10

Computer-Based Medical Decision Making: From MYCIN to VM

Lawrence M. Fagan, Edward H. Shortliffe, and Bruce G. Buchanan

We mentioned in the introduction to Chapter 5 that MYCIN provided a starting point for several additional research projects. The Ventilator Manager (VM) project of Larry Fagan had its beginnings in the MYCIN project but quickly diverged because of the dynamic nature of the intensive care unit (ICU) setting for which it was designed. MYCIN required a "snapshot" approach to patient assessment—temporal trends were poorly handled and advice was generally provided on the basis of a patient's situation at a single point in time. In dynamic settings like an ICU, however, decisions may be dependent on frequent sequential assessments of the patient's status. Fagan's work was accordingly also influenced by another earlier Stanford project known as HASP/SIAP (Nii et al., 1982). That system was not concerned with medical issues, but did develop iterative techniques for the ongoing analysis of signals.

Although Fagan was a graduate student at Stanford at the time, much of his work was based at Pacific Medical Center in San Francisco. The director of the postsurgical ICU there, Dr. John Osborn, had developed an elaborate monitoring system that was in routine use. However, the amount of data generated was sometimes overwhelming, particularly for physicians in training. It was clear that there was expertise involved in learning how to interpret the data, and the idea developed to build an expert system that could monitor the various physiological parameters and give advice accordingly. The following chapter discusses the resulting evolution from

From *Automedica*, 3: 97–106 (1980).

MYCIN to VM and explains how the differing requirements of the two clinical settings affected the ultimate design of the newer system. VM's specific approach is to use production rules for interpreting the physiological data, thereby permitting VM to aid in the management of patients being "weaned" from ventilators. This work is included here largely because of Fagan's insights regarding temporal reasoning in medicine. Particularly noteworthy is the development of techniques to allow the system to interpret patient data by comparing current findings with explicit expectations generated using production rules during earlier time periods.

10.1 Introduction

Since the early 1970s, researchers in computer-based medical reasoning have begun to recognize the potential benefits of applying symbolic reasoning techniques in clinical domains (see Chapter 3). One such research group is the Heuristic Programming Project at Stanford University. The first medical reasoning program developed by the project, known as the MYCIN system (Shortliffe, 1976), adopted symbolic processing techniques largely in response to a conviction that computer-based consultation systems, in order to be accepted by physicians, should be able to explain how and why a particular conclusion has been derived. Such systems should also be able to incorporate, organize, manipulate, and update large quantities of medical knowledge. Subsequently, a series of additional medical application programs using MYCIN's techniques has been created. In this paper we compare MYCIN, a program for infectious disease diagnosis and therapy, with a newer system, the Ventilator Manager (VM) program for measurement interpretation in the intensive care unit (ICU). Each of these programs uses a representation scheme, known as production rules (Davis and King, 1977), to encode the medical knowledge used for decision making. Each production rule is stated in the form "situation implies conclusion." Production rules may be chained together to form a line of reasoning leading from observed patient data to diagnostic and therapeutic conclusions. This report discusses the strengths of this form of knowledge representation and shows how production rules can be applied in two somewhat different clinical applications.

We begin by presenting the reasons that symbolic processing has been utilized for medical decision making. A brief discussion of the MYCIN program and a more detailed discussion of the VM program are included to demonstrate the use of the symbolic processing techniques. The design criteria for the two programs are compared. Differences in design criteria, plus experience with the MYCIN program, led to the extensions to the methodology described in the final section.

10.2 The Rationale for Using Symbolic Processing Techniques

There is increasing evidence that computer-based diagnosis and therapy programs will be accepted by physicians only if they meet a stringent set of design criteria. Several sets of design requirements have been suggested (Shortliffe et al., 1974) (see also Chapter 2). Although the overriding goal for any computer-based consultation program is, of course, that it be accurate, Gorry has suggested (see Chapter 2) that clinical decision systems should ideally have three additional capabilities: (1) the ability to maintain and manipulate a set of symbolic concepts, rather than mere numbers, (2) the ability to interact with clinicians using natural language, and (3) the ability to explain the reasoning process used to make conclusions. These goals were derived from his experience with a program that used decision analysis for the management of acute renal failure (Gorry et al., 1973). He concluded that detailed knowledge of medical concepts and the relationships between concepts would be required to reach reasonable conclusions reflecting a sense of the clinical context of the patient's problems. This could provide the program with a pragmatic view of the situation being analyzed. He encouraged the development of natural language communication in order to expedite the transfer of expertise, both from the expert to the program during the creation and expansion of the knowledge base and from the program to the user once the program becomes a clinical tool.

