Report 77-10
Stanford -- KSL

Scientific DataLink

A DSS for Diagnosis and Therapy.
Randall Davis,
Winter 1977

card 1 of 1

DATA BASE

A QUARTERLY NEWSLETTER OF SIGBDP VOLUME 8 NUMBER 3 WINTER 1977

PROCEEDINGS OF A CONFERENCE ON DECISION SUPPORT SYSTEMS

Edited by Eric D. Carlson

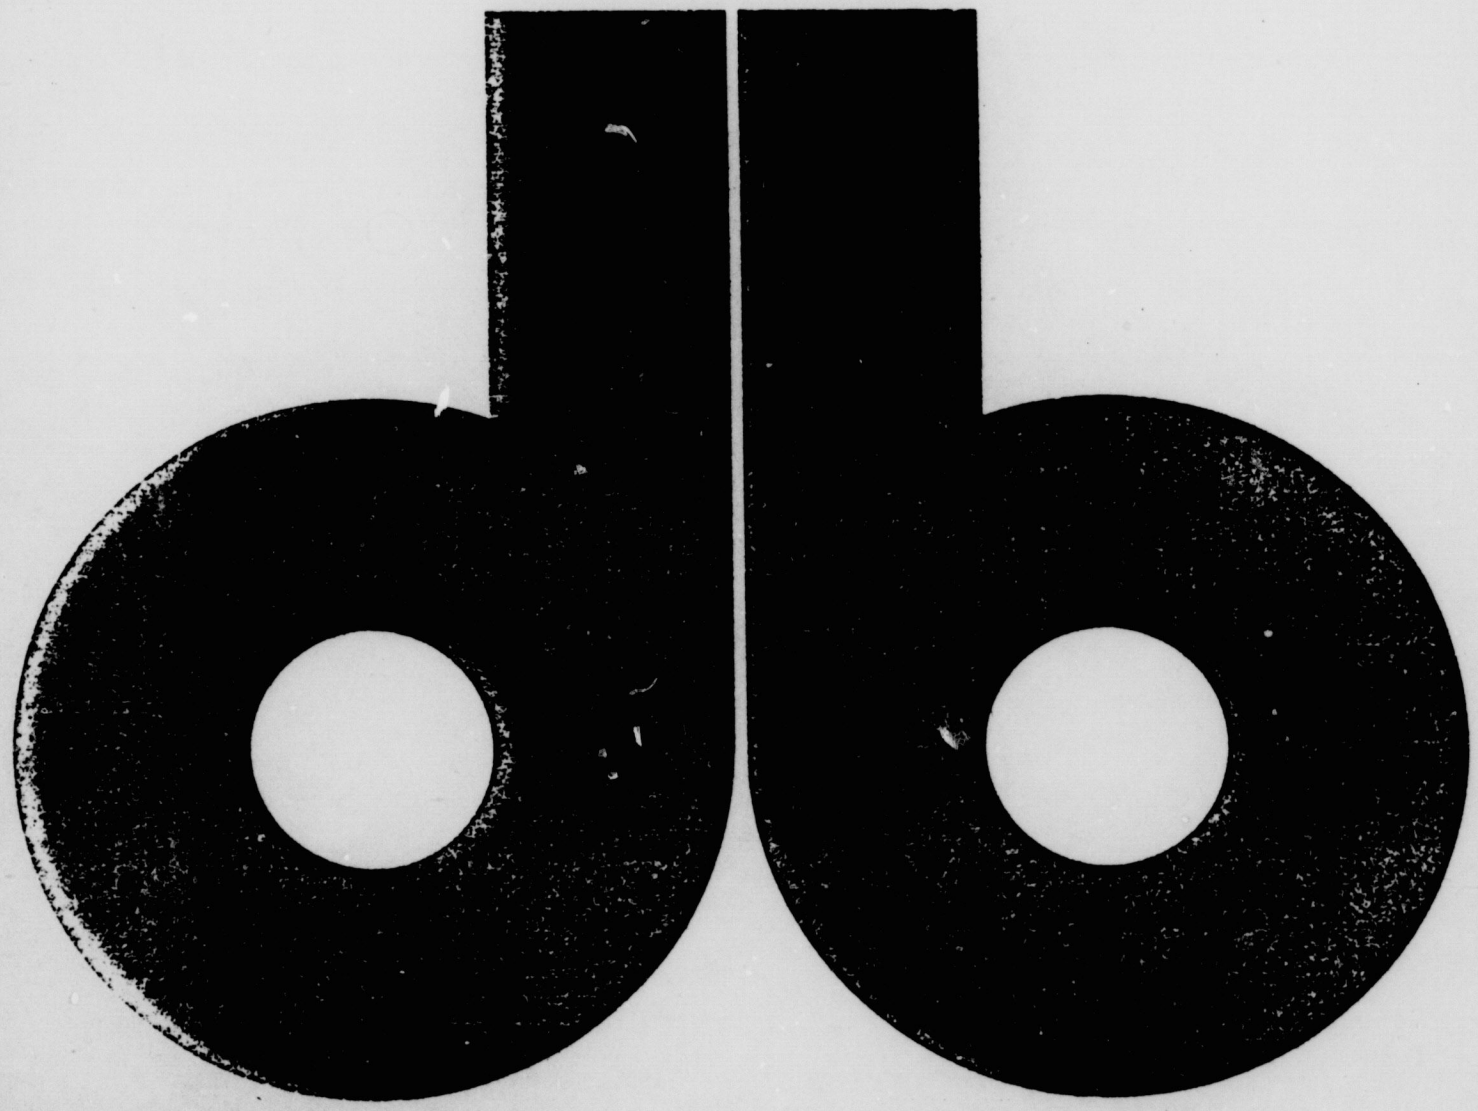

77-10

A DSS FOR DIAGNOSIS AND THERAPY

By Randall Davis
Computer Science Department
Stanford University

Abstract

This paper reviews an approach to the design and construction of a decision support system intended to function as a consultant on the question of diagnosis and therapy selection. It describes the system in terms of the nature of the decision problem involved, discusses factors which make the problem difficult and considers the design goals that have led to the construction of a system with several novel capabilities. Many of those capabilities result from representing domain-specific knowledge in the system in terms of numerous judgmental decision rules, and we examine a number of such rules. Examples of the system in operation are given to illustrate many of these issues, and performance is compared with previous approaches to automated medical decision making. Finally, we consider the domain independence and generality of the methodology and consider the potential impact the system may have as a tool for decision support.

Introduction

Over the past five years, a group of computer scientists and clinicians at Stanford University has developed MYCIN, a computer program designed to function as a consultant on problems of diagnosis and therapy selection.[1] Its current domain of application is infectious disease, and it displays an encouraging level of performance in dealing with cases of bacteremia (bacterial infections of the blood) and meningitis. In this paper we examine the system from the perspective of the medical decision-making process, considering the problems encountered and the solutions that have been developed to meet them.

