4

The Structure of the MYCIN System

William van Melle

A number of constraints influenced the design of the MYCIN system. In order to be useful, the system had to be easy to use and had to provide consistently reliable advice. It needed to be able to accommodate the large body of task-specific knowledge required for high performance, a knowledge base that is subject to change over time. The system also had to be able to use inexact or incomplete information. This applies not only to the absence of definitive laboratory data, but also to the medical domain itself (which is characterized by much judgmental knowledge). Finally, to be a useful interactive system, MYCIN needed to be capable of supplying explanations for its decisions and responding to physicians' questions, rather than simply printing orders.

The MYCIN system comprises three major subprograms, as depicted in Figure 4-1. The *Consultation Program* is the core of the system; it interacts with the physician to obtain information about the patient, generating diagnoses and therapy recommendations. The *Explanation Program* provides explanations and justifications for the program's actions. The *Knowledge-Acquisition Program* is used by experts to update the system's knowledge base.

This chapter is based on an article originally appearing under the title "MYCIN: A knowledge-based consultation program for infectious disease diagnosis" in *International Journal of Man-Machine Studies* 10: 313–322 (1978).

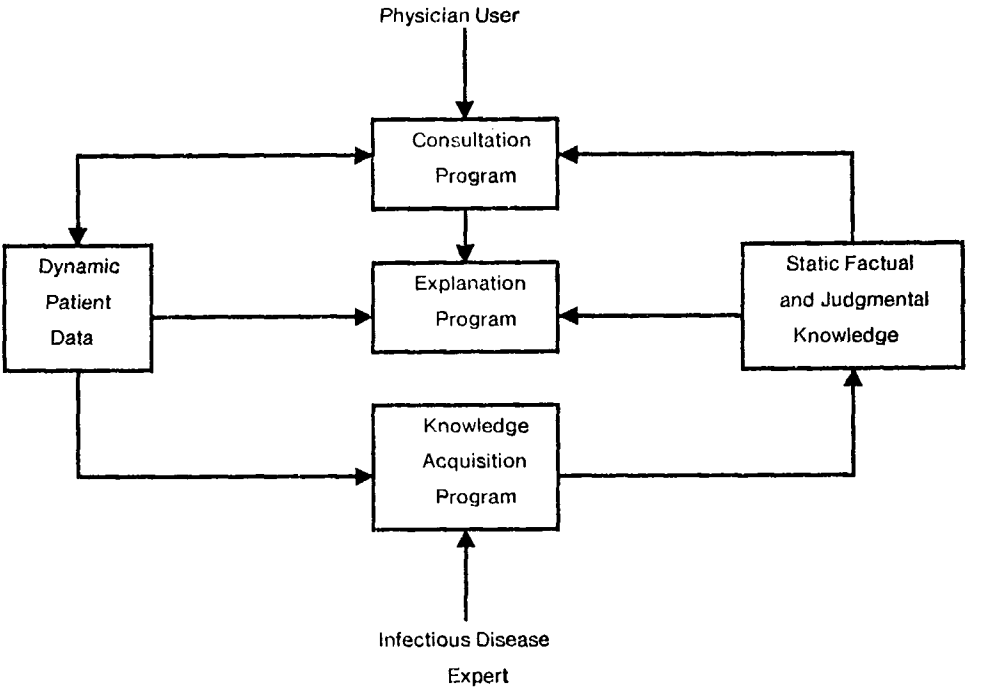

FIGURE 4-1 Organization of the MYCIN system. Arrows denote information flow between modules, knowledge base, and users.

4.1 The Consultation Program

Excerpts from a sample consultation are shown in Figure 4-2. (A complete annotated typescript is shown in the Appendix.) The doctor's responses are in upper-case letters following the double-asterisk prompt. Considerable effort has gone into the human engineering aspects of the system, to make the interaction as smooth as possible. If the user finds a question unclear, he or she can request a rephrasing with a list of sample responses. A large dictionary of synonyms allows the user great latitude in choice of responses, and simple typing or spelling errors are automatically corrected.

4.1.1 Knowledge Organization

Information about the patient and conclusions drawn during the consultation are represented as associative (object-attribute-value) triples. The

1) Patient's name: (first-last)
** **FRED BRAUN**
2) Sex:
** **M**
3) Age:
** **55**
4) Are there any cultures for Fred Braun which may be related to the present illness, and from which organisms have been grown successfully in the microbiology laboratory?
** **Y**

(Questions 1-3 request background patient data.)