These criteria imply that the same piece of knowledge must be used in many different ways. The knowledge should be represented in a fashion that does not limit the manner in which it can be used. In many programming languages, one part of a program cannot access or modify another part. Thus incorporating the knowledge directly into the program's procedures limits the possible utilization of that knowledge. Facts must be in a form that can be manipulated as easily as numerical data are manipulated in conventional programming tasks.

The subfield of computer science known as artificial intelligence (AI) (Winston, 1977) has concentrated on using computers for symbolic reasoning rather than for calculating with numbers. One goal of our project has been to determine the strengths and limitations of the production rule methodology drawn from AI. Production rules offer the advantage of containing a small "packet" of knowledge. These packets can be combined to create a knowledge base of facts and relations known to the system. Using current symbolic processing languages, these rules can be translated from an external English-like syntax into an internal form that can be examined and interpreted by a task-independent control program. Because they can be displayed in English for communication with the user

RULE209

IF:
1) The site of the culture is blood, and
2) There is significant disease associated with this occurrence of the organism, and
3) The portal of entry of the organism is GI, and
4) The patient is a compromised host
THEN:
It is definite (1.0) that bacteroides is an organism for which therapy should cover

FIGURE 10-1 Example of a MYCIN rule. This is the English translation of a rule used to determine which organism may be causing the patient's infection.

and because they also facilitate the development of simple techniques for understanding natural language, production rules have allowed us to respond effectively to the design criteria outlined above.

10.3 Overview of MYCIN

MYCIN selects antimicrobial therapy for patients with severe infections (Shortliffe, 1976). The program uses knowledge obtained from infectious disease specialists; this knowledge was captured in the form of heuristics or "rules of thumb" that relate microbiological data and clinical signs and symptoms to possible pathogenic organisms. The details of the MYCIN program have been outlined in several other publications (referenced below) and will be described only briefly here.

10.3.1 Knowledge Representation

The MYCIN system is built around a set of medical concepts, such as the surgical history of the patient and the identity of infecting organisms. Each of these concepts is called a *clinical parameter.* Relations between the clinical parameters are used to build production rules of the form "IF premise THEN action." The premise of the rule is formed by the conjunction of statements about clinical parameters, for example, "The age is greater than 8" or "The patient has had recent neurosurgery." The action portion of the rule states what conclusions can be drawn from the premise with an associated measure of certainty. The English translation of a MYCIN rule is shown in Figure 10-1.

To perform a consultation, the rules must be combined together to form a line of reasoning (see Chapter 5). MYCIN uses a goal-directed approach to integrate the knowledge, a process known as *backward chaining*. Starting with the top-level goal (i.e., to prescribe appropriate therapy), the program selects the set of rules that make a conclusion about this goal in their action part. The premise of each of these rules is evaluated to determine if a rule can be applied. If a fact needed to evaluate this premise is not available, then the program identifies other rules that make conclusions about the needed fact (or asks the user if no rules exist). In this manner, only the portion of the rule set that is relevant to the particular patient is examined. The number of questions asked is also minimized by this goal-directed search through the knowledge base.

The consultation program manipulates the rules as described above, but itself contains no knowledge about infectious diseases. The system also contains explanation and question-answering facilities that interact with both the knowledge in the rule set and an ongoing record of how rules were applied during a consultation (Scott et al., 1977). The definition and propagation of the measure of uncertainty *(certainty factor)* associated with each rule have also been a major area of concentration (Shortliffe and Buchanan, 1975). Evaluations (Yu et al., 1979a; 1979b) have shown that the performance of the system approaches that of a subspecialist in the two areas (bacteremia and meningitis) for which the knowledge base has been developed.

