

## **INFORMATION TO USERS**

The most advanced technology has been used to photograph and reproduce this manuscript from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book. These are also available as one exposure on a standard 35mm slide or as a 17" x 23" black and white photographic print for an additional charge.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

# **U·M·I**

University Microfilms International  
A Bell & Howell Information Company  
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA  
313/761-4700 800/521-0600

**Order Number 9009740**

**An approach for the incorporation of high level knowledge into  
clinical alarm systems**

**Henkind, Steven Joseph, Ph.D.**

**City University of New York, 1989**

**Copyright ©1989 by Henkind, Steven Joseph. All rights reserved.**

**U·M·I**  
300 N. Zeeb Rd.  
Ann Arbor, MI 48106

**AN APPROACH FOR THE INCORPORATION OF HIGH  
LEVEL KNOWLEDGE INTO CLINICAL ALARM  
SYSTEMS**

A

by

***STEVEN J. HENKIND***

**A dissertation submitted to the Graduate Faculty in Biomedical Sciences  
in partial fulfillment of the requirements for the degree of Doctor of  
Philosophy, The City University of New York.**


**1989**

© 1989

**STEVEN JOSEPH HENKIND**

**All Rights Reserved**

This manuscript has been read and accepted by the Graduate Faculty in Biomedical Sciences in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

4/18/89  
-----  
Date Chair of Examining Committee  


4/15/89  
-----  
Date Executive Officer  


Dr. Louis E. Teichholz  
-----

Dr. Malcolm C. Harrison  
-----

Dr. Anthony M. Benis  
-----

Dr. Ernest Davis  
-----

Supervisory Committee

The City University of New York

## Abstract

**AN APPROACH FOR THE INCORPORATION OF HIGH  
LEVEL KNOWLEDGE INTO CLINICAL ALARM  
SYSTEMS**

by

***STEVEN J. HENKIND*****Advisor: Professor Louis E. Teichholz**

Clinical alarm systems have been well accepted in medical practice but suffer from several major limitations; for example, these systems are generally not able to handle patient/disease specificity, temporal changes, multivariable combinations, and dynamic patterns. In order for an alarm system to overcome these limitations, it is necessary that the system utilize high level knowledge.

In this thesis we develop an approach for incorporating such knowledge into alarm systems. A fundamental aspect of our approach is the use of characteristic functions. In particular, we demonstrate that characteristic functions (a notion from fuzzy set theory) can serve as a powerful knowledge representation for knowledge based systems, even if fuzzy set operators are not utilized.

A prototype alarm system was developed using our approach; the system is specifically oriented to the detection of cardiac tamponade in cardiac surgical intensive care unit patients. This system exhibited promising behavior in limited tests.

**This dissertation is dedicated to the memory of  
my beloved grandmother, Sadie Barowsky.**

## Acknowledgments

Special thanks are extended to my advisor, Dr. Louis Teichholz, to the other members of my advisory committee, Dr. Anthony Benis, Dr. Malcolm Harrison, and Dr. Ernest Davis, and to Dr. Ronald Yager.

I would also like to extend my gratitude to Dr. Patrick Eggena, Dr. Robert Hummel, Dr. Roy Jurado, Dr. Lawrence Manevitz, Dr. Naomi Sager, and Dr. Herman Wyssbrod.

Finally, I would like to thank Barbara E. de Vries for her endless encouragement and love.

**Table of Contents**

<b>1. The Problem</b>	<b>1</b>
<b>2. Previous Work</b>	<b>10</b>
<b>3. Our Approach</b>	<b>35</b>
<b>4. The Program</b>	<b>56</b>
<b>5. Results</b>	<b>70</b>
<b>6. Discussion</b>	<b>81</b>
<b>7. Comments and Further Work</b>	<b>90</b>
<b>Sample Code</b>	<b>95</b>
<b>Details of Results</b>	<b>107</b>
<b>Appendix I</b>	<b>129</b>
<b>Appendix II</b>	<b>178</b>
<b>Appendix III</b>	<b>190</b>
<b>Appendix IV</b>	<b>205</b>
<b>Bibliography</b>	<b>219</b>

**CHAPTER 1****THE PROBLEM****1.1. The Cardiac Surgical Intensive Care Unit**

Cardiac surgery is relatively common. For example, approximately 200,000 coronary artery bypass graft (CABG) procedures are performed annually in the United States [Soch87].

The most common type of cardiac surgical procedure is the CABG. Also frequently performed are the replacement or reconstruction of defective mitral and aortic valves. Less frequently performed are the repair of other valves (e.g., tricuspid), the correction of congenital cardiac defects, the resection of aneurysms, and heart transplants. Recent experimental work has concentrated on the implantation of artificial hearts.

Cardiac surgery is quite complex: it requires large surgical teams and sophisticated instruments (e.g., heart-lung machines). A cardiac surgical team is generally composed of several surgeons, an anesthesiologist, surgical nurses, and technicians. The level of training among these individuals is high; for example, the typical cardiac surgeon has had a minimum of seven years of training beyond medical school.

The management of cardiac patients after surgery is also quite involved. Since these patients require complex care, they are generally placed into a specialized cardiac surgical intensive care unit (CSICU). These units are staffed by highly trained nurses and physicians. Typically, a nurse is assigned

to one, or at most two, patients (in a less intensive setting a nurse may be responsible for ten or more patients), while physicians circulate through the unit.

Much of the patient management is fairly routine; e.g., the changing of dressings, the administration of drugs, and so forth. The goal of this routine management is to stabilize the patient so that he<sup>1</sup> may be discharged from the unit.

The most specialized expertise of the unit personnel lies in their ability to handle the nonroutine: in particular, their ability to recognize and treat emergencies. Crises in the unit can occur suddenly, are often life-threatening, and must be promptly recognized and treated if the patient is to survive.

While patient management can be quite involved, it should be noted that "patients tend to follow particular postoperative courses, each one typically associated with the specific cardiac operation that has been performed" [Behr80]. These postoperative courses can be characterized in terms of specific physiologic progressions. For example, in CABG patients, the mean arterial pressure often becomes quite high during the first twelve hours or so (unless regulated by sodium nitroprusside), and then returns to a more normal range, whereas in patients who have undergone valve surgery, postoperative hypertension is much less common [Esta80].

Although patients can develop virtually any clinical problem after surgery, there are only a few complications which are commonly seen during the postoperative period. Among the complications that may be observed within a day or two after surgery are myocardial failure (low output syn-

drome), and cardiac tamponade [Frat75]. Other possible complications include pulmonary embolus and hypovolemia.

Some physicians believe that it is possible to detect clinical problems in postoperative patients by tracking the values of appropriate physiologic variables. Halpern and Levine [Halp85 pg. 145], for example, state that many of the typical clinical complications can be detected "... strictly by the numbers."

In order to provide the information necessary to manage patients, a very large amount of data is acquired in the CSICU. Vital signs, e.g., body temperature, are taken frequently, and the urine output, chest tube drainage (bleeding from the surgical incision), and other variables are measured periodically. Various invasive and noninvasive transducers, e.g., right atrial catheters, EKGs, are also used to gather data.

## **1.2. Computers in the Cardiac Surgical Intensive Care Unit**

In some CSICUs computers are used to collect data directly from transducers [Malo68]. Typically, the transducer signal is averaged or sampled over a preset time interval, and the data saved on tape or disk. The data may be recalled as desired, and hard-copy summaries may be produced for use by the staff and for inclusion in the patient's chart.

One of the first systems to automatically acquire physiologic data from patients was designed by Weil and his associates [Weil66] in the early 1960's. This system gathered cardiovascular and other data from patients in a specialized intensive care unit for the treatment of shock.

In 1968 Osborn et al [Osbo68] described a system that gathers cardiovascular and pulmonary data from patients in a cardiac surgical intensive care unit. The system currently in use at the Mount Sinai Hospital is based on this design.

Since the recording of data by hand is quite time consuming, automated data acquisition can lead to significant savings in time for the clinical staff [Shep79], hence increasing the time available for other aspects of patient care. Cullen and Teplick [Cull79], however, have claimed that automated data acquisition does not yield any significant time savings.

Computers can also play a somewhat more active role in patient management. For example, they may be used to regulate fluid or drug infusions based on physiologic measurements. Typically, such systems function as feedback loops. Sheppard et al [Shep68] describes a system that automatically regulates the infusion of blood in cardiac surgical patients. Sheppard [Shep80] and DeAsla et al [DeAs85] describe systems that control blood pressure by regulating the infusion of sodium nitroprusside. Linkens [Link86] provides a survey of some of the various on-line control systems that have been developed.

Probably the most common use for computers in the CSICU is the monitoring of patients. In particular, computers are used to generate alarms to alert the staff to emergencies in the patient's condition.

The motivation for such systems is straightforward: many clinical complications need to be caught early in order that they be effectively treatable. For example, Shoemaker et al [Shoe73] reports a mortality rate of 14.6% in tamponade cases (due to stab wounds) that are detected "early," and a mor-

tality rate of 60% for cases that are detected "late."

### 1.3. Conventional Computer-Based Alarm Systems

The basic paradigm that is used by most clinical alarm systems is as follows (Fig. 1.1): Each physiologic variable is viewed as running along a numeric scale. On this scale are two distinguished points—*Low* and *High*. These distinguished points are used as cutoff points (alarm limits): If the measured value of the variable falls out of the range between *Low* and *High*, then the value is considered "pathologic," and an alarm signal is activated. If the value falls between *Low* and *High*, then it is considered "normal," and no action is taken. The alarm limits are either preset in the alarm program, or entered by the physician.




Fig. 1.1. Conventional alarm paradigm.

Perhaps the earliest automated alarm system for the CSICU was the one developed by Osborn and his associates [Osbo68] in the late 1960's. It was based on the conventional alarm paradigm: in addition to acquiring and storing data, this system allowed the physician to set low and high alarm limits.

The conventional alarm paradigm is sometimes modified slightly. For example, in some systems a given variable may have to be out of limits for a minimum period of time before an alarm is triggered [Beni80]. In other systems there may be only one limit rather than two.

There have also been somewhat more substantial modifications to the alarm paradigm. For example, Warner [Warn69] developed a system which screens data by statistical methods and signals only if there has been a "statistically significant" change since the last data were recorded.

Some alarm systems are not based on the conventional paradigm. For example, monitors to detect cardiac arrhythmias are based on pattern recognition algorithms. Some alarm systems combine aspects of both the conventional paradigm and pattern recognition. For example, a system to detect an excessive number of premature ventricular contractions (*PVCs*) will use pattern recognition techniques to identify the *PVCs*, and a numerical threshold to determine whether the number that has been observed is excessive.