--------CULTURE-1--------
5) From what site was the specimen for CULTURE-1 taken?
** **BLOD**
=**BLOOD**
6) Please give the date and time when CULTURE-1 was obtained. (mo/da/yr time)
** **JUN 20, 1977**
The first organism isolated from the blood culture of 20-JUN-77 (CULTURE-1) will be referred to as:

(System corrects typing error.)

--------ORGANISM-1--------
7) Enter the laboratory-reported identity of ORGANISM-1:
**UNKNOWN
8) The stain (gram or Ziehl-Neelsen acid-fast) of ORGANISM-1:
** **NEG**
9) Is ORGANISM-1 rod or coccus (etc.):
** **ROD**
10) What is the form of the individual organisms (e.g. lancet-shaped for cocci, fusiform for rods, etc.)?
** **FUSIFORM**

(Typically identity is not yet known.)

(Preliminary lab results give some clues.)

(. . . more questions follow in order to gather sufficient information to infer the identity and significance of the infecting organisms . . .)

23) What is the suspected portal of entry of the gramneg rod from the blood culture of 20-JUN-77 (ORGANISM-1) into this sterile site?
** **GI (6)**
24) Is Fred Braun a burn patient?
** **N**
25) Is Fred Braun a compromised host (e.g. alcoholic, sickle-cell-disease, immunosuppressed)?
** **N**
26) Is Fred Braun febrile?
** **Y**

(User indicates the uncertainty of his response.)

(. . . after about 50–60 questions, MYCIN prints the diagnostic hypotheses on which therapy will be based . . .)

My therapy recommendations will be designed to treat for organisms that are either 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 BACTEREMIA

<ITEM-1> E.COLI [ORGANISM-1]
<ITEM-2> KLEBSIELLA [ORGANISM-1]
<ITEM-3> ENTEROBACTER [ORGANISM-1]
<ITEM-4> KLEBSIELLA-PNEUMONIAE [ORGANISM-1]

FIGURE 4-2 Excerpts from a MYCIN consultation. (Comments in italics are not part of the actual interaction.)

(. . . questions follow to evaluate possible therapy choices, and finally MYCIN prints its therapy recommendations . . .)

[REC-1] My preferred therapy recommendation is as follows:
In order to cover for items <1 2 3 4>:
Give: GENTAMICIN *(One drug covers 4 possible identities.)*
Dose: 119 mg (6.0 ml) q8h IV for 10 days [calculated on basis of 1.7 mg/kg]
Comments: Modify dose in renal failure.

FIGURE 4-2 continued

objects, known as *contexts* in MYCIN, are such things as individual cultures taken from the patient, organisms that grew out of them, and drugs the patient is currently receiving. Various attributes, termed *clinical parameters*, characterize these objects. Questions asked during the consultation attempt to fill in the values for relevant attributes of these objects. To represent the uncertainty of data or competing hypotheses, attached to each triple is a *certainty factor* (CF), a number between -1 and 1 indicating the strength of the belief in (or a measure of the importance of) that fact. A CF of 1 represents total certainty of the truth of the fact, while a CF of -1 represents certainty regarding the negation of the fact. While certainty factors are *not* conditional probabilities, they are informally based on probability theory (see Part Four). Some triples (with CF's) from a typical consultation might be as follows:

```
(IDENTITY ORGANISM-1 PSEUDOMONAS 0.8)
(IDENTITY ORGANISM-1 E. COLI 0.15)
(SITE CULTURE-2 THROAT 1.0)
(BURNED PATIENT-298 YES -1.0)
```

Here ORGANISM-1 is probably *Pseudomonas*, but there is some evidence to believe it is *E. coli*; the site of CULTURE-2 is (without doubt) the throat; and PATIENT-298 is known *not* to be a burn patient.

4.1.2 Production Rules

MYCIN reasons about its domain using judgmental knowledge encoded as production rules. Each rule has a *premise*, which is a conjunction of predicates regarding triples in the knowledge base. If the premise is true, the conclusion in the *action* part of the rule is drawn. If the premise is known with less than certainty, the strength of the conclusion is modified accordingly.

