Chapter 7

Conclusion

7.1 Summary

MYCIN is a large computer system developed (to the extent described in this text) over a two year period. The program's knowledge-base, and many aspects of system design, were contributed by collaborating physicians and computer scientists who met regularly throughout the two years. The project has recently expanded to include additional physicians and computer scientists who are contributing full-time efforts to the future expansion of MYCIN's capabilities.

This chapter summarizes the material that has been presented in this text. In this section, I reiterate the clinical problem for which MYCIN is designed to offer advice. I then briefly review how the program attempts to solve the problem. Section 7.2 then summarizes the limitations of the MYCIN approach. Section 7.3 discusses MYCIN's contribution to computer-based medical decision making, and I conclude in § 7.4 with consideration of the program's contribution to the field of AI.

7.1.1 THE CLINICAL PROBLEM

The principal goal of the MYCIN project has been to devise a computer-based system for assisting and educating physicians who need advice about appropriate antimicrobial therapy. The basis of rational infectious disease therapy is identification of the offending micro-organisms. Accurate identification is important because drugs that are highly effective against certain bacteria are often useless against others. The patient's clinical status and history, including such information as previous infections and treatment, provide valuable data to assist the physician with the identification task. However, bacteriological cultures that use specimens taken from the site of the patient's infection usually provide the most definitive identifying information.

Initial culture reports from a microbiology laboratory may become available within 12 hours from the time a clinical specimen is obtained from the patient. The information in these early reports often serves to classify the organism in general terms but does not permit precise identification. It may be clinically unwise to postpone therapy until identification of the infecting organism can be made with certainty, however–a process that usually requires 24 to 48 hours or longer. Thus, it is often necessary for the physician to estimate the range of possible organisms and to start appropriate treatment even before the laboratory is able to identify the offending organism and its antibiotic sensitivities.

As discussed in § 1.5.2, there is ample evidence that physicians often do not choose antimicrobial therapy wisely. Studies discussed in that chapter have shown that physicians may reach therapeutic decisions that differ significantly from those that would have been suggested by infectious disease experts. It is not uncommon for physicians to treat patients for whom experts believe no antimicrobial therapy is indicated. Furthermore, nonexperts sometimes choose a drug regimen designed to cover for all possibilities, prescribing either several drugs or one of the so-called "broad-spectrum" antibiotics, even though appropriate utilization of clinical clues might have led to a more rational (and often less toxic) therapy. Since professional resources are often overburdened in today's hospitals, a computer-based system that could serve effectively in a consultation role to the nonexpert–and gain his respect–would be highly useful. MYCIN has been designed to provide readily accessible advice and instruction that will help bridge this gap between practicing physicians and experts in infectious disease therapy.

MYCIN has also been developed with an awareness of the current lack of acceptance of computer-assisted decision making by the medical profession. We have attempted to analyze the reasons for the common opposition to such programs and to endow MYCIN with characteristics that will make it more acceptable. These points are discussed in detail in § 1.4.6 and in Chapter 2.

7.1.2 THE SOLUTION

The MYCIN system is offered as a solution to the two problems described in the previous section; i.e., it attempts to give good advice regarding antimicrobial selection and it attempts to do so in a way which will make the system acceptable to physicians. In order to solve both these problems, MYCIN has been designed with three principal capabilities in mind:

(1) an ability to give good advice;
(2) an ability to explain the basis for its advice;
(3) an ability to acquire new knowledge easily so that its advice can improve over time.

Thus, MYCIN consists of three subprograms, each of which addresses itself to one of these three goals.

Subprogram 1 is a Consultation System. This component uses information about a patient, plus MYCIN's knowledge of bacterial infections, in order to decide (a) whether the patient needs to be treated, (b) the likely identity of offending organisms, (c) the possible drugs for use against these organisms, and (d) the best drug or drugs for the particular patient in light of his current clinical condition. Information about the patient is entered by the physician in response to questions asked by MYCIN. Each question asks for the value of some clinical parameter used by the program when it makes decisions. If all such values were known, the patient's clinical status would be fully characterized. MYCIN's task, however, is efficiently to select those of the clinical parameters (currently 65 in number) that are needed for adequate consideration of a given patient. The program's current knowledge is stored in 200 isolated decision rules, each of which is invoked only if the program has reason to believe it may be useful. This efficient use of system knowledge is accomplished by a goal-oriented control structure which dynamically creates a reasoning network appropriate for the clinical problem under consideration. The details of Subprogram 1 are the subject of Chapter 3.

Subprogram 2 is an Explanation System. This component attempts to answer questions from the user both during and after a consultation session. Furthermore, it attempts to do so in terms that will convince the physician that it reaches decisions in much the same way that he does. The user may ask MYCIN to explain the reason for a question during the consultation or may demand explanations of decisions that the program has reached. In an effort to make such explanations easy to obtain, even by a novice user, Subprogram 2 has been given a limited ability to understand simple English. In addition, its responses to questions are expressed in English and require no knowledge of MYCIN's internal representation or control structure in order to be understood. The details of Subprogram 2 are described in Chapter 5.

