

Chapter 1

Introduction

1.1 Computer Applications in Medicine

1.1.1 PROBLEMS AND PROMISE

In the late 1960's, David Rutstein wrote a monograph entitled *The Coming Revolution In Medicine* [Rutstein, 1967]. His discussion was based on an analysis of several serious problems for the health professions:

- (1) modern medicine's skyrocketing costs;
- (2) the chaos of an information explosion involving both paperwork proliferation and large amounts of new knowledge that no single physician could hope to digest;
- (3) a geographic maldistribution of MD's;
- (4) increasing demands on the physician's time as increasing numbers of individuals began to demand quality medical care.

Rutstein concluded that technology provided a possible partial solution to several of these problems.

In subsequent years technology has indeed increased its influence in the medical sphere, but the problems listed above are still highly visible. Their ultimate solutions will undoubtedly involve a long process, only portions of which can be accomplished by technological innovation alone. Equally important are appropriate supportive legislation, at both state and federal levels, plus a gradual change in the attitudes of health personnel towards their training, their profes-

sional duties, and the technological environment that will increasingly surround them.

The attitudes of health personnel towards computers provide some of the greatest barriers to successful implementation of computer-based systems. A recent study [Startsman, 1972] used an open-ended questionnaire and factor analysis to provide information concerning the optimal interfacing of a computer-based information system with a medical staff. Results indicated that interns, nurses, and ancillary personnel expressed the least willingness to use data processing systems, while medical faculty, pre-clinical medical students, and medical record librarian students were most receptive. Although acknowledging that house staff attitudes may reflect the fast-paced environment in which preoccupation with the immediate physical needs of the patient is the norm, the authors point out that interns and residents comprise precisely the group for which many clinical computing systems should be oriented. Thus, since the study showed that familiarity with computers tends to dispel fears and breed interest, the authors suggest that health personnel should be exposed to data processing techniques during their educational years when they are apt to be most receptive to these kinds of innovation.

The most commonly expressed fears regarding computer applications in medicine involve loss of job (or job stature) due to 'replacement' by a computer, and presumed depersonalization of patient care due to machine intervention. In addition, some physicians are concerned about the legal ramifications in the use of, or failure to use, a computer-based facility [Hall, 1972]. Computers appear remarkably cold and sterile, particularly to individuals unfamiliar with their capabilities and limitations. "Scare" articles in professional journals also help reinforce attitudes of distrust [Eisenberg, 1974].

A group at Duke University Medical School has suggested that the key to physician acceptance of computer technology lies in a "practical demonstration that physicians or groups of physicians using [computers] have a clear advantage in practice over physicians who maintain the status quo" [Rosati, 1973]. Applications that can make such a demonstration convincingly, however, are difficult to imagine. Norms of practice already vary considerably, even within close geographic proximity, and mechanisms for measuring one clinician's "advantage" over another's have so far tended to emphasize economic considerations (e.g., length-of-stay and utilization review as a primary method for medical audit and quality assessment).

Introduction

The subject of economics also raises important questions regarding the cost of medical computing, another major impediment to acceptance of the technological innovation. Difficulty in quantifying the dollar-value of improved patient care quality has understandably frustrated economists who have tried to apply conventional theory to the unique medical marketplace. As a result, there are now specialists in medical economics who have proposed new conventions and analytical tools for considering questions of cost effectiveness and resource allocation within health care environments [Klarman, 1965]. The basic problem remains unsolved despite these efforts. One of the first questions a hospital administrator asks when a computer system is proposed is: how much will it cost? It is seldom easy to justify such systems as cost effective because the savings are buried in reduced length-of-stay data, in lowered lab or pharmacy charges for the patient, in "improved patient care," or in similar real but imprecise monetary measurements.

Finally, many computer innovations are proposed as time saving techniques for the physician. In an age when a doctor shortage and maldistribution is well recognized [Fein, 1967], such arguments can be highly compelling. By inference, however, any computer program that saves physician time must be doing a task that previously was done by the physician himself. The complex psychological and ethical issues involved here, both for the physician and the patient, will be discussed in greater detail when we describe computer-based clinical decision making in § 1.3.

1.1.2 (**) MEDICAL COMPUTING APPLICATION AREAS

The discussion in § 1.1.1 does not specify which computer applications are relevant to each point because almost all medical computing systems entail similar philosophical, ethical, and economic considerations. In this subsection I briefly describe the major areas of medical computing service and research. The categories are my own, and may therefore be nonexhaustive, but they should serve to give you a general feeling for the ways in which the so-called "computer revolution" is affecting the administration and the practice of medicine. General references on the subject of computer applications in medicine include Lindberg's volume from the University of Missouri [Lindberg, 1968], a comprehensive survey of medical computing in England [Abrams, 1970], a four-volume continuing series that sum-

MYCIN

marizes some of the work underway in the United States [Stacy, 1963, '65, '69, '74], and a survey article from the *New England Journal Of Medicine* [Barnett, 1968].

1.1.2-1 Business Applications

The most widely used and accepted computer-based applications involve hospital accounting systems. Business computing is perhaps the best developed of all computer applications, both because accounting uses have been a major concern of many computer firms since the industry was in its infancy, and because accounting problems are in general well-defined and thereby more straightforward to develop and implement. Automated accounting developed for the business world has required very little adaptation for hospital application. It is hardly surprising, then, that hospital accounting functions have been the first medical functions to be automated. Not only is this priority logical in light of the success and experience that general industry has acquired by using the computer for financial activities, but the application also demonstrates easily recognizable monetary benefits.

The need for computing systems to handle financial data and to print out forms has been heightened in recent years by the explosive rise in hospital rates and the concomitant need for increased and improved communication between the hospitals and third party payers or the government. The private physician has been faced with the same paperwork proliferation on a smaller scale. As a result, several service computing firms offer individual office-based financial packages to practitioners who find it difficult to maintain their patient care schedules, particularly with welfare cases, because processing all the paperwork by hand has become exceedingly tedious and time consuming.

It should be noted that much of the public opinion regarding computers is derived from direct contact at the financial level between the consumer and the computers that send him his bills. Thus a patient who is directed to sit at a console for an automated medical history may well think back to his last erroneous bank statement or computer-generated billing error and rebel at the thought that a similarly error-prone machine is about to take charge of his physical well-being. Physicians asked to read computer-generated summaries

Introduction

may also question the reliability of the information. Thus improved performance levels for business computer applications, both through increased machine reliability and utilization of well-trained and responsible systems personnel, may be a necessary first step towards improving the public image of computers and thus lowering the barriers of resistance to computing innovation in medicine. This trend is already underway and is being aided by the increasing number of young adults who have grown up in the computer age. The novelty and mysteriousness of computers have made them especially threatening to individuals who remember, for example, the hand-posted billing statements they received in the precomputer era.

A final important point regarding the introduction of financial computing into the doctor's office is that the related hardware and communications equipment will be increasingly familiar and accessible. The same computer terminal that is purchased for sending daily billing and insurance data from the office to a central financial computing service could presumably be used for connecting with a network of computer-based clinical resources. Thus, little or no additional capital outlay may be necessary for the future physician to interact with computer programs designed to help with the day-to-day practice of medicine. The challenge is, then, to develop good computer-based clinical tools so that the physician will take time to use them regularly (and be willing to pay for the associated computing charges) because they are of genuine help in his practice.

1.1.2-2 Biomedical Engineering

It is convenient to divide medical computing applications into two areas—those identifiable as biomedical engineering tasks and those more appropriately termed information processing or data handling. The primary component of biomedical engineering applications is the analysis of analog signals or the construction of sophisticated technologies for man-machine interaction. This is a vast field that includes such applications as medical computer graphics [Newton, 1973; Cox, 1967; Alderman, 1973], computer assisted pattern recognition from visual signals [Bahr, 1973; Neurath, 1966], computer analysis of real-time data [*Computers and Medicine*, 1973a; Harrison, 1971; Henry, 1968], and various kinds of patient monitoring.

Patient monitoring includes all applications in which computers

MYCIN

are used to process, or monitor, signals relayed by machines that measure the patient's physiological parameters. By far the largest subfield in this category is the development of programs that analyze electrocardiograms (EKG's) (in recent years, literally hundreds of articles on this subject have been published annually). The vastness of the literature reflects the well-recognized need for computer programs that can assist the physician with EKG analysis; this field is of particular value in medically underserved areas where the expertise of highly trained cardiologists may not be readily available. However, the size of the existing literature also suggests that the ultimate program for this purpose has not yet been created. Indeed, although several programs do very well at EKG analysis [Wartak, 1971; Carceres, 1964; Pryor, 1969; Wolk, 1972], none has yet achieved the accuracy of a good and experienced cardiologist. Similar work has also been done on the even more complex problems of electroencephalogram (EEG) analysis. Results in this field have so far been rather rudimentary and have tended to concentrate on the identification of abnormal spikes in the tracings from the various leads [Walter, 1968; Cox, 1972; Kellaway, 1973].

The phrase "patient monitoring", however, generally implies more than signal sampling and analysis [Warner, 1968; *American Medical News*, 1970; Felsenthal, 1973]. Also involved is the concept of a warning system, wherein a computer is programmed to sample a patient's physiologic parameters at specified intervals and to warn the nursing or medical staff if an abnormal or dangerous reading is noted. The ethical and legal implications of such systems are only gradually being worked out. Even more revolutionary will be systems in which the computer not only notes the abnormalities but takes corrective action by injecting a drug, altering a pacemaker setting, etc. Although such systems are often discussed, none has yet been implemented for ongoing service.

1.1.2-3 Multi-Phasic Health Testing

As health care critics have increasingly pointed out, the tendency of American medicine is to concentrate on crisis care, largely ignoring the need for improved preventive medicine. The health care and industrial communities have begun to counteract this tendency by screening large populations and identifying individuals with early

Introduction

or latent disease. "Multi-phasic health testing" (MPHT) is the common term for procedures whereby apparently healthy individuals are given a battery of screening tests to determine whether an individual may need further medical attention [Oszustowicz, 1972; Collen, 1964, 1965, 1966, 1969, 1971]. The various MPHT centers use computer technology to varying degrees, in most instances primarily for collecting data and for printing them in an organized fashion that facilitates their review by the staff physicians.

Many MPHT centers also use computers to obtain the patient's medical history. Automated history-taking has been developed primarily within the last decade [Grossman, 1968; Slack, 1966] and generally involves easy-to-use pushbutton display terminals. The patient sits at the 'scope for varying lengths of time, usually from 30 to 60 minutes (depending upon the complexity of his complaints), and answers the multiple choice questions by pushing the button beside the correct answer. The programs utilize branching logic so that more specific questions may be asked of patients from whom more detailed information seems relevant.

Such programs have also been used in hospital outpatient clinics. Summaries of the history are legibly printed by the computer for review by the physician when he sees the patient. He may then pursue in detail topics about which the computer has indicated an extensive history may be necessary. Another capability of automated history recording is that of asking the questions in one language and printing the medical history summary for the physician in another. Thus, the computer may serve as a useful translator in cases where, for example, the patient speaks only Spanish or French and the physician only English. Studies to evaluate such systems generally indicate that patients accept the automated history recording more readily than the physician does [Grossman, 1969, 1971]. The summary for the physician is gradually being improved, however, as the designers of these systems gain experience and acquire insight into the reasons for physician resistance.

1.1.2-4 Automated Medical Records

One of the great differences between modern medicine and the clinical practice of a century ago is that the care of a patient is shared, particularly in teaching institutions. Thus the medical record

MYCIN

that once served merely as a worksheet for the individual physician to jot down personal reminders now is an important means of communication for the physicians caring for the patient. Furthermore, the medical record now also serves as an important legal document.

Unfortunately, the medical record has not yet fully evolved to meet the demands of all these requirements. Charts are not usually standardized, are often poorly organized, and tend to be illegible. Redundancy of data is to be expected since health professionals using the medical record tend to duplicate information; they often have neither the time nor the practice to search the chart to see if the data have already been entered.

Recognizing the chaos that arises out of the conventional medical record system, several researchers have suggested new organization techniques and potential mechanisms for automation. Most notable, perhaps, is the Problem Oriented Medical Record (POMR) proposed by Weed [Weed, 1968, 1969a]. He developed the approach at Case Western Reserve, and in recent years has used computer technology to automate the system both there [Weed, 1969b] and at the University of Vermont. The POMR approach has also been advocated as an aid to medical audit [Weed, 1971], although recently questions have been raised regarding its usefulness for this purpose [Fletcher, 1974]. Nonetheless, the system has received wide attention [Bjorn, 1970; Collins, 1973; Esley, 1972; Feinstein, 1973; Goldfinger, 1973; Hurst, 1971a, 1971b, 1972, 1973; Mittler, 1972] and is now used routinely at several hospitals, particularly in the eastern United States. Only Weed's group has automated the POMR, although similar work has been undertaken at the Massachusetts General Hospital [Greenes, 1969, 1970a] where a computer-based clinical data management system has been utilized in the outpatient hypertension clinic, the coronary care unit, and for systemized input of radiology reports [Pendergrass, 1969; Bauman, 1972]. The important point to note regarding the computer systems of Weed and Greenes is that each is designed for use by the physician himself, both for data input and data retrieval. Thus, in accordance with our comments above, physician acceptance of such systems must remain a primary consideration during program development and implementation.