Nature of the Problem

MYCIN's fundamental task is to determine the identity and significance of organisms causing an infection and to select the appropriate drug(s) (if any) for treatment. A typical clinical situation begins with a patient showing signs of infection, and a specimen (of blood, urine and so on) is obtained and cultured to check for the presence of disease-causing bacteria. While cultures may show some evidence of bacterial growth within 12 hours, typically 24 to 48 hours are required for positive identification of the organisms. Treatment cannot often wait that long, so the physician must base a decision on whatever information is available. This typically includes several easily observable characteristics of the bacteria in the culture (overall shape, response to oxygen and so on), as well as the history of the patient (previous infections, other clinical evidence of infection or events that may make the patient particularly

susceptible to a particular type of bacteria).

Two fundamental characteristics of this information are central to the view of MYCIN as a system for medical decision making: the information is both *incomplete* and *inexact*. Incompleteness may arise from time constraints (as is the case with organism identity), gaps in the patient's medical records or simply because it is impractical to administer an exhaustive series of tests before initiating treatment. Inexactness is inherent in the domain because many test results are qualitative (rather than quantitative) and because important factors are often based on subjective impressions ("has the patient responded to previous therapy?").

A further complication arises from the fact that this inexactness affects not only the data on which decisions are based, but the decision criteria themselves. In common with other domains, there are relatively few statements that can be made with absolute certainty. Instead, the system must have some means of representing the fact that "X and Y *seem to suggest* Z," or "A and B *tend to rule out* C." Further sections below will make clear how the system deals with both the incompleteness of information and the two forms of inexactness.

One of the inherently difficult aspects of decision making in this domain is the presence of conflicting goals in therapy selection. On one hand, therapy should be broad enough to cover all organisms believed to be present, so the physician wants to prescribe enough drugs to take care of all reasonable possibilities. But every drug has side effects, a certain toxicity and the potential for interaction with other drugs. Thus the total number of drugs prescribed should be kept as small as possible. Then there are so-called wide-spectrum drugs that can act against a broad range of organisms, and hence reduce the number of drugs administered. But these may be less effective against an individual organism than a drug chosen specifically for that organism, so the physician wants to tailor the therapy more specifically. Yet the physician cannot simply choose the drug that would be most effective against each organism. This would often result in too many drugs and undesired ecological effects, since overuse of a drug promotes the proliferation of strains of bacteria resistant to it. Therapy selection is thus a difficult balancing act in which the physician must consider a number of goals simultaneously.

Suitability of the Domain

There are a number of reasons why infectious disease diagnosis and therapy was chosen as an appropriate domain in which to develop the system. There is, first, a constant need in the hospital for consultative advice of this sort. Infections are often developed as secondary effects of other events (surgery, burns, wounds), so the physician caring for the patient may not be an expert in infectious disease.

Second, as numerous recent studies have shown, there is a significant problem of antibiotic misuse. In a recent year, almost one out of every four people in the country was given penicillin, yet almost *90 percent* of those prescriptions were unnecessary [Kagan73]. One study [Ray76] examined the use of a common antibiotic (chloramphenicol) in 992 cases, and concluded that virtually *all* of the drug was prescribed inappropriately. The problem arises from a range of factors, including patient pressure for therapy when none may be necessary, the temptation on the part of the doctor to use a broad spectrum drug in lieu of a more precise diagnosis [Neu75] and the "antibiotic revolution" that started with the discovery of the first natural antibiotics and has led to a bewildering array of ever more new drugs.

Finally, antimicrobial therapy appears to be an especially suitable domain because the components of the decision-making process are more readily definable than in many other areas of medicine, and the consequences of the physician's decision can usually be assessed in terms of the direct therapeutic action. In our experience, this has proved to be a crucial factor in the development of the program. Before a decision support system can provide useful service, it must be possible to specify in detail the form of the decision criteria to be employed. One of the most difficult tasks in developing the MYCIN system has been the elucidation, testing and validation of the previously informal decision criteria used by physicians. This task would have been far more difficult in a domain where the results of therapeutic actions were less easily determined.

Previous Decision-Making Approaches

It will be useful to take a brief look at some of the other approaches that have been taken to the problem of medical decision making to motivate some of the features and capabilities of the MYCIN system. Three other approaches have received extensive attention in the literature:

1. Decision trees—as in [Meyer73], in which a sequence of decisions is structured in the form of a tree. Each node represents a particular question, and the answer determines which branch of the tree to follow to get to the next question. Final results are obtained by descending all the way to a leaf of the tree.
2. Bayesian techniques—as in [Warner64], in which extensive frequency data make it possible to use Bayes' theorem as a basis for diagnosis.
3. Decision analysis and utility theory—as in [Gorry68], in which there is associated with each piece of information a likely cost of obtaining it and a measure of the benefit to be derived from having it. Information is requested until the projected cost of asking another question (perhaps requiring another lab test or operative procedure) outweighs the benefit (in terms of a more precise diagnosis) to be obtained.

Each of these has a number of attractive aspects, but also encounters some limitations which provided the motivation for our work on a rule-based system. Decision trees, for example, offer simple, readily understandable procedures for diagnosing specific ailments. Problems occur, however, if they encounter unexpected data or if test results are unavailable. The representation of knowledge they offer can be somewhat inflexible as well, since the attempt to make changes deep down in the tree often requires consideration of all previous decisions made further up the tree.

The Bayesian technique offers an appealing generality and precision, since it is a domain-independent technique based on exact principles. Limitations here arise from the need for extensive amounts of frequency data concerning *a priori* and conditional probabilities. Where these data exist, the technique can be used quite effectively, but such figures may not often be available [Edwards72].

Techniques based on utility theory can present a well-motivated sequence of questions that appears to zero in

on the underlying ailment. Like the Bayesian approach, however, it requires extensive data on conditional probabilities of symptoms and disease.

Since none of these is intended to be a model of the reasoning process typically employed by clinicians, it can at times prove difficult for a clinician to discover the basis for the conclusions drawn by any of them. While they each present a compact encoding of knowledge that can provide an appealing efficiency to programs based on them, there is an unavoidable loss of comprehensibility to the physician using them. Reasoning which requires several distinct inferential steps by a clinician, for instance, might be expressed in a single value of a conditional probability in the Bayesian method.

One additional technique has received some attention lately, as other researchers (such as [Silverman75] and [Kulikowski73]) have developed sophisticated models of physiological processes. Where the system involved is sufficiently well-understood and isolatable (the glaucoma model in [Kulikowski73]), this can be a powerful approach. But both of these requirements are lacking in our domain, since infectious disease diagnosis and therapy selection involve a broad range of processes, many of which are only very imperfectly understood.

Design Goals

The limitations encountered in other approaches, along with our own estimation of the capabilities required for a useful medical decision making system provided a number of design goals for MYCIN.

