22

Extensions to the Rule-Based Formalism for a Monitoring Task

Lawrence M. Fagan, John C. Kunz, Edward A. Feigenbaum, and John J. Osborn

The Ventilator Manager (VM) program is an experiment in expert system development that builds on our experience with rules in the MYCIN system. VM is designed to interpret on-line quantitative data in the intensive care unit (ICU) of a hospital. After a major cardiovascular operation, a patient often needs mechanical assistance with breathing and is put in the ICU so that many parameters can be monitored. Many of those data are relevant to helping physicians decide whether the patient is having difficulty with the breathing apparatus (the ventilator) or is breathing adequately enough to remove the mechanical assistance. The VM program interprets these data to aid in managing postoperative patients receiving mechanical ventilatory assistance.

VM was strongly influenced by the MYCIN architecture, but the program was redesigned to allow for the description of events that change over time. VM is an extension of a physiologic monitoring system[1] and is designed to perform five specialized tasks in the ICU:

This chapter is a longer and extensively revised version of a paper originally appearing in *Proceedings of the Sixth IJCAI* (1979, pp. 260–262). Used by permission of International Joint Conferences on Artificial Intelligence, Inc.; copies of the *Proceedings* are available from William Kaufmann, Inc., 95 First Street, Los Altos, CA 94022.

[1]VM was developed as a collaborative research project between Stanford University and Pacific Medical Center (PMC) in San Francisco. It was tested with patient information acquired from a physiologic monitoring system implemented in the cardiac surgery ICU at PMC and developed by Dr. John Osborn and his colleagues (Osborn et al., 1969).

1. detect possible errors in measurement,
2. recognize untoward events in the patient/machine system and suggest corrective action,
3. summarize the patient's physiologic status,
4. suggest adjustments to therapy based on the patient's status over time and long-term therapeutic goals, and
5. maintain a set of case-specific expectations and goals for future evaluation by the program.

VM differs from MYCIN in two major respects. It interprets measurements over time, and it uses a state-transition model of intensive care therapies in addition to clinical knowledge about the diagnostic implications of data. Most medical decision-making programs, including MYCIN, have based their advice on the 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 to previous therapy are important for assessing the patient's situation.

Data are collected in different therapeutic situations, or *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. The correct interpretation of physiologic measurements depends on knowing which stage the patient is in. The goals for intensive care are also stated in terms of these clinical contexts. The program maintains descriptions of the current and optimal ventilatory therapies for any given time. Details of the VM system are given by Fagan (1980).

22.1 The Application

The intensive care unit monitoring system at Pacific Medical Center (Osborn et al., 1969) was designed to aid in the care of patients in the period immediately following cardiac surgery. The basic monitoring system has proven to be useful in caring for patients with marked cardiovascular instability or severe respiratory malfunction (Hilberman et al., 1975). Most of these patients are given breathing assistance with a mechanical ventilator until the immediate effects of anesthesia, surgery, and heart-lung bypass have subsided. The ventilator is essential to survival for many of these patients. Electrocardiogram leads are always attached, and patients usually have indwelling arterial catheters to assure adequate monitoring of blood pressure and to provide for the collection of arterial blood for gas analysis. The ventilator-to-patient airway is monitored to collect respiratory flows, rates, and pressures. Oxygen and carbon dioxide concentrations are also measured. All of these measurements are available at the bedside through the use of specialized computer terminals.

The mechanical ventilator provides total or partial breathing assistance (or ventilation) for seriously ill patients. Most ventilator therapy is with a type of machine that delivers a fixed volume of air with every breath, but a second type of machine delivers air at each breath until a fixed pressure is attained. Both the type and settings of the ventilator are adjusted to match the patient's intrinsic breathing ability. The "volume" mechanical ventilator provides a fixed volume of air under pressure through a tube to the patient. The ventilator can be adjusted to provide breaths at fixed intervals, which is called *controlled mandatory ventilation* (CMV), or in response to sucking by the patient, which is known as *assist mode.* Adjustments to the output volume or the respiration rate of the ventilator are made to ensure an adequate *minute volume* to the patient. When the patient's status improves, the mechanical ventilator is disconnected and replaced by a *T-piece* that connects an oxygen supply with the tube to the patient's lungs. If the patient can demonstrate adequate ventilation, the tube is removed (*extubation*). Often many of these clinical transitions must be repeated until the patient can breathe without assistance.

Three types of problems can occur in managing the patient on the mechanical ventilator:

1. changes in the patient's recovery process, requiring modifications to the life support equipment,
2. malfunctions of the life support equipment, requiring replacement or adjustment of the ventilator, and
3. failures of the patient to respond to therapeutic interventions within the expectations of the clinicians in charge.

22.2 Overview of the Ventilator Manager Program

The complete system (diagrammed in Figure 22-1) includes the patient monitoring sensors in the ICU, the basic monitoring system running on IBM 1800 and PDP-11 computers at the Pacific Medical Center, and the VM measurement interpretation program running on the SUMEX-AIM PDP-10 computer located at Stanford University Medical Center. Patient measurements are collected by the monitoring system for VM at two- or ten-minute intervals. Summary information, suggestions to the clinicians, and requests for additional information are generated at SUMEX for evaluation by research clinicians. The program's outputs are in the form of periodic graphical summaries of the major conclusions of the program and short suggestions for the clinician (as shown in Figures 22-2 and 22-3).