An alternative to both the traditional source-oriented record and the POMR is the time-oriented databank (TOD) introduced at Stan-

Introduction

ford Hospital [Fries, 1972]. The TOD System, like the POMR, is primarily a revision in the organization of the hard-copy record. Automation has been introduced only for off-line data entry and analysis. The TOD system emphasizes chronological organization of patient data so that flow-charted trends can be observed over time. Physician interaction with the computer is not yet a part of the TOD approach.

Several other groups have worked with automated records, most of which only peripherally involve the physician. The Kaiser Hospital System is particularly notable in the field [Davis, 1968; Collen, 1964], but other excellent work with both inpatient and outpatient records has also been done elsewhere in the United States [Grossman, 1973; Slack, 1967; Kiely, 1968] and abroad [Buckley, 1973]. Some investigators have looked for ways to automate records without sacrificing the conventional text format [Korein, 1963; Levy, 1964; Bross, 1969], while others have attempted to introduce structure to the records by using checklists or self-encoding forms [Yoder, 1966, 1969; Collen, 1971; Hall, 1967]. Finally, some observers have argued that it is premature to study the structure and optimization of patient data-handling without first assessing and improving the quality of the data themselves [Feinstein, 1970].

1.1.2-5 Laboratory and Pharmacy Systems

Unlike clinical parameters best known to the physician himself, patient data related to lab tests and administered drugs can be acquired from sources other than the doctor. Thus several systems have been developed to aid in the acquisition and control of laboratory and pharmacy data.

Chemistry laboratory systems are perhaps the most common clinical application of computers. Several excellent systems have been designed [Hamilton, 1973a, 1973b; Raymond, 1973; Katona, 1969] to accomplish one or more of the following tasks:

- (1) accept test orders, in some cases on-line from the wards;
- (2) generate schedules for the technicians who collect the appropriate samples from the patients;
- (3) generate worksheets for the technicians running the tests in the laboratory;

MYCIN

- (4) provide automatic accessioning for control and identification of samples;
- (5) accept test results on-line from various kinds of equipment;
- (6) accept other results from terminals in the laboratory;
- (7) provide rapid access to test results on any patient;
- (8) generate hard-copy records, in a variety of formats, for inclusion in the patient chart or for individual use by physicians.

Other programs suitable for inclusion in the category of laboratory systems are ones for reporting pathology lab diagnoses [Beckett, 1972], for analyzing antimicrobial sensitivity test results [Hulbert, 1973; Groves, 1974] or identification data [Mullin, 1970], for organizing and controlling large collections of laboratory specimens [Bachman, 1973], or for quality control in a microbiology laboratory [Petralli, 1970].

Pharmacy systems generally assist with label printing, inventory control, and maintenance of up-to-date patient drug profiles [Evans, 1971; Almquist, 1972]. One hospital has used such profiles to identify outpatients who are drug abusers [Maronde, 1972]. A novel pharmacy control system has been introduced at Stanford Hospital [Cohen, 1972, 1974] where new drug prescriptions are compared with the patient's drug profile and warnings for the physician are generated if a potential drug interaction is noted. Finally, the Kaiser Hospital System has reported a computer-based mechanism for monitoring the incidence of adverse drug reactions [Friedman, 1971].

1.1.2-6 Hospital Information Systems

A centralized computer that performs or oversees several of the automated functions described above is called a Hospital Information System (HIS). Since such systems tend to require massive computing facilities, commercial firms are particularly interested in such installations. An HIS usually involves an automated mechanism for patient admission and bed census [Hofmann, 1969] so that a computer-based record for each patient exists from the moment he enters the hospital. The patient record then serves as a focus for information flow. Laboratory and pharmacy data are centrally stored and the system transfers orders directly from the ward, where they are ordered, to the appropriate hospital service. Nursing personnel often use the system to post orders and to indicate when drugs have been

Introduction

administered or other patient care services have been performed. Physicians interact with ward terminals to varying extents, depending both upon the system design and the doctor's willingness to participate. A variety of additional services may also be performed by the central machine. Thus an HIS offers a variety of benefits to the various individuals who may use its data base [Barnett, 1968]:

. . . . To the physician, [HIS is] a system that will provide rapid, accurate, and legible communication of reports, better scheduling procedures and timely and precise implementation of activities ordered for patient care. To the nurse, HIS implies an operation to lighten the clerical load of communication functions, preparing requisitions and transcribing and charting. To the administrator, HIS is a means for using resources more effectively, for gathering the data necessary for appropriate management decisions and for ensuring that information necessary for the patient billing process is readily available and accurate. To the medical research investigator, HIS offers the potential for a data base of patient-care activities that is not only accurate but also organized and easily retrieved and analyzed.

Unfortunately this ideal picture of universal benefit and acceptance of an HIS has yet to be realized. The HIS at El Camino Hospital in Mountain View, California, has served as a model for other institutions considering such ventures. Initiated by Lockheed Aircraft but currently operated by the Technicon Corporation, this large system has surprised observers with its demonstrated cost effectiveness [Batelle Labs, 1973] but has been plagued by low user acceptance, particularly among physicians [*Computerworld*, 1973; *Computers and Medicine*, 1973b; Yasaki, 1973]. Suggested reasons for the problems encountered have been numerous. A 1971 article suggested several mechanisms for meeting resistance to hospital automation [Hofmann, 1971], some of which appear to have been overlooked by the El Camino planners. The need for eventual users of the system to participate in the planning process is particularly crucial, as is an effective feedback mechanism so that points of discontent can be overcome before they have a chance to grow. The need for thorough pre-implementation planning of the patient database for an HIS has also been recognized [Sauter, 1973]. Finally there are those who believe that any attempt to introduce a total hospital information system in a single step is doomed to failure from the outset. The alternate approach is to design the various

MYCIN

computer services as modules, perhaps on several small machines, and gradually to integrate them into a total system [Greenes, 1970b; Barnett, 1969; Hofmann, 1968].

1.1.2-7 Decision Support Systems

Computer programs to assist in clinical decision making are the subject of § 1.3. In that section some of the work that preceded the MYCIN system is discussed in detail. Here, we simply note that there are two kinds of clinical decisions that may be involved in such systems—the determination of the patient's diagnosis or the appropriate way to treat him. In some cases, treatment selection is straightforward once the proper diagnosis has been made. In others, treatment planning may be the most complex step in the decision making process.

1.1.2-8 Computer-Aided Instruction in Medicine

Computer-Aided Instruction (CAI) has become an accepted part of the educational process for many of today's younger students [Suppes, 1966b, 1969]. As the field has developed, students of the health professions have also begun to benefit from techniques developed by CAI researchers [Stolurow, 1970]. In medical education, a number of successful programs are available nationwide through a network supported by the National Library of Medicine [Wooster, 1973]. Several useful programs most of which avoid problems of natural language understanding, have been developed at Massachusetts General Hospital [Hoffer, 1973]. Ohio State University also has an extensive medical CAI facility [Weinberg, 1973]. Programs that play the role of a patient or otherwise enter into natural language discourse with the student include Cornell's ATS [Hagamen, 1973; Weber, 1972], and the CASE system at the University of Illinois [Harless, 1973a, 1973b]. A program that simulates the patient-physician encounter, with realistic simulation of the time required for the return of lab results, has also been reported [Friedman, 1973]. Little work has been done to evaluate the cost effectiveness of such systems, but a group at the University of California, San Francisco, has been sufficiently concerned with cost factors that

Introduction

they have developed a dedicated CAI system for use on inexpensive minicomputers [Kamp, 1973].

1.2 Artificial Intelligence

Although artificial intelligence (AI) has been defined in numerous ways, my preference is to acknowledge the intelligence of any machine that performs a task that a century ago would have been considered a uniquely human intellectual ability. This is a rather broad definition that thus encompasses a much wider group of machines and tasks than is usually ascribed to AI. The appeal of this definition, however, is the avoidance of arguments as to whether or not a specific machine should be classified as a product of the AI field. Furthermore, this definition implies that intelligence is a term that need not apply only to humans. However, one can argue that machine intelligence is not "artificial" at all, but is simply a variety of intelligence where the interplay of emotions, fatigue, and other "uniquely human" characteristics has been eliminated.

The more usual meaning of the term artificial intelligence encompasses a subset of the above definition in which (1) the machine is a digital computer or is controlled by a digital computer, and (2) the task involves symbolic reasoning ("thinking") rather than arithmetic calculations or information storage and retrieval. AI is therefore generally regarded as a subfield of computer science. The foundations of the field are often attributed to an article written by the late A. M. Turing [Turing, 1950], an English mathematician and logician who proposed an operational test of intelligence, the so-called Turing Indistinguishability Test. In addressing the question "Can machines think?", he suggests that, for all practical purposes, a machine is intelligent if an individual communicating with the machine (say by means of a teletype) is unable to decide whether he is interacting with a computer or with another human who is also using a teletype.

I shall not attempt to survey the field of artificial intelligence. Several excellent general texts are available that devote considerably more space and energy to such surveys than are available here [Feigenbaum, 1963; Minsky, 1968; Slagle, 1971; Nilsson, 1971]. Critics have also been moved to write entire volumes arguing against the AI field [Dreyfus, 1972]. The reader is therefore encouraged to

MYCIN

consult a recent survey paper [Nilsson, 1974] for a more thorough discussion of AI and for a comprehensive bibliography of the field. An earlier survey of the field also is available [Minsky, 1961]. In the rest of this section I shall follow Nilsson's categories [Nilsson, 1974] for organization of the AI field in an effort to give a brief overview of the kinds of problems with which AI is presently involved.

There are four basic AI methodologies that have been addressed by almost all workers regardless of their specific area of application. In addition, there are approximately eight application areas that encompass most of the work in AI. In the discussion below the eight application areas are listed and briefly described; then the four core topics common to most AI work are introduced.

1.2.1 (**) AREAS OF APPLICATION

1.2.1-1 (**) *Game-Playing*

Some of the best known work in AI involves the development of computer programs that can play highly complex games [Slagle, 1971]. Programs have been written to play checkers [Samuel, 1959, 1967], chess [Greenblatt, 1967], poker [Waterman, 1970], bridge [Berlekamp, 1963] and several other games that require complex strategies regarding a large number of alternative actions (moves). Such games must be contrasted with a contest such as tic-tac-toe in which the entire range of alternatives can be exhaustively analyzed by a computer and the machine can thereby be programmed never to lose a game.

1.2.1-2 (**) *Math, Science, and Engineering Aids*

There are few examples of applications in this category (the area into which MYCIN most appropriately falls). Such programs are perhaps best characterized as decision-support systems and, in general, are designed for noncomputer scientists. Some examples of these programs are discussed in § 1.3 and Chapter 3.

1.2.1-3 (**) *Automatic Theorem Proving*

We are all familiar with high school geometry problems in which the task is to use certain given information in order to prove

Introduction

something else about a geometrical figure. The proving of theorems from known axioms is a general problem area common to various other kinds of deductive logic. Some of the earliest AI programs dealt with this kind of theorem proving, and today the field involves some of the most sophisticated applications that have been developed [Nilsson, 1971; Chang, 1973]. This application area is thus closely related to several others (e.g., robot planning, automatic programming) in which theorem proving techniques are often used as the basic problem-solving methodology.

*1.2.1-4 (**) Automatic Programming*

Any computer science student who has slaved into the morning hours, trying to find mistakes in one of his programs, can testify to the "intelligence" required in order to write and debug computer programs that perform specified tasks. The idea of a computer that "figures out" how to program itself may seem absurd at first consideration, but considerable progress has been made in this area in recent years [Balzer, 1972]. For example, one approach to the problem is to give the computer some sample program inputs and the corresponding output data. The machine is then asked to create a program that will perform the required transformation.

*1.2.1-5 (**) Robots*

Science fiction films and modern television notwithstanding, a general purpose robot that walks, talks, and does what you ask it to do has yet to be developed. Work on robotics has involved AI researchers for over a decade, however, and several machines with limited capabilities have been developed [Rosen, 1972; Fikes, 1972; Coles, 1974]. In general this field involves more engineering technology than the other AI application areas because the electrical and mechanical problems in design of the robot itself are substantial. Some projects have limited themselves to computer-controlled arms with associated cameras for scene analysis [Feldman, 1971; Winston, 1972]. These "hand-eye" machines perform tasks in a fixed tabletop environment. Radio-controlled robots on wheels have also been developed [Hart, 1972] and are able to analyze their environment (by means of "on-board" television cameras) and to perform certain

MYCIN

limited tasks. Industry is particularly interested in progress in robotics, as is NASA because of the potential for the use of robots in space exploration. It should be emphasized, however, that the computer program that determines *how* the robot's task is to be accomplished and then sends appropriate signals to the robot's mechanical devices is an essential part of robot technology and underscores this field's association with the other AI application areas.