#### **1.4. Limitations of Conventional Alarm Systems**

There is little doubt that alarm systems based on the conventional paradigm have been of great value. Many systems have been built using the paradigm, and they are quite well accepted. Nevertheless, there are several limitations inherent in conventional alarm systems:

##### **(1) Patient/Disease Specificity:**

An alarm system may behave appropriately for some patients, but not for others. This limitation stems from the fact that expected physiologic values are highly dependent on the nature of the patient and his disease. For example, a patient with left ventricular hypertrophy is expected to have significantly higher left atrial pressures than a patient without

hypertrophy. If *Low* and *High* are preset to certain values, or set by hand to inappropriate values, then a system may generate false alarms or miss pathologic readings.

**(2) Temporal Changes:**

For a given patient, an alarm system may behave appropriately at one point in time, yet behave inappropriately at another point. This is due to the fact that the expected values of some physiologic variables change over time. For example, chest tube drainage is expected to be quite high immediately after surgery, but low several hours later. In a conventional alarm system the values of *Low* and *High* (whether preset, or set by hand) are fixed over time; hence, errors can result.

**(3) Multivariable Combinations:**

Conventional alarm systems generally do not take account of multivariable combinations. For example, suppose that the left and right atrial pressures are both moderately elevated. Then if the pressures are nearly equal (or equal) the patient may be in grave danger of impending tamponade, whereas if the pressures are not close to being equal there is less grounds for concern. A conventional alarm system does not handle these types of situations because each variable is viewed in isolation from the others.

**(4) Dynamic Patterns:**

Conventional alarm systems generally do not take account of the relative

ordering of events. Such orderings can be extremely significant. For example, if the cardiac index decreases, and then the left atrial pressure increases, the patient is probably in no danger; but if the atrial pressure increases, and then the cardiac index decreases, the patient may be in severe danger.

In actual practice the good clinician (nurse or doctor) is aware of the limitations of conventional alarm systems. By periodically adjusting alarm ranges, integrating information from multiple sources, and so forth, he or she is able to circumvent these limitations. Unfortunately, however, clinical personnel are not always fully experienced. Moreover, even the most highly trained individuals sometimes forget or make mistakes. Therefore, an alarm system that is not subject to the aforementioned limitations would seem to be a valuable supplement for patient care. *The fundamental hypothesis of this thesis is that it is possible to overcome the limitations of conventional computer-based alarm systems.*

There has been only a limited amount of previous work in this direction. Raison et al [Rais68] describes a system that bases its alarm criteria on the values of heart rate and arterial pressure. In this system both variables must be out of range in order for an alarm to be triggered. Another two variable alarm system has been constructed by Naghdy et al [Nagh84]. This system, which was specifically designed for a cardiac surgical intensive care unit, computes an aggregate score based on skin temperature and blood pressure. The VM system [Faga79, Faga84Com, Faga84Ext] uses a frame based approach to monitor patients being mechanically ventilated. There has also been work on a system that incorporates information from multiple sources

to assist in the monitoring and management of patients with cardiac arrhythmias [Long83Front].

In order to implement alarm systems based on the conventional paradigm, only a limited amount of information is needed. In particular, such a system needs only "know" the low and high limits. In order to overcome the limitations of conventional systems, however, much more information is required. For example, the patient's medical history is needed in order to handle patient/disease specificity, and the expected changes in physiologic variables are needed to handle temporal changes.

Information such as a medical history, expected physiologic changes, and so forth is generally referred to as *high level knowledge* (*knowledge*, for short) by computer scientists. The two types of knowledge that are of particular concern in this work are: (1) patient specific knowledge (e.g., the patient's medical history), and (2) domain knowledge (e.g., expected temporal changes of physiologic variables). The field of computer science that is specifically concerned with the handling and manipulation of knowledge is known as *artificial intelligence*.

#### Footnotes

- 1 In order to avoid excessive verbosity *he* is used instead of *he or she* throughout this manuscript.

**CHAPTER 2****PREVIOUS WORK****2.1. Artificial Intelligence**

Artificial intelligence (AI) [Barr81, Nils80, Wins84] has been an active field of research since the late 1950's. Although it is generally considered to be a field of computer science, it also encompasses aspects of linguistics, mathematics, philosophy, and psychology. A possible definition of artificial intelligence is the following: it is the study of how to get computers to do things that would normally be thought of as requiring reasoning ability or intelligence. At present, there is no universally accepted definition of artificial intelligence.

The following are two typical problems that a computer might be programmed to solve:

Problem 1) compute the sum  $921+7842$ ;

Problem 2) solve the integral  $\int \frac{dx}{x^2\sqrt{4+x^2}}$ .

The first problem is not considered to be an artificial intelligence type problem; the second is. Other problems that are the concern of artificial intelligence include the programming of a computer to understand visual or natural language inputs, the design of a program that is able to learn from its environment, and the development of expert systems (to be described).

Artificial intelligence differs from "mainstream" computer science in at least four respects.

First, the emphasis is on symbols rather than numbers. This is illustrated by the examples above.

Second, the emphasis is on reasoning rather than calculation. This, too, is illustrated by the examples above.

Third, there is a marked reliance on heuristics as well as algorithms. Heuristics can be thought of as rules of thumb. The correctness of a heuristic depends on the context in which it is used. For example, the integral

$\int \frac{dx}{x^2 \sqrt{4+x^2}}$  can be solved by making the substitution  $x = 2 \tan z$ . This is a

heuristic; for example, the substitution would fail on the integral

$\int \frac{dx \sqrt{x^3-x}}{x^2 \sqrt{4+x^2}}$ . On the other hand, the arithmetic problem  $921+7842$  can be

solved by invoking the well-known addition algorithm. Given *any* two integers this algorithm will produce the correct answer; hence, it is not a heuristic.

Finally, there is an emphasis on the use of knowledge rather than data. This will become more apparent throughout the remainder of this thesis.

Why are artificial intelligence techniques needed? In particular, since heuristics only sometimes work, whereas (correct) algorithms always work, why use heuristics?

First of all, for many problems there are not (and theoretically cannot be) general algorithms. For example, there is no general algorithm to solve an arbitrary indefinite integral (more precisely, there is an algorithm which will compute the closed form solution, if it exists, but the algorithm will not

halt if there is no solution). Therefore, mathematicians rely on clever substitutions and other heuristics when doing symbolic integration. Similarly, since there is no general algorithm for medicine, clinicians rely heavily on heuristics. For example, physicians make use of heuristics such as "give penicillin for gonococcal infections." This is a heuristic because some patients may be allergic to penicillin, some bacterial strains may be resistant, and so forth.

Even in cases where general algorithms exist, heuristics may still be needed. For example, there is a well-known algorithm to win at chess—simply compute all possible games and then play only those moves which lead to victory. The problem with this conceptually simple algorithm is that there are approximately  $10^{120}$  potential chess games. Even the fastest computers would need to run for the lifetimes of many universes to compute this number of games; thus, the algorithm is of little value. By the use of heuristics it is possible to very often, but not necessarily always, play a winning game.

Algorithms also generally perform poorly in the face of incomplete or uncertain information. Because of their very nature, heuristics often work better in the face of incompleteness or uncertainty. Note that clinical medicine is one activity where incomplete and uncertain information is the norm rather than the exception.

## 2.2. Expert Systems

Expert systems [Duda83, Haye83, Stef82] are AI systems that are

designed to handle real-world problems and exhibit performance that is as good, or better, than human experts. While this may sound highly ambitious, there are systems extant today that do, in fact, outperform human experts.

MACSYMA [Mose75], for example, does symbolic integration, solves highly complex equations, and performs other sophisticated mathematical operations. This system, which was developed at The Massachusetts Institute of Technology, is in wide use by mathematicians and scientific researchers.

DENDRAL [Linds80], which was developed at Stanford University, interprets mass spectographs and deduces molecular structures. It is the basis for numerous articles in the chemical literature.

RI (now called XCON) [McDe82] was developed by researchers from Carnegie Mellon University and the Digital Computer Corporation. It has been used to configure VAX computer orders for customers.

A great deal of work has gone into designing medical expert systems. Some of the more well known ones will be discussed in the next section.

Among the techniques that may be used to construct an expert system are *frames*, *semantic networks*, and the *logic programming*. The most commonly used technique, however, is the *production system*.

A typical *production system* consists of three components—a *rule base*, a *working memory*, and a *control structure* (Fig. 2.1.).

The *rule base* contains the knowledge that the system possesses. This knowledge is represented as a set of production rules. These rules are of the form "IF antecedent THEN consequent:"

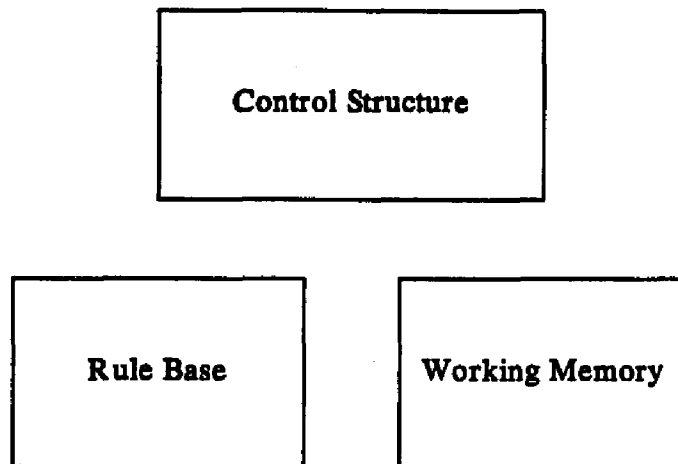


Fig. 2.1. Typical Production System

Example 1) IF  $x$  is a bacterium THEN  $x$  has no nuclear membrane;

Example 2) IF  $y$  is a medical student THEN  $y$  is overworked (.9).

In the examples above,  $x$  and  $y$  are symbols, and the (.9) is a reflection of uncertainty (i.e., only 9 out of every 10 medical students are overworked).

The *working memory* can be viewed as being a scratchpad where calculations take place, input and output are stored, and so forth.

The *control structure* interacts with the rule base and working memory and is responsible for the execution of the production system. The control structure is usually implemented by means of a *recognize-act* cycle:

*Typical recognize-act cycle:*

- a) Determine which production rules are currently applicable. If no rules are applicable then halt.

- b) If more than one rule is applicable, select which one to use ("conflict resolution").
- c) Use the selected rule ("fire the rule").
- d) Goto a.

