Report 79-31
Stanford -- KSL

Scientific DataLink

Representation of Dynamic Clinical Knowledge: Measurement Interpretation in the Intensive Care Unit. Lawrence M. Fagan, John C. Kunz, et al., Aug 1979

card 1 of 1

Heuristic Programming Project
Report No. HPP 79-31

August 1979

Representation of Dynamic Clinical Knowledge: Measurement Interpretation in the Intensive Care Unit

Lawrence M. Fagan, John C. Kunz, Edward A. Feigenbaum,
Stanford University, Computer Science Department

and

John J. Osborn
The Institutes of Medical Sciences, Pacific Medical Center

Reprinted from the Proceedings of the 6th International Joint Conference on Artificial Intelligence, Tokyo, Japan, August 1979.

Used by permission of the International Joint Conference on Artificial Intelligence, Inc.; copies of the Proceedings are available from Morgan Kaufmann Publishers, Inc., 95 First Street, Los Altos, CA 94022, USA.

REPRESENTATION OF DYNAMIC CLINICAL KNOWLEDGE: MEASUREMENT INTERPRETATION IN THE INTENSIVE CARE UNIT

Lawrence M. Fagan, John C. Kunz, Edward A. Feigenbaum
Heuristic Programming Project, Computer Science Department
Stanford University, Stanford, California, 94305

John J. Osborn
The Institutes of Medical Sciences, Pacific Medical Center
2200 Webster Street, San Francisco, California, 94115

This paper reports work in progress on a program to provide diagnostic and therapeutic suggestions about patients in the Intensive Care Unit (ICU). The Ventilator Manager program (VM) dynamically interprets the clinical significance of quantitative data from the ICU. This data is used to manage post-surgical patients receiving mechanical ventilatory assistance. An extension of a physiological monitoring system, VM (1) provides a summary of the patient's physiological status appropriate for the clinician, (2) recognizes untoward events in the patient/machine system and provides suggestions for corrective action, (3) suggests adjustments to ventilatory therapy based on a long-term assessment of the patient status and therapeutic goals, (4) detects possible measurement errors, and (5) maintains a set of patient-specific expectations and goals for future evaluation. The program produces interpretations of the physiological measurements over time, using a model of the therapeutic procedures in the ICU and clinical knowledge about the diagnostic implications of the data. These therapeutic guidelines are represented by a knowledge base of rules created by clinicians with extensive ICU experience.

1.0 INTRODUCTION

Many of the artificial intelligence approaches to medical decision making have concentrated on providing advice based on data available at one particular point in time. This formulation of the problem treats the diagnostic process as a static situation. In actual practice, the clinician receives additional information from tests and observations over time and must reevaluate hypotheses about the nature of the diagnosis and reevaluate the status and prognosis of the patient. The patient situation is affected by the progression of the disease process and the response to prior therapeutic interventions. Some of these aspects have been captured in current computer medicine work, particularly the digitalis therapy advisor [2]. This system generates conclusions over time pertaining to the size and timing of the next dose of digitalis. Both the MYCIN [4] and CASNET systems [8] allow the user to rerun a consultation as new patient information becomes available. The IRIS system [5] is capable of attaching a time specification to each conclusion created by the system, consisting of a numeric or symbolic time range when the conclusion was applicable.

One attempt to explore these issues is the development of the Ventilator Manager (VM) program, which provides diagnostic and therapeutic advice in the Intensive Care Unit. The input to VM are the values of 30 physiological measurements provided on a 2- or 10-minute basis by a automatic monitoring system [3]. The output is in the form of suggestions to clinicians and periodic summaries.

The clinical situation being modeled is the post-surgical progress of a patient in the Intensive Care Unit, concentrating on the status of his mechanical breathing assistance. The mechanical ventilator provides total or partial assist of ventilation for seriously ill patients. 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 called "controlled mandatory ventilation" or in response to sucking by the patient, known as "assist mode." Adjustments to the respiration rate or output volume of the ventilator are made to insure an adequate "minute volume" to

---

* This research is supported by the NIH under GM-24669, and utilizes the SUMEX-AIM computing resource (NIH RR-00785). We wish to thank Bruce Buchanan, Dianne McClung, Penny Nii, and Josh Rubin for their help on this project.

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 then the tube is removed ("extubation"). Often many of these clinical states must be repeated until the patient can breathe on his own.