*1.2.1-6 (**) Machine Vision*

Intimately related to robotics is the development of techniques for analyzing and understanding pictures, usually television pictures [Minsky, 1972; Duda, 1973]. For example, a robot arm that attempts to assemble an engine from parts placed in random locations on a table must be able to locate and recognize the pieces, regardless of their orientation. This problem of scene analysis also involves 3-dimensional perception, edge detection, and disambiguation of lines caused by shadows. Clearly a computer program that makes such judgments on the basis of electrical signals from a television camera is solving a complex intellectual problem.

*1.2.1-7 (**) Natural Language Systems*

Computer understanding of natural language [Schank, 1973; Simmons, 1970; Rustin, 1973], either spoken or written, has fascinated computer scientists ever since attempts were first made, in the 1950's, to write programs for translating from one human language to another (e.g., English to Russian). Researchers in this AI application area are closely involved with the field of linguistics, and have been forced to try to understand the nature of language itself. Their problems include analysis of syntax, disambiguation of words with multiple meanings, and analysis of the semantics of language, especially during a lengthy discourse when the over-all context determines the meaning of individual words. Understanding language typed into a machine by teletype has been extended recently to the development of programs that understand spoken words. The latter problem is similar to machine vision in that the program must first analyze electrical signals (in this case, from a microphone rather than a television camera) in order to determine what has been said. Then

Introduction

an attempt is made to understand the meaning of the words and to have the machine respond appropriately.

*1.2.1-8 (**) Information Processing Psychology*

Many AI researchers, in accordance with Turing's Indistinguishability Test, are concerned primarily with how well their programs perform the tasks for which they were designed; i.e., they do not necessarily care whether the program solves the problem in the same way that a human does. There are those who believe, however, that by attempting to create programs that solve problems in a manner similar to the workings of the mind, new insights into the psychology of human problem-solving can be discovered. Such work has taken several different forms [Newell, 1970; Schank, 1973; Lindsay, 1972] that interface with all seven of the other AI application areas I have discussed.

1.2.2 (**) AI METHODOLOGIES AND TECHNIQUES

Four core topics in artificial intelligence pervade all eight of the application areas discussed above. Section 1.3 and Chapter 3 describe how the MYCIN system has drawn on work in each of these areas.

*1.2.2-1 (**) Modeling and Representation of Knowledge*

Writers in the AI field are fond of citing examples of problems that seem exceedingly difficult until a simplified way of expressing the task is discovered. Consider a favorite such example—a 64-square checkerboard, 8 squares on each side, and a box of dominos. Each domino exactly covers two squares. Thus 32 dominos can be used to cover the entire checkerboard. You are asked to arrange 31 dominos on the board so that all squares are covered except the two squares in diagonally opposite corners.

Many people given this task would immediately begin trying to arrange dominos as requested. However, an individual who thinks about the problem in the right way will quickly announce that the task is impossible. The key here is to notice that the diagonally opposite squares on a square checkerboard are always the same color

MYCIN

Thus, performing the task would require covering 30 squares of one color and 32 squares of the other color. Since every domino must cover one square of each color, dominos arranged on the board must always cover as many squares of one color as the other. Hence the desired final state cannot be achieved (unless some dominos are cut in half).

A variety of modeling and representation schemes has been developed because it has been recognized that the representation of knowledge in the machine may be crucially important to the efficiency with which an AI program is able to perform. These approaches include use of the predicate calculus to represent facts and goals in problem-solving, semantic networks, production systems similar to the grammars that were first proposed by linguists, and procedural representations. The approaches that are most relevant to MYCIN are discussed in Chapter 3.

*1.2.2-2 (**) Reasoning, Deduction, and Problem Solving*

Since several AI applications involve the writing of programs that solve problems, the development of computer-based problem-solving techniques has been a central concern for many researchers in the field. The most common example used to describe the reasoning tasks involved in the so-called "monkey and bananas" problem. Consider a room containing a monkey, a box, and a bunch of bananas hanging from the ceiling. The distances are such that the monkey is unable to get the bananas unless he is standing on the box. The problem, then, is to write a program that derives a plan so that the monkey can get the bananas. Although the problem may at first seem absurdly simple, it must be remembered that computers have no "common sense" knowledge regarding boxes, monkeys, bananas or distances. The program must therefore be told that boxes may be pushed, that pushing has certain effects on a box and on the individual doing the pushing, that boxes may be climbed on, etc. An intelligent program then deduces, from this basic world knowledge, that the best plan is for the monkey to push the box under the bananas, to climb on the box, and finally to grasp the bananas.

This apparently trivial problem has served as the focus for innovative problem-solving techniques during the past decade. Numerous

Introduction

methodologies and representations of the problem have been suggested. Of course, many problems that are more difficult have been solved, but the puzzle of the “monkey and bananas” remains a convenient common ground for explaining suggested new approaches to computer-based reasoning.

*1.2.2-3 (**) Heuristic Search*

In many human problem-solving situations there are a large number of possible decisions or actions that may be taken. Imagine, for example, the large number of possible moves at most points during a game of chess or checkers. Since each action may in turn lead to several additional potential actions or responses, the number of possible decisions two or more steps into the future often becomes unmanageable. Humans therefore develop strategies for quickly discounting or eliminating possible actions that they can easily see are less desirable than the two or three best potential decisions. They can thus concentrate on the smaller number of actions, comparing their possible outcomes, and making a reasoned decision on the basis of the most rational alternatives. Programs for solving problems must be given similar strategies so that the machine’s computational power can be efficiently spent concentrating on a small number of possible actions. Despite the computer’s speed and computational powers, many human problems (such as selecting the best move in a game of chess) are so complex that thorough evaluation of each possible move can be shown to require a near-infinite amount of time! Any trick or strategy that can be used by a program in order to limit the number of alternative actions that it must investigate is known as a heuristic. Hence “heuristic search” is the name of the AI problem area in which researchers attempt to identify good strategies that adequately limit the number of alternatives that must be considered (but do not eliminate the alternative that would prove to be the best if all possibilities were thoroughly considered).

*1.2.2-4 (**) AI Systems and Languages*

A somewhat separate core topic is the development of computing systems and high-level languages for use by AI researchers [Bobrow, 1973]. Since AI applications typically require powerful capabilities

MYCIN

for symbol manipulation, the several common computer languages that emphasize numerical calculations are usually not adequate. Early AI languages emphasized list-processing [Newell, 1957; McCarthy, 1960], but in recent years newer languages have taken on some of the capabilities that originally were left to the applications programmer [Hewitt, 1969; Teitelman, 1974; Rulifson, 1972; Feldman, 1972]. These include search, pattern matching, and backtracking. MYCIN is written in one of these more recent programming languages, a descendant of LISP [McCarthy, 1962] called INTERLISP [Teitelman, 1974].

The brief overview given here has been intended to give you a sense of the kinds of problems and methodologies with which AI is centrally concerned. Perhaps now it is clear why the AI field holds much intuitive appeal for medical researchers who are examining the reasoning processes involved in clinical judgment, medical diagnosis, and the rational selection of appropriate therapy; this point is expanded on in § 1.3. Then, in § 1.4, the medical problem area for which the MYCIN system has been designed is introduced. Finally, § 1.5 introduces the program itself and gives an example of MYCIN's interactive decision making capabilities.

1.3 Computer-Assisted Medical Decision Making

1.3.1 MAJOR PROBLEM AREA

This section concentrates on an area of medical computing that was mentioned only briefly in § 1.1.2-7. Computer-assisted medical decision making fascinates numerous researchers, partly because analysis of human reasoning is itself challenging, but more importantly because modern medicine has become so complex that no individual can incorporate all medical knowledge into his decision making powers. The field has developed along several dimensions. Therefore it is somewhat difficult to devise an organizational structure for examining the work in this area. Three reasonable dimensions for classifying a computer-based system are:

Introduction

- (1) the program's mode of interaction;
- (2) the program's purpose;
- (3) the program's methodology.

I have chosen to summarize the field in terms of dimension (3), i.e., the various methodologies that have been utilized. The other two dimensions merit brief mention, however.

The decision making program's mode of interaction, like that of any computer program, is either on-line with the user (usually under some time-sharing monitor) or remote in a batch-processing or other off-line mode. The majority of such programs now operate on-line, interacting either directly with the decision maker or with someone who will transmit the computer's information to him. There is clearly more opportunity for discourse and explanation in such programs. An interactive system that gives advice in this fashion is often termed a "consultation program."

The "purpose" of a decision making program would provide a useful basis for classification of the field if there were not so much overlap among the categories. There are at least four kinds of programs along this dimension:

- (1) diagnostic programs
- (2) prognostic programs
- (3) treatment planning programs
- (4) educational programs

Programs specifically designed for educational purposes are mentioned in § 1.1.2-8. Any decision program has potential educational side-effects, however, particularly if it is able to explain the basis for its decisions. Similarly, programs for prognosis and treatment planning must in general make a partial diagnosis of the patient's problem (unless that information is provided by the user at the outset). As is described in § 1.4.1, MYCIN explicitly considers both diagnosis and treatment planning, and also has rules based upon patient prognosis that aid in therapy selection. Furthermore, as is explained in Chapter 2, educational capabilities have been an important design consideration during the current research. The MYCIN system is therefore an example of a system that encompasses all four of the "purpose"

MYCIN

categories I have named. Classification of decision making programs on the basis of these subcategories is hence not particularly useful.

You may well ask why I am so intent upon devising a classification scheme for the programs to be discussed in this section. One answer is that classification leads to structure and, in turn, to understanding. It is therefore the very basis of diagnosis itself [Jelliffe, 1973]. The reasoning processes used by a skilled diagnostician are usually poorly understood, even by the expert (see Chapter 4). Researchers attempting to devise computer-based approaches that parallel human decision making must first therefore assign structure to their problem area in some natural fashion. It is helpful to begin by analyzing the diagnostic process itself [Feinstein, 1967; Card, 1970a; Taylor, 1971] and then to seek a reasonable basis for its automation [Lusted, 1968; Gorry, 1970]. The methodology selected undoubtedly reflects both the specific clinical problem area and the researcher's own peculiar biases; in fact, the approaches are so numerous that national conferences have been held to communicate the new diagnostic techniques or applications developed [Jacquez, 1972]. Yet, two basic concepts underlying most methodologies are the use of some classification mechanism and, with very few exceptions [Ledley, 1973], the need for numerical techniques.

If the success of medical decision support programs is measured by user acceptance, however, the field has not produced more than a handful of truly useful programs. Croft has examined this field extensively [Croft, 1972] and suggests that attempts to develop new diagnostic models will be largely unsuccessful until three basic problems are solved:

- (1) lack of standard medical definitions;
- (2) lack of large, reliable medical data bases;
- (3) lack of acceptance of computer-aided diagnosis by the medical profession.

Croft explains the significance of the first two obstacles by observing that the more diseases a model is assigned to diagnose, the more difficult is the diagnostic task and, in turn, the less successful a program is apt to be in reaching correct decisions. Despite Croft's claim that model development should be set aside while his three listed obstacles are first overcome, one may argue that new diagnostic methodologies that pay more attention to the demands of the

Introduction

user are the only reasonable way to overcome point (3). MYCIN has been designed with this goal in mind (see Chapter 2).

Several attempts have been made to standardize medical definitions. These include the Standard Nomenclature of Diseases and Operations (SNDO), the International Classification of Diseases—Adapted (ICDA), and the Systemized Nomenclature of Pathology (SNOP). Few of these are used extensively in daily medical practice other than for certain reporting purposes. Brunjes has proposed an “anamnestic matrix” concept that would permit computer programs to handle nonstandardized input in a standardized fashion [Brunjes, 1971]. In addition, a British group that evaluated observer variation in history taking and examination found significant degrees of disagreement that were largely reduced when a system of agreed definitions was developed and utilized by the participating physicians [Gill, 1973]. MYCIN has avoided some of these problems by using a large synonym dictionary and by phrasing questions in a manner designed to maximize uniformity of user response (see § 3.3.2-2).

1.3.2 DATA RETRIEVAL AS A DECISION AID

The simplest kind of decision support system merely provides the data for others to make the complicated decisions that depend upon the retrieved information. Such systems generally rely on a computer-based information storage system that accumulates large amounts of data for several patients. Coded information may include physical parameters, diagnosis, treatment plan, and responses to therapy. Physicians may then request information on previous patients who match the current patient on the basis of one or more parameters. Detailed information on how other individuals with similar disease have responded to therapy may help the physician select the best treatment plan for his patient or better evaluate the prognosis for an individual with the particular constellation of symptoms [Feinstein, 1972]. Statistical programs may also provide correlation information that is difficult to deduce merely by looking at retrieved data [Fries, 1972]. A number of medical record systems have been designed with data retrieval requirements as an important consideration [Greenes, 1970a; Shortliffe, 1970; Karpinski, 1971; Feinstein, 1971].

MYCIN

1.3.3 DECISIONS BASED ON NUMERICAL COMPUTATIONS

A limited number of medical problem areas are now so well understood that they can be characterized by mathematical formulae. When the computations are complex, physicians are often tempted to take short cuts, making approximations on the assumption that this will compensate for the tendency to forget the formulae or their proper application. Thus computer programs that assist with calculations and their interpretation may be highly useful.