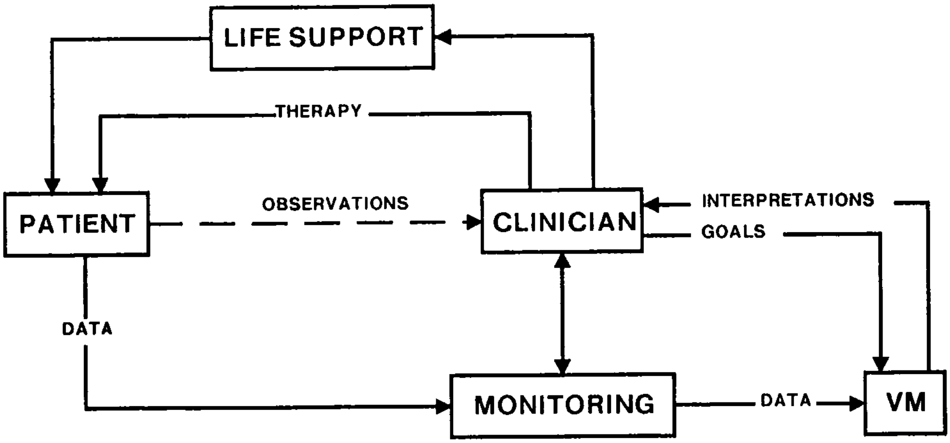

FIGURE 22-1 VM system configuration. Physiological measurements are gathered automatically by the monitoring system and provided to the interpretation program. The summary information and therapeutic suggestions are sent back to the ICU for consideration by clinicians.

```
Summary generated at time 15:40

All conclusions:
                                . . . | . . . . . | . . . . . | . . . . . | . . .
                                    1 2         1 3         1 4         1 5
        BRADYCARDIA[PRESENT]                = = = = = = = = = =     = =
        HEMODYNAMICS[STABLE]                                            = =
        HYPERVENTILATION[PRESENT]   = = = = = = = = = = = = = = = = = = =   = =
        HYPOTENSION[PRESENT]              = = = = = = =
        Goal Location               C C C C C C C C C C C C C / A A A A A A A A A A A
        Patient Location            V / C C C C C C C C C C C C C C C C C C C C C C C
                                . . . | . . . . . | . . . . . | . . . . . | . . .
                                    1 2         1 3         1 4         1 5
```

FIGURE 22-2 Summary of conclusions drawn by VM based on four hours of patient data. Current and optimal patient therapy stages are represented by their first letter: V = VOLUME, A = ASSIST, C = controlled mandatory ventilation, / = changing. A double bar (=) is printed for each ten-minute interval in which the conclusion on the left is made.

```
..  1640..
**  SUGGEST CONSIDER PLACING PATIENT ON T-PIECE IF
**  PA02 > 70 ON FI02 <= .4                    [measure of blood gas status]
**  PATIENT AWAKE AND TRIGGERING VENTILATOR
**  ECG IS STABLE

..  1650 .... 1700 .... 1710 .... 1720 .... 1730 .... 1740 .... 1750 ..
..  1800 ..
**  HYPERVENTILATION
**  PATIENT HYPERVENTILATING.
**  SUGGEST REDUCING EFFECTIVE ALVEOLAR VENTILATION.
**  TO REDUCE ALVEOLAR VENTILATION, REDUCE TIDAL VOLUME,
**  REDUCE RESPIRATION RATE, OR
**  INCREASE DISTAL DEAD SPACE TUBING VOLUME

..  1810 ..
**  SYSTEM ASSUMES PATIENT STARTING T-PIECE

..  1813 .... 1815 .... 1817 ..
**  HYPOVENTILATION

..  1819 .... 1822 ..
**  HYPOVENTILATION
```

FIGURE 22-3 Trace of program output. Format is ". . <time of day> . ." followed by suggestions for clinicians. Comments are in brackets.

22.2.1 Measurement Interpretation

Knowledge is represented in VM by production rules of the form shown in Figure 22-4.

The historical relations in the premise of a rule cause the program to check values of parameters for a period of time; e.g., HYPERVENTILATION is PRESENT for more than ten minutes. *Conclusions* made in the action part of the rule assert that a parameter has had a particular value during the time instance when the rule was examined. *Suggestions* are text statements, printed out for clinicians, that state important conclusions and a possible list of remedies. *Expectations* assert that specific measurements

```
IF:    Historical 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.
```

FIGURE 22-4 Format for rules in VM. Not every rule's action part includes conclusions, suggestions, and expectations.

should be within the specified ranges at some point in the future. Thus a rule examines the current and historical data to interpret what is happening at the present and to predict events in the future.

Additional information associated with each rule includes the symbolic name (e.g., STABLE-HEMODYNAMICS), 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 to apply the rule. The list of states is used to focus the program on the set of rules that are applicable at a particular point in time. Figure 22-5 shows a sample rule for determining hemodynamic stability (i.e., a measure of the overall status of the cardiovascular system).[2]

STATUS RULE: STABLE-HEMODYNAMICS
DEFINITION: Defines stable hemodynamics based on blood pressures and heart rate
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/minute in 15 minutes
    MEAN ARTERIAL PRESSURE is ACCEPTABLE
    MEAN ARTERIAL PRESSURE does NOT CHANGE by 15 torr in 15 minutes
    SYSTOLIC BLOOD PRESSURE is ACCEPTABLE
THEN
    The HEMODYNAMICS are STABLE

FIGURE 22-5 Sample VM rule.

The VM knowledge base includes rules to support five reasoning steps that are evaluated at the start of each new time segment:

1. characterize measured data as reasonable or spurious;
2. determine the therapeutic state of the patient (currently the mode of ventilation);
3. adjust expectations of future values of measured variables when the patient state changes;
4. check physiological status, including cardiac rate, hemodynamics, ventilation, and oxygenation; and
5. check compliance with long-term therapeutic goals.

Each reasoning step is associated with a collection of rules, and each rule is classified by the type of conclusions made in its action portion; e.g., all rules that determine the validity of the data are classed together.

---

[2]The complete rule set, from which this rule was selected, is included in the dissertation by Fagan (1980), which is available from University Microfilms, #AAD80-24651.

22.2.2 Treating Measurement Ranges Symbolically

Most of the rules represent the measurement values symbolically, using the terms 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 physiological measurements remains constant. For example, the rule shown in Figure 22-5 checks to see if the patient's heart rate is ACCEPTABLE. In the different clinical states, or stages of mechanical assistance, the definition of ACCEPTABLE changes. Immediately after cardiac surgery a patient's heart rate is not expected to be in the same range as it is when he or she is moved out of the ICU. Mentioning the symbolic value ACCEPTABLE in a rule, rather than the state-dependent numerical range, thus reduces the number of rules needed to describe the diagnostic situation.

The meaning of the symbolic range is determined by other rules that establish expectations about the values 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. (Physiologically, a patient will not be able to exhale all the CO_2 produced by his or her system, and so CO_2 will accumulate.) 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. A sample rule that creates these expectations is shown in Figure 22-6.