The following example traces a production system through several iterations of its recognize-act cycle:

*Sample execution of a production system:*

Initially:

Given that the rule base contains two rules:

P1: IF  $x$  has `pulmonic_stenosis` THEN  $x$  has `r_v_hypertrophy`  
P2: IF  $y$  has `r_v_hypertrophy` THEN  $y$  has `high_RAP`

and that working memory has one fact:

F1: Don has `pulmonic_stenosis`

Iteration 1:

Fact F1 in working memory matches rule P1.  
Therefore, P1 fires and produces a new fact.  
Working memory now contains two facts:

F1: Don has `pulmonic_stenosis`  
F2: Don has `r_v_hypertrophy`

Iteration 2:

The new fact matches rule P2.  
Therefore P2 fires and produces another fact.  
Working memory now contains three facts:

F1: Don has *pulmonic\_stenosis*  
F2: Don has *r\_v\_hypertrophy*  
F3: Don has *high\_RAP*

The first step of the recognize-act cycle is to determine which rules are currently applicable. This determination may be *data driven*, *hypothesis driven*, or *goal driven*.

In a *data driven* approach, rules are applicable when they match a new piece of data in working memory. This approach allows the system to be continually on the lookout for crucial data, and therefore responsive to its environment. The example above was data driven.

In a *hypothesis driven* (also called *model driven*) approach, rules are applicable when they can confirm or deny a current hypothesis of the system. By maintaining and revising hypotheses, the system can focus on what it considers to be of maximal importance. In a sense, the model driven approach uses incoming data for the verification of current hypotheses; thus, only data that is relevant to these hypotheses need be examined, and therefore some incoming data can be safely ignored (potentially useful if there is a large volume of incoming data).

In the *goal driven* approach, rules are applicable when they can help to satisfy a goal or subgoal of the system. This approach can be used to provide exhaustive reasoning, as well as some degree of system focus.

Other methods of determining rule applicability include the use of *meta-rules*, and the imposition of a structure on the rule base. Meta-rules [Davi80] are rules that describe the content of other rules. They can allow the system to "know what it knows," and therefore make a sophisticated determination of rule applicability. Various structures may be imposed on the rule base: for example, the rules may be grouped together into sets where each set is suitable for handling a particular type of subproblem. Another possible structure is to group the rules in a hierarchy [Saue83].

It should be noted that it is straightforward to implement an alarm system by means of production rules; e.g.,

IF variable out-of-range THEN issue alert.

### 2.3. Artificial Intelligence in Medicine

Since the early 1970's there has been a great deal of work applying artificial intelligence techniques to medicine [Clan84Read, Kuli80, Scho78, Shor79, Szol78, Szol82]. The emphasis has been on the construction of expert systems for diagnosis<sup>1</sup>, although a few systems have been developed to assist in therapeutics or patient monitoring. Five systems from the 1970's stand out as being particularly important, and have had a significant impact on later research on artificial intelligence in medicine (*AIM*)—MYCIN, CASNET, PIP, INTERNIST and The Digitalis Advisor.

### 2.3.1. MYCIN

MYCIN [Shor73, Shor76, Buch84, Cend84] was developed at Stanford University. Its original goal was to assist in the diagnosis and management of bacteremia (bacterial infections of the blood), although it was later adapted for meningitis as well as various non-medical problems. MYCIN is written in LISP.

Knowledge in the system is represented mainly in the form of production rules<sup>2</sup> [Davi77AI]. These rules are constructed from triples of the form (attribute, object, value).

MYCIN works in an interactive fashion, asking questions of the user as it makes its deductions. Inferences are made by a backward chaining production system<sup>3</sup> which produces a depth first search<sup>4</sup> of an and/or goal tree.<sup>5</sup> The mechanism is as follows: "if a condition in the IF (antecedent) part of a rule is decidable from the database, that is done; if the condition can be asserted by the THEN (consequent) part of some other rules, then they are applied; otherwise MYCIN asks the user" [Szol78]. *Certainty factors* are used to handle uncertain information [Shor75Amod], meta-rules can help guide the search and a context hierarchy is used to define and relate the various objects used during a given consultation session.

One of the prime considerations during the construction of MYCIN was that it should have a good user interface: the system has facilities for explanation of its reasoning, it can converse in a simple subset of English, and it can "learn" new rules by interacting with experts. The learning facility is provided by the TEIRESIAS [Davi77Inte, Shor75Comp] system, which was developed as an extension of MYCIN.

Davis et al [Davi77Prod] states that "experts approved of MYCIN's therapy recommendations in 72% of the evaluations." In a study of the MYCIN adaptation for meningitis, its "recommendations were uniformly judged preferable, or equivalent to those of five infectious disease experts who recommended therapy for the same patients" [Duda83]. See also [Yu84].

EMYCIN is an expert system shell that was created by abstracting out the control structure from MYCIN and incorporating facilities to handle antecedent rules [vanM79]. MYCIN itself was later re-implemented in EMYCIN [Buch84 pg. 298, Cend84 pg. 260]. PUFF [Aiki83], a system which evaluates pulmonary function data, was created using EMYCIN.

Other notable descendants of MYCIN are NEOMYCIN [Clan81], which is designed for teaching students, and ONCOCIN [Shor81], which is an assistant for oncology protocol management. Reference [Buch84] contains a thorough description of MYCIN and many of its descendants.

### 2.3.2. CASNET

CASNET [Weis77, Weis78Amod, Weis78Glau] was developed by researchers from Rutgers University and the Mount Sinai School of Medicine. Its domain of discourse is the diagnosis, prognosis, and treatment of glaucoma. It is written in FORTRAN and SNOBOL.

CASNET structures its information as a *casual associational network*, which is, in essence, a type of semantic network. Signs, symptoms, and test results are connected to pathophysiologic states by *associational links*. *Casual*

*links* connect various pathophysiologic states and serve as a model of the pathophysiology of the disease. Groups of pathophysiologic states are in turn linked to diseases by *classification links*.

Diagnostic inferences are made in three phases. In the first phase, initial data (signs, symptoms, and test results) are entered by the user. The system then examines the associational links and instantiates an initial configuration of pathophysiologic states, with all states either *confirmed*, *denied*, or *unknown*. This phase is data driven.

In the second phase, the casual links are used to guide inferencing and the gathering of additional information. The mechanism is as follows: Along any given casual path that does not contain a denied node, there may be nodes (*candidate nodes*) whose status is currently unknown, but which appear to be likely because of causal connections between themselves and confirmed nodes. These candidate nodes then become the target of further investigation: information is requested from the user which helps clarify the status of these nodes. The end result of the second phase is that a set of pathophysiologic states is identified which accounts for as many signs, symptoms and test results, as possible. This phase is primarily hypothesis driven.

In the third phase, the classification links are examined and they serve to provide a set of likely diseases.

Uncertainty is handled through the various phases by propagating *degrees of strength*. Therapy is determined by links between diseases and various therapy protocols.

CASNET is notable in that it has an explicit pathophysiologic model. It also has an ability (albeit limited) to handle time-varying data: it can

incorporate information from previous visits, and the pathophysiologic model itself is in a sense a description of the (idealized) time course of the disease.

Estimates of CASNET's accuracy range from 75% to 95% [Weis77]. This high level of performance is probably due to the fact that the domain of discourse is so limited—only glaucoma is dealt with—as well as the fact that glaucoma is quite well understood.

The EXPERT [Weis79] system, which is an expert system shell, was influenced by CASNET, but differs from it in several respects. It assists the user in creating a set of production rules (remember that CASNET represents knowledge as a network), and the inferencing mechanism is based partially upon a predefined ordering of production rules, and partially upon a hypothesis driven mechanism (although the consultations are primarily event driven). IRIS [Trig77] generalized the semantic net formalisms developed in CASNET.

### 2.3.3. PIP

PIP [Pauk76] (Present Illness Program) was developed by researchers from The Massachusetts Institute of Technology and and Tufts Medical Center. It deals with the diagnosis of patients presenting with edema, and is especially oriented to renal diseases. The system is written in LISP.

Knowledge in PIP is represented as frames about specific clinical entities; e.g., disease states, and pathophysiologic states. Each frame contains various objects including prototypical findings and *triggers*. The frames are linked into a network by various links representing causality, exclusion, and

so forth. *Daemons* are associated with each frame—they correspond to trigger findings (i.e., they are on the lookout for findings suggestive of their respective frames).

Initially, all frames are outside of working memory and considered to be *dormant*, and a number of the daemons are in working memory. To begin the consultation session, the user inputs some data which are placed into working memory. If a piece of data matches a daemon then the associated frame is moved into working memory and becomes *activated*. This process is data driven.

When a frame is activated, related frames (e.g., connected by a causal link) are pulled closer to working memory—they are made *semiactive*. This is accomplished by moving their daemons into working memory, so that they are then on the lookout for their trigger findings. The rationale for this process is that, clinically, certain entities are suggestive of other entities; for example, if it is suspected that a patient has recently had streptococcal pharyngitis, then acute glomerulonephritis should also be suspected.

*Binding scores* and *matching scores* are used to give a measure of fit between the current data about the patient, and active frames. If the fit exceeds a certain threshold then that active frame is *accepted* and the corresponding disease is considered to be present in the patient. On the other hand, if the scores are too low, an active frame may be moved to a semiactive state. The highest scoring active frame is used to guide further questioning and reasoning (e.g., look for prototypical findings of the disease). This process is hypothesis directed. After each new piece of information is input, scores are then re-evaluated, and processing continues. Thus, the program

alternates between being data driven and hypothesis driven until it has explained all reported findings.

We did not find any evaluations of PIP's performance in the literature that we reviewed.

#### 2.3.4. INTERNIST

INTERNIST [Mill82, Popl73, Popl75, Popl77, Popl82] was developed at the University of Pittsburgh. Its intended domain is vast—the diagnosis of diseases in all of internal medicine. The system has undergone many transitions over the years: versions include DIALOG, INTERNIST-I, INTERNIST-II, and most recently CADUCEUS. The discussion that follows pertains most closely to INTERNIST-I, however, much of the discussion is valid for later systems as well. INTERNIST is written mainly in LISP.

In most AIM systems, the domain of discourse is quite limited; e.g., bacteremia, glaucoma, renal disease. Therefore, these systems have only a limited hypothesis space to work within. On the other hand, internal medicine is a vast field, and the number of possible diseases is on the order of at least  $10^3$ . Therefore, the formation of a hypothesis set is one of the key concerns of INTERNIST. In most other AIM systems, the task is essentially to solve a known differential diagnosis; INTERNIST, on the other hand, must first formulate a differential diagnosis, and then solve it.

The data base in INTERNIST is quite large and is structured as follows:

- (1) Each disease has a *manifestation list* of signs and symptoms. Paired with each manifestation in the list is an estimate of how frequently the manifesta-

tion is seen in the disease (e.g., tachycardia is seen in 80% of hyperthyroid cases). (2) Each manifestation has an *evokes list* which lists the diseases in which the manifestation may be noted. Paired with each disease in the list is an estimate of how strongly suggestive the manifestation is of the disease (e.g., given tachycardia there is a 12% chance that the patient is hyperthyroid). The disease and manifestation lists together resemble a fully inverted file, but with the additional numerical estimates.

