

## Foreword

Edward Shortliffe's book represents excellent modern work in artificial intelligence (AI). Leaving aside for the moment its medical focus, the scheme viewed only as a computer system has considerable merit. Dr. Shortliffe broadens the usual definition of an AI program to include the requirement that such a program should justify its decisions: that is, defend its behavior. The MYCIN system can indeed show us the rules and assumptions upon which each decision is based. One advantage of this capability is that the user can evaluate the extent to which the system's reasoning resembles comparable human decision making, and in many cases the reader of this book will discover how frail a logical basis either machine or human has for many decisions.

In this vein one should also note that MYCIN is distinctly modern AI work in that it relies heavily upon what is called "content specific knowledge". That is, the program uses a strategy of generalized forms of representation for knowledge and relationships in order to seek and encode advice from experts in its problem domain—in this case infectious disease management. This approach is distinguished from the older and less successful ones which oblige themselves to "reason from first principles", invoking general concepts and processes rather than rules that have been particularized for the problem at hand. MYCIN then is in that class of AI work that assumes its problem task is difficult and important, and thus utilizes all knowledge helpful in obtaining reliable and credible automated decision making. The application of considerable resources of the system to remaining open-ended is a consequence of this strategy. MYCIN attempts to facilitate "fine tuning" and the acquisition of updated

## MYCIN

advice rules in agreement with the opinions of its expert consultants.

Viewed solely from a medical point of view, MYCIN must also be judged a significant accomplishment. The problem domain selected by Shortliffe, namely antibiotic medication for treatment of bacterial infections, is indeed medically exceedingly important. Often this function is performed by a subspecialist consultant in infectious disease. Alternatively the clinical decision to treat with an antibiotic may be based upon the general knowledge of the clinician (internist, family practitioner, surgeon, or pediatrician) in consultation with the laboratorian (pathologist or microbiologist). It is definitely not necessary for MYCIN to command the array of knowledge reflected by these medical specialists. Only a small subset of that knowledge is needed: namely that which relates to the treatment of bacterial infections with antibiotics. As in the case of the human decision maker, MYCIN's knowledge and rules must be up-to-date with respect to microbiological nomenclature and the names and types of pharmacological preparations. As with the human information processors, MYCIN's ultimate decisions can be no more valid than the initial medical observations provided. Basically these are observations of the patient, laboratory isolations of pathogenic microorganisms, and results of antibiotic sensitivity tests. MYCIN, like the physician, will inevitably be led astray by incorrect information about what organism actually was recovered from a patient specimen, or other such errors.

MYCIN shares one more difficulty with its human counterpart: it must deal with the uncertainty that results from missing data. In infectious disease management, some observations simply cannot be recreated. Once antibiotic treatment begins, for example, certain isolations are impossible. Probabilistic operation is natural for a computer program, more so than for most humans. On the other hand, dealing with the conditional statements that relate the probable conclusion to real world events, a task more pertinent to MYCIN's function, is difficult. An example of such a relationship is the appearance of contaminating bacteria in cultures from a patient wound. It requires considerable sophistication for a program to recognize a contaminant. Consider the fact that the hospital laboratorian will ordinarily not feel competent to make the distinction. At least, he or she will feel obliged to report the questionable isolate. That MYCIN can even attempt such things is a tribute to the



## Foreword

maturing AI field and serves to set it aside completely from ordinary hospital automation. MYCIN is therefore a significant step forward for AI research, and is directed at the solution of a substantial medical problem.

When one examines in detail the proposal for the scheme to be used in actual hospital and office practices, he or she sooner or later hits upon cost as a potential obstacle. The high cost of operation of artificial intelligence programs written in flexible, high-level languages such as LISP is easily rationalized but not easily avoided. INTERLISP, an interactive version of the same language with additional accommodations, is perhaps 50 times less efficient than FORTRAN, which is perhaps ten times less efficient than an inefficient assembler. Yet LISP permits the construction of lists of relationships describing most naturally the real world information structures necessary for AI programs. LISP also readily permits the alteration or respecification of the lists to reflect a new expert's contribution. Using this capability, MYCIN's formal certainty factors and combining functions permit new knowledge to be incorporated even without explicit commitments as to how or when it will be used. INTERLISP permits easy access to the endlessly evolving structures in order that the computer programs be kept up to date with the changing world. Operational expenses are not related directly to simple machine efficiency. Yet even so, the desirable features of a system such as INTERLISP are expensive.

One could imagine such a system "finished" one day, ready for "hard coding" in an efficient language in an efficient environment. Technically this could easily be accomplished, with a dramatic reduction in operating costs. Based upon past experience with laboratory automation, such a change would be rash because of the quicksilver nature of the medical laboratories. No week will ever pass for a system like MYCIN—or any other medical laboratory computer program—in which changes are not required to up-date the program with changes in practice. The changes need not be conceptual; MYCIN is actually prepared to accept conceptual changes, whereas no ordinary laboratory program can do this. In fact most programs cannot tolerate so much as the change of a field type or an additional digit on a patient number. The majority of the endless changes will be ordinary: new drugs, new names for old drugs, new preparations or formulations, changes in bacteriological nomenclature, new isola-

## MYCIN

tion or subculture techniques, new combinations of results to be reported together or separately, or new types of specimens in new preservatives or holding solutions to be treated in new ways by the laboratory workers.