22.2.3 Therapy Rules

Therapy rules can be divided into two classes: the long-term therapy assessment (e.g., when to put the patient on the T-piece), and the determination of response to a clinical problem, such as hyperventilation or hypertension. The two rules shown in Figure 22-7, for selecting T-piece therapy and for responding to a hyperventilation problem, demonstrate several key factors in the design of the rule base:

- use of a hierarchy of physiological states,
- use of the program's determination of patient's clinical state,
- generation of conditional suggestions.

The abstracted hierarchy of states, such as hemodynamic stability, is important because it makes the rules more understandable. Since the definition of stability changes with transition to different clinical stages, as described above, rules about stability are clearer if they mention the concept rather than the context-specific definition. It is important for the program to determine what state the patient is in, since the program is

```
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 22-6 Portion of an initializing rule. This type of rule establishes initial expectations of acceptable and ideal ranges of variables after state changes. Not all ranges are defined for each measurement. EXPIRED pCO2 is a measure of the percentage of carbon dioxide in expired air measured at the mouth.

designed to avoid interrupting the activities in the ICU to ask questions of the physicians or nurses. Its design is thus different from the design of a one-shot consultation system such as MYCIN. A physician will change the mode of assistance from CMV (where the machine does all the work of breathing) to ASSIST (where the machine responds to a patient's attempts to breathe). The VM program has to know that this transition is normal and to determine when it occurs in order to avoid drawing inappropriate conclusions. The advice that VM offers is often conditional. Unlike other consultation programs such as MYCIN, VM attempts to avoid a dialogue with the clinician. When the appropriateness of a suggestion depends on facts not known to VM, it creates a conditional suggestion. The clinician can check those additional facts and make an independent determination of the appropriateness of the suggestion.

22.2.4 Selecting Optimal Therapy

The stages of ventilatory therapy are represented in VM by a finite state graph (see Figure 22-8). The boxed nodes of the graph represent the values associated with the parameters "PatientLocation," specifying the current state, and "GoalLocation," specifying alternative therapies. The arcs of the graph represent transition rules and therapy rules. Thus goals are expressed as "moves" away from the current therapeutic setting, and each

```
THERAPY-RULE: THERAPY.A-T
DEFINITION: DEFINES READINESS TO TRANSITION FROM ASSIST MODE TO T-PIECE
COMMENT: If patient has stable hemodynamics, ventilation is acceptable, and patient has been
         awake and alert enough to interact with the ventilator for a period of time then transition
         to T-piece is indicated.
APPLIES TO: ASSIST
IF
  HEMODYNAMICS ARE STABLE
  HYPOVENTILATION NOT PRESENT
  RESPIRATION RATE ACCEPTABLE
  PATIENT IN ASSIST FOR > 30 MINUTES
THEN
  THE GOAL IS FOR THE PATIENT TO BE ON THE T-PIECE
  SUGGEST CONSIDER PLACING PATIENT ON T-PIECE IF
    Pa02 > 70 on FI02 <= 0.4
    PATIENT AWAKE AND TRIGGERING VENTILATOR
    ECG IS STABLE

