1975.

[6] Shortliffe, E.H., and Buchanan, B.G. A model of inexact reasoning in medicine Math. Biosci. 23,351-379 (1975).

[7] Shortliffe, E.H., Davis, R., Buchanan, B.G., Axline, S.G., Green, C.C., and Cohen, S.N. Computer-based consultations in clinical therapeutics: explanation and rule-acquisition capabilities of the MYCIN system. Comput. Biomed. Res. 8,303-320 (1975).

[8] Green, B.F., Wolf, A.K., Chomsky, C., and Laughery, K. BASEBALL: an automatic question-answerer (1961). In Computers and Thought (E.A. Feigenbaum and J. Feldman, eds.), pp.207-216, McGraw Hill, San Francisco, 1963.

[9] Colby, K., Weber, S., and Hilf, F. Artificial paranoia. Artificial Intelligence 2,1-25 (1971).

[10] Quillian, M.R. Semantic memory. In Semantic Information Processing (M. Minsky, ed.), pp.227-270, MIT Press, Cambridge, MA. (1968).

[11] Shortliffe, E.H. MYCIN: A Rule-Based Computer Program For Advising Physicians Regarding Antimicrobial Therapy Selection. Doctoral dissertation, Stanford University, October 1974. Also available as AIM-251, Stanford Artificial Intelligence Laboratory, Stanford University, November 1974.