The data base also contains causal links between diseases, and a disease hierarchy. The hierarchy is not strict. For example, infectious mononucleosis is included under both liver disease and diffuse lymphadenopathy.

The program begins execution by accepting a set of initial data items from the user. Then, by using the evokes and manifestations lists, various weighting schemes, and "a partitioning algorithm that mimics the process of multiple set intersection," [Popl82 pg. 188] it establishes a hypothesis problem set—i.e., a set of diseases to consider. This process of problem set formation is data driven.

INTERNIST then uses its knowledge base, and various heuristic strategies, to guide further questioning in order to refine the problem set. For example, INTERNIST might attempt to differentiate between two diseases in the problem set by asking about a symptom that is known to be associated with only one of the diseases. This process of refinement is hypothesis driven. Each time an additional piece of information is acquired the system recomputes the problem set; thus, the system alternates between data driven control and hypothesis driven control until it accounts for all pathologic findings. During this processing it may diagnose one or more diseases.

INTERNIST has facilities for handling uncertain data. It is also worth noting that the control strategies of PIP and INTERNIST bear some resemblance to one another, but while PIP bases its questioning on the most likely hypothesis at any given time, INTERNIST bases its questioning on a set of current hypotheses.

INTERNIST currently handles over 500 diseases. In a performance study its behavior "appeared qualitatively similar to that of the hospital clinicians" [Mill82]. Shortcomings include an inability to filter out relatively trivial findings, and a lack of user friendliness.

#### **2.3.5. The Digitalis Advisor**

The Digitalis Advisor [Gorr78, Swar77] was developed by researchers from The Massachusetts Institute of Technology and and Tufts University. It serves as an advisor for the administration of digitalis therapy. The Advisor is written in OWL (which is a language implemented in LISP).

Knowledge is represented by a procedural hierarchy. For example, a top level procedure is BEGIN-THERAPY. This procedure in turn can call CHECK-SENSITIVITIES which may then call CHECK-SENSITIVITIES-DUE-TO-POTASSIUM and so forth. The system also contains a detailed model of the pharmacokinetics of digitalis.

In order to give therapeutic advice, the Advisor forms a patient specific model. The model incorporates both general knowledge about digitalis, and specific knowledge about the patient (obtained during an initial consultation session). This model is then used to generate initial therapy advice.

One of the key features of the Advisor is that it is designed to use feedback information on the patient's response to therapy: "The program classifies the patient's therapeutic response into three states: *none*, *partial*, and *complete*; it classifies his toxic response into three other states: *none*, *partial*, and *definite*. Thus, there are nine possible 'therapeutic-toxic' states, and the program associates a particular set of actions with each of them. It makes new therapeutic recommendations by first determining the patient's current state and then suggesting steps to move him to a 'better' state" [Gorr78]. Thus, the system asks questions about the patient's clinical response, and then uses this information to generate new therapy recommendations. It also updates the patient specific model in response to the new information.

It is notable that the Digitalis Advisor has some ability (albeit limited) to handle temporal information: the patient specific model evolves over a series of consultations, and the system uses feedback information on the patient's response. In limited trials the program appeared to perform quite well [Gorr78, Long80].

### 2.3.6. Second Generation AIM

MYCIN, CASNET, PIP, INTERNIST, and the Digitalis Advisor are considered to be *first generation* AIM systems. One characteristic of these systems is that the bulk of their reasoning is performed with "shallow" knowledge; e.g., symptom-disease correlations, rather than causal or other forms of "deeper" knowledge (this is, perhaps, less true of CASNET than the other four systems).

Much current (*second generation*) AIM research is aimed at constructing systems that reason at deeper levels. For example, there is work on constructing systems that do pathophysiologic reasoning [Long83MED, Pati81]. These systems are still in the prototypical stage, and hence have not been extensively tested.

Another focus of current AIM research is the design of appropriate knowledge representations for complex medical information. For example, the MDX system [Gome81] models medical knowledge as a "collection of essentially decoupled conceptual structures, each with an embedded problem-solving mechanism."

### 2.3.7. AIM Monitoring Systems

There has been a limited amount of work applying AI techniques to monitoring. For example, Gallanti et al [Gall85] report on an AI based system for monitoring a thermal power plant, and Scarl et al [Scar85] report on an AI system for monitoring oxygen loading in the space shuttle. At least two clinical monitoring systems have been constructed using AI techniques: BABY, and VM.

The BABY system [Mich83, Rode84] is designed to monitor infants in an intensive care unit; we were not able to gather any additional information about this system.

The VM system [Faga79, Faga84Comp, Faga84Exte] was developed at Stanford University. Its aim is to assist in the management of post cardiac surgical patients who are receiving mechanical ventilatory assistance. It

receives values of approximately thirty physiologic variables every two or ten minutes (these values are supplied by Osborn's automated monitoring system [Osbo68]). VM periodically summarizes the patient's status and provides advice about therapy.

Knowledge is represented in two ways—as a state network and as production rules. The state network consists of four possible ventilation modes ("therapeutic states")—*volume*, *assist*, *CMV*, and *t-piece*. Each of these four states contains information about expected physiologic values and so forth (the states can be thought of as frames). Production rules are of four types: *transition rules*, *initialization rules*, *status rules*, and *therapy rules*. Transition rules govern when the system makes a transition between states. Initialization rules define what the system "expects" in a newly entered state (it is not clear from the references whether these expectations can be customized for the individual patient). Status rules summarize the patient's status, and therapy rules provide advice.

VM's inference mechanism is based on the EMYCIN approach, but it is data driven rather than goal driven. VM also has a strong hypothesis driven aspect: data (and production rules) are interpreted in the light of the patient's current status.

VM treats the value of certain variables in a symbolic fashion; e.g., the exact meaning of *heartrate acceptable* depends on the therapeutic state. VM has a limited ability to handle temporal data; e.g., transition rules can mention temporal intervals, and expectations can have associated time intervals. However, "VM's reasoning about time is limited to adjacent time intervals, being concerned only with the previous state and the next state" [Haye83 pg.

98].

VM does not have a facility for handling uncertain information [Faga84Exte]. The literature that we reviewed does not provide any evaluations of VM's performance.

#### **2.4. Current AI Techniques and the Alarm System Limitations**

In previous sections we described four limitations of conventional alarm systems, and also described some current techniques from artificial intelligence. In this section we discuss why these techniques, by themselves, do not seem to be adequate to overcome the alarm limitations. It is assumed that knowledge is represented as production rules (e.g., an alarm system implemented by means of a production system), but the discussion generalizes to other knowledge representations as well.

##### **2.4.1. Patient/Disease Specificity**

Recall that expected physiologic values are dependent on the nature of the patient and his disease; e.g., a patient with left ventricular hypertrophy is expected to have a significantly higher left atrial pressure (*LAP*) than a patient without hypertrophy. Similarly, a patient with right ventricular hypertrophy is expected to have a significantly higher right atrial pressure (*RAP*) than a patient without hypertrophy.

One method for handling patient/disease specificity is to have multiple sets of production rules; e.g., rules for a patient with hypertrophy, rules for

a patient without hypertrophy, and so forth. The problem with this method is the potentially explosive growth in the size and complexity of the rule base.

A somewhat better solution is to retain only one set of production rules, but to incorporate *symbolic ranges* into these rules. For example, the antecedent of a rule might be "IF LAP is *normal*" where *normal* is interpreted with reference to the particular patient. One way of implementing this is by means of a lookup table; e.g., "nonhypertrophy: *normal* = 4 to 14; hypertrophy: *normal* = 10 to 16 ...."

Several AIM systems have used symbolic ranges. For example, VM [Faga84Exte] used them to interpret data in the context of one of the four possible ventilator settings. MYCIN used symbolic ranges as a way to optimize storage [Shor84]. Long [Long83MED] discusses the use of symbolic ranges in a pathophysiologic reasoning system. See also [KahnM86].

Note that symbolic ranges require the use of cutoff points (i.e., low and high limits); of course, many other techniques also use cutoff points. Unfortunately, cutoff points can lead to rather disturbing behavior. One difficulty is that these points do not capture the idea that some values are more pathologic than others. For example, suppose that for the given patient an upper limit of 14 is selected for LAP. Then an observed LAP of 15 and an observed LAP of 29 would be considered to be equally bad since both exceed the cutoff; in reality the latter value is far worse. In addition, the discontinuous behavior at the cutoff points is disturbing; e.g., an LAP of 13.99 is considered to be completely normal, whereas an LAP of 14.01 is considered to be completely abnormal.

### 2.4.2. Temporal Changes

Recall that expected values of some physiologic variables change over time; e.g., chest tube drainage is expected to be high immediately after surgery, but low several hours later.

One method for handling temporal changes is to define time intervals during which the rules are applicable [Russ86]. Another technique is to use symbolic ranges where the symbolic ranges are applicable during specified time intervals; e.g., "0 to 60 minutes: *CT-normal* = 40 to 250; 61 to 120 minutes: *CT-normal* = 30 to 100 ...." Both of these approaches require cut-off points for time (as well as for the measured value of the variable), and thus suffer from the difficulties discussed above.

### 2.4.3. Multivariable Combinations

Recall that, in some cases, it is important to look at combinations of variables; e.g., if the left and right atrial pressures are elevated and approximately equal, then the patient may be in danger.

It is possible to extend the conventional paradigm to handle multivariable combinations. Typically, this is done by defining Boolean combinations of variables; e.g., "LAP is high & RAP is high ...." Besides the usual cutoff problems there are also problems of exponential combinatorics: e.g., given  $n$  variables there are  $2^n - (n + 1)$  multivariable combinations (combinations of 2 or more variables).

#### 2.4.4. Dynamic Patterns

Recall that the order in which events take place can be of significance; e.g., a decrease in cardiac index preceding an increase in LAP is probably of little significance; but an increase in LAP preceding a decrease in cardiac index is grounds for concern. It should be noted that event ordering is only one facet of the general problem of temporal reasoning; see [KahnK77] for a more general discussion.

The determination of the relative order of events is a more subtle problem than might be thought. The standard approach is to (1) assume that each event has a discrete time of occurrence, e.g., [Koya81]; (2) determine the times of occurrence; and then (3) order the events. The times of occurrence are determined by the variable exceeding some threshold,  $T$ ; e.g., (Fig. 2.2.).




Fig. 2.2: Time of occurrence is given by the threshold.

This approach suffers from some rather serious problems. First, it is susceptible to "spikes" (i.e., spurious values) in the data (Fig. 2.3.). Second, the approach runs into the usual cutoff problems (e.g., the variable approaches but does not exceed the threshold) (Fig. 2.4.). There are other




Fig. 2.3. Spike in the data.




Fig. 2.4. Just below the threshold.

difficulties as well; e.g., what if the cutoff changes over time?