THERAPY-RULE: THERAPY.VENTILATOR-ADJUSTMENT-FOR-HYPERVENTILATION
DEFINITION: MANAGE HYPERVENTILATION
APPLIES TO: VOLUME ASSIST CMV
IF
  HYPERVENTILATION PRESENT for > 8 minutes
  COMMENT wait a short while to see if hyperventilation persists
  V02 not low
THEN
  SUGGEST PATIENT HYPERVENTILATING.
    SUGGEST REDUCING EFFECTIVE ALVEOLAR VENTILATION.
    TO REDUCE ALVEOLAR VENTILATION, REDUCE TV BY 15%, REDUCE RR, OR
      INCREASE DISTAL DEAD SPACE TUBING VOLUME
```

FIGURE 22-7 Two therapy rules. The first (THERAPY.A-T) suggests a T-piece trial; the second resolves a hyperventilation problem.

possible move corresponds to a decision rule. The overall clinical goal, of course, is to make the patient self-sufficient, specifically, to remove the mechanical breathing assistance (extubate) as soon as is practical for each patient. The knowledge base is linked to the graph through the APPLIES TO statement specified in the introductory portion of each rule.

The mechanism for deriving and representing therapy decisions in VM takes into account the relationship between VM's suggestions and actual therapy changes. Computer-generated suggestions about therapy changes are decoupled from actual changes due to: (1) additional information to the clinician suggesting modification to or disagreement with VM's suggestion; (2) sociologic factors that delay the implementation of the therapy decisions (e.g., T-piece trials have been delayed due to concern about disturbing a patient in the next bed); or (3) variation of criteria among clinicians for making therapy decisions. Because of the discrepancy between computer-generated goals and actual therapy, VM cannot assume that the patient is actually in the stage that the program has determined

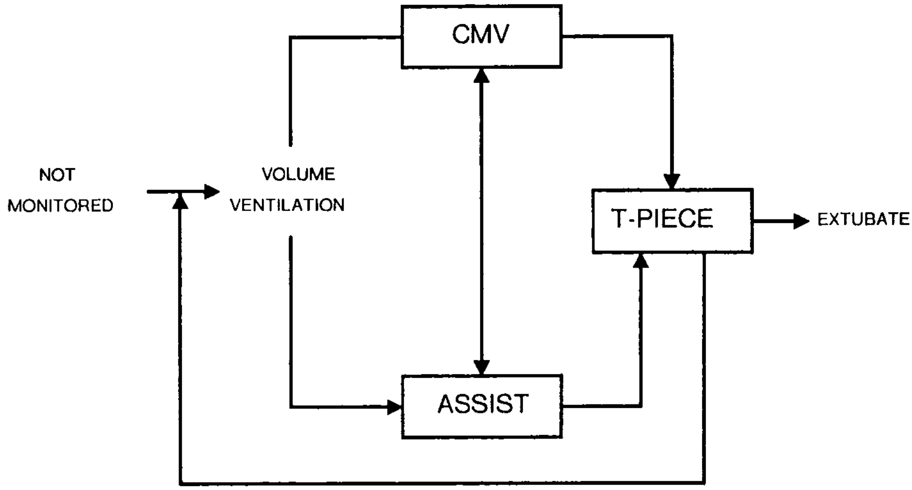

FIGURE 22-8 Therapy state graph.

is optimal. Transition rules in VM allow the program to notice changes in a patient's state. They reset the description of the context, then, so that the data will be interpreted correctly. However, when the therapy rules are evaluated, the program may determine that the previous state is still more appropriate.

Two models can be created for representing the period of time between the suggestion of therapy (a new goal) and its implementation. The first model is that therapy goals are the same as last stated, unless explicitly changed. It assumes that once a new therapeutic goal is established, the goal should persist until either the therapy is initiated or the goal is negated by a rule. This model is based on the common clinical practice of continuing recently initiated therapy even if the situation has changed. This clinical practice, which might be termed *hysteresis*,[3] is used to avoid frequent changes in treatment strategy—i.e., avoid oscillation in the decision-making process. While clinicians acknowledge this behavior, they find it hard to verbalize rules for rescinding previous therapy goals. This hysteresis has also been evident in the formulation of some of the therapy rules. The rule that suggests a switch from assist mode to T-piece is stated in terms of ACCEPTABLE limits; the rule for aborting T-piece trials (back to assist mode or CMV) is stated in terms of VERY HIGH or VERY LOW limits. This leaves a "grey area" between the two decision points and precludes fluctuating between decisions.

[3]Hysteresis is "a lag of effect when the forces acting on a body are changed" (*Webster's New World Dictionary*, 1976).

The second model for representing therapeutic goals requires that the appropriate goal be asserted each time the rule set is evaluated. If no therapy rules succeed in setting a new goal, the goal is asserted to be the current therapy. This scheme ignores the apparent practices of clinicians, but represents a more "conservative" approach that is consistent with the rule-writing strategy used by our experts. This model is potentially sensitive to minor perturbations in the patient measurements, but such sensitivity implies that a borderline therapy decision was originally made.

22.3 Details of VM

22.3.1 Parameters

The knowledge in VM is based on relationships among the various parameters of the patient, such as respiration rate, sex, and hyperventilation. The program assigns values to each of these parameters as it applies its knowledge to the patient data: the respiration rate is high, the sex of the patient is male, and hyperventilation is present. In a changing domain, the values associated with each parameter may vary with time, for example, "hyperventilation was present for one-half hour, starting two hours ago." Not all parameters have the same propensity to change over time; a classification is given in Figure 22-9.

Parameters are represented internally by using the property list notation of LISP. The property list contains both static elements (e.g., the list of rules that use the parameter in the premise) and dynamic elements (e.g., the time when the parameter was last updated). The static elements are input when the parameter is described or calculated from the contents of the rule set. The dynamic elements are computed as the program interprets patient data. Figure 22-10 lists the properties associated with parameters (although not every parameter has every property).

Figure 22-11 shows a "snapshot" of the parameters RR (respiration rate) and HEMODYNAMICS taken after 120 minutes of data have been processed. Associated with values assigned to parameters (e.g., RR LOW or HEMODYNAMICS STABLE) are lists of intervals when those conclusions were made. Each interval is calculated in terms of the elapsed time since patient data first became available. Thus, in the example, the hemodynamics were stable from 2–8 minutes into the program, momentarily at 82 minutes, and in the interval of 99–110 minutes of elapsed time.

The properties USED-IN, CONCLUDED-IN, and EXPECTED-IN are used to specify how the parameters are formed into a network of rules. These pointers can be used to guide various strategies for examining the rules—e.g., find and evaluate each rule that uses respiration rate or each rule that concludes hemodynamic status.

Constant

Examples: surgery type, sex
Input: once
Reliability: value is good until replaced

Continuous

Examples: heart rate, blood pressure
Input: at regular intervals (6–20 times/hour)
Reliability: presumed good unless input data are missing or artifactual

Volunteered

Examples: temperature, blood gases
Input: at irregular intervals (2–10 times/day)
Reliability: good for a period of time, possibly a function of the current situation

Deduced

Examples: hyperventilation, hemodynamic status
Input: calculated whenever new data are available
Reliability: a function of the reliability of each of the component parameters.

FIGURE 22-9 Classification of parameters.

The UPDATED-AT and GOOD-FOR properties are used to determine the status of the parameter over time, when it was last given a value and the time period during which a conclusion made about this parameter can reasonably be used for making future conclusions. The GOOD-FOR property can also be a pointer to a context-dependent rule.

DEFINITION: free form text describing the parameter
USED-IN: a list of the names of rules that use this parameter to make conclusions
CONCLUDED-IN: names of rules where this parameter is concluded
EXPECTED-IN: names of rules where expectations about this parameter are made
GOOD-FOR: length of time that a measurement can be assumed to be valid; if missing, then must be recomputed, input if possible, or assumed unknown
UPDATED-AT: the last time any conclusion was made about this parameter

FIGURE 22-10 Properties associated with parameters.

```
RR
--
  DEFINITION: (RESPIRATION RATE)
  USED-IN: (TRANSITION.V-CMV TRANSITION.V-A TRANSITION.A-CMV
            TRANSITION.CMV-A STATUS.BREATHING-EFFORT/T THERAPY.A-CMV
            THERAPY.A-T THERAPY.T-V ABNORMAL-EC02)
  EXPECTED-IN: (INITIALIZE.V INITIALIZE.CMV INITIALIZE.V-RETURN
                INITIALIZE.A INITIALIZE.T-PIECE)
  GOOD-FOR: 15                [information is good for 15 minutes]
  UPDATED-AT: 82              [last updated at 82 minutes after start]
  LOW: ((72 . 82) (52 . 59))  [concluded to be LOW from 52-59 minutes and 72-82
                              minutes after start]