A typical rule is shown in Figure 4-3. The predicates (such as SAME) are simple LISP functions operating on associative triples, which match the declared facts in the premise clause of the rule against the dynamic data known so far about the patient. $AND, the multi-valued analogue of

```
RULE035

PREMISE:  ($AND (SAME CNTXT GRAM GRAMNEG)
                (SAME CNTXT MORPH ROD)
                (SAME CNTXT AIR ANAEROBIC))
ACTION:   (CONCLUDE CNTXT IDENTITY BACTEROIDES TALLY .6)

IF:  1) The gram stain of the organism is gramneg, and
     2) The morphology of the organism is rod, and
     3) The aerobicity of the organism is anaerobic
THEN: There is suggestive evidence (.6) that the identity
      of the organism is bacteroides
```

FIGURE 4-3 A MYCIN rule, in both its internal (LISP) form and English translation. The term CNTXT appearing in every clause is a variable in MYCIN that is bound to the current context, in this case a specific organism (ORGANISM-2), to which the rule may be applied.

the Boolean AND function, performs a minimization operation on CF's. The body of the rule is actually an executable piece of LISP code, and "evaluating" a rule entails little more than the LISP function EVAL. However, the highly stylized nature of the rules permits the system to examine and manipulate them, enabling many of the system's capabilities discussed below. One of these is the ability to produce an English translation of the LISP rule, as shown in the example. This is possible because each of the predicate functions has associated with it a translation pattern indicating the logical roles of the function's arguments.

It is intended that each rule be a single, modular chunk of medical knowledge. The number of rules in the MYCIN system grew to about 500.

4.1.3 Application of Rules—The Rule Interpreter

The control structure is a goal-directed backward chaining of rules. At any given time, MYCIN is working to establish the value of some clinical parameter. To this end, the system retrieves the (precomputed) list of rules whose conclusions bear on this goal. The rule in Figure 4-3, for example, would be retrieved in the attempt to establish the identity of an organism. If, in the course of evaluating the premise of one of these rules, some other piece of information that is not yet known is needed, MYCIN sets up a subgoal to find out that information; this in turn causes other rules to be tried. Questions are asked during the consultation when rules fail to deduce the necessary information. If the user cannot supply the requested information, the rule is simply ignored. This control structure results in a highly focused search through the rule base.

4.1.4 Advantages of the Rule Methodology

The modularity of rules simplifies the task of updating the knowledge base. Individual rules can be added, deleted, or modified without drastically affecting the overall performance of the system. And because each rule is a coherent chunk of knowledge, it is a convenient unit for explanation purposes. For example, to explain why the system is asking a question during the consultation, a first approximation is simply to display the rule currently under consideration.

The stylized nature of the rules is useful for many operations. While the syntax of the rules permits the use of any LISP function, there is a small set of standard predicates that make up the vast majority of the rules. The system contains information about the use of these predicates in the form of function *templates*. For example, the predicate SAME is described as follows:

function template:	(SAME CNTXT PARM VALUE)
sample function call:	(SAME CNTXT SITE BLOOD)

The system can use these templates to "read" its own rules. For example, the template shown here contains the standard tokens CNTXT, PARM, and VALUE (for context, parameter, and corresponding value), indicating the components of the associative triple that SAME tests. If the clause above appears in the premise of a given rule, the system can determine that the rule needs to know the site of the culture, and that the rule can only succeed if that site is, in fact, blood. When asked to display rules that are relevant to blood cultures, MYCIN will be able to choose that rule.

An important function of the templates is to permit MYCIN to precompute automatically (at system generation time) the set of rules that conclude about a particular parameter; it is this set that the rule monitor retrieves when the system needs to deduce the value of that parameter.

The system can also read rules to eliminate obviously inappropriate ones. It is often the case that, of a large set of rules under consideration, several are provably false by information already known. That is, the information needed to evaluate one of the clauses in the premise has already been determined, and that clause is false, thereby making the entire premise false. By reading the rules before actually invoking them, many can be immediately discarded, thereby avoiding the deductive work necessary in evaluating the premise clauses that precede the false one (this is called the *preview mechanism*). In some cases this means the system avoids the useless search of one or more subgoal trees, when the information thereby deduced would simply be overridden by the demonstrably false premise.

Another more dramatic case occurs when it is possible, on the basis of information currently available, to deduce with certainty the value of some parameter that is needed by a rule. This is the case when there exists a

chain of one or more rules whose premises are known (or provable, as above) with certainty and that ultimately conclude the desired value with certainty. Since each rule in this chain must have a certainty factor of 1.0, we term such a chain a *unity path*; and since a value known with certainty excludes all other potential values, no other rules need be tried. MYCIN always seeks a unity path before trying a set of rules or asking a question; typically, this means "commonsense" deductions are made directly, without asking the user "silly" questions or blindly invoking all the rules pertaining to the goal. Since there are usually few rules on any potential unity path, the search tends to be small.