The most fundamental of these, of course, is consistently high performance, and this in turn has a number of implications. It means, for instance, the ability to deal with a large and constantly changing body of technical knowledge. Large amounts of task-specific knowledge seem to be required for high performance, and it is not often possible to specify this knowledge in one, or even a small number of attempts. We rely instead on what may be called an incremental approach to competence. This in turn means that the system's collection of medical knowledge (its "knowledge base") is subject to significant changes over time. Each modification must therefore be a reasonable task, or the program will soon begin to stagnate. A flexible knowledge base also means that the system is inherently dynamic in character. It will be easy to modify it to take into account regional variations in practice, new results which arise from progress in medical research or changes in drug-resistance patterns.

The second goal is the ability to handle an interactive dialog. The system should not be a "black box," printing a collection of orders for the user to follow, but should instead be capable of supplying coherent explanations of its results. (This was perhaps the major motivation for the selection of a symbolic reasoning paradigm, rather than one which, for example, relies on statistics.) Giving the program the ability to explain its results offers a number of useful features. It can lead to greater acceptance by the user population, since the system's conclusions need not be accepted blindly. Instead the user can examine and determine the basis for each of them. An explanation facility also makes it possible for an expert in the field to check the validity of the system's reasoning process, and (as we will see) this can be a significant aid in improving the system's performance. Finally, such explanations can also have an educational influence for the user who does not have extensive experience in the domain.

The desire to provide interactive dialogs means that the system will require extensive human engineering features designed to make interaction simple for someone unaccustomed to computers. Examples below will demonstrate that this has motivated a number of features in the system.

Offering decision support in a medical domain brings up its own particular set of additional goals [Shortliffe74]. Speed, access and ease of use gain additional emphasis, since a physician's time is typically limited. The program also has to fill a need well-recognized by the clinicians who would actually use the system, since the lure of pure technology is usually insufficient. In addition, it has to be designed with an emphasis on its supportive role as a tool for the physician, rather than as a replacement for his own reasoning process.

Finally, the system has to be able to handle the problem noted earlier, of dealing with both incomplete and inexact information. It must be at least flexible enough to continue to function despite the lack of any particular piece (or pieces) of information, and at best, it might (like people) be able to use the very lack of information as a diagnostic clue.

System Design and Operation

Figures 1 and 2 show the initial and final parts of a sample dialog with MYCIN (a complete consultation with additional background information is given in the final section). User type-in is in boldface capitals; italicized comments at the right are for explanatory purposes and are not part of the actual dialog. A great deal of attention had been paid to making the interaction easy and natural. Questions have been carefully worded, for example, so as to set up expectations of responses without making them as uninteresting as multiple-choice questions. Currently the dialog is directed primarily by the system; patient information can be entered only by answering each question in turn.

There are several options available to make the interaction more informative. The user can, for instance, ask for a rephrasing of the question or a display of some (or all) of the acceptable responses (as in question 15). If a requested item of information is unavailable, he can respond with UNKNOWN (U, or UNK for short). The user can also modify any answer of which he is unsure by attaching a number between 0 and 1, indicating a degree of certainty (as, for example, in question 13). The system also has the ability to correct minor typing mistakes (as in question 7), to insure that unfamiliarity with computer terminals or lack of typing ability does not present a problem.

System Organization

Figure 3 offers an overview of MYCIN; arrows indicate information flow. The system is made up of four programs (center boxes) and two information bases (double-walled boxes at the sides). The *consultation program* uses the rules in the *knowledge base* as the basis for the consultation. A *question answering program* handles simple natural language inquiries about either the specific consultation or about the program's general knowledge of infectious disease. The *explanation program* makes it possible for the system to explain the motivation behind the request for an item of information or the basis for drawing any particular conclusion. The *knowledge acquisition program* provides a link between the knowledge

```
--------PATIENT-1--------
   1) Patient's name: (first-last)
   **FRED SMITH                                    {background data}
   2) Sex:
   **MALE
   3) Age:
   **55
   4) Have you been able to obtain positive cultures from a site at which Fred
      Smith has an infection?
   **YES
      --------INFECTION-1--------
      5) What is the infection?
      **PRIMARY-BACTEREMIA
      6) Please give the date and approximate time when signs or symptoms of the
         primary-bacteremia (INFECTION-1) first appeared. (mo/da/yr)
      ** MAY 5, 1975
      The most recent positive culture associated with the primary-bacteremia
      (INFECTION-1) will be referred to as:
         --------CULTURE-1--------
         7) From what site was the specimen for CULTURE-1 taken?
         **BLODD
= BLOOD
         8) Please give the date and time when this blood culture (CULTURE-1)
            was obtained. (mo/da/yr)
         ** MAY 9, 1975
         The first significant organism from this blood culture (CULTURE-1)
         will be referred to as:
            --------ORGANISM-1--------
            9) Enter the identity  of ORGANISM-1:
            **UNKNOWN          {typically the identity is not known and must be inferred}
            10) Is ORGANISM-1 a rod or coccus (etc.):
            **ROD
            11) The gram stain of ORGANISM-1:
            **GRAMNEG
            12) Have there been POSITIVE cultures from sites which are NOT
                associated with infections about which you are seeking advice?
            **NO
            13) Did ORGANISM-1 grow aerobically?
            **YES (.8)
            14) Is the patient's illness with ORGANISM-1 a hospital-
                acquired infection?
            ** ?
            Was the infection with ORGANISM-1 acquired while the
            patient was hospitalized?
            Expected responses are: YES NO
            Enter HELP for list of user options.
            ** Y
              ...
```

{more questions follow, to gather sufficient information to infer possible identities of the infecting organisms}

Figure 1

```
My therapy recommendation will be based on the following probable infection(s)
and potential causative organism(s):

INFECTION-1 is PRIMARY-BACTEREMIA                    {possible identities}

<Item 1>  The identity of ORGANISM-1 may be PSEUDOMONAS-AERUGINOSA
<Item 2>  The identity of ORGANISM-1 may be KLEBSIELLA-PNEUMONIAE
<Item 3>  The identity of ORGANISM-1 may be E.COLI
<Item 4>  The identity of ORGANISM-1 may be BACTEROIDES-FRAGILIS
<Item 5>  The identity of ORGANISM-1 may be ENTEROBACTER
<Item 6>  The identity of ORGANISM-1 may be PROTEUS-NON-MIRABILIS

   46) On a scale of 0 to 4, where higher numbers indicate increasing severity
       how would you rate Fred Smith's degree of sickness?
   **3
   47) Does Fred Smith have a clinically significant allergic reaction to any
       antimicrobial agent?
   **NO
   48) Fred Smith's weight in kilograms (or <number> POUNDS):
   **70
                                  {note that 2 drugs cover for 6 potential organisms}
[Rec 1] My preferred therapy recommendation is as follows:
   In order to cover for Items <1 2 3 5 6>:
        Give: GENTAMICIN
        Dose: 119 mg (1.7 mg/kg) q8h IV [or IM] for 10 days
        Comments: Modify dose in renal failure
   In order to cover for Item <4>:
        Give: CLINDAMYCIN
        Dose: 595 mg (8.5 mg/kg) q6h IV [or IM] for 14 days
        Comments: If diarrhea or other GI symptoms develop, patient
            should be evaluated for possible pseudomembranous colitis.

Do you wish to see the next choice therapy?
                                  {alternative choices are available}
** YES

Which, if any, of the above drugs do you wish eliminated from consideration?
**CLINDAMYCIN

[Rec 2] Next best therapy recommendation:
   In order to cover for Items <2 3 4 5 6>:
        Give: CHLORAMPHENICOL
        Dose: 1 gm (15 mg/kg) q6h IV for 14 days
        Comments: Monitor patient's white count
   In order to cover for Item <1>:
        Give: GENTAMICIN
        Dose: 119 mg (1.7 mg/kg) q8h IV [or IM] for 10 days
        Comments: Modify dose in renal failure
```

Figure 2

base and an expert who wants to teach the program new things about the domain.