The difficulties with this approach stem largely from the *assumption* that events have a discrete time of occurrence. In fact, for many things it is not possible to specify an exact time; e.g., at precisely which point in time does a person's cardiac function become inadequate?

It should be noted that, frequently, what is of clinical significance is not the exact times of occurrence of events, but rather their relative order of occurrence. In the next chapter we shall propose a method that determines the order of events, without the need to determine exact times of occurrence. Long [Long83AAAI] has also noted that time relationships can be of more

significance than exact times of occurrence; however, his method (defining times of occurrence by specifying temporal intervals) still suffers from the problems described above.

### Footnotes

- 1 Earlier attempts at computer-assisted diagnosis were based on statistical techniques. They are reviewed in [Cro72], [Rog79], and [Ward78].
- 2 Production rules turned out to be inadequate for representing certain types of knowledge: in particular, therapy selection knowledge is represented in a much more procedural fashion [Clan84Deta].
- 3 The control mechanism is actually much more complex than simple backward chaining. For example, Davis et al [Davi77Prod] states that "it makes good sense to ask the user first about positive cultures (those showing bacterial growth) before asking about negative cultures ... this design choice was embedded in the ordering of a list buried in the system code" (this choice was later implemented by means of meta-rules). Davis [Davi80] also notes that "in MYCIN subgoal ordering is hardwired in by the choice of a depth-first search and the ordering of preconditions in the premise of a rule." As yet another example, the weights on certain rules needed to be adjusted in order to force their consideration [Buch84 pg. 217].
- 4 The search is not exhaustive; for example,  $CF = .2$  (the .2 was "selected empirically" [Shor84]) is used as a threshold to constrain search, as is  $CF = 1$  [Buch84 pg. 63].
- 5 Clancey [Clan79] notes that the tree is never more than five levels deep.

**CHAPTER 3****OUR APPROACH****3.1. The Fundamental Strategy**

In the first chapter we noted that conventional alarm systems suffer from several limitations. In addition, we noted that although virtually any clinical problem can develop after cardiac surgery, there are only a few complications which are typically seen during the postoperative period.

The fundamental strategy of our approach<sup>1</sup> is to build a *recognizer* for each of these typical complications. The function of each recognizer is to detect its respective complication, and to issue an appropriate alert.

This strategy differs from diagnosis: A diagnostic system seeks to determine the disease (or diseases) which account for the patient's signs and symptoms. A recognizer, on the other hand, seeks only to detect a specific complication. Note that a diagnostic system has a large hypothesis space, whereas a recognizer has only one hypothesis. This greatly reduces the combinatorics (but with a potential loss in generality).

Each recognizer seeks to detect a specific complication. We view the detection of clinical complications from two perspectives: the *Set-Subset Perspective*, and the *Prototypical-Sequence Perspective*.

The basic idea of the set-subset perspective is that for each complication there is a set of signs which characterize that complication (e.g., cardiac tam-

ponade is characterized by abnormal chest tube drainage, low urine output, low cardiac index, low mean arterial pressure, and elevated and approximately equal left and right atrial pressures). At any point in time, only a subset of the possible signs may be present. The more of these signs that are present, the more strongly the complication should be suspected. When the set of observed signs is identical with the set of signs that characterize the complication, then it is virtually certain that the complication is present.

The basic idea of a prototypical sequence is that for each complication there are sequences of events which characterizes that complication. For example, in cardiac tamponade, abnormal chest tube drainage is followed by rising and converging left and right atrial pressures and a decreasing cardiac index, and then finally by a decreasing mean arterial pressure. The more closely a patient follows such a sequence, the more strongly the complication should be suspected. If a patient exactly follows a prototypical sequence, then it is virtually certain that he has the complication.

Since single variables and multivariable combinations must be interpreted with reference to the particular patient and the particular point in time, the set-subset perspective addresses the first three alarm limitations (patient/disease specificity, temporal changes, and multivariable combinations). The prototypical-sequence perspective addresses the fourth alarm limitation (dynamic patterns).

### **3.2. Parameterized Characteristic Functions**

### 3.2.1. Using Characteristic Functions to Represent Clinical Knowledge

*Characteristic functions* are one of the fundamental tools in our approach (these functions are reviewed in section 4 of Appendix I). Characteristic functions are functions which define the degree of membership of an object in a set. They can serve as a reasonably natural way of representing the "goodness" or "badness" of the measured value of a physiologic variable.


Example:

In the typical patient the left atrial pressure runs in the range from 4 to 14 mm Hg. Define the function  $\mu_{LAP}(l)$  where

$$\mu_{LAP}(l) = \begin{cases} 1 & \text{if } l < 1 \\ 1 - ((l-1)/3) & \text{if } 1 \leq l < 4 \\ 0 & \text{if } 4 \leq l \leq 14 \\ (l-14)/3 & \text{if } 14 < l \leq 17 \\ 1 & \text{if } l > 17, \end{cases}$$

and where  $l$  is the measured value of the left atrial pressure,  $\mu_{LAP}(l) = 0$  means that the measured value is completely normal,  $\mu_{LAP}(l) = 1$  means that the measured value is completely abnormal, and  $0 < \mu_{LAP}(l) < 1$  means that the measured value is somewhat abnormal. Note that by defining the function in this way it is possible to represent the fact that some measured values are more pathologic than others (rather than classifying all values as either "normal" or "pathologic"). In addition, discontinuous behavior at the distinguished points (4 and 14) is avoided. See Fig. 3.1.<sup>2</sup>

Since the expected values of a physiologic variable depend upon the particular situation (e.g., preexisting pathology, length of time since surgery, etc.) it is not possible to define a single characteristic function for each vari-

Fig. 3.1.  $\mu_{LAP}(l)$ .

able. Hence, we allow the characteristic functions to contain parameters.

**Example:**

Suppose that in the given patient the left atrial pressures are expected to run in the range from  $L$  to  $H$ . Define the function  $\mu_{LAP}^*(l)$  where

$$\mu_{LAP}^*(l) = \begin{cases} 1 & \text{if } l < L - 3 \\ 1 - ((l - (L - 3)) / 3) & \text{if } L - 3 \leq l < L \\ 0 & \text{if } L \leq l \leq H \\ (l - H) / 3 & \text{if } H < l \leq H + 3 \\ 1 & \text{if } l > H + 3. \end{cases}$$

See Fig. 3.2. This function exhibits the same type of behavior as  $\mu_{LAP}(l)$  but it can be "customized" for the given situation by setting the parameters  $L$  and  $H$ . In particular, if  $L$  and  $H$  are set appropriately for the given patient, then patient/disease specificity can be handled. In addition, if  $L$  and  $H$  are appropriately adjusted over time, then temporal changes can be handled.

It should be noted that the characteristic function for a given variable may be different in different complications; for example, both low LAP and




Fig. 3.2.  $\mu_{LAP}^*(l)$ .

high LAP are of clinical concern, but low LAP is suggestive of hypovolemia, whereas high LAP is suggestive of tamponade or cardiac failure.

### 3.2.2. The Set-Subset Perspective

In the following sections we describe in more detail how the set-subset perspective can handle patient/disease specificity, temporal changes, and multivariable combinations.

#### 3.2.2.1. Patient/Disease Specificity

One difficulty in conventional alarm systems is that appropriate values for the distinguished points *Low* and *High* differ from individual to individual (patient/disease specificity). In actual practice this limitation can be overcome because the good clinician knows that there are certain expected values

(i.e., default values) for the typical patient, but that these values must be modified depending upon the particular pathology.

Usually, the number of pathologic conditions that can alter the expected values is fairly small. For example, in a typical patient normal left atrial pressure runs in the range from 4 to 14 mm Hg. However, if the patient has ventricular hypertrophy, or ventricular dilatation (or one of a few other conditions), then the expected range is 10 to 16 mm Hg. This sort of information can be modeled by production rules; e.g.,

```
IF pathology present
THEN set pathologic range
ELSE set normal range.
```

Our approach for handling patient/disease specificity is to define a characteristic function for each of the measured physiologic variables. In addition, a production system is used to determine appropriate values for *Low* and *High*, and these values are then used to instantiate the characteristic functions.

Note that certain pathologic conditions may be fuzzy (e.g., degree of hypertrophy). These conditions can be handled by means of fuzzy initialization rules (e.g., greater degree of hypertrophy -> higher expected range).

#### 3.2.2.2. Temporal Changes

Recall that some physiologic variables change over time. It is usually possible to describe expected temporal curves for these variables. For example, in most cases, the chest tube drainage is initially quite high, and then

drops off over the next few hours (Fig. 3.3.). Note that these sorts of curves really represent expected ranges rather than expected values.




Fig. 3.3. Expected chest tube drainage.

While all patients do not follow such curves, most patients who are being properly managed and who are on the road to recovery do; thus, deviation from a curve can be viewed as an indication of pathology.

Our approach for handling temporal changes is to define a characteristic function for each of the time-dependent variables. In addition, a temporal curve is defined for each such variable, and these curves are used to parameterize the respective characteristic functions. Note that, in some cases, it is necessary to shift or modify the expected curves depending upon the particular patient (e.g., the expected chest tube drainage is a function of the patient's weight); this can be done by a production system. As before, fuzzy initialization rules could be used, if needed.

### 3.2.2.3. Multivariable Combinations

An important limitation of most conventional alarm systems is that they

do not take account of multivariable combinations. One difficulty in overcoming this limitation is that the number of combinations is exponential: given  $n$  variables, there are  $2^n - (n + 1)$  multivariable combinations.

It might seem desirable that the alarm system examine all multivariable combinations. Due to the rapid growth of  $2^n$ , however, a system that attempts to do so is doomed to be extremely unwieldy for even small values of  $n$ .

Fortunately, however, for each typical complication it is possible to identify a limited number of combinations which are of clinical significance. For example, in cardiac tamponade the LAP-RAP difference is important, whereas many other combinations (e.g., right atrial pressure and urine output) are of little or no significance. By selecting only the clinically significant combinations the exponential combinatorics can be greatly reduced.

Thus, our approach for handling multivariable combinations is to define a characteristic function for each relevant combination. This can be done in a manner analogous to that used for single variables. Note that the multivariable characteristic functions may depend upon the particular clinical complication; for example, a small LAP-RAP difference is characteristic of tamponade, whereas a large LAP-RAP difference is characteristic of myocardial failure [Engl70].

#### **3.2.2.4. Combining the Evidence**

For each complication, characteristic functions are defined for the single variables. In addition, the important multivariable combinations are

selected, and characteristic functions are defined for those combinations. Thus, the physiologic measurements from the patient are converted into  $\mu$ -values<sup>3</sup> which reflect the "goodness" (equivalently "badness") of the single variables, and multivariable combinations. Each of these  $\mu$ -values is a piece of evidence concerning the condition of the patient.