HEMODYNAMICS
------------
  DEFINITION: (HEMODYNAMICS)
  CONCLUDED-IN: (STATUS.STABLE-HEMODYN/V,A,CMV)
  USED-IN: (THERAPY.CMV-A THERAPY.A-T THERAPY.T-PIECE-TO-EXTUBATE)
  GOOD-FOR: NIL              [this is a derived parameter so reliability is based on
                             other parameters]
  UPDATED-AT: 110            [last updated at 110 minutes after start]
  STABLE: ((99 . 110) (82 . 82) (2 . 8))
```

FIGURE 22-11 "Snapshot" of parameters RR (respiration rate) and HEMODYNAMICS after 120 minutes of elapsed time.

22.3.2 Measurements

Over 30 measurements are provided to VM every 2 to 10 minutes.[4] The interval is dependent on the situation; shorter intervals are used at critical times as specified by the clinician at the bedside. It is not appropriate to store this information using the interval notation above, since most measurements change with every new collection of data. Predefined intervals, e.g., respiration rate from 5–10, 10–15, and 15–20 breaths/minute, could be used to classify the data, but meaningful ranges change with time. Instead, symbolic ranges such as HIGH and LOW are calculated from the measurements as appropriate. A large quantity of data is presented to the program, in contrast to typical knowledge-based medical systems. About 5000 measurement values per patient are collected each day (30 measurements per collection with 6–8 data collections/hour). Patients are monitored from a few hours to a few weeks, with the average about 1.5 days. While this amount of information could be stored in a large-scale program such as VM, only the most recent information is used to make conclusions. The program stores in memory about one hour's worth of data, independent of the time interval between measurement collections (the remainder

[4]Clinicians can select the default sample rate: fast (2 minutes) or slow (10 minutes). An extra data sample can be taken immediately on request.

of the data are available on secondary storage). Technically, this storage is accomplished by maintaining a queue of arrays that contain the entire collection of measurements that vary over time. The length of the queue is adjusted to maintain an hour's worth of data. Schematically, the measurement storage might be represented as follows:

Elapsed time	69	59	58	. . .	09
Respiration rate	9	9	10	. . .	9
Systolic blood pressure	141	154	153	. . .	150
.					
.					
.					
Clock time	1230 [*current time*]	1220	1219	. . .	1130

Throwing away old measurements does not limit the ability of the program to utilize historical data. The *conclusions* based on the original data, which are stored much more compactly, are maintained throughout the patient run. Thus the numerical measurement values are replaced by symbolic abstractions over time.

One current limitation is the program's inability to reevaluate past conclusions, especially when measurements are taken but are not reported until some time later. One example of this is the interpretation of blood gas measurements. It takes about 20–30 minutes for the laboratory to process blood gas samples, but by that time the context may have changed. The program cannot then back up to the time that the blood gases were taken and proceed forward in time, reevaluating the intervening measurements in light of the new data. The resolution of conflicts between expectations and actual events may also require modification of old conclusions. This is especially true when forthcoming events are used to imply that an alternative cause provides a better explanation of observed phenomena.

22.3.3 Rules

Rules used in VM have a fixed structure. The premise of a rule is constructed from the conjunction or disjunction of a set of clauses. Each clause checks relationships about one or more of the parameters known to the program. Each of these relationships, such as "the respiration rate is between an upper and lower limit," will be tested to determine if the premise is satisfied. If the clauses are combined conjunctively and each clause is true, or combined disjunctively and at least one clause is true, then the rule is said to "succeed." As explained in Section 22.3.6 on uncertainty in VM, no probabilistic weighting scheme is currently used in the rule evaluation (although the mechanism is built into the program).

When a rule succeeds, the action part of the rule is activated. The action portion of each rule is divided into three sections: conclusions (or interpretations), suggestions, and expectations. The only requirement is that at least one statement (of any of the three types) is made in the action part of the rule. The first section of the action of the rule is composed of the conclusions that can be drawn from the premise of the rule. These conclusions (in the form of a parameter assuming a value) are asserted by the program to exist at the current time and are stored away for producing summaries and to provide new facts for use by other rules. When the same conclusion is also asserted in the most recent time when data are available to the program, then the new conclusion is considered a *continuation* of the old one. The time interval associated with the conclusion is then extended to include the current time. This extension presumes that the time period between successive conclusions is short enough that continuity can be asserted.

The second section of the action is a list of suggestions that are printed for the clinician. Each suggestion is a text string to be printed that summarizes the conclusions made by the rule.[5] Often this list of suggestions includes additional factors to check that cannot be verified by the program—e.g., the alertness of the patient. By presenting the suggestions as conditional statements, the need to interact with the user to determine the current situation is minimized. The disadvantage of this method is that the program maintains a more nebulous view of the patient's situation, unless it can be ascertained later that one of the suggestions was carried out.

The last section of the action part of the rule is the generation of new expectations about the ranges of measurements for the future. Expectations are created to help the program interpret future data. For example, when a patient is first moved from assist mode to the T-piece, many parameters can be expected to change drastically because of the stress as well as the altered mode of breathing. When the measurements are taken, then, the program is able to interpret them correctly. New upper and lower bounds are defined for the acceptable range of values for heart rate, for example, for the duration of time specified. The duration might be specified in minutes or in terms of a context (e.g., "while the patient is on the T-piece").

MYCIN does not place any constraints on the types of conclusions made in the action part of the rule, although most rules use the CONCLUDE function in their right-hand sides. For example, MYCIN calls a program to compute the optimal therapy as an action part of a rule (Chapter 6). The basic motivation behind imposing some structure on rules was to act as a mnemonic device during rule acquisition. The same advantage is found in framelike systems with explicit component names—e.g., CAUSED-BY, MUST-HAVE, and TRIGGERS in the Present Illness Program (Szolovits and Pauker, 1978).

---

[5]Not every conclusion has a corresponding suggestion, particularly when the conclusion denotes a "normal" status—e.g., hemodynamic stability.

A rule is represented internally by a property list with a fixed set of properties attached to the name of the rule:

RULEGROUP	Defines type or class of the rule; in this case, the rules that deduce the status of the patient
DEFINITION	Free text that defines the main idea of the rule
COMMENT	The collected comments from the external form of the rule
NODE	All of the contexts for which this rule makes sense (currently limited to the values associated with the patient's therapeutic setting)
EVAL	Specifies the methods of evaluation; ALL OF for conjunction, ONE OF for disjunction, X% for requirement of a fixed percentage of verified premise clauses
ORIGLHS	A copy of the external notation of the premise of the rule, used in explanations and tracing
FILELOCATION	The description of the location on a file of the original text of the rule
M	The translated premise of the rule, a list of calls to premise functions (M stands for match)
I	The list of interpretations (conclusions) to be made
S	The list of suggestions to be printed out
E	The list of expectations to be made

The actual processing of a rule is carried out by a series of functions that test conditions, make interpretations, make suggestions, or create expectations. Each of these functions has a well-defined semantic interpretation and provides the primitives for encoding the knowledge base.

The translation between an external format, e.g., RESPIRATION RATE > 30, and the corresponding internal format, (MCOMP RR > 30), is made by the same parsing program used in EMYCIN.[6] The MCOMP function is given a parameter name (RR), a relation (less than, greater than, or equal to), and a number with which to compare it. The execution of the MCOMP function returns a numerical representation of TRUE, FALSE or UNKNOWN, based on the current value of respiration rate. Figure 22-12 demonstrates the external and internal representations of a typical rule.

22.3.4 Premise Functions

One goal of the VM implementation is to create a simple set of premise functions that are able to test for conditions across time. Many of the static premise functions have been adapted from the MYCIN program; e.g.,

[6]The parsing program was written by James Bennett, based on work by Hendrix (1977).