The system is written in INTERLISP (an advanced version of the LISP programming language), runs on a DEC PDP-10 with approximately 130,000 words of nonshared code and is fast enough to permit convenient real-time interaction. The consultation program (with the required data structures) requires approximately 50,000 words, the explanation program 10,000, the question answering program 15,000, the rule acquisition program 20,000 and the rules themselves approximately 8,000 (the remainder includes a dictionary, information on drug properties and various system utility features).

As described below, each of these four interrelated programs is relevant to some of the design goals outlined above.

The primary source of medical knowledge in the system is a set of some 400 decision rules like the one below, each with a premise and an action (Figure 4).

Many of the system's unique and important capabilities are made possible by encoding knowledge in rules like the

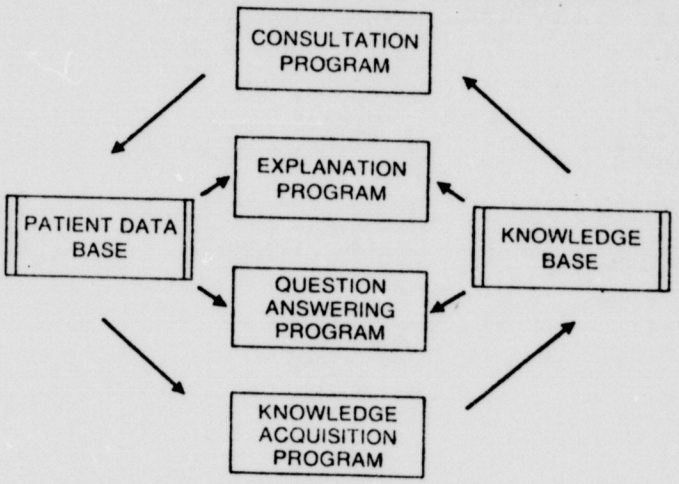

Figure 3

The rules and judgmental knowledge

The primary source of medical knowledge in the system is a set of some 400 decision rules like the one below, each with a premise and an action (Figure 4).

English version

```
If    1) the gram stain of the organism is gram negative, and
      2) the morphology of the organism is rod, and
      3) the aerobicity of the organism is anaerobic,
then there is suggestive evidence (.7) that the identity of the organism is
     Bacteroides.
```

Internal format

```
PREMISE          ($AND(SAME CNTXT GRAM GRAMNEG)
                       (MEMBF CNTXT MORPH ROD)
                       (SAME CNTXT AIR ANAEROBIC))

ACTION           (CONCLUDE CNTXT IDENT BACTEROIDES TALLY .6)
```

Figure 4

one above. Such rules form modular "chunks" of knowledge about the domain, represented in a form that is comprehensible to a clinician.

The consultation system, for instance, uses the collection of rules to make conclusions about the patient. If it is attempting to determine the identity of an organism responsible for a particular infection, it retrieves the entire list of rules which, like the one above, conclude about identity. It then attempts to ascertain whether the conclusion of the first rule is valid by evaluating in turn each of the clauses of the premise. Thus, for the rule above, the first thing to find out is gram stain. If this information is already available in the data base, the program retrieves it. If not, determination of gram stain becomes the objective of a new rule, and the program retrieves all rules which conclude about it and tries to use each of them to obtain the value of gram stain. If, after trying all the relevant rules, the answer still has not been discovered, the program asks the user for the relevant clinical information which will permit it to establish the validity of the premise clause. Thus, the

rules "unwind" to produce a succession of goals, and it is the attempt to achieve each goal that drives the consultation.

We noted earlier the necessity for dealing with inexact decision criteria. The number in the conclusion of each rule (called the "certainty factor" or CF) is the mechanism used to capture this inexactness. For the rule in Figure 4, the evidence cited in the premise is strongly indicative (.7 out of 1) of the conclusion, but cannot justify it with absolute certainty.

The methods for combining CFs are embodied in a model of what we might call "approximate implication." While CFs are derived from and are related to probabilities, they are distinctly different (for a detailed review of the concept, see [Shortliffe75]). Evidence confirming a hypothesis is collected separately from that which disconfirms it, and the truth of the hypothesis at any time is the algebraic sum of the current evidence for and against it. This is an important aspect of the truth model, since it makes plausible the simultaneous existence of evidence in favor and against the same hypothesis. We believe this is an important characteristic of any model of inexact reasoning.

Explanations

Normally during a consultation the system takes the initiative, asking questions of the user. If one of those questions seems inappropriate, however, the user can interrupt and use the explanation facilities to find out what's going on. By examining the chain of reasoning that prompted the question, he can find out if the reasoning was motivated by plausible considerations. Several examples are shown below, in a segment extracted from a consultation (additional examples can be found in [Davis76] and [Davis77]).

To determine the motivation for any question, the user can ask why it was asked, and can follow the reasoning chain as far as he likes:

---

18) In this series, how many of the cultures were positive for the e.coli (ORGANISM-1)? *{typical system prompt}*

**WHY

{request from user to determine motivation for question}

[i.e. WHY is it important to determine the number of cultures which were positive for the organism?]

{system's rephrasing of the WHY question, and answer}

[1.0] This will help to determine whether there is significant disease associated with this occurence of the organism.

It has already been established that

[1.1] the culture was taken from a sterile source

[1.2] the number of cultures in the series is greater than 2

Therefore, if

[1.3] the number of cultures which were positive for ORGANISM-1 is greater than 2

**WHY

Then there is suggestive evidence (.5) that there is significant disease associated with this occurrence of the e.coli (ORGANISM-1). [RULE049]

{repeated WHY to determine relevance of previous explanation}

[i.e. WHY is it important to determine if there is significant disease associated with this occurrence of the organism?]

[2.0] This will help to determine a diagnosis for which the infection should be treated.

It has already been established that

[2.1] the infection is not one of: primary-bacteremia, meningitis

---

[2.2] for some current culture of the patient it is true that the site of the culture is blood

Therefore, if

[2.3] there is significant disease associated with this occurrence of the organism

Then

it is definite (1.0) that the diagnosis for which the infection should be treated is secondary-bacteremia

[RULE103]

Figure 5

He can also determine the basis for any of the conclusions the system has drawn by inquiring how the conclusion was made:

**HOW 1.1

{the 1.1 refers to the information in the line labelled [1.1] in Figure 5}

[i.e. HOW was it determined that the culture was taken from a sterile source?]