An important issue is how to combine these pieces of evidence. There are at least two considerations in this regard: (1) It must be possible to account for the degree of severity; for example, a very elevated LAP (e.g.,  $\mu_{LAP}(l) = .95$ ), should count more heavily than a moderately elevated LAP (e.g.,  $\mu_{LAP}(l) = .70$ ). (2) It must be possible to account for the relative weightings of variables and multivariable combinations; for example, the cardiac index is more significant than the mean arterial pressure.

Various numerical calculi have been developed to deal with the combination of evidence; some of these calculi are reviewed in Appendix I. One of the conclusions of that appendix is that there is not one calculus that is "the best," and that it is necessary to carefully choose (or construct) a calculus that meets the needs of a given problem.

Since our notion of characteristic functions is borrowed from fuzzy set theory, it is natural to look to that theory for suitable operators with which to combine the evidence. In the framework of fuzzy set theory, an intersection operator is required since a recognizer might, for example, look for abnormal chest tube drainage, *and* low urine output, *and* low cardiac index, etc.

Recall that the standard fuzzy set definition of intersection is the minimum operator (*min*). This operator is clearly not appropriate for the needs of our system since one small value will overwhelm many large values

(e.g.,  $\min(.01, .9, .9, .9, \dots) = .01$ ). Other definitions for intersection have been proposed, e.g., [Yage80], but these also do not seem appropriate.

The Bayesian and Dempster-Shafer calculi do not satisfy the needs of our system (e.g., they do not offer a convenient method for representing the relative severity of a physiologic reading). The MYCIN and EMYCIN calculi also are not appropriate. Accordingly, we have chosen to develop our own method for combining evidence.

The method that we use is the following: Given that  $n$  single variables are monitored, and that  $\mu_1, \dots, \mu_n$  are the respective characteristic functions for those variables. Given that  $m$  multivariable combinations are monitored and that  $\mu_{n+1}, \dots, \mu_{n+m}$  are the respective characteristic functions for those combinations. Define the function  $WS$  (weighted sum) where

$$WS = \alpha_1 \mu_1 + \dots + \alpha_{n+m} \mu_{n+m},$$

and where  $\sum_{i=1}^{i=n+m} \alpha_i = 1$ . This function can account for degrees of severity

(e.g.,  $WS$  will be greater if  $\mu_{LAP} = .95$ , than if  $\mu_{LAP} = .75$ ). In addition, the  $\alpha$ 's are weights which can be modified to account for the relative importance of the variables and multivariable combinations. In practice, the  $\alpha$ 's should be adjusted to tune the performance of the the system.

It should be noted that each of the variables may be measured on a different numerical scale (or in different units entirely!); thus it would be problematic to construct a function  $WS$  which combines the raw values of the variables. This, then, is one advantage of using characteristic functions—the functions *normalize* the values, thus allowing for the meaningful combination of information.

Note that if *WS* is to be used to trigger a binary alarm (e.g., *on* or *off*) then it is necessary to select a cutoff point, *C* where  $WS > C$  will trigger the alarm. This is discussed in further detail in chapter six.

### **3.2.2.5. Summary of the Set-Subset Perspective**

For the given complication the important single variables and multivariable combinations are selected. The characteristic functions for those single variables and multivariable combinations are constructed. In order to handle patient/disease specificity, parameters are instantiated by a production system. In order to handle temporal changes, temporal curves are defined; a production system may be used to adjust the curves. The scoring function *WS* is determined (i.e., the  $\alpha$ 's are adjusted). *WS* computes a number in the unit interval [0,1] where 0 means that the complication is definitely not present, 1 means that the complication is definitely present, and a number between 0 and 1 has an intermediate interpretation.

### **3.2.3. The Prototypical-Sequence Perspective**

The standard technique for determining the order of events is to determine a discrete time of occurrence for each event and then to order these times. As we discussed, this technique has several difficulties; these difficulties stem from the (problematic) assumption that events have a discrete time of occurrence.

### 3.2.3.1. Outline of a Sequencing Technique

The technique which we describe below *directly* determines the order of events without determining times of occurrence. Since the technique does not determine times of occurrence, it is not necessary to make the assumption that events occur at a distinct point (or interval) in time.

Outline of the technique:

- 1) At each point in time map the measured value of each variable into its corresponding  $\mu$ -value ("degree of badness") by means of appropriate characteristic functions; i.e., for each of the variables  $v_i$  compute  $\mu_{v_i}$ .
- 2) Compute the sum over time of the  $\mu$ -values for each variable; i.e., compute  $\int \mu_{v_i}$  for each of the variables  $v_i$ .
- 3) Determine the order of events by sorting the  $\int \mu_{v_i}$  in numerical order; e.g., if  $\int \mu_{v_i} > \int \mu_{v_j}$  then variable  $v_i$  "went bad" before variable  $v_j$  went bad.

### 3.2.3.2. Examples and Elaboration of the Sequencing Technique

The following example should serve to clarify this technique: Suppose that we are given the physiologic variable,  $A$ , where

$$\mu_A = \begin{cases} 0 & \text{if } A < 10 \\ (A - 10)/10 & \text{if } 10 \leq A \leq 20 \\ 1 & \text{if } A > 20. \end{cases}$$

Now suppose that the value of  $A$  from time  $t_1$  to  $t_2$  is given by Fig. 3.4.

Then the corresponding  $\mu$ -values for  $A$  ( $\mu_A$ ) will be as depicted in Fig. 3.5.




Fig. 3.4. Value of  $A$  from time  $t_1$  to time  $t_2$ .




Fig. 3.5.  $\mu_A$  from time  $t_1$  to time  $t_2$ .

Suppose that we are also given the physiologic variable,  $B$ , in which  $\mu_B = \mu_A$  (substituting  $B$  for  $A$ ). Then if the value of  $B$  from time  $t_1$  to time  $t_2$  is given by Fig. 3.6., then  $\mu_B$  is given by Fig. 3.7. In Fig. 3.8.  $A$  and  $B$  are superimposed on the same graph. Fig. 3.9. superimposes  $\mu_A$  and  $\mu_B$  on the same graph. Note that, in Fig. 3.9., the area under curve  $\mu_B$  is greater than the

area under curve  $\mu_A$ ; i.e.,  $\int_{t_1}^{t_2} \mu_B > \int_{t_1}^{t_2} \mu_A$ ; therefore, event  $B$  precedes event  $A$ .




Fig. 3.6. Value of  $B$  from time  $t_1$  to time  $t_2$ .




Fig. 3.7.  $\mu_B$  from time  $t_1$  to time  $t_2$ .

In this particular example, we have constructed  $A$  and  $B$  on the same numerical range; due to this construction it is also true that  $\int_{t_1}^{t_2} B > \int_{t_1}^{t_2} A$ . However, most physiologic variables do not run along the same range; e.g., a typical right atrial pressure might be in the range 4 to 12 mm Hg, whereas a typical left atrial pressure might be in the range 4 to 14 mm Hg. Moreover, many physiologic variables are measured in entirely different units; e.g., cardiac pressures are measured in mm Hg, whereas cardiac output is meas-




Fig. 3.8.  $A$  and  $B$  on the same graph.




Fig. 3.9.  $\mu_A$  and  $\mu_B$  on the same graph.

ured in liters/min.

Note that if variable  $A$  and variable  $B$  have different numerical ranges

then it is meaningless to compare  $\int_{t_1}^{t_2} A$  and  $\int_{t_1}^{t_2} B$ . However, the conversions

from  $A$  to  $\mu_A$  and from  $B$  to  $\mu_B$  are in effect normalizations, and these nor-

malizations make the comparison of  $\int_{t_1}^{t_2} \mu_A$  and  $\int_{t_1}^{t_2} \mu_B$  meaningful; i.e., even if

$A$  and  $B$  have different numerical ranges, then given that  $\int_{t_1}^{t_2} \mu_B > \int_{t_1}^{t_2} \mu_A$  we can

say that event *B* precedes event *A*. Moreover, such normalizations allow for the comparison of variables that are measured in entirely different units.

As was noted earlier, there are several difficulties with the standard approach for determining the order of events. One difficulty is due to spikes in the data; i.e., a spurious high value which exceeds the threshold will lead to an erroneous time of occurrence, hence the event sequence may be incorrect. Since our technique computes an integral, momentary spikes in the data are essentially "averaged out" (see Fig. 3.10.).




Fig. 3.10. In the upper diagram *A* crosses the threshold, *T*, before *B* because of the "spike" (i.e., spurious value); therefore, it is incorrectly determined that event *A* precedes event *B*. In the lower diagram  $\int \mu_B > \int \mu_A$ , therefore it is correctly determined that event *B* precedes event *A* (the spike in the data is "averaged out").

Another problem with the standard approach lies in the use of arbitrary thresholds; e.g., even if a variable is infinitesimally less than its corresponding threshold, the event is determined not to have taken place, hence the event sequence may be incorrect. Our technique does not suffer from this limitation (see Fig. 3.11.).




Fig. 3.11. In the upper diagram  $A$  crosses the threshold,  $T$ , before  $B$  (even though  $B$  is extremely close), therefore it is incorrectly determined that event  $A$  precedes event  $B$ . In the lower diagram  $\int \mu_B > \int \mu_A$ , therefore it is correctly deduced that event  $B$  precedes event  $A$ .

The reason that our technique does not suffer from the problems of the standard approach is that *it does not determine exact times of occurrence*. In the standard approach these times must be computed as an intermediate result; it is because our technique does not require this additional computa-

tional step that it does not suffer from the same difficulties as the standard approach.

Certain additional observations are in order about the technique. First, in many real applications, variables are measured periodically rather than continuously; e.g., LAP may be measured and averaged over 10 minute intervals. Our technique can easily handle such variables by interpolating between successive values. Moreover, as the sampling frequency increases, so does the accuracy of the technique.

Second, note that in those cases where variables make sudden transitions from "good" to "bad" (i.e., stepwise transitions) the technique can in fact be used to determine an exact time of occurrence (if desired): if the variable is measured during the period  $t_1$  to  $t_2$ , then the time of occurrence is given by  $t_2 - \int \mu_v$ .

Our technique does have several limitations; these limitations are discussed in chapter six.

### 3.2.3.3. Summary of the Prototypical-Sequence Perspective

The sequence of events that the patient follows (the *actual sequence*) is determined; e.g., by means of our technique. The actual sequence is then matched against the prototypical sequence; the closer the match, the stronger the evidence is that the patient has the complication.