Subprogram 3 is a Rule-Acquisition System designed for use by experts in infectious disease therapy. The capabilities of this system component are currently incomplete, but it is possible for an expert to teach MYCIN certain simple rules which are then incorporated into the system's knowledge base for use in future consultations. An expert is encouraged to use Subprograms 1 and 2 in an effort to identify problems with MYCIN's knowledge of infectious disease therapy. Subprogram 3 then permits him to enter new rules or to modify old ones which he has found to be inadequate. The rule-acquisition procedure, like Subprogram 2, attempts to understand knowledge statements expressed in English so that the expert need not learn a computer language nor details of MYCIN's implementation. Subprogram 3 is the subject of § 6.3.

7.2 Limitations of MYCIN's Approach

One of MYCIN's principal limitations is shared by all clinical consultation programs: it requires that the physician take the initiative in asking for an interactive session. Despite attention that may be paid to "human engineering" issues during the development of such programs, physicians seldom choose to use computers for tasks they feel they can do themselves [Croft, 1972; Startsman, 1972]. Evidence that the medical profession may be failing to perform a job optimally, such as I presented in § 1.5.2, may not be sufficiently compelling for the individual practitioner to seek out a program's advice. Programs which can monitor physician prescribing habits automatically and generate warnings when appropriate [Cohen, 1974] may thus be more likely to influence physicians and the general quality of patient care.

The MYCIN approach is also limited by its complex control structure and representation scheme. These are probably unnecessary for medical inference tasks in which uncertainty and incomplete knowledge are uncommon occurrences. Until the MYCIN formalism has been tested on other medical decision problems, however, it will be difficult to predict what proportion might be handled more directly by other techniques. Furthermore, the efficiency advantages of approaches relying more heavily on statistical theory must be weighed against the importance of the natural mechanism for explanation and knowledge modification that may be achieved through MYCIN's flexible representation scheme.

Another limitation results because MYCIN currently requires more computing power and memory (a Digital Equipment Corporation PDP-10 with 256K of core) than is reasonable to expect to find in most hospital computing environments. Although the PDP-10 has been a powerful research machine for development of the MYCIN system, we are hopeful that the program may eventually be translated and simplified for smaller machines that may more reasonably be purchased or shared by hospitals. At present the system functions solely in a research environment, and long-range cost estimates are therefore not feasible.

MYCIN's explanation capabilities are limited somewhat by the current state of the art in computer-based processing of natural language. Although the program can understand simple, straightforward questions (Chapter 5), automated language understanding has not progressed enough to allow any program to handle unrestricted discourse.

Finally, it should be noted that MYCIN's reasoning mechanism must necessarily become more complex as the number of rules (i.e., amount of knowledge) increases. The present program exhaustively searches all rules that seem to be useful and does not distinguish those that are more useful than others. We have identified several heuristics for limiting the search space for a consultation (§ 3.7), and these must be implemented as MYCIN's knowledge grows or else the problem of excessive use of computing time will become overriding.

7.3 Contribution to Computer-Based Medical Decision Making

MYCIN has several novel attributes that distinguish it from other programs for medical decision making. Foremost among these is its ability to reason with informal judgmental knowledge acquired from experts. Although the system makes no attempt explicitly to model the psychological processes of a clinical decision maker, its modular decision rules and the certainty factor quantification scheme permit a physician's intuitions to be coded without major difficulty. Thus MYCIN's decisions need not depend upon the diagnostic algorithms, physiologic models, nor the statistical analyses that pervade much of the field (§ 1.3). The MYCIN formalism is therefore potentially applicable to decision making in the large number of clinical problem areas for which pathophysiology is poorly understood and statistical data are incomplete or nonexistent.

It should be noted that the MYCIN approach does not rule out applications for which reliable data become available. The formal certainty factor definitions and combining functions permit probabilistic information and judgmental knowledge to be used in unison. Furthermore, extensions to MYCIN may permit causal links to be coded in rule form so that the present control structure need not be modified. Although MYCIN may not provide the "best" solution for decision making in every clinical problem area, it may well serve as a useful adjunct to alternative techniques in most medical decision making applications.

Another important contribution of MYCIN's approach is its ability to reach decisions based upon whatever information is available at the time of the consultation. As is true for human consultants, MYCIN gives more reliable advice as more comprehensive information becomes available. Explicit decision trees or decisions based upon clinical algorithms tend to require pieces of information in a fixed order; if a datum is unavailable, the physician must wait for the appropriate test result before completing the consultation session. Of course, there are times when so little information is available that MYCIN cannot reach a reasonable decision. In general, however, MYCIN makes the best decision it can on the basis of current data and the user is encouraged to return for more definitive advice as further information becomes available. In a problem area such as the treatment of infectious disease, interim decisions while awaiting further data are often in the best interests of the acutely ill patient.