[3.1] RULE001 was used to conclude that there is suggestive evidence (.7) that the culture was taken from a sterile source.

[3.2] RULE022 was used to conclude that there is strongly suggestive evidence (.7) that the culture was taken from a sterile source.

**HOW 3.1

{another request from the user}

[i.e. HOW was RULE001 used?]

It has already been established that

[4.1] the site of the culture is one of: those sites from which the sterility of cultures depends upon the method of collection

[4.2] the method of collection of the culture is one of: the recognized methods of collection associated with the site of the culture, and

[4.3] it is not known whether care was taken in collecting the culture

Therefore

there is strongly suggestive evidence (.8) that the culture was taken from a sterile source

[RULE022]

Figure 6

First, we are dealing with a domain in which deduction, and deduction in the face of uncertainty, is a primary task. The system's fundamental approach to explanation is thus to display some recap of its internal actions, a trace of its reasoning. The success of this technique is predicated on the claim that the system's basic approach to the problem is sufficiently intuitive that a summary of those actions is at least a reasonable basis from which to start. While it would be difficult to prove the claim in any formal sense, there are several factors which suggest its plausibility.

The use of production rules in an if/then format seems therefore to be a natural way of expressing things about the domain, and the display of such rules should be comprehensible. Second, the use of such rules in a backward-chaining mode is, we claim, a reasonably intuitive scheme. *Modus ponens* is a well-understood and widely (if not explicitly) used mode of inference. Thus, the general form of the representation and the way it is employed should not be unfamiliar to the average user. More specifically, however, consider the source of the rules. They have been given to us by human experts who were at-

tempting to formalize their own knowledge of the domain. As such, they embody accepted patterns of human reasoning, implying that they should be relatively easy to understand, especially for those familiar with the domain. As such, they will also attack the problem at what has been judged an appropriate level of detail. That is, they will embody the right size of "chunks" of the problem to be comprehensible.

We are not, therefore, recapping the binary bit level operations of the machine instructions for an obscure piece of code. We claim instead to be working with primitives and a methodology whose substance, level of detail and mechanism are all well-suited to the domain and to human comprehension precisely because they were provided by human experts. This approach seems to provide what may plausibly be an understandable explanation of system behavior.

By way of contrast, we might try to imagine how a program based on a statistical approach might attempt to explain itself. Such systems can, for instance, display a disease which has been deduced and a list of relevant symptoms with prior and posterior probabilities. No more informative detail is available, however. When the symptom list is long, it may not be clear how each of them (or some combination of them) contributed to the conclusion. It is more difficult to imagine what sort of explanation could be provided if the program were interrupted with interim queries while in the process of computing probabilities. The problem, of course, is that statistical methods are not good models of the actual reasoning process (as shown in psychological experiments of [Edwards68] and [Tversky-74]), nor were they designed to be. While they are operationally effective when extensive data concerning disease incidence are available, they are also for the most part, shallow, one-step techniques which capture little of the ongoing process actually used by expert problem solvers in the domain.[2]

We have found the presence of even the current basic explanation capabilities to be extremely useful, as they have begun to pass the most fundamental test: It has become easier to ask the system what it did than to trace through the code by hand. The continued development and generalization of these capabilities is one focus of our present research.

Question Answering

The use of a rule-based representation of knowledge makes it possible for the system to explain the basis for its clinical recommendations. For example, if the clinician asks, "How did you decide that the organism might be an e.coli?" the program answers by displaying the rules which were actually used and explaining, if requested, how each of the premises of the rules was established. (Several examples are shown below, additional examples can be found in [Shortliffe76] and [Scott76].) This is something which the clinician can readily understand, and it provides a far more comprehensible and acceptable explanation that would be possible if the program were to use a simple statistical approach to diagnosis.

Note that the system is capable of answering questions concerning actions the program *failed* to take (as in question [d] below). For questions such as [e] below, the search through the knowledge base is directed by a simple parsing of the question into a request for a set of rules, with constraints on premise and/or action contents. The retrieval of relevant rules is guided primarily by pre-established (but automatically generated) lists which indicate premise and action contents.

(a) ** HOW DID YOU DECIDE THAT ORGANISM-1 MIGHT BE AN E.COLI?

I used RULE021 to conclude that the identity of ORGANISM-1 is e.coli. This gave a cumulative CF of (.47).

Then I used RULE084 to conclude that the identity of ORGANISM-1 is e.coli. This gave a cumulative CF of (.55).

Finally, I used RULE003 to conclude that the identity of ORGANISM-1 is e.coli. This gave a cumulative CF of (.74).

(b) ** DID YOU USE RULE 163 TO FIND OUT ANYTHING ABOUT ORGANISM-1?

RULE163 was tried in the context of ORGANISM-1, but it failed because it is not true that the patient has had a genito-urinary tract manipulative procedure (clause 3).

(c) ** DID RULE 116 TELL YOU ANYTHING ABOUT INFECTION-1?

No. RULE116 would be relevant to the primary-bacteremia when applied to ORGANISM-1.

But clause 2 ("the gram stain of the organism is not known") was already known to be false for ORGANISM-1, so the rule was never tried.

(d) ** WHY DIDN'T YOU CONSIDER STREPTOCOCCUS AS A POSSIBILITY FOR ORGANISM-1?

The following rule could have been used to determine that the identity of ORGANISM-1 was streptococcus: RULE033

But clause 2 ("the morphology of the organism is coccus") was already known to be false for ORGANISM-1, so the rule was never tried.

(e) ** DO YOU RECOMMEND CLINDAMYCIN FOR BACTEROIDES?

Yes.

RULE060

If the identity of the organism is bacteroides

Then I recommend therapy be chosen from among the following:

CHLORAMPHENICOL, CLINDAMYCIN, RIFAMPIN, ERYTHROMYCIN, CARBENICILLIN TETRACYCLINE, LINCOMYCIN

(f) ** IS BLOOD A STERILE SITE?

Yes, blood is one of those sites that are normally sterile.

(g) ** WHAT ARE THE ORGANISMS ASSOCIATED WITH THE GASTRO-INTESTINAL TRACT?

The likely pathogens associated with the gastro-intestinal tract are: ENTEROCOCCUS, CLOSTRIDIUM-GANGRENE, BACTEROIDES, KLEBISELLA, PSEUDOMONAS, E.COLI, ENTEROBACTER, PROTEUS.

Figure 7

Since the field of infectious disease therapy is both large and constantly changing, it was apparent from the outset that the program would have to deal with an evolving knowledge base. The size of the domain makes writing a complete set of rules in a single try an impossible task, so the system was designed to facilitate an incremental approach to competence. In addition, new research in the domain produces new results and modifications of old principles, so that a broad scope of knowledge-base management capabilities is clearly necessary.