In addition, a number of other computer-based systems are beginning to deal with drug-drug interactions, and some with the ways in which patient status (e.g., renal failure) relates to choice and quantity of the administered drug (including antibiotics). Eventually MYCIN must either talk to such programs, or incorporate their pertinent knowledge. An additional future linkage for the MYCIN system, and an additional advantage in remaining in INTERLISP, is the obvious need to interface with hospital data acquisition systems. MYCIN is undertaking to render consultation; however there is the other world of acquisition, quality control, and reporting of the basic laboratory observations and measurements. Since these other systems are typically rigid, it will be up to the advanced systems like MYCIN to remain adaptable so that interchange of information between the two can take place in the future.

The general argument that highly flexible and accommodating programs are necessary simply to support interaction with medical subject matter experts is a moot point. If there must be a general rule it should be that a system should be capable of taking advice however an expert wants to give it: on-line, off-line, from rules, from tables, from graphs. Many systems, for instance the author mode of computer-aided instruction systems, have been "improved" by compulsory on-line case-building to the point of becoming burdensome. Such approaches fundamentally shift the program updating task from the program to the terminal user. In the case of MYCIN, however, the point is moot because the expert probably needs interaction with the system in order to discover what rule needs challenging and to see the full consequences of a change before it is made. In other words, the user needs the interaction more than the program does. Consequently, in order for the program to continue to grow—at least in wisdom—the relatively expensive high-level interaction may be desirable almost indefinitely.

These reasons suggest that the advantages to MYCIN outweigh the higher costs of operation. Artificial intelligence programs were made possible by flexible programming languages. The lack of flexibility, not the costs, has signaled the eventual death of almost all previous



## Foreword

laboratory systems that used traditional programming techniques. MYCIN should be evaluated—sink or swim—as a relatively costly but exquisite machine tool.

The future importance to medicine of AI programs is twofold. First, it is possible that this approach to knowledge acquisition is easier—at least for the medical people—than the older programming methods that relied upon simple trees or likelihood tables. A MYCIN-like approach could explore in time a more complex problem domain and provide a means serially to examine the state of knowledge in a field not yet wholly rationalized, but such a system could be capable of accepting fairly radical correctives during its subsequent development. MYCIN and its derivatives will want and need changes to assumptions and rules, as long as the changes move it toward more valid decision making. An AI program such as MYCIN can be expected to facilitate the extension of computer-based automation into new medical areas simply because of its ability to evolve a knowledge representation for the field.

Secondly, and perhaps even more exciting is the potential for MYCIN and similar systems to enhance human understanding of the problem domain itself. This advantage is characteristic of good modeling work in general. One expects major gains in perception to emerge from the initial task of formalizing one's knowledge. The difficulty then (with respect to medicine) lies in the selection of the proper medical problem domain. Minimally, the domain must have some structure, that is, a valid fundamental model of the data generating process. To take an outlandish example, a system such as MYCIN based upon idle ramblings about vapors, phlogistin, and bad air could not make valid decisions, although it might successfully imitate the decisions of physicians of 200 years ago. On the other hand, in the present case we do have the advantage of a microbial model of infectious disease which gives at least the possibility that MYCIN can make a valid decision.

Another requirement for successful AI work in the medical domain is the existence of a reasonable amount of knowledge about the relevant decision making—however organized—in the heads of the expert practitioners. In infectious disease management, there are many explicit rules relating diagnosis to the outcome of drug therapy. In contrast, the action of the individual antibiotics is poorly conceptualized. The testing of bacterial isolates by impregnated discs

## MYCIN

for “antibiotic sensitivity” is at best a kind of standardized artifact that results in notoriously poor *in vivo* versus *in vitro* correlation. For example, *Salmonella typhosa*, the bacterium that causes typhoid, is by standard laboratory testing “sensitive” to most antibiotics in the test battery, yet experience has shown that successful treatment must be with chloramphenicol (or perhaps ampicillin). Fortunately experts know this and can cite the proper behavioral rules even though they cannot give explanations; unfortunately there are scores of such exceptions, and they are not always known to the non-expert. It is in this type of practical problem in clinical medicine that MYCIN can be useful.

This discrepancy between theory and practice in treatment with antibiotics is related to the well-known scientific quest to deduce pharmacological function from chemical structure. While the chemically complex antibiotic molecules might not be the obvious area in which to anticipate such a discovery, examination of common clinical outcomes and MYCIN relational rules might well contribute to this kind of investigation.

Ultimately this work must be evaluated according to whether it results in drug selections that are associated with improved clinical outcomes, and whether new understanding of biological relationships results from the development of the programs. Indeed there is an excellent chance that MYCIN can improve drug selection for the non-expert. With respect to the possibility of gaining new understanding of infectious disease relationships, MYCIN has a vast advantage over ordinary laboratory automation systems. The traditional systems attempt merely to automate present procedures. MYCIN attempts to formalize the representation of knowledge of infectious disease. Doctor Shortliffe’s book explains quite clearly how he goes about this process. His book is as much a pleasure to read as his system is a pleasure to contemplate.

University of Missouri  
Columbia, Missouri

DONALD A. B. LINDBERG, M.D.  
*Director, Information Science Group*