10.4 Overview of VM

The VM program is designed to interpret on-line quantitative data in the intensive care unit (ICU). These data are used to manage postsurgical patients receiving mechanical ventilatory assistance. VM is an extension of a physiologic monitoring system (Osborn et al., 1969) and is designed to perform five specialized tasks in the ICU: (1) to detect possible measurement errors, (2) to recognize untoward events in the patient/machine system and suggest corrective action, (3) to summarize the patient's physiologic status, (4) to suggest adjustments to therapy based on the patient's status over time and long-term therapeutic goals, and (5) to maintain a set of patient-specific expectations and goals for future evaluation by the program. The program produces interpretations of the physiologic measurements over time, using a model of the therapeutic procedures in the ICU and clinical knowledge about the diagnostic implications of the data.

Most medical decision-making programs, including the MYCIN system described above, have based their advice on data available at one particular time. In actual practice, the clinician receives additional information from tests and observations over time and reevaluates the diagnosis and prognosis of the patient. Both the progression of the disease and the response

```
STATUS RULE: STABLE-HEMODYNAMICS
DEFINITION: Defines stable hemodynamics based on blood pressures and heart rates
APPLIES to patients on VOLUME, CMV, ASSIST, T-PIECE
COMMENT: Look at mean arterial pressure for changes in blood pressure and
systolic blood pressure for maximum pressures
IF
  HEART RATE is ACCEPTABLE
  PULSE RATE does NOT CHANGE by 20 beats/min. in 15 min.
  MEAN ARTERIAL PRESSURE is ACCEPTABLE
  MEAN ARTERIAL PRESSURE does NOT CHANGE by 15 torr in 15 min.
  SYSTOLIC BLOOD PRESSURE is ACCEPTABLE
THEN
  The HEMODYNAMICS are STABLE
```

FIGURE 10-2 Sample VM interpretation rule. The meaning of ACCEPTABLE varies with the clinical context—for example, the type of ventilatory assistance. VOLUME, CMV, ASSIST, and T-PIECE refer to types of ventilation therapies.

to prior therapeutic interventions are important for assessing the patient's situation.

Data are collected in different therapeutic contexts. In order to interpret the data properly, VM includes a model of the stages that a patient follows from ICU admission through the end of the critical monitoring phase. Correct interpretation of physiologic measurements depends on knowing which stage the patient is in. The goals for patient management are also stated in terms of these clinical contexts. The program maintains descriptions of the current and optimal ventilatory therapies for any given time.

Knowledge is represented in VM by production rules of the following form:

```
IF:    Relations about one or more parameters hold
THEN:  1) Make a conclusion based on these facts,
       2) Make appropriate suggestions to clinicians, and
       3) Create new expectations about the future values of parameters
```

Additional information associated with each rule includes the symbolic name, the rule group (e.g., rules about instrument faults), the main concept (definition) of the rule, and all of the therapeutic states in which it makes sense. Figure 10-2 shows a sample rule for determining hemodynamic stability.

The VM knowledge base includes rules to support five reasoning steps that recur whenever a new time segment begins: (1) characterizing measured data as reasonable or spurious; (2) determining the therapeutic state of the patient (currently the mode of ventilation); (3) adjusting expectations of future values of measured variables when patient state changes; (4) checking physiologic status, including cardiac rate, hemodynamics, ventilation, and oxygenation; and (5) checking compliance with long-term

```
INITIALIZING RULE: INITIALIZE-CMV
DEFINITION: Initialize expectations for patients on controlled mandatory ventilation (CMV) therapy
APPLIES to all patients on CMV
IF ONE OF:
  PATIENT TRANSITIONED FROM VOLUME TO CMV
  PATIENT TRANSITIONED FROM ASSIST TO CMV
THEN EXPECT THE FOLLOWING:
```

		[--------- acceptable range	---------]			
			[---- ideal	----]		
	very low	low	min	max	high	very high
Mean pressure	60	75	80	95	110	120
Heart rate		60			110	
Expired pCO2	22	28	30	35	42	50

FIGURE 10-3 Portion of an initializing rule. This rule establishes initial expectations of acceptable and ideal ranges of variables. Not all ranges are defined for each measurement. pCO2 is a measure of the percentage of carbon dioxide in expired air measured at the mouth.

therapeutic goals. Each reasoning step is associated with a collection of rules sorted by the type of conclusions made in the action portion of the rule, for example, all rules that determine the validity of the data.