One such clinical problem area is the classification and management of electrolyte and acid-base disorders. The relationship of blood pH to variables such as kidney function and electrolyte levels is well characterized by formulae that utilize the numerical values of blood gas and other laboratory tests. Bleich has written a program that assists the physician with evaluation of such problems [Bleich, 1969, 1971, 1972], and a similar program has been reported by Schwartz [Schwartz, 1970]. These systems were designed primarily to assist physician users. Their developers, therefore, faced many of the same problems of user acceptance and human engineering that have been encountered during the design of MYCIN. Both programs take advantage of time-shared systems with flexible storage mechanisms that permit not only the calculation of patient parameters but also the presentation of useful information regarding the patient's status. Possible etiologies are listed and literature references are given so that the physician may pursue the topic if necessary. A similar program that evaluates the respiratory status of patients in a respiratory care unit, and makes therapeutic recommendations, has also been described [Menn, 1973].

Another problem area in which numerical calculations using well-defined formulae are the primary concern is the customization of drug doses once the agent to be used has been selected. Several examples of programs in this field involve the selection of a digoxin regimen for a patient with heart disease [Sheiner, 1972; Jelliffe, 1972; Peck, 1973]. There is also a program that helps physicians decide on insulin doses for diabetics [Bollinger, 1973]. These systems depend upon a pharmacokinetic model of the body's absorption, metabolism, distribution, and excretion of the drug in question. Inputs to the programs are various clinical parameters for the patient that are then used to calculate the dosage regimen needed to achieve optimal blood levels of the therapeutic agent.

Introduction

1.3.4 PROBABILISTIC APPROACHES TO DECISION MAKING

Most computer-based decision making tools for medical practitioners are based upon statistical decision theory. The methods used range from simple binary decision trees to conditional probability, discriminant analysis, and clustering techniques.

Explicit decision trees offer advantages in that they clearly represent, when diagrammed, an algorithmic approach to diagnosis. Such diagrams, if memorized or easily accessible, may be useful in visualizing a particular patient's status and the clinical parameters that should be checked in order to further define his diagnostic (or prognostic) category. The trees are nondynamic, however, and therefore cannot adjust easily to unexpected findings or to unavailable test results. Furthermore, modification of the trees when they are found to be incomplete or inaccurate can be highly complex due to the subtle interrelationships within such reasoning networks. There are several examples of programs that are at least partially dependent upon tree-structured decision pathways [Warner, 1972a; Sletten, 1973; Brodman, 1966; Button, 1973; Koss, 1971; Meyer, 1973].

By far the most commonly used statistical technique employed for computer-based medical decision making is Bayes' Theorem in its various forms. It is generally utilized as a first-order approximation to conditional probability under the assumption that the patient's signs and symptoms are jointly independent. In Chapter 4, I discuss the theory in some detail and explain why we chose to reject Bayesian analysis as the basis for MYCIN's decision model. When comprehensive patient data are available, however, Bayes' Theorem offers both excellent results and a methodology that lends itself to automation.

In 1964, Warner *et al.* introduced a computer program that aided in the diagnosis of congenital heart disease [Warner, 1964]. Data had been gathered for several hundred patients with congenital cardiac malformations. As a result, all the conditional probabilities needed for the use of Bayes' Theorem could be computed. The program accordingly classified new patients with an accuracy similar to that of cardiologists.

Four years later, Gorry and Barnett presented a program that used the same patient data to give results of similar accuracy [Gorry, 1968a]. However, their program used a modification of Bayes' Theorem (see § 4.2) that permitted diagnoses to be reached in a

MYCIN

sequential fashion. The system was therefore able to suggest the laboratory or physical tests that were most valuable at each step in the decision process. Using a selection function that considered both the current degree of certainty regarding a diagnosis and the cost of additional testing (in terms of money, time delay, and physical pain or inconvenience), the program attempted to minimize the number of tests while maximizing its diagnostic accuracy.

Bayesian programs continue to pervade the literature on computer-based diagnosis. Recent reports from several countries in addition to the United States have presented computer programs using Bayesian analysis both for diagnosis [Gledhill, 1972; Knill-Jones, 1973] and for screening patients who have given automated medical histories [Warner, 1972b]. The technique has been shown to be highly useful in cases where adequate data are available.

Nordyke *et al.* presented an interesting study using Bayes' Theorem and two other mathematical techniques for the diagnosis of thyroid disease [Nordyke, 1971]. Having previously reported a pattern recognition approach to the problem [Kulikowski, 1970], the authors compared both Bayes' Theorem and pattern recognition to a linear discriminant model. "Pattern recognition" is a general term, the interpretation of which depends upon the application area being discussed. In medical diagnosis, the term usually describes a method that "... attempts to extract the most characteristic features of each diagnostic category, rather than trying to discriminate directly between categories. A patient is then classified into the category with which his data shares the most features..." [Nordyke, 1971]. One variation of this technique may be characterized mathematically using a feature extraction procedure that specifies data vectors that may be subjected to cluster analysis. The linear discriminant model, on the other hand, is an attempt to consider the effects of correlation (or second-order interdependence) between characteristics. The discriminant used in the thyroid study is described in detail in the Nordyke paper.

The data used by Nordyke *et al.* were extracted from the records of 2405 patients who had been seen over a six year period for evaluation of thyroid disease. Their results showed that although the pattern recognition technique performed best in identifying ill patients on the basis of historical data alone, it produced an inordinate number of false positives. Bayes' Theorem, on the other hand, gave

Introduction

comparatively better diagnostic accuracy as more physical findings and laboratory test results became available. Their report therefore concludes:

Because each of the methods uses the characteristics of a patient differently, some taking advantage of discriminating information at a given stage better than others, it would seem that a combination of these would be best for a sequential diagnostic procedure. . . . However, since the simpler Bayes method provides comparable results at the pre-laboratory stage of diagnosis, it might prove the most effective clinical aid.

Another technique used for sequential decision making is the Shannon entropy formula [Shannon, 1949]:

$$\text{entropy} = - \sum_i p(X_i) \log p(X_i)$$

Here $p(X_i)$ is the probability that X_i is true (e.g., that the patient has disease D_i). Steps in the sequential process are selected so as to maximize the entropy of the set of possible diagnoses. Several programs have successfully used this selection function [Mullin, 1970; Gleser, 1972], but it should be noted that entropy too is dependent on good probabilistic information.

All the methodologies discussed so far are examples of techniques utilized in the field of decision analysis [Raiffa, 1968]. The last programs for discussion in this subsection are those that encompass several of the techniques—conditional probabilities, decision trees, utility measures, and selection functions for sequential decision making. Ginsberg's program for diagnosis and management of patients with pleural effusions is an excellent example of this kind of eclectic approach [Ginsberg, 1968, 1970]. In addition, one of the early workers with Bayesian diagnostic programs [Gorry, 1968a, 1968b], has gradually broadened his approach to include several additional facets of decision theory. In joint papers published in the *American Journal of Medicine*, he and his coworkers presented a comprehensive look at decision theory as applied to medical diagnosis [Schwartz, 1973], and reported a program that uses the techniques to evaluate the etiology of acute renal failure [Gorry, 1973b]. Although neither their techniques nor their results are unique, their presentation is lucid and complete. It has generated positive commentary [Jelliffe, 1973] at a time when, as I have remarked before,

MYCIN

the acceptance of computers by physicians is in need of reasoned support.

1.3.5 ARTIFICIAL INTELLIGENCE AND MEDICAL DECISIONS

There are relatively few examples of AI programs used for medical decision making. Since 1970, however, a small number of researchers, most of whom have had experience rooted in the traditional approaches described in § 1.3.4, have begun to consider AI techniques. Notable among these is G.A. Gorry, then from M.I.T. He became aware that the purely statistical programs have had three failings that are major impediments to physician acceptance of the systems. First, the programs have no real “understanding” of their problem area. Gorry explains this point as follows [Gorry, 1973a]:

There are several approaches to inferring renal function and assessing whether it is stable or changing. This determination is very important in diagnosis and in choosing management strategies. From the experts, it is possible to obtain the procedure by which they infer a value for renal function. Further, many statements about the interpretation of changes in renal function can be made. To capture the knowledge embodied in these statements, some computer realization of the concept of renal function must be developed.

AI, with its emphasis upon representation of knowledge, offers a natural environment for examining the kind of “concept formation” that Gorry feels is needed.

The second problem is that, even if the traditional programs have been given an understanding of their problem area, they have no mechanism for discussing their knowledge with the user. Physicians are often uninspired by programs that produce a diagnosis and a four-decimal-place probability estimate without being able to answer questions about how the conclusion was reached. Furthermore, physicians attempting to give the programs new information have shared no common language with the computer. Gorry therefore calls for the development of natural language interfaces to permit discourse between physicians and diagnostic programs. Once again AI provides a natural environment for examining this requirement.

The third problem, closely related to the first two, is the need for programs that can explain (i.e., justify) their advice. This capability requires that a program both understand its reasoning processes and be able to generate explanations in a language that is easily under-

Introduction

stood by the physician. Gorry's group has therefore worked on developing knowledge representations and language capabilities that will heighten the acceptability of a system such as their acute renal failure program [Gorry, 1973b]. In Chapter 2, where the design criteria for the MYCIN system are discussed, the similarities between our desiderata and those of Gorry are readily apparent.

The system requirements discussed by Gorry entail more than a natural language "front end" in combination with a statistically-based program. As discussed in Chapter 5, efficient knowledge representation is generally the foundation for man-machine discourse in natural language. Isner's medical knowledge system, for example, has demonstrated the need for an efficient representation scheme, plus a program with problem-solving skills, if a computer system is to communicate with minimally trained users [Isner, 1972]. I do not mean to suggest that statistical theory has no place in AI research. Several AI programs have used traditional numerical techniques [Good, 1970] but have also utilized data structures that facilitate utilization of knowledge in ways that are not possible if system information is stored solely in probability tables. Our own mathematical decision model is introduced in Chapter 3 and discussed in detail in Chapter 4.

Problem-solving techniques from AI also hold a natural appeal for certain researchers in computer-based medical decision making. The various AI methodologies will not be surveyed here because those most pertinent to MYCIN are discussed in Chapter 3. Four medical projects warrant comment in this context, however.

The first is the theory formation system of Pople and Werner [Pople, 1972] that does not attempt diagnosis as such, but does make inferences on the basis of model behavior. The program uses an alternative to deduction and induction—abductive logic [Pople, 1973]. A convincing argument can be made that abduction is the basis for medical diagnosis. Consider, for example, the three statements:

- (1) If a person has pneumonia, then he has a fever;
- (2) John has pneumonia;
- (3) John has a fever.

Deductive logic allows us to derive (3) from (1) and (2); i.e., "since people with pneumonia have fevers, and since John has pneumonia,

MYCIN

John must have a fever." Induction, on the other hand, uses one or more observations of people for whom (2) and (3) hold in order to infer that (1) is true; i.e., "since I have observed several people with pneumonia, all of whom have fevers, it is perhaps generally true that people with pneumonia have fevers." Abduction is the remaining combination, namely using (1) and (3) to infer (2); i.e., "since people with pneumonia have fevers, and since John has fever, *perhaps* it is true that John has pneumonia." Clearly, the last example parallels a clinical diagnosis on the basis of a patient's symptomology.

Pople and Werner use the abductive model as the basis of a program for inferring neuroanatomical explanations of the behavior of human neurons in response to central stimulation. The system also includes a simulator that tests hypotheses by modeling them and seeing whether the observed responses are duplicated. The problem, of course, is that the word "perhaps" is not quantified in our explanation of abduction above. It is therefore unclear how to select between two competing hypotheses that are both abductively supported by the same observation(s). In fact, Bayes' Theorem and the other numerical methods discussed in § 1.3.4 are attempts to solve precisely this problem, although the term "abduction" does not generally appear in the formulation of those techniques.

An Italian group has recently proposed a more quantitative problem-solving approach that uses AI techniques and addresses itself specifically to medical diagnosis [Gini, 1973]. Their central concern, as has been true for several other researchers, is sequential test selection for effective diagnosis, but they propose a model based upon state-transition networks. Having defined operators for transition from one state in the network to another, they present an algorithm for creating a dynamic ordering of the operators on the basis of their "promise." The algorithm interfaces with a heuristic mechanism for obtaining a diagnosis, i.e., for finding a set of tested symptoms that match a particular disease definition. It is probably wise to reserve judgment about the approach until this model has been automated in a computer program, but it initially appears to offer little advantage over other programs (*cf.* pattern recognition) that have attempted to define diseases as sets of symptoms.