```
STATUS RULE:  STATUS. STABLE-HEMODYNAMICS
DEFINITION:  Defines stable hemodynamics based
        on blood pressures and heart rate
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/minute
        in 15 minutes
  MEAN ARTERIAL PRESSURE is ACCEPTABLE
  MEAN ARTERIAL PRESSURE does NOT CHANGE by 15
        torr in 15 minutes
  SYSTOLIC BLOOD PRESSURE is ACCEPTABLE
THEN
  The HEMODYNAMICS are STABLE

RULEGROUP:  STATUS-RULE
DEFINITION:  ((DEFINES STABLE HEMODYNAMICS BASED)
              (ON BLOOD PRESSURES AND HEART RATE))
COMMENT:  ((LOOK AT MEAN ARTERIAL PRESSURE FOR)
            (CHANGES IN BLOOD PRESSURE AND SYSTOLIC)
            (BLOOD PRESSURE FOR MAXIMUM PRESSURES))
NODE:  (VOLUME CMV ASSIST T-PIECE)
EVAL:  (ALL OF)
ORIGLHS:  ((HEART RATE IS ACCEPTABLE)
           (PULSE RATE DOES NOT CHANGE BY 20 BEATS/MINUTE IN 15
              MINUTES)
           (MEAN ARTERIAL PRESSURE IS ACCEPTABLE)
           (MEAN ARTERIAL PRESSURE DOES NOT CHANGE BY 15 TORR
              IN 15 MINUTES)
           (SYSTOLIC BLOOD PRESSURE IS ACCEPTABLE))
FILELOCATION:  (<Puff/VM>VM.RULES;18 12538 13143)
M:  ((MSIMP HR ACCEPTABLE NIL)
     (FLUCT PR CHANGE 20 (0.0 15) NOT)
     (MSIMP MAP ACCEPTABLE NIL)
     (FLUCT MAP CHANGE 15 (0.0 15) NOT)
     (MSIMP SYS ACCEPTABLE NIL))
I:   ((INTERP HEMODYNAMICS = STABLE NIL))
```

FIGURE 22-12 External and internal representations for a rule in VM.

MCOMP encompasses the functions of GREATERP, LESSP, numeric EQUAL, and their negations in MYCIN. Most of the functions listed below test the value of a parameter within a time interval and return TRUE, FALSE or UNKNOWN. As mentioned earlier, they reference concepts, such as HIGH value or STABLE value, that are defined by rules at each stage. Each function is composed of the following program steps: (a) find out the value of the parameter mentioned in the time period mentioned (otherwise, use the current time), (b) make the appropriate tests, and (c) negate the answer, if required. Table 22-1 lists the premise functions.

TABLE 22-1

Class	*Function*	*Format*	*Example*	*Action*
Testing measurement values	FLUCT	FLUCT [parameter, direction, amount, time-range, negation]	THE MEAN ARTERIAL PRESSURE DOES NOT CHANGE BY 15 TORR IN 15 MINUTES	Calculates slopes or ranges from the table of measurements; direction can be rises, drops, or changes
	STABILITY	STABILITY [parameter, time-range, negation]	RESPIRATION RATE IS NOT STABLE FOR 20 MINUTES	Calculates the stability of the measurement by comparing the distance between the average and the maximum and minimum values over the time range
	MSIMP	MSIMP [parameter, relation, negation]	RESPIRATION RATE IS NOT HIGH	Compares the current value of a parameter with one of the expectation types; relation can be very low, low, abnormally low, acceptable, high, very high, or abnormally high
	MCOMP	MCOMP [parameter, relation, compared to, negation]	HEART RATE IS NOT >150 HEART RATE IS < THE IDEAL-MINIMUM EXPECTATION LIMIT	Compares the current value of the parameter with a numeric cutoff or a symbolic range; relation can be greater than, less than or equal to; compared-to can be any number, ideal-minimum, ideal-maximum, very low, etc.
	BETWEEN	BETWEEN [parameter, lower-limit, upper-limit, negation]	PULSE RATE IS NOT BETWEEN 40 AND 60 BEATS/MINUTE	Compares the current value of the parameter with a numeric range

Testing derived conclusions	TIMEEXP	TIMEEXP [parameter, value, relation, time-range, negation]	PATIENT IS ON CMV FOR GREATER THAN 30 MINUTES	Checks that the parameter has had a value for the specified period of time; the time range may be in the past, e.g., 20 to 40 minutes ago; relation can be greater than, less than, or equal to
	TRANSITION	TRANSITION [parameter, from-node, to-node]	PATIENT TRANSITIONED FROM ASSIST MODE TO THE T-PIECE	Checks for state changes; parameter is usually PATIENT or GOAL; from-node and to-node are equipment configurations, e.g., ASSIST MODE.
Interpretation	INTERP	INTERP [parameter, value, negation]	THERE IS HYPERVENTILATION	Makes conclusions; updates property list of parameter to show changes
Suggestion	SUGGEST	SUGGEST [arbitrary-text]	SUGGEST PUTTING THE PATIENT ON THE T-PIECE IF PATIENT IS AWAKE	Prints suggestion when rule succeeds
Expectation	EXPECT	EXPECT [parameter, limit-type, limit, expectation-type, time-range]	THE ACCEPTABLE RANGE FOR SYSTOLIC BLOOD PRESSURES IS 110-150 (while the patient is on volume ventilation)	Sets up one of three levels of expectations: for all time, for the duration of the current context, or for a fixed interval; limit-type can be upper limit, lower limit, range, or table; expectation-type can be default, context, or timed
Miscellaneous	EQUATION	EQUATION [variables, equation]	FN(TIDAL VOLUME, WEIGHT): TIDAL VOLUME > 3 * WEIGHT	Evaluates equation and returns Boolean value
	EXECUTE	EXECUTE [rule-group]	EXECUTE STATUS-RULES	Evaluates each of the rules in the named group according to the method specified in Section 22.3.5

22.3.5 Control Structure

A simple control structure is used to apply the knowledge base to the patient data. This method starts by the execution of the goal rule, which in turn evaluates a set of rules corresponding to each level of abstraction in order: first, data validation, followed by context checking and expectation setting, determination of physiological status, and finally, therapeutic response, if necessary. From the group of rules at each level of reasoning, each rule is selected in turn. The current context as determined by the program is compared against the list of applicable contexts for each rule (the NODE property). The premise portions of acceptable rules are examined. If the parameter mentioned in a premise clause has not yet been fully evaluated, an indexing scheme is used to select the rules within this rule set that can make that conclusion. Using this method avoids the necessity of putting the rules in a specific order. The rule is added to a list of "used" rules, and the next unexamined rule is studied. The list of evaluated rules is erased each time the rule set is evaluated. When a rule succeeds, the action part of the rule is used to make interpretations, print suggestions, and set expectations.

Most rules attempt to explain the interpretation of measurements that have "violated" their expectations. Thus, for the portion of the rules that mention an "out-of-range" measurement value in their premise or that are based on the conclusions of these rules, the following strategy could be used: compare all measurements against the current expectations, and forward chain only those measurements with values that require explanation. This method is not useful when the rule specifies that several normal measurements imply a normal situation, e.g., determining hemodynamic stability. These "normal" rules would have to be separated and forward- or backward-chained as appropriate.

22.3.6 Uncertainty in VM