10.4.1 Treating Measurement Ranges Symbolically

Most of the rules represent the measurement values symbolically, using the term ACCEPTABLE or IDEAL to characterize the appropriate ranges. The actual meaning of ACCEPTABLE changes as the patient moves from state to state, but the statement of the relation between the physiologic measurements remains constant. The use of symbolic statements (e.g., "HEART RATE is ACCEPTABLE") allows for the exposition of common principles of physiologic interpretation in different contexts. In addition, it minimizes the number of rules needed to describe the complexity of the diagnostic situation.

The meaning of the symbolic range is determined by rules that establish expectations about the value of measured data. For example, when a patient is taken off the ventilator, the upper limit of acceptability for the expired carbon dioxide measurement is raised. The actual numeric calculation of "expired pCO2 high" in the premise of any rule will change when the context switches (removal from ventilatory support), but the statement of the rules remains the same. An example of a rule that creates these expectations is shown in Figure 10-3.

10.4.2 Rule Interpretation

The VM rule interpreter is based on the MYCIN interpreter. The major changes include (1) forward-chaining (data-driven) rule invocation as opposed to backward chaining, (2) checking to see that information acquired in a previous time frame is still valid for making conclusions, and (3) cycling through appropriate parts of the rule set each time new information is available.

A data-driven approach is necessary to take advantage of the small set of measurement values available in each time frame. This means that the reasoning process works forward from the available information as opposed to working backward from a goal and obtaining information as necessary. Because of the demanding nature of the ICU environment, the system must acquire and interpret data with minimal staff intervention.

Each of the rule groups corresponding to the five reasoning steps mentioned above is considered in order. Each rule is examined to determine if it applies to the current context. The premise of the rule is examined to determine validity, and the appropriate conclusions are recorded by the program, as well as expectations on the future ranges of measurement values. Suggestions to clinicians are also printed out.

Often the examination of the rule premise requires the utilization of a value acquired earlier, for example, the temperature measurement, which is volunteered to the patient-monitoring system on an episodic basis. The reliability of the stored value is determined by evaluating either a time constant (for variables that predictably change over time) or a rule (for cases in which the assessment of a value's reliability is dependent on context-specific information). Associated with each parameter in the system is a specific mechanism for determining its reliability over time. If a measurement is concluded to be spurious or outdated, then it is treated as if it were unknown, requiring alternative methods for determining the status of the patient. The rule invocation process is repeated each time a new set of measurements is available (currently every 2 to 10 minutes).

Identical conclusions made in contiguous time frames are represented by the interval specified by the times of the first and last assertion. A list of these intervals summarizes the history of a particular conclusion. The evaluation of a rule clause such as "Patient hyperventilating for the past 30 minutes" is made by direct examination of the time intervals stored along with the conclusions, as opposed to looking at the original measurements. Expectations are associated with the appropriate measurement and are classified by duration and type, such as the upper limit of the acceptable range. Expectations can persist for a fixed interval, such as "for 20 minutes starting in 10 minutes," or for the duration of one or more clinical situations, for example, "while the patient is on the ventilator."

10.5 Comparison of Design Goals for MYCIN and VM

MYCIN was designed to serve in the ward setting as an expert consultant for antimicrobial therapy selection. A typical interaction might take place after the patient has been diagnosed and preliminary cultures have been drawn but little microbiological data are available. In critical situations, a tentative decision about therapy must often be made pending actual culture results. In return for assistance in making this decision, the clinician is asked to spend the small amount of time required to seek a consultation. As we have discussed, there are numerous challenges involved in the effort to motivate clinicians to use such a resource. The environment of the intensive care unit is quite different, however. Continuous surveillance and evaluation of the patient's status are required. The problem is one of making therapeutic adjustments over a long period of time, many of which are minor, such as adjusting the respiratory rate on the ventilator. The main reasons for interacting with VM would be to obtain status information or to investigate an unusual event. The program must therefore be able to interpret measurements with minimal human participation. When an interaction does take place, for example, when an unexpected event is noted by the program, it must be terse and concise.

This difference in the timing and style of the user/machine interaction has considerable impact on system design. For example, the VM system must (1) presume that the clinician's input into the system will be brief, (2) use historical data to determine the clinical situation, (3) be able to provide advice at any point in the hospital course of the patient, (4) be able to follow up on the outcomes of previous therapeutic decisions, and (5) be able to provide summaries of conclusions made over time. VM's environment thus differs from MYCIN's in that typed natural language input is an unlikely modality for communication with the clinician.