A fundamental assumption in the approach we use is that the expert teaching the system can be debriefed, thus transferring his knowledge to the program. That is, presented with any conclusion he makes during a consultation, the expert must be able to state a rule indicating all relevant premises for that conclusion. The rule must, in and of itself, represent a valid chunk of clinical knowledge.

There are two reasons why this seems a plausible approach to knowledge acquisition. First, clinical medicine appears to be at the correct level of formalization. That is, while relatively little of the knowledge can be specified in precise algorithms (at a level comparable to, say, elementary physics), the judgmental knowledge that exists is often specifiable in reasonably firm heuristics. Second, on the model of a medical student's clinical training, we have emphaized the acquisition of new knowledge in the context of debugging (although the system is prepared to accept a new rule from the user at any time). We expect that some error on the system's part will become apparent during the consultation, perhaps through an incorrect organism identification or therapy selection. Tracking down this error by tracing back through the program's actions is a reasonably straightforward process which presents the expert with a methodical and complete review of the system's reasoning. He is obligated to either approve of each step or to correct it. This means that the expert is faced with a sharply focused task of adding a chunk of knowledge to remedy a specific bug. This makes it far easier for him to formalize his knowledge than would be the case if he were asked, for example, "Tell me about bacteremia."

Knowledge acquisition dialogs are somewhat extended and are not included here for the sake of brevity. The interested reader may consult [Davis76] and [Davis77] for a detailed review and discussion.

The fundamental design and implementation of MYCIN does not restrict its use to medical domains. This is due in part to the widespread applicability of the concept of diagnosis and therapy as a problem-solving method and in part to the way in which the system was implemented. Many problems that can be viewed as the discovery and correction of errors can be viewed as diagnosis and therapy, whether the domain is medicine, repair of machines or program debugging. This means that the fundamental approach to problem solving is potentially widely applicable.

The implementation of the system was kept strongly domain-independent from the outset, primarily to permit extension to other areas of infectious disease. The current knowledge base can thus easily be augmented or removed entirely and replaced with another.

It has proven possible, for instance, to build additional knowledge bases for such disparate fields as chemotherapy for psychiatry and auto mechanics. In one of the first such efforts [vanMelle74], a small part of an auto repair manual was rewritten as production rules and inserted in place of the bacteremia knowledge base. What resulted was a very simple but fully functional consultant capable of diagnosing and curing problems in a part of an auto electrical system. More recently, a pilot system for psychiatric diagnosis is being assembled. While it has

currently only some 50 rules, it is fully functional and displays primitive but encouraging performance. In both systems, all of the established explanation facilities work as designed without modification.

Finally, the basic methodology developed in designing and implementing MYCIN has provided a basis for a number of other systems. The work described in [Hart75], for instance, deals with a system oriented toward repair of electromechanical devices, while [Anderson75] describes a system for the creation of intelligent terminals. Both of these share significant points of methodology with MYCIN.

There are, naturally, some domains that might be less profitable to explore. One of the interesting lessons of the auto repair system was that domains with little inexactness in the reasoning process—those for which algorithmic diagnostic routines can be written—are not particularly appropriate for this methodology. The precision in these domains means that little use is made of the certainty factor mechanism, and many of MYCIN's more complicated (and computationally expensive) features go unused.

Nor is it reasonable to expect to be able to write rules for an arbitrary domain. As knowledge in an area accumulates, it becomes progressively more formalized. There is a certain state of this formalization process when it is appropriate to use rules of the sort shown above. Earlier than this the knowledge is too unstructured, later on it may (like the auto repair system) be more effective to write straightforward algorithms.

It is also possible that the knowledge in some domains is inherently unsuited to a rule-like representation, since rules become increasingly awkward as the number of premise clauses increases. Dealing with a number of interacting factors may be difficult for any representation, but given the reliance here on rules as a medium of communication of knowledge, the problem becomes especially significant.

Impact as a DSS

We expect that the primary long-term impact of MYCIN itself will be in providing consultative support in primary health care centers where such expertise is currently in short supply. While such a step is still a long way off, with a sufficiently large knowledge base it may prove possible to provide near-human-level performance. With the advent of phone line networks for communication to computers, this service can quickly and easily be made available to hospitals and clinics in rural areas that are typically undersupplied with expert medical care. Given the flexibility of the system, modifications could easily be introduced to deal with regional variations in accepted medical practice or seasonal variations in infection susceptibility.

It is not hard to see the advantages to the busy clinician that result. Imagine the utility of having the services of a recognized expert in the field available 24 hours a day for consultation and advice. There is the potential of handling more cases with greater accuracy and reduced likelihood of serious error.

More generally, we believe that our basic methodology for handling decision making is capable of offering assistance in a wide a range of domains. Its fundamental power lies in the ability to deal with a certain sort of complexity. To see this, consider two of the factors that may make decision making difficult: interconnectedness and size. Problems whose subparts are tightly interconnected are difficult because they are not decomposable: connections between subparts are sufficiently rich that no small number of parts can be considered independently. The difficulty here lies in trying to keep track of too many interrelated items at once.

But even where problems can be decomposed and subproblems solved separately, difficulty can arise because there is simply too much to do. The decision-making capabilities in MYCIN are currently well-suited to problems of this character. Given a sufficiently large and diverse collection of rules, the system can deal with all necessary factors. Its primary virtue then would lie in its exhaustive consideration of all factors and an encyclopedic approach to the task. The presence of the explanation and question-answering facilities means that the system also has the ability to explain its conclusions to any desired level of detail and may thus become a useful tool in the attempt to understand the problem.

Finally, the attempt to write a computer program to perform a complex task often highlights (in intense detail) exactly how imprecise is our understanding of the processes by which people perform the task. More positively stated, the computer provides an excellent environment in which to elucidate and verify the sources of knowledge required. This has been our experience with MYCIN, as an interesting side effect of the efforts to create the system has been the progressive formalization of knowledge about infectious disease diagnosis and therapy. We speculate that this may be one of the most enduring effects of the system when it is applied to other domains, providing experts in other fields with a medium in which to formalize, state and test the principles on which their decisions are made.

Detailed Example

As part of our effort to monitor performance development in the system, we recently undertook a large-scale evaluation, using real cases selected from current admissions to the Stanford Medical Center. To give a feeling for how MYCIN would function in a true-to-life clinical situation, one of those cases is reviewed below in detail. The background and recent medical history is presented first to supply the context and motivate the need for a consultation (parenthetical comments are added to make the terminology clear to nonmedical readers, along with a summary at the end to explain the major decisions to be made in the case). The consultation then follows and the results are compared with the initial course of treatment.