As I have described (§ 1.2.1-8), there is a large subfield of AI in which investigators are motivated by an interest in psychology. A psychologist from Duke University has reported a fascinating program based upon this approach to medical diagnosis [Wortman,

Introduction

1972]. He views diagnosis as “. . . a search through a hierarchically organized memory composed of diseases, disease categories, categories of categories, etc. . . along with a parallel hierarchy containing the heuristic decision rules for evaluating these categories.” After asking a neurologist to “think aloud” while solving clinical problems, Wortman analyzed the resulting protocols and wrote a program that attempted to mimic the neurologist’s approach to cerebellar disease diagnosis. Not only did the program perform as well as the expert in subsequent tests (correctly diagnosing the disease in 19 of 20 sample cases), but it also generated protocols that closely resembled those of the neurologist himself. It is important to note, however, that the program’s performance was also based on the expert’s subjective probabilities relating cerebellar symptomology to each of the 16 selected diseases that were the subject of the experiment. As a result, Wortman’s information processing approach still relies upon the availability of data that reflect the preferences of the expert being modeled. MYCIN also needs such information. AI does not necessarily offer a means for avoiding numerical representation of data relationships, but does suggest new and potentially powerful methods for analyzing the problem domain and selecting relevant knowledge. It will be fascinating to observe Wortman’s future work to see if his success continues as the range of possible diagnoses increases and the clinical problem areas are expanded.

Noteworthy work combining AI techniques and mathematical models of disease has been progressing at Rutgers University for the last several years. Like some of the investigators discussed in § 1.3.3, the Rutgers researchers have sought clinical problem areas that could be well-characterized by mathematical models. Envisioning tiered levels of modeling addressed to various degrees of detail, they assert that an appropriate representation scheme provides an important basis for the design of diagnostic strategies [Amarel, 1972]. Their concern reflects a basic agreement with Gorry in his claim that a diagnostic program needs to “understand” the decisions that it reaches [Gorry, 1973a].

The problem area they have selected for testing their approach is the diagnosis and management of glaucoma. This is an ocular disease that may be characterized both by causal relationships over time and mathematical formulae reflecting fluid resistance and flow [Kulikowski, 1971]. They represent disease states in a network based on causal links reflecting various weights (e.g., “always,” “almost al-

MYCIN

ways," "sometimes," "never," etc.). The network provides the basis of a consultation program for ophthalmologists who need help in evaluating a patient's status [Kulikowski, 1972a]. Working in close collaboration with an ophthalmologist, the group has also written programs that permit an expert interactively to modify nodes in the causal network or to add new information to the inferential structure [Kulikowski, 1972b]. The result is a dynamic program that has created considerable interest among clinical professionals to whom it has been presented at a national meeting of ophthalmologists [Kulikowski, 1973]. The causal network and mathematical model lend themselves well to the development of novel strategies for test selection during the consultation process [Kulikowski, 1972c]. Furthermore, the group's agreement with Gorry's call for programs that can explain their decisions [Gorry, 1973a] is reflected in the program's ability to present a "parse" of those portions of the network that explain the patient's current clinical state [Kulikowski, 1974]. Although certain of the program's human engineering features currently leave much to be desired (the organization of questions during a consultation and the motivation for individual queries appear somewhat confusing to this observer), the glaucoma system represents a pleasing blend of mathematical and AI techniques that holds great promise for those medical problem areas that can be adapted to this kind of causal modeling.

It is unfortunately the case that most human disease states are not sufficiently well understood to be characterized by well-defined mathematical formulae. Even causal relationships are seldom understood. MYCIN is a program that attempts to use AI techniques to model decision making in ill-defined areas such as these. After all, experts *do* reach decision when such medical problems arise, and they can usually offer theoretical arguments for making the judgments that they do. Our goal has been to capture such judgmental knowledge and to create a program that uses the information effectively and in a way that is acceptable to the physicians for whom it is designed. These considerations are described in detail in Chapter 2.

1.3.6 PHILOSOPHICAL OBSERVATIONS

Although medical professionals often demonstrate great resistance to computing innovation, obstacles to acceptance are greatest when

Introduction

the application demands "hands-on" use at a computer terminal or when the program appears to take over intellectual functions, transcending housekeeping or simple "number crunching" chores. Decision making systems must therefore overcome huge barriers, not only because they usually demand interaction by the professional and are attacking a problem that demands intelligence, but also because the user of the program is in most cases the physician himself. Of all health professionals, the physician is perhaps most pressed for time and most wedded to a self-image that has been ingrained since medical school. Schwartz has discussed this last point [Schwartz, 1970]:

Physicians as a group have traditionally cherished their ability to learn and retain large numbers of facts, to formulate a differential diagnosis and to carry on decision making activities. Introduction of the computer into these processes could well be viewed by the doctor as devaluating his hard-won medical education and as undermining his intellectual contribution to medical care. This loss of self-esteem would, of course, be exacerbated if the patient were to find in the transfer of many intellectual functions from man to machine a basis for viewing the doctor with diminished admiration and respect. Such loss of status could have serious social, economic, and political consequences for a profession that has historically enjoyed eminence in the public mind.

Concern regarding the attitudes of patients is not without foundation. I recently heard a group of individuals agree that, all other things being equal, if they had to choose between a doctor who used computer-based consultation programs and one who did not, they would select the physician who was "intelligent enough" to make decisions for himself.

And what of today's medical consultants? How will they react if they are made to feel that their professional expertise is no longer in demand because a computer program has intruded into their clinical problem area? The potential economic implications for both the consultant and the practicing physicians are enormous. Not only may the programs infringe directly on the physician's duties, but, by providing decision support for individuals less highly trained than physicians, may contribute to a reorganization of responsibilities among allied health personnel.

Concerns are also often voiced regarding the effect of such programs on medical education [Schwartz, 1970]. It is not uncommon

MYCIN

to hear the suggestion that such programs will remove the motivation for both doctors and medical students to think or read since they will always know that there is a computer program to help them out if there is something they do not know. Schwartz even suggests that the kind of student attending medical school could change because the primary focus of medical training might become the management of a patient's emotional needs.

Partially because the public image of computers has grown to encompass visions of massive data banks monitoring the daily lives of the public, physicians often express concern that computers capable of making decisions will be used to monitor their medical practice. In an age when federal legislation is already threatening the sacred privacy of the individual physician entrepreneur, technical innovations that could potentially automate the peer review process are especially threatening (see, for example, the discussion of MYCIN's possible extension into the monitoring arena, § 6.6).

Finally there are enormous legal questions that remain essentially unanswered at present. Who is culpable if a physician follows a computer's advice and the patient's condition worsens, especially under circumstances when a panel of experts agree that an alternate therapy would have been preferable? Must program designers assume legal responsibilities for their system's mistakes, or does the physician assume ultimate responsibility when he follows a program's advice?

I have proposed a sufficient number of potentially serious questions that you may have begun to wonder whether research in computer-based medical decision making should be encouraged to continue at all! Let us step back for a moment, however, to ask how many of the itemized concerns are valid and how many are the result, rather, of misunderstanding on the part of physicians and the public or of poor public relations efforts on the part of system designers.

Perhaps the most important point to note initially is that many of the programs have been developed in response to a well-demonstrated need. Despite the availability of expert consultants in university environments, the expertise of specialists is either unavailable or over-taxed in many parts of the country. As a result, local physicians are often forced to make decisions that are less than optimal. Furthermore, even experts may find it difficult adequately

Introduction

to incorporate their experience with several thousand patients into coherent diagnostic strategies. In this sense programs with access to large data bases are potentially useful for physicians at all levels of experience.

Secondly, developers of decision support programs must make it clear, both from their system design and from the tone and content with which they present their work to the medical community, that computer programs for medical decision making are meant to be *tools* for the physician, not crutches to replace his own clinical reasoning. There is no reason that a computer-based consultation need be any more threatening than a chest xray or a battery of tests from the clinical chemistry laboratory. If a consultation program prods the physician to consider a diagnosis or treatment that might otherwise have slipped his mind, it has done a service both to him *and* to the patient. Patient education on this point is therefore similarly important. An effort must be made to inform the public that, since certain clinical problems are highly complex, the medical care they receive may be better if their physician seeks the unique capabilities of a computer rather than forging headlong into a diagnostic or therapeutic decision that is based solely upon his current knowledge. After all, few patients object to their physician seeking the advice of a human consultant.

The concern regarding the effect of such programs on medical education may be answered by pointing out that consultation systems, if properly designed, have considerable educational side-effects (see Chapter 2). The physician can therefore become more familiar with the problem area and its important considerations after each consultation session. The result is a growing body of knowledge that may gradually decrease the physician's need for the program's advice. A consultation program's success could in fact be measured in part by the tendency for physicians to become decreasingly reliant upon the system.

What of the specialist's concern that consultation programs will take over his role? There is some basis for this worry because computer-based consultations are likely to be less expensive than consultations with human experts. However, it is likely that most physicians will prefer the advice of fellow doctors when the experts are readily available. The greatest contribution of computer programs is therefore apt to arise at odd hours when consultants are not

MYCIN

accessible (even the specialists may welcome programs that can assume their roles at 4 a.m.!) or in rural or other nonuniversity environments where the expertise simply does not exist. Furthermore, in an era when the shortage of doctors and their maldistribution are reaching crisis proportions [Fein, 1967; Schwartz, 1970], computer innovation that encourages reallocation of health care responsibilities among medical personnel may perhaps be viewed more as a social boon than an economic threat to physicians.

Even the concerns regarding automated monitoring of physicians' habits may be largely overinflated. In § 6.6 a model is proposed for prospective peer review monitoring that could avoid the threats of retrospective punitive actions on the part of utilization review and medical audit committees. The latter practices are abhorrent to many physicians and partially account for organized medicine's opposition to recent legislation that sets up mandatory peer review mechanisms.

Finally, the questions of legal responsibility are difficult ones to answer since the judicial precedents are not yet well established [Hall, 1972]. However, it seems likely that if the consultation programs are designed to serve as decision tools rather than replacements for the physician's own reasoning processes, the responsibility for accepting or rejecting the computer's advice will probably rest with the physician himself. A more complicated problem arises if a physician diagnoses or treats incorrectly after failing to use a computer program that was readily available to him. Despite the legal questions raised, the potential benefits of decision making programs seem sufficiently large that unanswered judicial concerns should not be allowed to interfere with progress in the field.

1.4 Antimicrobial Selection

1.4.1 (**) NATURE OF THE DECISION PROBLEM

An antimicrobial agent is any drug designed to kill bacteria or to arrest their growth. Thus the selection of antimicrobial therapy refers to the problem of choosing an agent (or combination of agents) for use in treating a patient with a bacterial infection. The terms "antimicrobial" and "antibiotic" are often used interchangeably, although the latter actually refers to any one of a number of drugs that are

Introduction

isolated as naturally occurring products of bacteria or fungi. Thus, the well-known penicillin mold is the source of an antibiotic, penicillin, that is used as an antimicrobial. Some antibiotics are too toxic for use in treating infectious diseases but are still used in research laboratories (e.g., dactinomycin) or in cancer chemotherapy (e.g., daunomycin). Furthermore, some antimicrobials (such as the sulfonamides) are synthetic drugs and are therefore not antibiotics. There are also semi-synthetic antibiotics (e.g., methicillin) that are produced in chemical laboratories by manipulating a naturally occurring antibiotic molecule. Throughout this text, I shall not rely upon this formal distinction between "antimicrobial" and "antibiotic" but will, rather, use the terms as though they were synonymous. The following list of commonly used antimicrobial agents will introduce you to the names of several of these agents. The list includes many of the generic drugs (i.e., nonbrand names) with which the MYCIN system is familiar:

ampicillin	ethambutal	penicillin
bacitracin	gentamicin	polymyxin
carbenicillin	INH	rifampin
cephalothin	kanamycin	streptomycin
chloramphenicol	methicillin	sulfisoxazole
clindamycin	nalidixic-acid	tetracycline
colistin	nitrofurantoin	vancomycin
erythromycin	PAS	

This list does not include the several nonbrand name antimicrobials that are chemically related to the generic drugs above but that have some distinctive feature such as a different preferred route of administration.

The name MYCIN is taken from the common suffix shared by several of the antimicrobial agents. It reflects the central concern of the program, namely the selection of an appropriate therapeutic regimen for a patient with a bacterial infection. MYCIN does not yet consider infections caused by viruses or pathogenic fungi, although these other kinds of organisms cause significant diseases that may be difficult to distinguish clinically from disorders with bacterial etiology.

Antimicrobial selection would be a trivial problem if there were a single nontoxic agent effective against all bacteria capable of causing

MYCIN

human disease. However, drugs that are highly useful against certain bacteria are often not the least effective against others. The identity (genus) of the organism causing an infection is therefore an important clue for deciding what drugs are apt to be beneficial for the patient. The following list summarizes the organisms with which MYCIN is familiar. Subtypes are specified only in those cases where the subdivisions have important therapeutic implications:

arizona	mycobacterium-tb
bacteroides	neisseria-gonorrhea
borrelia	neisseria-meningitidis
brucella	neisseria-species
citrobacter	pasteurella
clostridium-botulinum	peptococcus
clostridium-species	proteus-mirabilis
clostridium-tetani	proteus-non-mirabilis
corynebacteria-diphtheriae	providence
corynebacteria-species	pseudomonas
diplococcus-pneumoniae	salmonella
e.coli	serratia
edwardsiella	shigella
enterobacter	staphylococcus-coag+
fusobacterium	staphylococcus-coag-
hafnia	streptobacillus
hemophilus-influenzae	streptococcus-alpha
hemophilus-non-influenzae	streptococcus-anaerobic
herellea	streptococcus-beta(group-A)
klebsiella	streptococcus-beta(non-group-A)
listeria	streptococcus-gamma
mima	streptococcus-group-D
moraxella	streptococcus-microaerophilic
mycobacterium-atypical	treponema
mycobacterium-balnei	vibrio
mycobacterium-leprae	

Selection of therapy is a four-part decision process. First, the physician must decide whether the patient has a significant bacterial infection requiring treatment. If there is significant disease, the organism must be identified or the range of possible identities must be inferred. The third step is to select a set of drugs that may be appropriate. Finally, the most appropriate drug or combination of

Introduction

drugs must be selected from the list of possibilities. Each step in this decision process is described below.