Avoiding explicit decision trees has provided other advantages besides an ability to operate solely on the basis of current information. Most important among these is MYCIN's ability to incorporate new knowledge without explicitly being told how or when it will be useful. The program's control structure for dynamic reasoning (Chapter 3) automatically utilizes any rule-based knowledge that appears to be relevant. Storing knowledge in rules has also facilitated an ability to explain why questions are asked and to justify the basis for the program's therapeutic recommendations.

Finally, MYCIN has been designed to be more than merely an interesting theoretical approach to medical decision making in this therapeutic problem area. From the outset we have stressed the goal of eventually implementing the program for ongoing use by physicians. We have sought to understand why such programs have met resistance in the past, and we have in turn implemented a number of features, including a comprehensive explanation capability, designed to heighten MYCIN's acceptability to physicians. Although the program is not yet sufficiently knowledgeable for ongoing clinical use (§ 6.4.1), physicians who have used the system have uniformly indicated that they believe the program *can* become sufficiently reliable and will hence be used by the clinicians for whom it has been designed.

7.4 Contribution to AI

MYCIN's mechanisms for representing and utilizing judgmental knowledge also heighten its interest for computer scientists working in AI. Unlike formal problem-solving systems based on axiomatic knowledge, MYCIN suggests an approach for modeling the kinds of inexact reasoning that typify many real-world problems. AI researchers have recognized the need for some way to combine the attributes of decision theory with those of machine problem-solving [Feldman, 1974], and MYCIN provides what is perhaps the first general approach to this problem. Certainty factors are potentially applicable to a number of AI application areas. For example, conversations with

AI researchers have revealed that tasks such as identifying objects in machine vision or phonemes in speech understanding are typified by the kind of indecision that CF's are designed to handle.

Although neither MYCIN's goal-oriented control structure nor its dependence upon rule-based knowledge is unique (see Chapter 3), no other AI system has used its knowledge in quite the same way. As I have emphasized, MYCIN's formalism is domain independent and thus may prove useful for AI researchers who wish to automate other tasks that are dependent upon the heuristics of individuals. Furthermore, the use of rules with CF=1, or with a certainty factor derived from reliable statistical data, provides a mechanism for coding theorems, real-world data, and definitional information. This formal knowledge may then be used simultaneously with the informal knowledge that is representative of the intuitive inexact reasoning that typifies much of human problem-solving.

MYCIN has also been developed with more attention to human engineering than is typical of much of the AI field. The goal has been to develop mechanisms for interacting with medical professionals who are not only unfamiliar with AI but have often never used computers before. MYCIN's rules have therefore served as a highly useful representation scheme since they can be individually retrieved in order to explain why questions have been asked or to justify aspects of the program's advice. As AI applications for use by scientists and other individuals become more common, MYCIN may well suggest some useful guidelines for interactions with novice computer users.

Another lesson to be learned from MYCIN is that a single programmer, working full-time for two years with a powerful interactive language such as INTERLISP, can create an AI program that serves a useful purpose. Observers often bemoan the current state of the art in AI, asserting that it will be years before machines can perform problem-solving tasks at a level approximating that of humans. MYCIN has shown, however, that if researchers are willing to accept the current limitations of the AI field, and to select real-world goals that are compatible with those limitations, a useful system can be developed using techniques for representation and control that would not have been available if it were not for prior work in AI. MYCIN's question-answering (QA) skills are an example of this point. The techniques used for natural language understanding are dependent upon several simplifying assumptions that ignore syntax, semantics, and the psychology of language. The last two of these are perhaps the principal barriers to further AI progress in the field of linguistics (Chapter 5). The limited QA capabilities that result, however, are in general satisfactory for the application area in which they are to be used. Although it would clearly be preferable if the program could participate in free form discourse, MYCIN has shown that a useful interim solution can be developed once the current limitations of the field have been accepted.

Finally, MYCIN has contributed to the AI field by providing evidence that suggests that current AI techniques may be adequate for assisting professionals with an important real-world problem. There has been a tendency for theoretical AI work to concentrate on tasks that are often described as "toy problems." Although such problems are generally nontrivial and AI researchers can themselves appreciate the challenges involved, the relative paucity of AI programs that deal with real-world tasks has not always benefitted the image of the field. Although MYCIN's effectiveness as a clinical tool has not yet been fully demonstrated, the preliminary evaluations mentioned in § 6.4 make us optimistic about its future. We are therefore pleased to be able to offer MYCIN as an example of a way in which current AI technology can potentially contribute to the betterment of public health through improved care for patients with infections.