2.0 Knowledge representation

The availability of new measurement data requires updated interpretations based on the changing values and trends. As the patient setting changes—e.g., as a patient starts to breathe on his own during removal (weaning) from the ventilator—the same measurement values can lead to different interpretations. In order to properly interpret data collected during changing therapeutic contexts, the knowledge base includes a model of the stages that a patient follows from admission to the unit through the end of the critical monitoring phase. Recognition and utilization of the appropriate patient context is an essential step in determining the meaning of most physiological measurements. The goals for patient management are also stated in terms of these clinical contexts. The program maintains a description of the current and optimal ventilatory therapies for any given time. Figure 1 shows a summary of conclusions generated by the program on a periodic basis (currently one hour).

```
  Current conclusions:
  HYPOTENSION PRESENT for 41 MINUTES
  HYPERVENTILATION PRESENT for 33 MINUTES
  SYSTOLIC BLOOD PRESSURE LOW for 46 MINUTES

Conclusions: {time of day} |.....|.....|.....|.
                           13    14    15    16
  {physiological states}
HEMODYNAMICS - STABLE                 ====
HYPERVENTILATION - PRESENT     =     == == ====
HYPOTENSION - PRESENT            ====  == =====
TACHYCARDIA - PRESENT          ==== ==
  {patient context}
Patient is ASSIST                        =====   ==
Patient is CMV                   ====     ===
Patient is VOLUME              ==
Patient is NOT-MONITORED   ====
  {therapeutic goals}
Goal is CMV                    ================
Goal is VOLUME             ====
                           |.....|.....|.....|.
                           13    14    15    16
```

Figure 1. Summary prepared by VM based on data from a patient in the ICU. Comments are in {}.

Knowledge is represented in VM by production rules [4,1,7]. The rules are of the form:

```
IF      facts about measurements
        and/or previous conclusions are true
THEN
        1) Make a conclusion based on
               these facts;
        2) Make appropriate suggestions to
               clinicians; and
        3) Create new expectations about
               future acceptable ranges
               for measured variables.
```

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

```
STATUS RULE: STABLE-HEMODYNAMICS
DEFINITION: Defines stable hemodynamics for
        most settings
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 2. Sample VM Interpretation Rule.

2.1 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 the physiological measurements remains constant. The use of symbolic statements (e.g., "heart rate is acceptable") allows for a simple representation of common clinical practice and the exposition of common principles of physiological 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 the context-dependent 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 CO2 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 rule that creates these expectations is shown in Figure 3.

IF PATIENT TRANSITIONED FROM ASSIST TO CMV
THEN EXPECT THE FOLLOWING

		[acceptable range]				
	very		[ideal]			very
	low	low	min	max	high	high
MEAN PRESSURE	60	75	80	95	110	120
HEART RATE		60			110	
EXPIRED CO2	22	28	30	35	42	50

Figure 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.

The VM knowledge base includes rules to support the following reasoning steps: (1) characterize measured data as reasonable or spurious; (2) determine therapeutic state of the patient (currently the mode of ventilation); (3) establish expectations of future values of measured variables; (4) check physiological status, including cardiac rate, hemodynamics, ventilation, oxygenation; and (5) check compliance with long-term therapeutic goals.

3.0 RULE INTERPRETATION

The VM rule interpreter is based on the EMYCIN interpreter [6,4]. The major changes are: forward chaining (data-driven) rule invocation as opposed to backward chaining, checking to see that information acquired in a previous time frame is still valid for making conclusions, and the ability to cycle through the rule set each time new information is available.

A data-driven approach is used to take advantage of the small set of measurement values available in each time frame. Because of the nature of the ICU environment, one can assume almost no dialogue will take place with clinicians when they are using the system. Thus, conclusions must be based on the available data. Each of the rule groups corresponding to the reasoning steps mentioned above are considered in order. It is necessary to separate out the reasoning steps since one part of the reasoning chain may conclude that the clinical context has changed, affecting the interpretation of the more "abstract" reasoning steps.

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 over time. 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 conclusions as opposed to looking at the original measurements. Expectations are associated with the appropriate measurement and are classified by type—e.g., the upper limit of the acceptable range—and duration. Expectations can persist for a fixed interval, such as "for twenty minutes starting in ten minutes," or for the duration of one or more clinical contexts.

4.0 REFERENCES

[1] Davis, R. and King, J. "An Overview of Production Systems." in Machine Intelligence 8: Machine Representations of Knowledge, ed. Elcock, E. and Michie, D., John Wylie, 1977.

[2] Gorry, G., Silverman, H. and Pauker, S. "Capturing clinical expertise: a computer program that considers clinical responses to digitalis." Amer. J. Med 64 (1978) pp. 452-460.

[3] Osborn, J., Beaumont, J., Raison, J. and Abbott R. "Computation for Quantitative On-Line Measurement in an Intensive Care Ward," Computers in Biomedical Research 3 1969.

[4] Shortliffe, E. Computer-based medical consultations: Mycin, American Elsevier, New York, 1976.

[5] Trigoboff, M. and Kulikowski, C: "IRIS: A System for the Propagation of of Inferences in a Semantic Net." In Proc. IJCAI-77. MIT, Cambridge, Mass., August, 1977.

[6] van Melle, W. "A Domain-Independent Production Rule System for Consultation Programs." In Proc. IJCAI-79. Tokyo, Japan, August, 1979.

[7] Waterman, D. and Hayes-Roth, F. "An Overview of Pattern-Directed Inference Systems." In Pattern-Directed Inference Systems, ed. Waterman, D. and Hayes-Roth, F. Academic Press, 1978, pp. 3-22.

[8] Weiss, S., Kulikowski, C., and Safir, A. "A Model-Based Method for Computer-Aided Medical Decision-Making." Artif. Intell. 11 (1978).