In July 1975, C.R., a 52-year-old white male, was admitted to the hospital with acute symptoms of diarrhea, fever, nausea and vomiting. He also had a two-month history of weight loss with pain in the lower abdomen. During this hospitalization the patient had two positive blood cultures for *Salmonella* taken prior to antibiotics. Stool and urine cultures were negative. Physical exam indicated an asymptomatic abdominal aortic aneurysm 5 cm in diameter and 7 cm in length. Barium enema in August 1975 showed a perforated appendiceal mass (a ruptured appendix). Patient underwent partial colectomy (partial removal of the colon) with cholecystectomy (removal of the gall bladder). Cholelithiasis (gall stones) was noted with chronic cholecystitis (infection of the gall bladder) at that time. Cultures of gall bladder stones revealed no *Salmonella*.

During this hospitalization patient received ampicillin iv (intravenously), 12 gm daily for 28 consecutive days.

Subsequent blood cultures postantibiotic were all negative (the medication appears to have taken care of the infection). Patient was observed in the hospital for one week after discontinuation of the antibiotics without return of fever or GI (gastro-intestinal) symptoms. Patient was discharged in good condition and arrangements made for elective readmission for resection of the abdominal aneurysm.

Discharge diagnoses were Salmonella sepsis—resolved, aortic abdominal aneurysm and anemia of chronic disease.

While at home, patient complained of occasional sweats without fever. However, patient had a good appetite and denied any abdominal discomfort. Patient did have occasional episodes of diarrhea after his discharge. In addition patient had intermittent nonspecific lower back pain.

C.R. returned to the hospital for the resection operation on October 20, 1975. He was in no acute distress. Blood pressure was 120/80 (normal), pulse 120 (very high), respiration 20 (high), temperature 101.5 (feverish), weight 75 kg. Abdomen has 10 X 15 cm expansile mass in mid-abdomen (large bulge has developed).

Of immediate concern to the surgeon was the increasing size of the aneurysm. Ultrasound was done and confirmed the increase. The decision was made to operate, cephalothin 2 gm iv and gentamicin 80 mg iv were given preoperatively. An aneurysmectomy was in progress when three of the three blood cultures drawn the previous day were reported positive for gram negative rods in aerobic and anaerobic bottles.

The patient's first admission was for GI problems and weight loss. Physical exam indicated a ruptured appendix and the aortic aneurysm. Appendix was removed; the gall bladder was also removed because of the gall stones and its chronic inflamation. Patient was given ampicillin, and all signs of infection disappeared. At discharge, arrangements were made for an elective readmission to take care of the aneurysm.

When readmitted for the aneurysm operation, aneurysm was noted to have gotten much worse (a potentially life-threatening situation, since it may burst). Patient was also feverish (a sign of infection).

There were thus two conflicting problems: the aneurysm may burst, and so must be taken care of as soon as possible, but the operation is a serious one, and it is not desirable to undertake it on someone who may have an active infection. The aneurysm problem was judged too serious to be delayed, and the surgeons chose two standard and powerful antiobiotics (cephalothin and gentamicin) to be given preoperatively to offset the possible infection.

The report of all three blood cultures turning positive makes it clear that the patient does, indeed, have an infection. Therapy must be started immediately, or the patient, who has survived a major operation but been weakened by it, may succumb to the infection. At this point we ran a consultation on MYCIN, supplying the system with all of the facts given above.

Consultation

```
MYCIN  31-OCT-75 ...

   22-OCT-75 12:00

   --------PATIENT-248--------
   1) Patient's name: (first-last)
   ** C.R.
   2) Sex:
   ** MALE
   3) Age:
   ** 52
   4) Have you been able to obtain positive microbiological information about
      a possible infection of C.R.?
   ** YES

      --------INFECTION-1--------
      5) What is the infection?
      ** ENDARTERITIS (.6)
      {An infected artery wall, since the clinician believes the aneurysm was infected. Since
      he is not sure, he uses a certainty factor to indicate his degree of belief.}
      6) Please give the date and approximate time when signs or symptoms of
         the endarteritis (INFECTION-1) first appeared. (mo/da/yr time)
      ** 21 JULY 1975
      The most recent culture yielding positive microbiological information
      about the endarteritis (INFECTION-1) will be referred to as:

         --------CULTURE-1--------
         7) From what site was the specimen for CULTURE-1 taken?
         ** BLOOD
```

8) Please give the date and time when this blood culture (CULTURE-1) was obtained. (mo/da/yr time)
** 21 OCT 1975
The first significant organism from this blood culture (CULTURE-1) will be referred to as:

--------ORGANISM-1--------
9) Enter the identity of ORGANISM-1:
** UNKNOWN
10) Is ORGANISM-1 a rod or coccus (etc.):
** ROD
11) The gram stain of ORGANISM-1:
** GRAMNEG
12) Have there been POSITIVE cultures yielding organisms about which you will NOT be seeking advice?
** YES

{The previous history of salmonella may be relevant, but since it occurred several months ago, it no longer needs treatment.}

--------CULTURE-2--------
13) From what site was the specimen for CULTURE-2 taken?
** BLOOD
14) Please give the date and time when this blood culture (CULTURE-2) was obtained. (mo/da/yr time)
** 21 JULY 1975
The first significant organism from this blood culture (CULTURE-2) will be referred to as:

--------ORGANISM-2--------
15) Enter the identity of ORGANISM-2:
** SALMONELLA
16) Any other significant organisms isolated from this blood culture (CULTURE-2)?
** NO
17) Have there been any other POSITIVE cultures yielding organisms about which you will NOT be seeking advice?
** NO
18) Did ORGANISM-1 grow in the aerobic bottle?
** YES
19) Did ORGANISM-1 grow in the anaerobic bottle?
** YES
20) What is the suspected portal of entry of ORGANISM-1 into this sterile site?
** GI (.7)
21) Is the patient's illness with ORGANISM-1 a hospital-acquired infection?
** NO
22) Is C.R. a burn patient?
** NO
23) Has C.R. had any gu surgery or manipulation of the urinary tract (eg. indwelling catheter, cystoscopy, etc.)?
** NO
24) Does C.R. have 'pseudomonas-type' hemorrhagic skin lesions (ecthyma gangrenosum)?
** NO
25) Do you consider C.R. to be a compromised host in the sense that s/he has increased susceptability to infection (alcoholic, diabetic, uremic, or immunosuppressed)?
** NO

26) Including this blood culture (CULTURE-1), how many times during this series of cultures was this site cultured?
** 3
27) In this series, how many of the cultures from this site were positive for ORGANISM-1?
** 3
28) Any other significant organisms isolated from this blood culture (CULTURE-1)?
** NO
29) Did any additional cultures yield positive microbiological information which you believe is associated with the endarteritis (INFECTION-1)?
** NO
30) Does C.R. have any other infection about which you have obtained positive microbiological information?
** NO
31) Do you suspect an infection at a site which you have been unable to culture, or from which all cultures were negative?
** NO
32) Has the lab reported the in-vitro antibiotic sensitivities of the Salmonella (ORGANISM-2)?
** UNKNOWN
33) Is C.R. currently receiving therapy with any antimicrobial agent?
** NO
34) During the present illness, has C.R. been treated with an antimicrobial agent which is no longer being administered?
** NO