The ability to read rules opens the way to the writing of rules that manipulate other rules. We term such rules *meta-rules* (see Part Nine); they are used to make deductions not about the medical entities of the domain but about strategies to be used by the system. Whenever the rule interpreter is about to invoke a list of rules to establish some goal, it first applies any meta-rules associated with that goal. These meta-rules can reorder or prune the rule list to make the search more suitable for the given case.

4.2 Explanation Capability

A major subprogram of MYCIN is a general *Question-Answering* (QA) *Module*, which answers simple English-language questions concerning the system's decisions in a particular consultation or about the system's general knowledge. A limited set of commonly desired explanations is also provided in a command style during the consultation by the *Reasoning Status Checker*. This special module explains the system's current line of reasoning and motivation for the questions asked. (A more detailed study of explanation in MYCIN is included in Part Six of this book.) The QA Module is entered automatically at the end of each consultation and is also available to the user during the consultation if the need arises. The user is encouraged to ask questions about any aspect of MYCIN's performance.

Because of the relatively unambiguous vocabulary used in the medical domain, and because of our desire to minimize the delay in answering questions, we have forgone any complicated parsing techniques in the QA Module. Rather, we use a simple pattern-matching approach, relying on the strong implications of certain keywords in the user's questions. The words in the translation of a clinical parameter, and values of that parameter, are tagged in the system dictionary as implying that parameter. By suitably weighing the implications of each word in the sentence, the system determines one or more parameter/value pairs being asked about. Pattern matching is used to determine the general type of the question, which indicates what is being asked about those parameter/value pairs. Typical

**** IS ORGANISM-1 A STREPTOCOCCUS?**
No. The identity of ORGANISM-1 is e.coli (.74) bacteroides (.7) klebsiella-pneumoniae (.4) pseudomonas-aeruginosa (.22) enterobacter (.21) proteus-non-mirabilis (.19) serratia (.03) salmonella (.02)

****WHY DID YOU RULE OUT STREPTOCOCCUS AS A POSSIBILITY FOR ORGANISM-1?**
RULE033 could have been used to determine that the identity of ORGANISM-1 is streptococcus. However, clause 2 of RULE033 ["the morphology of the organism is coccus"] was already known to be false for ORGANISM-1, so the rule was never tried.

**** HOW DO YOU KNOW THAT CULTURE-1 WAS FROM A STERILE SOURCE?**
I used RULE189 to conclude that this blood culture was taken from a sterile source. This gave a cumulative CF of (1.0). The last question asked before the conclusion was made was 7.

**** WHAT DOES GU MANIPULATION TELL YOU ABOUT THE IDENTITY OF AN ORGANISM?**
The rules listed below use:
whether the patient has had a genito-urinary manipulative procedure
to conclude about:
the identity of the organism

156, 163, 190
Which do you wish to see?
**** 156**

RULE156
IF: 1) The site of the culture is blood, and
2) The gram stain of the organism is gramneg, and
3) The morphology of the organism is rod, and
4) The portal of entry of the organism is urine, and
5) The patient has not had a genito-urinary manipulative procedure, and
6) Cystitis is not a problem for which the patient has been treated
THEN: *There is suggestive evidence (.6) that the identity of the* organism is e.coli

**** IS BLOOD A STERILE SITE?**
Yes. Blood is one of those sites that is normally sterile.

**** WHAT ARE THE NONSTERILE SITES?**
Those sites that are normally nonsterile are: CERVIX, CUTANEOUS-ULCER, LOCHIA, NOSE, SKIN, STOOL, THROAT, URETHRA, VAGINA.

**** IS BACTEROIDES AEROBIC?**
No, the aerobicity of bacteroides is anaerobic.

FIGURE 4-4 Interaction with the QA Module. The user's questions follow the double-asterisk prompt.

questions (see Figure 4-4) seek the value of a certain parameter, how this value was concluded, and how this parameter was used. Questions about the value of a given parameter of an object are answered by simply displaying the current hypothesis regarding it. To explain how the value of a parameter was concluded, MYCIN retrieves the list of rules that were successfully applied and prints them, along with the conclusions drawn. Alternatively, if the user supplied the value by answering a previous question, this is noted. More general questions about how information is used or concluded are answered by retrieving the relevant rules from the rule base.