Although the MYCIN certainty factor mechanism (Chapter 11) is incorporated into the VM structure, it has not been used. Most of the representation of uncertainty has been encoded symbolically in the contents of each rule. Rules conclude that measurement values can be spurious (under specified conditions), and the interpreter prohibits using such aberrant values for further inferences. Any value associated with a measured parameter that was concluded too long ago is considered to be unknown and, therefore, no longer useful in the reasoning mechanism. This is meant to be a first approximation to our intuition that confidence in an interpretation decays over time unless it is reinforced by new observations.

Uncertainty has been implicitly incorporated in the VM knowledge base in the formulation of some rules. In order to make conclusions with a higher level of certainty, premise clauses were added to rules that correlated strongly with existing premise clauses—e.g., using both mean and systolic blood pressures. The choice of measurement ranges in several therapy rules also took into account the element of uncertainty. Although the experts wanted four or five parameters within the IDEAL limit prior to suggesting the transition to the next optimal therapy state, they often used the ACCEPTABLE limits. In fact, it would be unlikely that all measurements would simultaneously fall into IDEAL range. Therefore, incorporating these "grey areas" into the definition of the symbolic ranges was appropriate. There are at least two possible explanations for the lack of certainty factors in VM rule base: (1) on the wards, it is only worthwhile to make an inference if one strongly believes and can support the conclusion; and (2) the measurements available from the monitoring system were chosen because of their high correlation with patients' conditions.

As mentioned elsewhere, the PUFF and SACON systems also did not use the certainty factor mechanism. The main goal of these systems was to classify or categorize a small number of conclusions as opposed to making fine distinctions between competing hypotheses. This view of uncertainty is consistent with the intuitions of other researchers in the field (Szolovits and Pauker, 1978, p. 142):

> If possible, a carefully chosen categorical reasoning mechanism which is based on some simple model of the problem domain should be used for decision making. Many such mechanisms may interact in a large diagnostic system, with each being limited to its small subdomain. . . . When the complex problems need to be addressed—which treatment should be selected, how much of the drug should be given, etc.—then causal or probabilistic models are necessary. The essential key to their correct use is that they must be applied in a limited problem domain where their assumptions can be accepted with confidence. Thus, it is the role of categorical methods to discover what the central problem is and to limit it as strongly as possible; only then are probabilistic techniques appropriate for its solution.

22.3.7 Representation of Expectations in VM

Representing expectations about the course of patient measurements is a major design issue in VM. In the ICU situation, most of the expectations are about the typical ranges (bounds) associated with each physiological measurement. Interpreting the relationship between measurement values and their expectations is complicated particularly at the discontinuities caused by setting numeric boundaries. For example, on a scale of possible blood pressure values ranging from 50 to 150, how much difference can there be between measurement values of 119 and 121, in spite of some boundary at 120? However, the practice of setting specific limits and then treating values symbolically (e.g., TOO HIGH) appears to be a common educational and clinical technique. The ill effects on decision making of

Symbolic value	*Interpretation*
IDEAL	The desired level or range of a measurement
ACCEPTABLE	The limits of acceptable values beyond which corrective action is necessary—bounds are high and low (similar for rate)
VERY UNACCEPTABLE	Limit at which data are extremely out of range—e.g., on which the definition of severe hypotension is based
IMPOSSIBLE	Outside the limits that are physiologically possible

FIGURE 22-13 Representing expectations using symbolic bounds.

setting specific limits are minimized by the practice of using multiple measurements in coming to specific conclusions. One alternative to using symbolic ranges would be to express values as a percentage of some predefined norm. This has the same problems as discrete numeric values, however, when the percentage is used to draw conclusions. When it was important clinically to differentiate how much an expectation was exceeded, the notion of alternate ranges (e.g., VERY HIGH) was utilized. For the physiological parameters, several types of bounds on expectations have been established, as shown in Figure 22-13. In VM these limits are not static; they are adapted to the patient situation. Currently, the majority of the expectation changes are associated with changes in ventilator support. These expectations are established on recognition of the changes in therapy and remain in effect until another therapy context is recognized. A more global type of expectation can be specified that persists for the entire time patient data are collected. A third type of expectation type corresponds to a perturbation, or local disturbance in the patient's situation. An example of this is the suctioning maneuver where a vacuum device is put in the patient's airway. This disturbance has a characteristic effect on almost every measurement but only persists for a short period of time, usually 10–15 minutes. After this time, the patient's state is similar to what it was in the period just preceding the suction maneuver. It is possible to build a hierarchy out of these expectation types based on their expected duration; i.e., assume the global expectation unless a specific contextual expectation is set, provided a local perturbation has not recently taken place.

Knowledge about the patient could be used to "customize" the expectation limits for the individual patient. The first possibility is the use of historical information to establish *a priori* expectations based on type of surgery, age, length of time on the heart/lung machine, and presurgical pulmonary and hemodynamic status. The second type of customization could be based on the observation that patient measurements tend to run within tighter bands than the *a priori* expectations. The third type of expectation based on transient events can be used to adjust for the effects of temporary intervention by clinicians. This requires expert knowledge about the side effects of each intervention and about the variation between different classes of patients to these temporary changes.

22.3.8 Summary Reports