*1.4.1-1 (**)* Is The Infection Significant?

The human body is normally populated by a wide variety of bacteria. Organisms can invariably be cultured from samples taken from a patient's skin, throat, or stool. These normal flora are not associated with disease in most patients and are, in fact, often important to the body's homeostatic balance. The isolation of bacteria from a patient is therefore not presumptive evidence of significant infectious disease.

Another complication is the possibility that samples obtained from normally sterile sites (such as the blood, cerebrospinal fluid, or urinary tract) will become contaminated with external organisms either during the collection process itself or in the microbiology laboratory where the cultures are grown. It is therefore often wise to obtain several samples and to see how many contain organisms that may be associated with significant disease.

Because the patient does have a normal bacterial flora and contamination of cultures may occur, determination of the significance of an infection is usually based upon clinical criteria. Does the patient have a fever? Is he coughing up sputum filled with bacteria? Does he have skin or blood findings suggestive of serious infection? Is his chest x-ray normal? Does he have pain or inflammation? These and similar questions allow the physician to judge the seriousness of the patient's condition and often explain why the possibility of infection was considered in the first place.

*1.4.1-2 (**)* What is the Organism's Identity?

There are a variety of laboratory tests that allow an organism to be identified. The physician obtains a sample from the site of suspected infection (e.g., a blood sample, an aspirate from an abscess, a throat swabbing, or a urine collection) and sends it to the microbiology laboratory for culture. There the technicians first attempt to grow organisms from the sample on an appropriate nutritional medium. Early evidence of growth may allow them to report the morphological and staining characteristics of the organism. However, complete

MYCIN

testing of the organism so that a definite identity is determined usually requires 24–48 hours or more.

The problem with this identification process is that the patient may be sufficiently ill at the time when the culture is first obtained that the physician cannot wait two days before he begins antimicrobial therapy. Early data regarding the organism's staining characteristics, morphology, growth conformation, and ability to grow with or without oxygen may therefore become crucially important for narrowing down the range of possible identities. Furthermore, historical information about the patient and details regarding his clinical status may provide additional useful clues as to the organism's identity.

*1.4.1-3 (**)* What are the Potentially Useful Drugs?

Even once the identity of an organism is known with certainty, its range of antimicrobial sensitivities may be unknown. For example, although a pseudomonas is usually sensitive to gentamicin, an increasing number of gentamicin-resistant pseudomonads are being isolated. For this reason the laboratory will often run *in vitro* sensitivity tests on an organism they are growing, exposing the bacterium to several commonly used antimicrobial agents. This sensitivity information is reported to the physician so that he will know those drugs that are likely to be effective *in vivo*.

Sensitivity data do not become available until one or two days after the culture is obtained, however. The physician must therefore often select a drug on the basis of his list of possible identities plus the antimicrobial agents that are statistically likely to be effective against each of the identities. These statistical data are available from many hospital laboratories (e.g., 82% of E.coli isolated at Stanford Hospital are sensitive *in vitro* to kanamycin) although, in practice, physicians seldom use the probabilistic information except in a rather intuitive sense (e.g., "Most of the E.coli infections I have treated recently have responded to kanamycin").

*1.4.1-4 (**)* Which Drug is Best for This Patient?

Once a list of drugs that may be useful has been considered, the best regimen is selected on the basis of a variety of factors. These

Introduction

include not only the likelihood that the drug will be effective against the organism, but a number of clinical considerations. For example, it is important to know whether the patient has any drug allergies or whether the drug is contraindicated because of his or her age, sex, or kidney status [Kovnat, 1973]. If the patient has meningitis or brain involvement, does the drug cross the blood-brain barrier? Since some drugs can be given only orally, intravenously (IV), or intramuscularly (IM), the desired route of administration may become an important consideration. The severity of the patient's disease may also be important, particularly for those drugs for which use is restricted on ecological grounds [Finland, 1970; Rose, 1968] or which are particularly likely to cause toxic complications. Furthermore, as the patient's clinical status varies over time and more definitive information becomes available from the microbiology laboratory, it may be wise to change the drug of choice or to modify the recommended dosage regimen.

1.4.2 EVIDENCE THAT ASSISTANCE IS NEEDED

The "antimicrobial revolution" began with the introduction of the sulfonamides in the 1930's and penicillin in 1943. The beneficial effects that these and subsequent drugs have had on mankind cannot be overstated. However, as early as the 1950's it became clear that antibiotics were often being misused. A study of office practice involving 87 general practitioners [Peterson, 1956] revealed that antibiotics were given indiscriminately to all patients with upper respiratory infections by 67% of the physicians, while only 33% ever tried to separate viral from bacterial etiologies. Despite attempts to educate physicians regarding this kind of inappropriate therapy, similar data are reported even today [Kunin, 1973].

Antibiotic misuse has recently received wide attention [Scheckler, 1970; Roberts, 1972; Kunin, 1973; Simmons, 1974; Carden, 1974]. The studies have shown that very few physicians go through the methodical decision process that I described in § 1.4.1. In the outpatient environment antibiotics are often prescribed without the physician having identified or even cultured the offending organism [Kunin, 1973]. In 1972, the FDA certified enough of the commonly used antibiotics (2,400,00 kg) to treat two illnesses of average duration for every man, woman, and child in the country. Yet it has

MYCIN

been estimated that the average person has an illness requiring antibiotic treatment no more often than once every 5 to 10 years [Kunin, 1973]. Part of the reason for such overprescribing is the patient's demand for some kind of prescription with every office visit [Muller, 1972]. It is difficult for many physicians to resist such demands, so improved public education is one step toward lessening the problem.

However, antibiotic use is widespread among hospitalized patients as well. Studies have shown that, on any given day, one third of the patients in a general hospital are receiving at least one systemic antimicrobial agent [Roberts, 1972; Scheckler, 1970; Resztak, 1972]. The monetary cost to both patients and hospitals is enormous [Reimann, 1966; Kunin, 1973]. Simmons and Stolley have summarized the issues as follows [Simmons, 1974]:

- (1) Has the wide use of antibiotics led to the emergence of new resistant bacterial strains?
- (2) Has the ecology of "natural" or "hospital" bacterial flora been shifted because of antibiotic use?
- (3) Have nosocomial (i.e., hospital-acquired) infections changed in incidence or severity due to antibiotic use?
- (4) What are the trends of antibiotic use?
- (5) Are antibiotics properly used in practice?
 - Is there evidence that prophylactic use of antibiotics is harmful, and how common is it?
 - Are antibiotics often prescribed without prior bacterial culture?
 - When cultures are taken, is the appropriate antibiotic usually prescribed and correctly used?
- (6) Is the increasingly more frequent use of antibiotics presenting the medical community and the public with a new set of hazards that should be approached by some new administrative or educational measures?

These authors, after stating the issues, proceed to cite evidence that indicates that each of these questions has frightening answers. They show that the effects of antibiotic misuse are so far-reaching that the consequences may often be worse than the disease (real or imagined) being treated!

Our principal concern is with the fifth question, i.e., whether or not physicians are rational in their prescribing habits and, if not, why not? Roberts and Visconti examined these issues in 1,035 patients

Introduction

consecutively admitted to a 500-bed community hospital [Roberts, 1972]. Of 340 patients receiving systemic antimicrobials, only 36% were treated for infection. The rest received either prophylactic therapy (56%) or treatment for symptoms without verified infection (10%). A panel of expert physicians and pharmacists evaluated these therapeutic decisions and only 13% were judged to be rational whereas 66% were assessed as clearly irrational. The remainder were said to be questionable.

Of particular interest were the reasons that therapy was judged to be irrational in those patients for whom some kind of antimicrobial therapy was warranted. This group consisted of 112 patients—50.2% of the 223 patients who were treated irrationally. It is instructive to list the reasons that were cited, along with the corresponding percentages indicating how many of the 112 patients were involved:

Antimicrobial contraindicated in patient	7.1
Patient allergic	2.7
Inappropriate sequence of antimicrobials	26.8
Inappropriate combination of antimicrobials	24.1
Inappropriate antimicrobial used to treat condition	62.5
Inappropriate dose	18.7
Inappropriate duration of therapy	9.8
Inappropriate route	3.6
Culture and sensitivity needed	17.0
Culture and sensitivity indicate wrong antibiotic being used	16.1

The percentages sum to more than 100% because each therapy may have been judged inappropriate for more than one reason. Thus 62.5% of the 112 patients who required antimicrobial therapy but were treated irrationally were given a drug that was inappropriate for the patient's clinical condition. This observation reflects the need for improved therapy selection in patients requiring therapy. This is precisely the decision task with which MYCIN is designed to assist.

The hospital at which Roberts and Visconti conducted their study is certainly not the only institution at which physicians tend to prescribe antimicrobials inappropriately. Macaraeg *et al.* have also reported serious disagreement between some of the practices and opinions of hospital physicians and those of infectious disease experts practicing at the same institution [Macaraeg, 1971]. Recent

MYCIN

review articles [Kunin, 1973; Simmons, 1974] have cited additional studies that have shown similar data.

Now that a need for improved continuing medical education in antimicrobial selection is recognized, there are a variety of valid ways to respond. One is to offer appropriate post-graduate courses for physicians. Another is to introduce surveillance systems for the monitoring and approval of antibiotic prescriptions within hospitals [Edwards, L., 1972; Kunin, 1973]. In addition, physicians should be encouraged to seek consultations from infectious disease experts when they are uncertain how best to proceed with the treatment of a bacterial infection. Finally, an automated consultation system that can substitute for infectious disease experts when they are unavailable or inaccessible could provide a valuable component of the solution to the therapy selection problem. The computer program described in the remainder of this report is an attempt to fill that need.

1.5 MYCIN System

1.5.1 SYSTEM'S ORGANIZATION

MYCIN is an evolving computer program that has been developed to assist physicians who are not experts in the field of the antimicrobials with the decision task discussed in § 1.4.1. Work on the system began early in 1972 when it was recognized that the Stanford community could provide the professional and computing resources necessary for attempting a partial solution to the problem of antibiotic misuse that was discussed in § 1.4.2. The project has involved both physicians, with expertise in the clinical pharmacology of bacterial infections, and computer scientists, with interests in artificial intelligence and medical computing.

The computing techniques used in the development of MYCIN were formulated over several months as the collaborators met in weekly meetings and discussed representative case histories of patients with infections. It was decided to concentrate initially on the process of selecting therapy for patients with bacteremia (i.e., bacteria in the blood). This remains our primary focus to date. As patients with bacteremia were discussed by the clinicians, the project members tried to identify the semi-formal decision criteria that were

Introduction

being used. It gradually became clear that these criteria, once defined, can be expressed as rules that reflect the knowledge of the experts. Thus, MYCIN was developed as a program that could efficiently utilize such rules in an attempt to model the decision processes of the experts from whom they were obtained.

The discussion in § 1.4.1 pointed out that there are four parts to the process of selecting antimicrobial therapy. MYCIN must accordingly follow each of these steps when giving advice to a physician. To reiterate, decision rules have been sought that allow the program to do the following:

- (1) decide whether the patient has a significant infection;
- (2) determine the likely identity of the offending organism;
- (3) decide what drugs are apt to be effective against this organism;
- (4) choose the drug that is most appropriate given the patient's clinical condition;

Approximately 200 such decision rules have been identified to date. This corpus of rules is termed the "knowledge-base" of the MYCIN system.

System knowledge must be contrasted with MYCIN's "data base." MYCIN uses two kinds of data when it gives advice. Information about the patient under consideration is termed "patient data." These data are entered by the physician in response to computer-generated questions during the consultation. "Dynamic data," on the other hand, are the data structures created by MYCIN during the consultation—the deductions it makes and an ongoing record of why these conclusions were reached. This distinction between MYCIN's knowledge-base and data base should be understood because the terms are used in their respective senses throughout this text.

The program itself consists of three subcomponents, each of which performs a specialized task. Subprogram 1 is the Consultation System, that portion of MYCIN that asks questions, makes conclusions, and gives advice. Subprogram 1 is the subject of Chapter 3.

Subprogram 2 is the Explanation System, the component of MYCIN that answers questions from the user and attempts to explain its advice. The need for such a capability is discussed in Chapter 2, and Chapter 5 explains the implementation details of the explanation capability.

MYCIN

Subprogram 3, the most recent addition to MYCIN, is the Rule-Acquisition System. This module permits experts to teach MYCIN new decision rules or to alter pre-existing rules that are judged to be inadequate or incorrect. Chapter 3 also discusses the need for this kind of capability. Since this subprogram presently exists only in preliminary form, its current capabilities and plans for future extensions are discussed in § 6.3 in the chapter describing future work.