A consultation program should also be able to model the changing medical environment so that the program can interpret the available data in the appropriate context. Of course, areas like infectious diseases often have critical points where a consultation is most necessary. In the development of the meningitis section of the MYCIN knowledge base, the concept of "partially treated meningitis" (prior treatment with an antibiotic) was handled quite distinctly from the untreated case, even though the laboratory findings might be identical.

It was also necessary for VM to contain knowledge that could be used to evaluate the results of its therapeutic advice, just as a human consultant follows a case over a period of time. This is complicated by the fact that the user of the system may not follow the recommended therapy regimen. If the patient does not react as expected to the given therapy, then the program has to determine what alternative therapeutic steps may be required.

10.6 Extending the MYCIN Design

The VM program has been used as a test-bed to investigate methods for increasing the capabilities of symbolic processing approaches by extending the production rule methodology. The main area of investigation has been in the representation of knowledge about dynamic clinical settings. There are two components to representing a situation that changes over time: (1) providing the mechanism for accessing and evaluating data in a new time frame, and (2) building a symbolic model to represent the ongoing processes in the medical environment.

Another aspect of VM development has been to experiment with more general extensions to production rules based on observations of the use of the MYCIN system. These changes can be described by two research directions: (1) expanding the level of detail in the knowledge base, and (2) increasing the global structure of the knowledge base. The problem of designing an advice-giving program with limited user/machine interaction has also been explored.

10.6.1. Representing Knowledge About Dynamic Clinical Settings

With VM we have begun to experiment with mechanisms for providing MYCIN-like systems with the ability to represent the dynamic nature of the diagnosis and therapy process. The original MYCIN system was designed to produce therapeutic decisions for one critical moment in the patient's hospital course. This was extended with a "restart mechanism" that allows for selectively updating those parameters that might change in the interval between consultations. MYCIN can start a new consultation with the updated information, but the results of the original consultation are lost. In VM three requirements are necessary to support the processing of new time frames: (1) examining the values of historical data and conclusions, (2) determining the validity of those data, and (3) combining new conclusions with previous conclusions.

New *premise functions,* which define the relationships about parameters that can be tested when a rule is checked for validity, were created to examine the historical data. Premise functions used in MYCIN include tests to see if: (a) any value has been determined for a parameter, (b) the value associated with a parameter is in a particular numerical range, or (c) there is a particular value associated with a parameter. VM includes a series of time-related premise functions. The first function examines trends in input data over time, for example, "The mean arterial pressure does not rise by 15 torr in 15 minutes." A second function determines the stability of a series of measurements by examining the variation of measurements over a specific time period. Other functions examine previously deduced conclusions, as in "The patient has been on the T-piece for greater than 30 minutes" or "The patient has never been on the T-piece." Functions also exist for determining changes in the state of the patient, for example, "The patient has transitioned from assist mode to the T-piece." When VM is required to check whether a parameter has a particular value, it must also check to see if the value is "recent" enough to be useful.

The notion that data are reliable for only a given period of time is also used in the representation of conclusions made by the program. When the same conclusion is made in contiguous time periods (two successive evaluations of the rule set), then the conclusions are coalesced. The result is a series of intervals that specify when a parameter assumed a particular value. In the MYCIN system this information is stored as several different parameters. For example, the period during which a drug was given is represented by a pair of parameters corresponding to the starting and ending times of administration. In MYCIN, if a drug was again started and stopped, a new entity—DRUG-2—would have to be created. The effect of the VM representation is to aggregate individual conclusions into "states" whose persistence denotes a meaningful interpretation of the status of the patient.

10.6.2 Building a Symbolic Model

A sequence of states recognized by the program represents a segmentation of a time line. Specifying the possible sequences of states in a dynamic setting constitutes a symbolic model of that setting. The VM knowledge base contains a model of the ventilatory therapies. This model is used in three ways by the program: (1) to limit the number of rules examined by the program, (2) to provide a basis for comparing actual therapy with potential therapies, and (3) to provide the basis for the adjustment of expectations used to interpret the incoming data.