Summary reports are also provided at fixed intervals of time, established at the beginning of the program. Summaries include: (1) a description of current conclusions (e.g., PATIENT HYPERVENTILATING FOR 45 MINUTES); (2) a graph with time on one axis (up to six hours) and recent conclusions on the other; and (3) a similar graph with time versus measurements that are beyond the expected limits. (Figure 22-2 shows a portion of a sample summary report.)

The summary report is based on several lists generated by the program. The first list is composed of parameter-value pairs concluded by the program. This list is extended by the INTERP function called from the action portion of rules. The second list includes pairs of measurement types and symbolic ranges (e.g., RESPIRATION RATE—HIGH). This list is augmented during the process of comparing measurement values to expected ranges. These lists are built up from the start of the program and are not reset during new time intervals. The graphs are created by sorting the lists alphabetically, and then collecting the time intervals associated with each parameter-value pair. The conclusion and expectation graphs cover the period from six hours ago until the current time, with a double bar (=) plotted for each ten-minute period that the conclusion was made.

The number of items in each graph is controlled by the number of currently active pathophysiological conclusions subject to a static list of parameters and values that are omitted (for example, some intermediate conclusions are not plotted). When the rule base is extended into other problem areas of ICU data interpretation, new sets of rules may have to be created to select which of the current conclusions should be graphed.

The graph of "violated" expectations presents a concise display of the combination of measurements that are simultaneously out of range. Most of these conclusions have been fed into rules that determine the status of the patient. Patterns that occur often, but fail to trigger rules about the status of the patient, become candidates for the development of new rules.

22.4 Summary and Conclusions

VM uses a simple data-directed interpreter to apply a knowledge base of rules to data about patients in an intensive care unit. These rules are arranged according to a set of levels ranging from measurement validation to therapy planning, and are currently formulated as a categorical system. Interactive facilities exist to examine the evaluation of rules while VM is monitoring data from a patient, and to input additional test results to the system for interpretation.

22.4.1 Representing Knowledge About Dynamic Clinical Settings

In 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. As mentioned in the introduction, MYCIN was designed to produce therapeutic decisions for one critical moment in a 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. One function examines trends in input data over time—e.g., 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—e.g., THE PATIENT HAS TRANSITIONED FROM ASSIST MODE TO THE T-PIECE. When VM is required to check if 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 when 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.

22.4.2 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 stages involved in 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 used 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.

22.4.3 Comparison of MYCIN and VM Design Goals

MYCIN was designed to serve on a hospital ward 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 before very much microbiological data are available. In critical situations, a tentative decision about therapy must often be made on partial information about cultures. In return for assistance, the clinician is asked to provide answers to questions during a consultation.

The intensive care unit is quite different from the static situation addressed by MYCIN, however. Continuous monitoring and evaluation of the patient's status are required. The problem is one of making therapeutic adjustments, many of which are minor, such as adjusting the respiratory rate on the ventilator, over a long period of time. The main reasons for using VM are to monitor status 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, e.g., when an unexpected event is noted, the program must be concise in its warning. VM's environment differs from MYCIN's in that natural language is an unlikely mode of communication.

This difference in the timing and style of the user-machine interaction has considerable impact on system design. For example, the VM system must be able to:

1. reach effective decisions on the presumption that input from a clinician will be brief,
2. use historical data to determine a clinical situation,
3. provide advice at any point during the patient's hospital stay,
4. follow up on the outcomes of previous therapeutic decisions, and
5. summarize conclusions made over time.

A consultation program should also be able to model the changing medical environment so that the program can interpret the available data in context. Areas such as that of infectious disease require an assessment of clinical problems in a variety of changing clinical situations, e.g., "patients who are severely ill but lack culture results," "patients after culture data are available," "patients after partial or complete therapy," or "patients with acquired superinfection."

It is also necessary that VM contain knowledge that can be used to follow a case over a period of time. This is complicated by the fact that the user of the system may not follow the therapy recommended. VM then has to determine what actions were taken and adjust its knowledge of the patient accordingly. Also, 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.

During the implementation of the VM program, we observed many types of clinical behavior that represent a challenge to symbolic modeling. One such behavior is the reluctance of clinicians to change therapies frequently. After a patient meets the criteria for switching from therapy A to therapy B, e.g., assist mode to T-piece, clinicians tend to allow the patient's status to drop below optimal criteria before returning to therapy A. This was represented in the knowledge base by pairs of therapy selection rules (A to B, B to A) with a grey zone between the two criteria. For example, ACCEPTABLE limits might be used to suggest going from therapy A to therapy B, whereas VERY HIGH or VERY LOW limits would be used for going from B to A. If the same limit were used for going in each direction, a small fluctuation of one measurement near a cutoff value would provide very erratic therapy suggestions. A more robust approach makes decisions in such situations based on the length of time a patient has been in a given state and on the patient's previous therapy or therapies.

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 of representing a situation that changes over time: (1) providing the mechanism for accessing and evaluating data in each new time frame, and (2) building a symbolic model to represent the ongoing processes and transitions in the medical environment.