My therapy recommendations will be designed to treat for organisms that are either very likely or, although less likely, would have a significant effect on therapy selection if they were present. It is important to cover for the following probable infection(s) and associated organism(s):

INFECTION-1 is ENDARTERITIS with BACTEREMIA

<Item 1> E.COLI [ORGANISM-1]
<Item 2> SALMONELLA [ORGANISM-1]
<Item 3> KLEBSIELLA-PNEUMONIAE [ORGANISM-1]
<Item 4> PSEUDOMONAS-AERUGINOSA [ORGANISM-1]
<Item 5> ENTEROBACTER [ORGANISM-1]
<Item 6> PROTEUS-NON-MIRABILIS [ORGANISM-1]

35) On a scale of 0 to 4, where higher numbers indicate increasing severity, how would you rate C.R.'s degree of sickness?
** 3
36) Does C.R. have a clinically significant allergic reaction to any antimicrobial agent?
** NO
[Considering GENTAMICIN for use against INFECTION-1...]
[Considering CHLORAMPHENICOL for use against INFECTION-1...]

37) C.R.'s weight in kilograms (or <number> POUNDS):
** 75

[Rec 1] My preferred therapy recommendation is as follows:
In order to cover for Items <1 3 4 5 6>:
Give: GENTAMICIN
Dose: 128 mg (1.7 mg/kg) q8h IV (or IM) for 10 days

```
     Comments: Modify dose in renal failure
In order to cover for Item <2>:
     Give: CHLORAMPHENICOL
     Dose: 563 mg (7.5 mg/kg) q6h IV for 14 days
     Comments: Monitor patient's white count
```

Figure 8

Commentary

Faced with the necessity of operating, the choice was made to administer two powerful antibiotics (cephalothin and gentamicin) which are the treatment choice for the sort of gram negative rod infection the patient appeared to have. However, after the operation and upon further consideration during the consultation, it became clear that the current infection was likely to be a recurrence of the *Salmonella*, and neither of the drugs chosen is effective against this organism. MYCIN recommends a combination of gentamicin and chloramphenicol, with the latter drug chosen specifically to combat the *Salmonella*.

In the actual case, the same determination was made, the patient's therapy was altered appropriately and he recovered.

Notes

1. The system was developed by a group that has included Stanley Cohen MD, Stanton Axline MD, Jan Aikins MS, Bob Blum MD, Bruce Buchanan PhD, Bill Clancey BS, Randall Davis PhD, Larry Fagan BS, Frank Rhame MD, Carlisle Scott MS, Ted Shortliffe MD PhD, Bill van Melle BS, Sharon Wraith BS and Victor Yu MD. The system's name is derived from the suffix common to many antibiotic drug names.
2. However, the reasoning process of human experts may not be the ideal model for all knowledge-based problem solving systems. In the presence of reliable statistical data, programs using a decision-theoretic approach are capable of performance surpassing those of their human counterparts. In domains like infectious disease-therapy selection, however, which are characterized by judgmental knowledge, statistical approaches may not be viable. This appears to be the case for many medical decision-making areas. See [Gorry73] for further discussion of this point.

References

1. [Anderson75] Anderson R H, Gillogly J J. RAND intelligent terminals agent: design philosophy, RAND R-1809-ARPA, The RAND Corporation, Santa Monica, CA.
2. [Davis77] Davis R, Buchanan B G, Shortliffe E H, Production rules as representation for a knowledge-based consultation program, *Artificial Intelligence*, January, 1977 (to appear).
3. [Davis76] Davis R, Applications of meta level knowledge to the construction, maintenance, and use of large knowledge bases, AI Memo 283, July 1976.
4. [Edwards68] Edwards R, Conservatism in human information processing, in Kleinmuntz (Ed.), *Formal Representation of Human Judgment*, Wiley (1968), pp 17–52.
5. [Edwards72] Edwards W N = 1, Diagnosis in unique cases, *Computer Diagnosis and Diagnostic Methods*, (Jacquez, ed.), pp 139–151, C C Thomas, Springfield, Illinois, (1972).
6. [Gorry68] Gorry G A, Barnett G O. Experience with a model of sequential diagnosis, *Computers and Biomedical Research;* 1 (1968) pp 490–507.
7. [Gorry73] Gorry G A, Kassirer J P, Essig A, Schwartz W B. Decision analysis as the basis for computer-aided management of acute renal failure, *American Journal of Medicine*, 55 (1973) pp 473–484.
8. [Hart75] Hart P E. Progress on a computer-based consultant, AI Group Technical Note 99, January 1975, Stanford Research Institute, Menlo Park, CA.
9. [Kagan73] Kagan B M, Fannin S L, Bardie F. Spotlight on antimicrobial agents—1973, *JAMA*, 226, #3 (October 1973) pp 306–310.
10. [Kulikowski73] Kulikowski C A, Weiss S, Saifr A. Glaucoma Diagnosis and Therapy by Computer, Proceedings of Annual Meeting of Assn. for Research in Vision and Ophthalmology, May 1973.
11. [Meyer73] Meyer A V, Weissman W K. Computer analysis of the clinical neurological exam, *Computers and Biomedical Research*, 3(1973) pp 111–117.
12. [Neu75] Neu H C, Howry S P, Testing the physician's knowledge of antibiotic use, *NEJM*, 293:1291–5, 18 December 1975.
13. [Ray76] Ray W A, Federspiel C F, Schaffner W. Prescribing of chloramphenicol in ambulatory practice, *Ann Internal Med*, 84:866–870 (March 1976).
14. [Scott76] Scott A C, Clancey W, Davis R, Shortliffe E H. Explanation capabilities of knowledge-based production systems, (submitted to the *American Journal of Computational Linguistics*).
15. [Shortliffe74] Shortliffe E H, Axline S G, Buchanan B G, Cohen, S N. Design considerations for a program to provide consultations in clinical therapeutics. Presented at San Diego Biomedical Symposium 1974 (February 6–8, 1974).
16. [Shortliffe75] Shortliffe E H, Buchanan B G. A model of inexact reasoning in medicine, *Mathematical Biosciences* 23 (1975) pp 351–379.
17. [Shortliffe76] Shortliffe, E H. *Computer-Based Medical Consultations: MYCIN*, American Elsevier, 1976.
18. [Silverman75] Silverman H. A digitalis therapy advisor, MAC TR-143 (January 1975) Project MAC, Mass Inst Tech.
19. [Tversky74] Tversky A, Kahneman D. Judgment under uncertainty: heuristics and biases, *Science*, 185 (18 September 1974) pp 1129–1131.
20. [van Melle74] van Melle W. Would you like advice on another horn, MYCIN project internal working paper, December 1974, Stanford University, Stanford California.
21. [Warner64] Warner H R, Toronto A F, Veasy L G. Experience with Bayes' theorem for computer diagnosis of congenital heart disease, *Annals New York Academy of Science*, 115 (1964) pp 558–567.