Figure 1-1 is an overview of the three subprograms and the way in which they access MYCIN's knowledge and data. The heavy arrows indicate the system's flow of control between the subprograms, while the light arrows represent information flow between program components and MYCIN's knowledge and data.

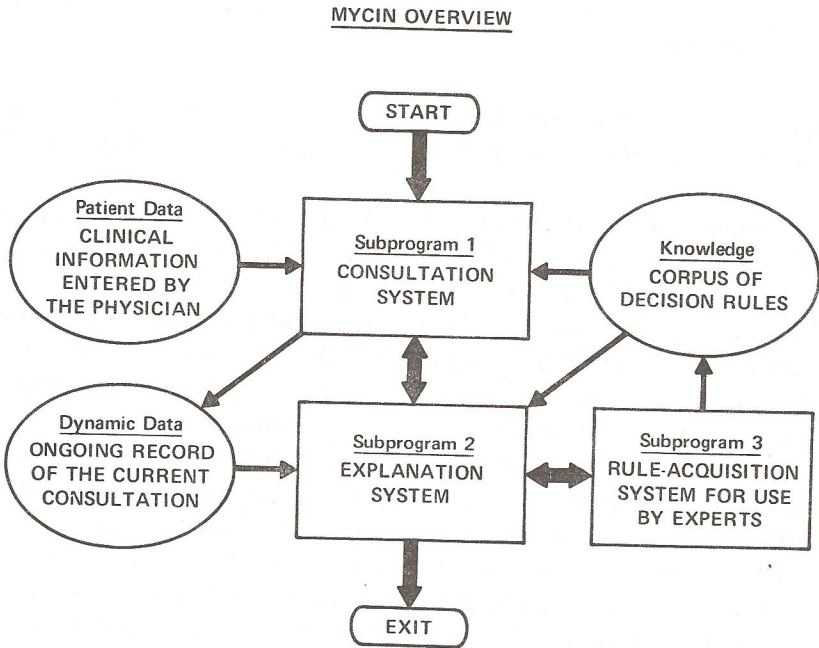


Figure 1-1: Diagram demonstrating the flow of control and the flow of information within the MYCIN system. The three subprogram components are enclosed in boxes. Control passes from one subprogram to another as shown by the heavy arrows. Light arrows indicate program access to information used by the system. The program's knowledge-base is contained in the corpus of rules shown on the right. The way in which the Consultation System uses such rules is described in Chapter 3. [Reproduced from *Computers and Biomedical Research* [Shortliffe, 1975b] with permission of the publishers.]

Introduction

The physician begins an interactive session by starting the Consultation System (Subprogram 1). When MYCIN asks questions, the physician enters patient data as indicated in Figure 1-1. MYCIN uses its knowledge-base to process this information and to decide what question to ask next. Whenever a conclusion is made, MYCIN saves the information in its dynamic data structure. If the physician wants to interrupt the consultation in order to ask questions, he may enter the Explanation System (Subprogram 2). After the question-answering session, he returns to Subprogram 1 and the consultation proceeds from the point of digression.

When MYCIN is through asking questions, it gives its therapeutic recommendation, and control then automatically passes to Subprogram 2. At this point the physician may ask questions regarding the consultation and how MYCIN reached its decisions. This feature forces MYCIN to justify its conclusions and permits the physician to reject the program's advice if he feels that some step in the reasoning process has been unsound.

Subprogram 3 is an option available to experts with whom the system is familiar. If an expert (when using Subprogram 2) notes an invalid, incomplete, or missing rule, he may enter the Rule-Acquisition System in order to teach MYCIN the new information. This new knowledge is then incorporated into the corpus of rules so that it will be available to Subprogram 1 during future consultation sessions. As noted above, this feature currently exists only in rudimentary form.

Throughout all three subprograms there are a variety of features designed to heighten MYCIN's acceptability to physicians. For example, the system is quite tolerant of spelling or typographical errors, and Subprograms 2 and 3 permit the physician to communicate with MYCIN in the language of clinical medicine rather than some specialized computer language. The need for these kinds of human-engineering considerations is discussed in Chapter 2, and the details are described in the chapters that explain each of the subprograms.

1.5.2 SAMPLE CONSULTATION SESSION

This chapter closes with a sample consultation session that illustrates MYCIN's consultation and explanation capabilities. The following comments will help you follow the sample interactive session:

MYCIN

- (1) MYCIN asks questions and makes comments in mixed upper- and lower-case letters;
- (2) MYCIN answers questions and displays decision rules using upper-case letters only;
- (3) The physician types words in upper-case letters and these phrases *always* follow MYCIN's prompt character (**), the double asterisk (note the asterisks distinguish words typed by the physician from upper-case phrases printed by MYCIN);
- (4) Explanatory comments in italics have occasionally been inserted in square brackets at the extreme right margin of the page.

The entire sample consultation session, including the question-answering interaction, required approximately 20 minutes at a computer terminal.

*****Welcome from MYCIN*****

Instructions?: (Yes or No)

**YES

I am a computer program named MYCIN that has been designed to advise you regarding an appropriate choice of infectious disease therapy. I understand that you have a patient from whom a possibly positive culture (called CULTURE-1) has been obtained. Please answer the following questions, terminating each response with RETURN.

If you are not certain of your answer, you may modify the response by inserting a certainty factor (a number from 1 to 10) in parentheses after your response. Absolute certainty (10) is assumed for every unmodified answer.

If you wish to change an answer to a previous question, enter CHANGE followed by the number(s) of the question(s) to be altered. Try to avoid going back, however, because the process requires reconsidering the patient from the beginning and it may therefore be slow.

Note that you may also enter UNK (for UNKnown) if you do not know the answer to a question, ? if you would like to see some examples of recognized responses, ?? if you want to see all recognized responses, the word RULE if you would like to see the decision rule which has generated the question being asked, the word WHY if you would like to see a more detailed explanation of the question, or the letters QA if you would like to interrupt the consultation in order to ask questions regarding the current

Introduction

status of the program's reasoning. If you are ever puzzled about what options are available to you during a consultation, enter the word HELP and a list of options will be listed for you.

Sample Response

Did ORGANISM-1 grow in clumps, chains, or pairs?

**CHAINS (7) PAIRS (3)

SUMMARY OF USER OPTIONS:

UNK	- answer not known
?	- request for examples of recognized responses
??	- request for list of all recognized responses
RULE	- request for current decision rule
QA	- request for enter question-answering mode
WHY	- request for detailed explanation of the question being asked
CHANGE XXX	- go back and re-request answer to question number XXX
STOP	- Halts the program without completing the consultation
HELP	- Prints this list

[At this point the Consultation System (Subprogram 1) begins.]

7-JUN-74 16:27:27

-----PATIENT-1-----

1) Patient's name: (first-last)

**JEREMIAH SAMPLE

2) Sex:

**MALE

3) Age:

**60

I will refer to the most recent positive culture for which you desire therapeutic advice as:

-----CULTURE-1-----

4) From what site was the specimen for CULTURE-1 taken?

**BLODD

=BLOOD

5) How many days ago was this blood culture (CULTURE-1) obtained?

**2

MYCIN

I will refer to the first offending organism from this blood culture (CULTURE-1) as:

-----ORGANISM-1-----

6) Enter the identity (genus) of ORGANISM-1:

**UNK

7) Staining characteristics of ORGANISM-1 (gram):

**?

WHAT IS THE STAIN OF ORGANISM-1?

Expected responses are: ACIDFAST GRAMPOS GRAMNEG

Enter HELP for list of user options.

**GRAMNEG

8) Is ORGANISM-1 a rod or coccus (etc.):

**ROD

9) Were any organisms that were significant (but no longer require therapeutic attention) isolated within the last approximately 30 days?

**YES

-----CULTURE-2-----

10) From what site was the specimen for CULTURE-2 taken?

**SPUTUM

11) How many days ago was this upper-resp culture (CULTURE-2) obtained?

**4 1/2 DAYS

PLEASE USE DECIMAL NUMBERS FOR FRACTIONS

4 1/2 DAYS NOT A RECOGNIZED RESPONSE

WHAT IS THE NUMBER OF DAYS SINCE THE SPECIMEN FOR THIS UPPER-RESP CULTURE WAS ACQUIRED?

Expected responses are: NUMBER OF DAYS

Enter HELP for list of user options.

**4.5

I will refer to the prior organism from this upper-resp culture (CULTURE-2) as:

-----ORGANISM-2-----

12) Enter the identity (genus) of ORGANISM-2:

**PSEUDOMONAS

13) Any other significant organisms isolated from this upper-resp culture (CULTURE-2)?

**YES

Introduction

-----ORGANISM-3-----

14) Enter the identity (genus) of ORGANISM-3:

**KLEBSIELA

=KLEBSIELLA

15) Any other significant organisms isolated from this upper-resp culture (CULTURE-2)?

**YES

-----ORGANISM-4-----

16) Enter the identity (genus) of ORGANISM-4:

**ENTEROB

=ENTEROBACTER

17) Any other significant organisms isolated from this upper-resp culture (CULTURE-2)?

**NO

18) Any other significant earlier cultures from which pathogens were isolated?

**NO

19) What is the suspected non-sterile portal of entry of ORGANISM-1 into this sterile site?

**RULE

RULE121

IF: 1) THE SITE OF THE CULTURE IS ONE OF: THOSE SITES THAT ARE NORMALLY STERILE, AND
2) THE PORTAL OF ENTRY OF THE ORGANISM IS THE SAME AS THE SITE OF A PRIOR ORGANISM WITH THE SAME IDENTITY AS THIS CURRENT ORGANISM
THEN: THERE IS FURTHER SUGGESTIVE EVIDENCE (.3) THAT EACH OF THEM IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS THIS CURRENT ORGANISM

**GU

20) Did ORGANISM-1 grow aerobically?

**YES

21) Did ORGANISM-1 grow anaerobically?

**NO

22) Was any chemotherapy recently begun specifically to combat ORGANISM-1?

**QA

MYCIN

[The physician indicates that he wants to enter the Explanation System (Subprogram 2) for a question-answering (QA) session.]

Type HELP for the list of question options.

** IQ WHAT DO YOU THINK THE IDENTITY OF ORGANISM-1 IS AT
** THIS POINT?
THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS (.333) KLEB-
SIELLA (.333) ENTEROBACTER (.333)

**RETURN

[Having asked the question of interest, the physician returns to the Consultation System (Subprogram 1).]

22) Was any chemotherapy recently begun specifically to combat ORGANISM-1?

**NO

23) Is Jeremiah Sample (PATIENT-1) a burn patient?

**RULE

RULE047

IF: 1) THE SITE OF THE CULTURE IS BLOOD, AND
2) THE IDENTITY OF THE ORGANISM IS NOT KNOWN WITH CERTAINTY, AND
3) THE STAIN OF THE ORGANISM IS GRAMNEG, AND
4) THE MORPHOLOGY OF THE ORGANISM IS ROD, AND
5) THE PATIENT HAS BEEN SERIOUSLY BURNED

THEN: THERE IS WEAKLY SUGGESTIVE EVIDENCE (.4) THAT THE IDENTITY OF THE ORGANISM IS PSEUDOMONAS

**NO

24) Is the patient's illness with ORGANISM-1 a hospital-acquired infection?

**YES

25) Do you consider Jeremiah Sample (PATIENT-1) to be a compromised host? (alcoholic, immunosuppressed, etc.):

**NO

26) Has Jeremiah Sample (PATIENT-1) had any gu surgery or

Introduction

manipulation of the urinary tract? (e.g., indwelling catheter, cystoscopy, etc.)

****WHY**

I AM TRYING TO FIND OUT ABOUT THE IDENTITY OF THE ORGANISM.

I ALREADY KNOW THAT:

THE PORTAL OF ENTRY OF THE ORGANISM IS UPPER-GU

THE MORPHOLOGY OF THE ORGANISM IS ROD

THE STAIN OF THE ORGANISM IS GRAMNEG

THE SITE OF THE CULTURE IS BLOOD

THEREFORE, IF:

1) THE PATIENT HAS NOT HAD A GENITO-URINARY MANIPULATIVE PROCEDURE, AND

2) URINARY-TRACT-INFECTION 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

(RULE156)

[The WHY option demonstrated here is briefly mentioned in Chapter 2. Its implementation details are not described in this text, however.]

****YES**

27) Does the patient have 'pseudomonas-type' hemorrhagic skin lesions (ecthyma gangrenosum)?

****YES**

28) Was the patient treated for the injection with the enterobacter (ORGANISM-4)?

****NO**

29) Do you feel that the finding of the enterobacter (ORGANISM-4) at this normally non-sterile site is suggestive of disease in this patient?

****NO**

30) Do you believe that the enterobacter (ORGANISM-4) is a contaminant (or normal flora in the case of non-sterile sites)?

****NO**

31) Was the patient treated for the infection with the klebsiella (ORGANISM-3)?

****NO**

MYCIN

32) Do you feel that the finding of the klebsiella (ORGANISM-3) at this normally non-sterile site is suggestive of disease in this patient?

**NO

33) Do you believe that the klebsiella (ORGANISM-3) is a contaminant (or normal flora in the case of non-sterile sites)?