### **3.2.4. Integrating the Two Perspectives**

Both the set-subset perspective and the prototypical-sequence perspective provide information regarding the presence of the clinical complication. This information may be combined: e.g., the alarm is activated only if *WS* is sufficiently high *and* the patient has exhibited a close match to a prototypical sequence. Alternatively, the information may be viewed as being complementary: e.g., the alarm is activated if *WS* is sufficiently high *or* (nonexclusive) the patient has exhibited a close match to a prototypical sequence.

In our domain the two perspectives are probably best viewed as being complementary: If a patient has all the characteristic signs, then even if he has not followed a prototypical sequence, it is likely that he has the complication. Similarly, if a patient has followed a prototypical sequence, then even if he does not have all the signs, it is likely that he has the complication. In other domains, however, this may not be the case.

### **3.3. A Pragmatic Consideration**

Several authors [Draz80, Shep79] have noted that in order for a computer system to be accepted into an intensive care unit, it must be well integrated into the existing routine (e.g., it can *not* require substantial modifications to the established clinical protocol).

An important consideration in this regard is that the system can expect only a very limited amount of manual input: since the staff is already quite busy there is little time to answer queries from the computer. Therefore, the system must acquire most of its data automatically. Similarly, the system

should produce only a limited amount of output in order that the staff not be overwhelmed and ignore important information [Conn86, Osbo68].

### **3.4. Why Not Pathophysiologic Reasoning**

Much current AIM research is concerned with pathophysiologic reasoning. Although the incorporation of pathophysiologic reasoning mechanisms might seem to be desirable, we have specifically chosen *not* to use this approach for several reasons.

First, the physiology is complex and not completely understood [Cull79]. For example, pulsus paradoxus is considered by many to be a nearly sure sign of cardiac tamponade—yet after many years of research its mechanism has still not been definitively established (Appendix III). In addition, the physiologic mechanisms are influenced by factors that may have unpredictable results. For example, the effects of many of the medications administered in the unit can be quite variable depending upon the dosage, individual response, etc. [Long86] (e.g., dopamine can both increase and decrease renal perfusion). In some cases, factors which are essentially intangible can affect the physiology; for example, pain and nausea can affect hemodynamics [Russ74].

Second, even if the physiology were well understood it is not clear that we would be able to build an adequate computer model. For example, the laws of mechanics are quite simple and well understood (certainly more so than human physiology), yet, at this point in time, the construction of a robust physical reasoning system is still an open problem in artificial intelli-

gence.

Finally, even if it were possible to build an adequate physiologic model, a system based on such a model would not be pragmatically acceptable due to the tremendous amount of manual input that would be required (e.g., time, route, and dosage of all drugs administered). This particular difficulty might be overcome, however, if a reliable speech input system were available.

#### Footnotes

- 1 Early in the course of this research we investigated two alternative approaches - 1) modeling the alarm system as a production system, and 2) modeling the alarm system as a discrete finite automata. Although both of these formalisms have great theoretical power, they turned out to be quite unwieldy when applied to the construction of our alarm system; hence, we rejected these approaches.
- 2 Note that the figures in this manuscript are imprecise and have not been drawn to scale.
- 3 We will call the values computed by characteristic functions  $\mu$ -values.

**CHAPTER 4****THE PROGRAM****4.1. The Scope of the Program**

In order to test our approach, we designed a program to detect the presence of cardiac tamponade. This complication was selected since it is life-threatening, but potentially curable if recognized early enough [Shoe73, Week76].

Currently a computer system is in use in the CSICU [Jura77, Jura82]. This system, which is based on Osborn's [Osbo68] design, automatically acquires values of selected physiologic variables at ten minute intervals (although it can also sample at one minute intervals). In addition, it accepts values of other variables which are measured by hand. The values (both automatically and manually acquired) are stored and available for use by other programs. Thus, our program is designed to use, as input, data that has been gathered by the existing program.

The existing program measures (or derives) over twenty physiologic variables. Our program uses only six of these variables: chest tube drainage (*CT*), urine output (*UO*), cardiac index (*CI*), mean arterial pressure (*MAP*), left atrial pressure (*LAP*), and right atrial pressure (*RAP*).

The patient population which the program is designed to monitor has been restricted as follows: The patients are adults between the ages of twenty-one and seventy-five (since the very young and the very old often

exhibit markedly atypical postoperative behavior). The patients have undergone mitral valve, aortic valve, or CABG surgery. The patients are not grossly obese ( $> 350$  lbs), nor are they markedly cachectic ( $< 80$  lbs), since such patients often exhibit atypical postoperative behavior. Most patients in the CSICU satisfy these restrictions.

In addition, it is assumed that patients who are predisposed to heavy bleeding, have a fragile aorta, or who have undergone a CABG procedure, are administered sodium nitroprusside in order to control the arterial pressure. The program is intended to monitor patients only until extubation (typically around 24 hours) since after that point patients become more active and hence have variable hemodynamics.

## **4.2. Description of the Six Physiologic Variables**

### **4.2.1. Chest Tube Drainage**

The chest tube drainage is measured by hand approximately every hour. This process could be automated. In some cases (e.g., abnormal bleeding patterns) more frequent measurements may be appropriate.

Both the flow rate and the cumulative volume are of clinical significance; markedly low or high values of either of these is grounds for concern. In general, chest tube drainage is high initially, and then tapers off during the postoperative period. Chest tube drainage is a function of body weight (larger individuals have higher bleeding rates).

Certain factors may predispose the patient to higher than normal chest tube drainage. Among these are previous cardiac surgery, multiple CABGs, a left internal mammary artery (*LIMA*) CABG, excessive bleeding in the operating room, excessive time on cardiopulmonary bypass, and a preexisting coagulopathy.

#### **4.2.2. Urine Output**

The urine output is measured by hand approximately every hour. This process could be automated. In some cases more frequent measurements may be appropriate (e.g., poor renal function).

The flow rate is of clinical significance (cumulative volume is generally disregarded); a low flow rate is grounds for concern. In general, urine output is high initially and then decreases during the postoperative period. Urine output is a function of body weight (larger individuals have higher flow rates).

Significant preoperative cardiac dysfunction and preoperative renal dysfunction may predispose the patient to low urine output.

#### **4.2.3. Cardiac Index**

Cardiac index is measured by hand using a thermodilution catheter. This process could be automated. The frequency of measurements is at the discretion of the physician, but, as a minimum, readings should be taken every few hours. If the patient appears to be doing poorly, then more fre-

quent measurements are essential.

A cardiac index greater than 2.4 is considered acceptable; a value below 1.6 is, in most cases, incompatible with survival.

Certain conditions may predispose to low cardiac index. Among these are mitral valve surgery, significant preoperative cardiac dysfunction, significant preoperative pulmonary hypertension, a low preoperative cardiac index, and mitral valve pathology.

#### 4.2.6. Mean Arterial Pressure

The mean arterial pressure is automatically measured by an indwelling arterial catheter. It is the thirty second average of the arterial pressure waveform.

A low MAP is of concern; high pressures can also be dangerous, (especially in patients who have fragile aortas or who are predisposed to bleeding problems). In the typical patient the MAP ranges between 60 and 90 mm Hg. Those patients who were hypotensive preoperatively tend to run a little lower, i.e., 55-80 mm Hg; those patients who were hypertensive preoperatively tend to run a little higher, i.e., 60-95 mm Hg. In many patients the MAP may be somewhat low during the first three postsurgical hours.

Patients who have had a CABG often develop dangerous hypertension, hence, may have their MAP regulated by sodium nitroprusside; for the purposes of this program we assume that *all* CABG patients are so regulated (MAP kept between 60 and 110 mm Hg). In some patients (i.e., those with fragile aortas and those predisposed to heavy bleeding) even modest

hypertension is dangerous; we assume that these patients are regulated by nitroprusside (MAP kept between 55 and 75 mm Hg).

#### **4.2.4. Left Atrial Pressure**

The left atrial pressure is automatically measured by a left atrial catheter. The value is the thirty second average of the a, c, and v waves.

In the typical patient the left atrial pressure is usually in the range from 4 to 14 mm Hg. The LAP may be somewhat volatile during the first post-surgical hour.

Certain factors lead to higher than normal expected values of LAP: aortic surgery, mitral surgery, left ventricular hypertrophy, left ventricular dilation, aortic valve problems, mitral valve problems, and left sided congestive heart failure.

#### **4.2.5. Right Atrial Pressure**

The right atrial pressure is automatically measured by a right atrial catheter. The value is the thirty second average of the a, c, and v waves.

In the normal patient the right atrial pressure is usually in the range from 4 to 12 mm Hg. The RAP may be somewhat volatile during the first post-surgical hour.

Certain factors lead to higher than normal expected values of RAP: pulmonary hypertension, right ventricular hypertrophy, right ventricular dilation, tricuspid valve problems, and right sided congestive heart failure. In

addition, patients with left sided pathology need to be forward loaded, and hence will have a higher than normal expected RAP.

#### **4.3. Post-Surgical Cardiac Tamponade**

Cardiac tamponade is a potentially life-threatening complication in which excessive bleeding into the pericardial space leads to compromised cardiac function. Tamponade can occur either "early" (within hours after surgery) or "late" (days after surgery) [Hill82]; we are concerned with early tamponade. Engleman et al [Engl70] have reported an incidence of 4 early tamponade cases in 310 post-surgical patients. The complication can be corrected by surgical means if it is detected early enough [Shoe73, Week76].

Among the signs that characterize tamponade are abnormal chest tube drainage (either high or low), low urine output, low cardiac index, low mean arterial pressure, and elevated and approximately equal left and right atrial pressures. Many references claim that the phenomenon of pulsus paradoxus is nearly always found in tamponade. However, in Appendix III we conclude that the phenomenon is widely misunderstood; hence our system does not utilize information about pulsus paradoxus. A prototypical scenario for a patient with cardiac tamponade is the following: (1) abnormal chest tube drainage is followed by (2) rising and converging left and right atrial pressures and a decreasing cardiac index, and then finally (3) a decreasing mean arterial pressure. Cardiac tamponade is reviewed more thoroughly in Appendix IV.

#### **4.4. Details of the Program**

The program is written in Franz Lisp. It is approximately 2,000 lines long, excluding comments and blank lines. The program could be made somewhat shorter by recoding certain portions into a language more suitable for numerical computations; e.g., C.

The program is broken into three sections: the *questioning module*, the *deductive module*, and the *monitoring module*. The questioning and deductive modules can be run before the patient returns from surgery; the monitoring module must run on-line and in real-time during the post-surgical period.

##### **4.4.1. The Questioning Module**

The questioning module gathers the necessary background information about the patient (i.e., the relevant medical history). The module gathers the information by posing a series of questions to the nurse or physician (*the user*). The questions require either a short answer response, or a multiple choice type response.

This module contains a list of approximately forty questions which may be posed to the user. Since not all questions are necessary or appropriate for a given individual, the actual questions posed vary from patient to patient (e.g., if the patient has not had a CABG, then it is inappropriate to ask whether he has had a LIMA CABG).