34) Was penicillinase added to this blood culture (CULTURE-1)?
**** WHY**

[i.e. WHY is it important to determine whether penicillinase was added to CULTURE-1?]

[3.0] This will aid in determining whether ORGANISM-1 is a contaminant.
It has already been established that
[3.1] the site of CULTURE-1 is blood, and
[3.2] the gram stain of ORGANISM-1 is grampos
Therefore, if
[3.3] penicillinase was added to this blood culture
then
there is weakly suggestive evidence (.4) that ORGANISM-1 is a contaminant [RULE039]

**** WHY**

[i.e. WHY is it important to determine whether ORGANISM-1 is a contaminant?]

[4.0] . . . In order to determine whether there is significant disease associated with this occurrence of ORGANISM-1.
It has already been established that
[4.1] this blood culture was taken from a sterile source
Therefore, if
[4.2] ORGANISM-1 is not a contaminant
then
there is strongly suggestive evidence (.9) that there is significant disease associated with this occurrence of ORGANISM-1
[RULE032]

**** WHY**

[i.e. WHY is it important to determine whether there is significant disease associated with this occurrence of ORGANISM-1]

[5.0] . . . In order to determine the organisms which require therapy.
It has already been established that
[5.1] the infection is known, and
[5.2] the identity of ORGANISM-1 is known
Therefore, if
[5.3] there is significant disease associated with this occurrence of ORGANISM-1
then
it is definite (1.0) that ORGANISM-1 is one of the organisms which require therapy
[RULE090]

**** HOW 5.2**

[i.e. HOW was it established that the identity of ORGANISM-1 is known?]

Question 9 was asked, but you weren't sure; so the following was used:

[5.5] RULE084 was used to conclude that the identity of ORGANISM-1 is streptococcus-group-a (.29).

FIGURE 4-5 Use of the Reasoning Status Checker during the consultation to explain MYCIN's line of reasoning.

As shown in Figure 4-5, the Reasoning Status Checker is invoked by the HOW and WHY commands. At any time during the consultation, when

the user is asked a question, he or she can delay answering it and instead ask why the question was asked. Since questions are asked in order to establish the truth of the premise of some rule, a simple answer to WHY is "because I'm trying to apply the following rule." Successive WHY questions unwind the chain of subgoals, citing the rules that led to the current rule being tried.

Besides examining the current line of reasoning, the user can also ask about previous decisions, or about how future decisions might be made, by giving the HOW command. Explaining how the truth of a certain clause was established is accomplished as described above for the general QA Module. To explain how a presently unknown clause might be established, MYCIN retrieves the set of rules that the rule interpreter would select to establish that clause and selects the relevant rules from among them by "reading" the premises for applicability and the conclusions for relevance to the goal.

4.3 Knowledge Acquisition

The knowledge base is expanded and improved by acquiring new rules, or modifications to old rules, from experts. Ordinarily, this process involves having the medical expert supply a piece of medical knowledge in English, which a system programmer converts into the intended LISP rule. This mode of operation is suitable when the expert and the skilled programmer can work together. Ideally, however, the expert should be able to convey his or her knowledge directly to the system.

Work has been undertaken (see Part Three) to allow experts to update the rule base directly. A rule-acquisition routine parses an English-language rule by methods similar to those used in parsing questions in the QA Module. Each clause is broken down into one or more object-attribute-value triples, which are fitted into the slots of the appropriate predicate function template. This process is further guided by *rule models* (see Chapter 28), which supply expectations about the structure of rules and the interrelationships of the clinical parameters.

One mode of acquisition that has received special attention is acquiring new rules in the context of an error. In this case, the user is trying to correct a localized deficiency in the rule base; if a new rule is to correct the program's faulty behavior, it must at the very least apply to the consultation at hand. In particular, each of the premises must evaluate to TRUE for the given case. These expectations greatly simplify the task of the acquisition program, and also aid the expert in formulating new rules.

One difficult aspect of rule acquisition is the actual formulation of medical knowledge into decision rules. Our desire to keep the rule format

simple is occasionally at odds with the need to encode the many aspects of medical decision making. The backward chaining of rules by the deductive system is also often a stumbling block for experts who are new to the system. However, they soon learn to structure their knowledge appropriately. In fact, some experts have felt that encoding their knowledge into rules has helped them formalize their own view of the domain, leading to greater consistency in their decisions.