Attached to each rule in VM is a list of the clinical situations in which the rule makes sense. When rules are selected for evaluation, this list is examined to determine if the rule is applicable. This provides a convenient filter to increase the speed of the program. A set of rules is utilized to specify the conditions for suggesting alternative therapeutic contexts. Since these rules are examined every few minutes, they serve both to suggest when the patient's condition has changed sufficiently for an adjustment in ventilatory therapy and to provide commentary concerning clinical maneuvers that have been performed but are not consistent with the embedded knowledge for making therapeutic decisions. The model also provides mechanisms for defining expectations about reasonable values for the measured data. Much of the knowledge in VM is stated in terms of these expectations, and they can be varied in response to changes in the patient's situation.

RULE 236
(This rule applies to organisms from positive cultures, and is tried in order to find out about the infection that requires therapy or whether there is significant disease associated with this occurrence of the organism.)

IF:
1) The site of the culture is urine, and
2) The method of collection of the culture is voided, and
3) The colony count (in thousands) of the organism is greater than or equal to 100

THEN:
1) There is suggestive evidence (0.5) that the infection that requires therapy is cystitis, and
2) There is suggestive evidence (0.7) that there is significant disease associated with this occurrence of the organism

Author: Yu.
Comments: This definition of significance differs from E. Kass's original definition (Am. J. Med., 18:764, 1955) where two consecutive cultures are required. However, for practical purposes, if the patient is symptomatic, physicians generally start treatment on the basis of only one culture.
Created: 19 May 1977, 13:43.
Last edited: 1 June 1977, 11:50.

FIGURE 10-4 Example of a MYCIN rule with justificatory comments.

10.6.3 Expanding the Level of Detail in the Knowledge Base

Those who implement production rule systems often assume that the knowledge to be represented will be broken into small pieces corresponding to individual rules. What would happen in MYCIN if this assumption were violated? At one extreme there would be a single rule that weighed all of the clinical inputs in order to conclude the presence or absence of a single organism, say *E. coli,* but this would be too large and complicated to understand. The other extreme would be to base the deductive steps on the most minute details of physiologic knowledge, for example, knowledge of the cell wall properties of each species of bacteria. Explanation and modification would be very difficult in either situation. The approach taken in the development of MYCIN has been between these two extremes. Although no fixed criteria have been established, an examination of the rule set shows that intermediate steps have been left out when they appeared to be definitional in nature. Since the major performance requirements of a consultation system, that is, reaching correct hypotheses, revolve around propagation of the uncertainty associated with each piece of knowledge, definitional facts affect the outcome primarily by providing "commonsense" domain knowledge. Currently each of MYCIN's rules is augmented with a free-text justification or rationale that discusses some of the intervening steps that were used in formulating the particular content of that rule. The text justifications are available to the user if the basis for the knowledge in a rule is not clear from the translation of the rule itself (Figure 10-4).

Our representation of medical knowledge has been particularly stereotyped so that the programs we write can examine and manipulate the knowledge in many different ways. For example, in the middle of a MYCIN consultation the user can ask for an explanation of why a particular question is asked, resulting in the description of the chain of reasoning leading to the current rule under consideration (Davis, 1976; Scott et al., 1977). However, because a rule's justification is stored as unformatted text, it is unavailable for dissection and manipulation by the program as it gives explanations. It has become clear to us that the development of more formal mechanisms for encoding the basic knowledge that underlies a single rule (in a form that a computer can manipulate) will improve the educational and explanatory features of the program by providing an additional level of detail that can be explored and utilized programmatically. The detailed justifications could also be used for consistency checking since they represent the same knowledge but are stated in terms of "first principles." The requirements for augmenting the knowledge base in this way for the purpose of tutoring medical students have been described by Clancey (1979a).

The approach taken in VM is to introduce additional rules that are often definitional in nature (e.g., the rule in Figure 10-2 that defines hemodynamic stability). We have found that these additional rules act to form a convenient method for introducing abstract concepts into the rule base. This, in turn, has provided a basis for separating out the portion of the knowledge that was independent of the current context, for example, the physiology, from the knowledge that must adjust to the changing medical situation.