**UNK

34) Was the patient treated for the infection with the pseudomonas (ORGANISM-2)?

**NO

35) Do you feel that the finding of the pseudomonas (ORGANISM-2) at this normally non-sterile site is suggestive of disease in this patient?

**YES (8)

36) Has the lab reported the in-vitro antibiotic sensitivities of the pseudomonas (ORGANISM-2)?

**NO

37) Has Jeremiah Sample (PATIENT-1) had any operations that you feel may be relevant to the current infection?

**YES

-----OPERATION-1-----

38) How long ago was OPERATION-1 performed?

**11 DAYS

39) Were any antibiotics given during OPERATION-1?

**YES

-----DRUG-1-----

40) What drug?

**CEFALOTHIN

=CEPHALOTHIN

41) Were any other antibiotics given at the time of OPERATION-1?

**NO

42) Any other significant operative procedures?

**NO

43) Including this blood culture (CULTURE-1), how many cultures were obtained from this normally sterile site on the same day?

**RULE

RULE044

IF: 1) THE SITE OF THE CULTURE IS ONE OF: THOSE
SITES THAT ARE NORMALLY STERILE, AND

Introduction

- 2) THE NUMBER OF CULTURES IN THE SERIES INCLUDING THE CULTURE IS GREATER THAN 1, AND
3) THE NUMBER OF CULTURES CONTAINING THE ORGANISM IS GREATER THAN 1

THEN: THERE IS EVIDENCE (CF= THE QUOTIENT OF THE NUMBER OF CULTURES CONTAINING THE ORGANISM DIVIDED BY THE NUMBER OF CULTURES IN THE SERIES INCLUDING THE CULTURE) THAT THERE IS SIGNIFICANT DISEASE ASSOCIATED WITH THIS OCCURRENCE OF THE ORGANISM

**4

44) How many of the cultures were positive for ORGANISM-1?

**2

[Note that the certainty factor (CF) for RULE044 is calculated from the answers to questions 43 and 44. CF's are explained in Chapter 4.]

45) Any other organisms isolated from this blood culture (CULTURE-1) for which you would like a therapeutic recommendation?

**NO

46) Does Jeremiah Sample (PATIENT-1) show any evidence of infection at a site inaccessible to culture (e.g., brain, bone, liver, etc.)?

**NO

47) Were any additional positive specimens recently drawn from Jeremiah Sample (PATIENT-1) which revealed one or more organisms for which you wish therapeutic advice?

**NO

[MYCIN has now asked all questions necessary in order to decide (a) whether ORGANISM-1 is significant, and (b) the likely identities of ORGANISM-1.]

My therapy recommendation will be based on the following possible identities of the organism(s) that seem to be significant:

- <Item 1> The identity of ORGANISM-1 may be PSEUDOMONAS
- <Item 2> The identity of ORGANISM-1 may be KLEBSIELLA
- <Item 3> The identity of ORGANISM-1 may be ENTEROBACTER
- <Item 4> The identity of ORGANISM-1 may be E.COLI
- <Item 5> The identity of ORGANISM-1 may be PROTEUS (Subtype unknown)

MYCIN

[MYCIN now begins asking questions to assist with drug selection.]

48) On a scale of 0 to 4, where higher numbers indicate increasing severity, how would you rate Jeremiah Sample (PATIENT-1)'s degree of sickness?

**4

49) Does Jeremiah Sample (PATIENT-1) have a clinically significant allergic reaction to any antimicrobial agent?

**NO

[Considering GENTAMICIN for use against ORGANISM-1 . . .]

[Rec 1] My preferred therapy recommendation is as follows:

In order to cover for Items <1> <2> <3> <4> <5>:

Give the following in combination:

1 GENTAMICIN

Dose: 1.7 MG/KG Q8H - IV OR IM

Comments: MODIFY DOSE IN RENAL FAILURE

2 CARBENICILLIN

Dose: 25 MG/KG Q2H - IV

Comments: MODIFY DOSE IN RENAL FAILURE

Do you wish to see the next choice?

**YES

[Considering COLISTIN for use against ORGANISM-1 . . .]

[Rejecting SULFISOXAZOLE for use against ORGANISM-1 . . .]

[Rec 2] Next best therapy recommendation:

In order to cover for Items <1> <2> <3> <4>:

Give: COLISTIN

Dose: 1.25 MG/KG Q6H - IM

Comments: MODIFY DOSE IN RENAL FAILURE

Do you wish to see the next choice?

**NO

[The Consultation System (Subprogram 1) is now finished. Control therefore automatically passes to the Explanation System (Subprogram 2).]

Please ask any simple-English questions regarding the above consultation session or the overall knowledge of the consultation system. When discussing the consultation above, please refer to

Introduction

organisms, cultures, etc. by their numeric names (e.g., ORGANISM-1). The names of rules need not be printed in full when using the PR option ('PR 37' is okay), but use the whole name in other contexts (e.g., RULE037). Questions may be specified by the numbers that preceded them above. Enter the word STOP when you have no further questions. Certain standard question types may also be preceded by special codes that will speed up the answers to your questions.

Type HELP for the list of question options.

**HELP

QUESTION-ANSWERING (QA) OPTIONS:

- HELP - PRINTS THIS LIST
- EQ - REQUESTS AN EXPLANATION OF THE SPECIFIED QUESTION(S) FROM THE CONSULTATION
- IQ - PREFIX TO A QUESTION WHICH ASKS ABOUT INFORMATION ACQUIRED BY THE PROGRAM DURING THE CONSULTATION
- NO PREFIX - THIS QUESTION QUERIES CONTENTS OF DECISION RULES IN THE SYSTEM
- PR - REQUESTS THAT SPECIFIED RULE(S) BE PRINTED
- STOP - ESCAPE FROM EXPLANATION SYSTEM
- RA - ENTRY TO RULE-ACQUISITION MODULE FOR RECOGNIZED EXPERTS

** IQ WHAT IS THE FINAL DECISION REGARDING THE IDENTITY OF
** ORGANISM-1?

THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS (.357)
KLEBSIELLA (.177) ENTEROBACTER (.166) E. COLI (.137)
PROTEUS (.118) SERRATIA (.042)

** WHAT MADE YOU THINK ORGANISM-1 MIGHT BE A PROTEUS?

I USED RULE163 TO CONCLUDE THAT THE IDENTITY OF
ORGANISM-1 IS PROTEUS. THIS GAVE A CUMULATIVE
C.F. OF .277. THE LAST QUESTION ASKED BEFORE THE
CONCLUSION WAS MADE WAS 26.

I USED RULE084 TO CONCLUDE THAT THE IDENTITY OF
ORGANISM-1 IS PROTEUS. THIS GAVE A CUMULATIVE
C.F. OF .119. THE LAST QUESTION ASKED BEFORE THE
CONCLUSION WAS MADE WAS 24.

** PR 84 163

MYCIN

RULE084

- IF: 1) THE SITE OF THE CULTURE IS ONE OF: THOSE
 2) THE LIST OF LIKELY PATHOGENS ASSOCIATED
 3) THIS ORGANISM AND THE MEMBERS THAT YOU
 4) ARE CONSIDERING AGREE WITH RESPECT TO THE
 5) FOLLOWING PROPERTIES: GRAM MORPH AIR
- THEN: THERE IS STRONGLY SUGGESTIVE EVIDENCE (.8)
 THAT EACH OF THEM IS THE IDENTITY OF THE
 ORGANISM

RULE163

- IF 1) THE MORPHOLOGY OF THE ORGANISM IS ROD, AND
 2) THE STAIN OF THE ORGANISM IS GRAMNEG, AND
 3) THE IDENTITY OF THE ORGANISM IS NOT KNOWN
 4) WITH CERTAINTY, AND
 5) THE PATIENT HAS HAD A GENITO-URINARY
 6) MANIPULATIVE PROCEDURE
- THEN: 1) THERE IS WEAKLY SUGGESTIVE EVIDENCE (.36)
 2) THAT THE IDENTITY OF THE ORGANISM IS
 3) PSEUDOMONAS, AND
 4) THERE IS WEAKLY SUGGESTIVE EVIDENCE (.18)
 5) THAT THE IDENTITY OF THE ORGANISM IS PROTEUS,
 6) AND
 7) THERE IS WEAKLY SUGGESTIVE EVIDENCE (.1)
 8) THAT THE IDENTITY OF THE ORGANISM IS SERRATIA,
 9) AND
 10) THERE IS WEAKLY SUGGESTIVE EVIDENCE (.18)
 11) THAT THE IDENTITY OF THE ORGANISM IS
 12) ENTEROBACTER, AND
 13) THERE IS WEAKLY SUGGESTIVE EVIDENCE (.18)
 14) THAT THE IDENTITY OF THE ORGANISM IS
 15) KLEBSIELLA

** HOW DID YOU DECIDE THAT PSEUDOMONAS WAS THE MOST
** LIKELY IDENTITY OF ORGANISM-1?

I USED RULE163 TO CONCLUDE THAT THE IDENTITY OF
ORGANISM-1 IS PSEUDOMONAS. THIS GAVE A

Introduction

- CUMULATIVE C.F. OF .651. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 26.
- I USED RULE084 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS. THIS GAVE A CUMULATIVE C.F. OF .455. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 24.
- I USED RULE050 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS. THIS GAVE A CUMULATIVE C.F. OF .352. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 24.
- I USED RULE040 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS. THIS GAVE A CUMULATIVE C.F. OF .833. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 27.
- I USED RULE021 TO CONCLUDE THAT THE IDENTITY OF ORGANISM-1 IS PSEUDOMONAS. THIS GAVE A CUMULATIVE C.F. OF .19. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 21.

** PR 21

RULE021

- IF: THERE IS A PRIOR ORGANISM WHICH MAY BE THE SAME AS THE CURRENT ORGANISM
- THEN: IT IS DEFINITE (1) THAT THESE PROPERTIES - IDENTITY - SHOULD BE TRANSFERRED FROM A PRIOR ORGANISM WHICH MAY BE THE SAME AS THE ORGANISM TO THIS CURRENT ORGANISM

** IQ DID YOU DECIDE ORGANISM-2 WAS THE SAME AS

** ORGANISM-1?

ORGANISM-4 (.19) ORGANISM-3 (.19) ORGANISM-2 (.19) IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS ORGANISM-1

** HOW DID YOU DECIDE THAT ORGANISM-2 WAS THE SAME AS

** ORGANISM-1?

I USED RULE005 TO CONCLUDE THAT THE PSEUDOMONAS IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS ORGANISM-1. THIS GAVE A CUMULATIVE C.F. OF .19. THE LAST QUESTION ASKED BEFORE THE CONCLUSION WAS MADE WAS 21.

I USED RULE004 TO CONCLUDE THAT THE PSEUDOMONAS IS A PRIOR ORGANISM WITH THE SAME IDENTITY AS

MYCIN

ORGANISM-1. THIS GAVE A CUMULATIVE C.F. OF .133.
THE LAST QUESTION ASKED BEFORE THE CONCLUSION
WAS MADE WAS 18.

** PR 4 5

RULE004

IF: 1) THE IDENTITY OF THE ORGANISM IS NOT KNOWN
WITH CERTAINTY, AND
2) THIS CURRENT ORGANISM AND PRIOR ORGANISMS
OF THE PATIENT AGREE WITH RESPECT TO THE
FOLLOWING PROPERTIES: MORPH GRAM
THEN: THERE IS WEAKLY SUGGESTIVE EVIDENCE (.4) THAT
EACH OF THEM IS A PRIOR ORGANISM WITH THE SAME
IDENTITY AS THE ORGANISM

RULE005

IF: THIS CURRENT ORGANISM AND A PRIOR ORGANISM
WITH THE SAME IDENTITY AS THE ORGANISM AGREE
WITH RESPECT TO THE FOLLOWING PROPERTIES: AIR
THEN: THERE IS WEAKLY SUGGESTIVE EVIDENCE (.2) THAT
EACH OF THEM IS A PRIOR ORGANISM WITH THE SAME
IDENTITY AS THE ORGANISM

** EQ 48

48 QUESTION 48 WAS ASKED IN ORDER TO FIND OUT JEREMIAH
SAMPLE'S DEGREE OF SICKNESS (ON A SCALE OF 4) IN AN
EFFORT TO EXECUTE RULE068.

** PR 68

RULE068

IF: 1) THE THERAPY UNDER CONSIDERATION IS
GENTAMICIN, AND
2) THE IDENTITY OF THE ORGANISM MAY BE
PSEUDOMONAS, AND
3) THE PATIENT'S DEGREE OF SICKNESS (ON A SCALE
OF 4) IS GREATER THAN 3, AND

Introduction

4) [A - THE PATIENT IS NOT ALLERGIC TO ANTI-BIOTICS, OR B - CARBENICILLIN IS NOT A DRUG TO WHICH THE PATIENT IS ALLERGIC]

THEN: RECORD THE FOLLOWING AS A MORE APPROPRIATE THERAPY THAN GENTAMICIN: THE COMBINATION OF GENTAMICIN WITH CARBENICILLIN

** STOP

50) Do you wish advice on another patient?

**NO

THANK-YOU

* * * * *