10.6.4 Increasing the Structure of the Knowledge Base

In addition to the need for more highly formalized justifications associated with each rule, we have observed the potential value of a more global organization of the rule base. In the development of a set of rules for the treatment of meningitis, we identified a situation in which a series of very similar rules were used to represent a "case analysis" of patient findings. The development of the meningitis knowledge base also included the need to represent *default decision rules* that applied to the majority of the patients considered but could still be customized for individual patient histories. The problem was broken up into a master rule that would make a preliminary set of conclusions and more specific rules that could modify the preliminary conclusions in response to unusual items from the patient's history. These more specific rules, therefore, cannot be understood without first considering the default rule. Two different methods can be used to handle this dependence. The first would be to rewrite each of the specific rules in order to incorporate all of the information in the default rule (and the default rule would then have to be changed to specifically exclude each of the special cases). Then each of the rules would be more complex but somewhat more independent. However, it would be difficult to relate the differences in conclusions based on one special situation versus another. An alternative solution would be to recognize the inherent structure in the segment of knowledge that has been distributed across several rules. A technique used in other symbolic processing approaches (Pople, 1977) (see also Chapter 6) is to promote prototypical situations (and their exceptions) as the basic unit of knowledge representation. Information in these systems is often organized around individual diagnoses and groups together all of the knowledge pertaining to a particular disease. This method has the disadvantage that the size of each of these prototypical units, known as *frames* or *schemas* (Minsky, 1975), can become too large to comprehend. These organizational structures can also be used to provide for a more coherent consultation by supplying a larger context for the question-asking mechanism.

We have experimented with another representation for structuring the rule base: the creation of a rule set containing knowledge about the medical knowledge of the system (*meta-knowledge*) (Davis, 1976). These *meta-rules* can be used as "strategy rules" to order the application of rules in the knowledge base. They provide a heuristic mechanism for taking into account the facts that some information may be more relevant for making a specific conclusion and that other rules, although potentially applicable, can likely be ignored.

Another use for a global structure overlaid on the knowledge base would be to provide for anatomical models. Reggia (1978) suggests that this would have been useful in the development of a production rule system for neurological localization. Aikins (1979) has explored the combination of production rules and frames using the MYCIN methodology for the interpretation of pulmonary function tests.

The designers of future rule-based systems should consider some of the above methods for providing a global structure for the knowledge. Not all rules can be considered independently, and when rules are related, the connections should be available for manipulation by the computer.

10.6.5 Handling Limited User Input

In the intensive care unit, the lack of communication is partly solved by the availability of a large mass of on-line computer-processed data. Another approach to solving the communication problem is to display for the clinician conditional conclusions that require clinical observation before being carried out. For example, rather than asking whether a patient is sweating, VM might display a recommendation such as "If the patient is diaphoretic I suggest . . . , otherwise. . . ."

One additional solution to the problem of limited user-to-machine communication would be to anticipate the key questions that might be posed by the clinician at the bedside and provide a "menu" of likely questions for exploring the conclusions generated by the program. During the development of part of the meningitis knowledge base, the MYCIN program was modified to generate automatically the answers to a few key questions specified in advance by the medical expert. Such a key question for the ICU setting is "What is the status of ventilatory therapy?" The program, by the evaluation of several of the rules, can produce the following type of explanation: "Before transition to the T-piece can be suggested, hemodynamic stability must be present, which requires systolic blood pressure to be acceptable (current systolic blood pressure value is 170)."

10.7 Summary

Several years of experience with the MYCIN program have led to an understanding of additional requirements for symbolic processing approaches to medical decision making. These include extending the knowledge base beyond the facts necessary for high performance, providing an organizing structure for a large number of production rules, and extending the decision-making aids to include assistance throughout the patient's clinical course. For decision aids in the intensive care unit or other equally dynamic situations, programs cannot depend on interaction with the clinical users. Furthermore, they must handle data that are changing over time, but might be missing or spurious. They must also be able to provide tracking of the patient's status during the course of the underlying disease or in response to therapeutic intervention. A more complete description of the VM program can be found in Fagan (1980).

ACKNOWLEDGMENTS

We wish to thank Jonathan King for his comments on an earlier draft of this paper. The following people have also contributed to the work described here: J. S. Aikins, S. G. Axline, W. J. Clancey, S. N. Cohen, R. Davis, E. A. Feigenbaum, J. C. Kunz, J. J. Osborn, B. J. Rubin, A. C. Scott, S. M. Wraith, and V. L. Yu. This work has been supported by the National Institutes of Health (general medical sciences grant GM-24669) and by the Bureau of Health Sciences (research and evaluation grant HS-01544). Dr. Shortliffe is the recipient of a research career development award (LM-00048) from the National Library of Medicine. Computing resources have been provided by the SUMEX-AIM facility, under NIH grant RR-00785.