Simple error checking is performed to make sure that the responses to the questions are reasonable. The user may request elaborations to clear up

questions that he feels are imprecise. In addition, the responses are echoed back to the user, and the user may then change his responses if not satisfied. Another feature of the questioning module is that it screens out patients who do not meet the program's criteria (e.g., patients who are too old, grossly obese, etc.).

Each question is composed of six components (the sixth one being optional). The first component is the actual query; e.g., "Were any of the CABGs a LIMA?". The second component is a shortened form of the query (for experienced users); e.g., "LIMA?".

The third component is an elaboration of the query. The elaboration may contain a definition so as to avoid problems of lexical imprecision (see Appendix II), or it may give some insight as to why the question has been posed, and how the response will be used (e.g., "Often patients with a LIMA CABG have heavy bleeding postop").

The fourth component is a piece of code that determines whether the response to the question is legal (e.g., (memq response '(yes no y n))). It should be noted that the error checking is quite rudimentary; it is essentially "syntactic" rather than "semantic."

The fifth component is an atom that stores the response (assuming that it is legal).

The sixth component is a piece of code that determines if the query should be posed to the user; e.g., (or (eq CABG 'yes) (eq CABG 'y)). This piece of code may check the responses to previous questions in order to determine whether the current question is appropriate—this allows the sequence of questions to be varied depending upon the particular patient.

Another potential method for gathering the background information about the patient would be to accept "free form" input from the user. We have rejected this approach because question asking is faster and easier to implement. In addition, question asking guarantees that all the necessary information is gathered.

#### 4.4.2 The Deductive Module

The deductive module instantiates the parameters for the characteristic functions, modifies the temporal curves as needed, and so forth. This module uses, as input, information that has been acquired by the questioning module.

The heart of the deductive module is a set of production rules. Currently there are approximately 120 rules. Fuzzy initialization rules were not utilized, although they could have been easily incorporated into the rule base.

The rules are partitioned into lists (each rule is an element of one and only one list). Each variable and multivariable combination has associated with it two of these lists: a *silent action list*, and a *nonsilent action list*. The silent action list contains rules that perform actions which are of no concern to the user (e.g. the initialization of certain atoms). The nonsilent action lists contain rules that perform actions which are of concern to the user (e.g., the instantiation of parameters for the characteristic functions). The rules on the nonsilent lists contain English language descriptions; these descriptions are output to the user as the production rules are fired. In addition, some of the

rules contain elaborations which describe the purpose of the rule in greater detail.

Each list of rules constitutes a rule base. Certain observations are in order about these rule bases: (1) The rules are commutative in the sense of Nilsson [Nils pg. 35] (intuitively, with certain constraints, the order of rule firings does not make any difference to the final outcome). (2) The rules will be fired only once (if at all). (3) There are only a few multistep reasoning chains.

The control structure (i.e., rule interpreter) executes as follows: move down each list a rule at a time—if a rule antecedent is non-nil then fire the consequents, otherwise don't; continue until the end of the list. This scheme is simple, but adequate in view of the above observations. Since certain rules depend upon deductions made by other rules, it is necessary that each list be properly ordered (in general, the ordering needs to be no more complex than a topological sort). Note that we have ordered the rules in the lists to make the sequence of deductions seem as logical as possible to the user.

Using information acquired by the questioning module, and knowledge encoded in the production rules, the deductive module makes various deductions concerning the patient. Two particularly important types of deductions are *expectations* and *warnings*.

Expectations are deductions about the *expected values* of variables (e.g., "EXPECT: rap to be high because of pulmonary hypertension"). Expectations are used by the program to instantiate parameters for the characteristic functions (e.g., if RAP is expected to be normal, then the parameters for  $\mu_{RAP}$  are 4 and 12, whereas if RAP is expected to be high, then the parame-

ters for  $\mu_{RAP}$  are 8 and 16).<sup>1</sup> Note that, in some patients, several factors may lead to the same expectation; (e.g., "EXPECT: rap to be high because of pulmonary hypertension," and "EXPECT: rap to be high because of right ventricular hypertrophy"). The program handles these sorts of situations by assigning even higher values for the parameters (e.g., 8 and 18 for  $\mu_{RAP}$ ).

Warnings are deductions about an *increased probability* of certain problems (e.g., "WARNING: because the patient is a reop, may have heavy bleeding"). Warnings about a particular variable are *not* used to instantiate the parameters for that variable since, for example, a bleeding rate of 500 ml/hr is dangerous regardless of whether or not the patient is predisposed to heavy bleeding. Warnings, however, may be used to instantiate parameters for other variables; e.g., if the patient is predisposed to heavy bleeding then the MAP should be kept fairly low. This is because, clinically, such information influences patient management; e.g., if the patient is at risk for heavy bleeding then clinical protocol dictates that the MAP should be kept low (e.g., by means of sodium nitroprusside). Warnings might also be used to increase the sampling rate of a variable (e.g., if there is an increased probability of abnormal bleeding, then CT should be measured with increased frequency). Note that increased sampling rates have not yet been implemented in our test program.

The distinction between expectations and warnings is important: Expectations are what one expects to see if the patient's progress is *satisfactory*. Warnings, on the other hand, refer to conditions that may arise if the patient's progress is *unsatisfactory*. In addition, warnings refer to conditions that are not necessarily expected to occur (although these conditions have a

greater than normal probability).

The program outputs information about all expectations, warnings, and other important deductions. This output allows the user to better understand the program's behavior. In addition, the output can serve as useful reminders for the less experienced staff.

The questioning and deductive modules can be run before the patient returns from surgery. The staff is not overly busy during this period, and, therefore, the amount of time necessary to respond to the questions and examine the deductions (in practice, only a few minutes) is pragmatically acceptable.

#### 4.4.3. The Monitoring Module

The monitoring module accepts, as input, values of physiologic variables (the values having been acquired automatically, or entered by hand). Using the parameters and temporal curves established by the deductive module, the monitoring module computes  $\mu$ -values for each variable and multivariable combination, and computes  $WS$ . In addition,  $\int \mu$  is computed for each variable and multivariable combination.

Each reading consists of a value, and a time at which that value was measured. The times are needed because, for variables which are measured by hand, the data may be entered at a different time than when it was measured. In addition, the times are needed because the  $\mu$ 's are computed with respect to the time of measurement, and not with respect to the current clock time. For convenience, readings are stored in lists, with each variable having

its own list.

For each variable the module does the following: Determine if a new reading has been entered. If a new reading has been entered then screen the value to make sure that it is reasonable. If the value is reasonable then the reading is concatenated onto its respective list, and  $\mu$  and the  $\int \mu$  are computed. If the value is not reasonable, then the reading is discarded, and the user is so informed. Note that the screening process is quite rudimentary; e.g.,  $1 \leq CI \leq 15$ .

After each of the single variables has been processed the multivariable combinations (i.e., LAP-RAP) are handled in a similar fashion. Finally,  $WS$  for the clinical complication is computed.

Note that certain variables are measured more frequently than others (e.g., LAP is measured every ten minutes, whereas CI may be measured as infrequently as every few hours).  $WS$  uses the most recently acquired value of each variable; hence, different values may have been measured at different points in time (e.g.,  $\mu_{MAP}$  may be computed for a value of MAP that is a few seconds old, whereas  $\mu_{CI}$  may be computed for a value of CI that is an hour old). The characteristic functions are computed with respect to the time of measurement, *not* the clock time (e.g., if the UO was last measured at 0500, and it is now 0600, then  $\mu_{UO}$  is computed with respect to time 0500).

The module is driven by a clock which is incremented every ten minutes. The above tasks are repeated during each clock cycle.

At present the module is not on-line. Accordingly, certain aspects of the program have been simulated. In particular, the readings come from

lists, and the clock is faster than real-time.

Selected portions of the program can be found in the section of this thesis entitled *Sample Code*.

### Footnotes

- 1 The user is given the opportunity to suggest values for the parameters. If the user's and the computer's suggested values are "close" (i.e., within 2 for LAP and RAP, and within 5 for MAP) then the user's value is automatically selected; if the user's suggested values differ significantly from the computer's, then the user decides which values the system will utilize.

## CHAPTER 5

## RESULTS

**5.1. The Scope of the Tests**

Two experiments were performed using the program. The purpose of the first experiment was to assess the soundness of our approach; in particular, to assess whether the initialization deductions (e.g., parameter instantiations),  $\mu s$ ,  $WS$ , and the  $\int \mu s$  exhibit behavior that seems to be "reasonable" and in accord with intuition. The limited nature of this experiment is emphasized—it was *not* intended to be a rigorous evaluation of a fully implemented alarm system. For example, no attempt was made to select a value  $C$  (where  $WS > C$  triggers an alarm), no matching of the actual and prototypical sequences was done, and  $WS$  and the sequence information were not integrated. In addition, the test data was relatively "clean"; e.g., it contained only a few erroneous values.

The purpose of the second experiment was to assess whether our approach might be suitable for use in the CSICU.

**5.2. Experiment I: Performance Tests**

The program was run on six sets of patient data. Four of these sets were used to *tune* the program; the other two sets were used to *test* the program. The data sets were relatively "clean" (i.e., contained few erroneous

values).

An initial set of production rules, temporal curves, characteristic functions, and  $\alpha$ s was provided by a senior physician from the CSICU (Dr. A. M. Benis) after we described our approach to him. The program was then run on the tuning data sets. Three of the four tuning sets consisted of simulated data (constructed by Dr. Benis). Two of these three were simulated tamponade cases; the other was a simulated non-tamponade case. The fourth tuning set was based on real data from a tamponade case (certain missing values needed to be "filled in" due to gaps in the clinical records).

The program was tuned in an interactive fashion: it was run on the tuning data sets (Dr. Benis acted as the user), the results were analyzed (by Dr. Benis and myself), changes were made, the program was run again, and so forth. We ran the program approximately ten times (on each of the data sets). During this tuning process the production rules, temporal curves, characteristic functions, and  $\alpha$ s were all adjusted.

The program was then run off-line on the test data sets. The two sets consisted entirely of real data: a tamponade case, and a non-tamponade case.

The characteristic functions used for testing were as follows:

$$\mu_{CT}(x,y) = \max(CT_F(x) + CT_C(y)) \text{ where}$$

$$CT_F(x) = \begin{cases} 1 & \text{if } x \leq l_{CTF}/2 \\ (l_{CTF} - x)/(l_{CTF} - (l_{CTF}/2)) & \text{if } l_{CTF}/2 < x \leq l_{CTF} \\ 0 & \text{if } l_{CTF} < x \leq h_{CTF} \\ (x - h_{CTF})/(2 \cdot h_{CTF} - h_{CTF}) & \text{if } h_{CTF} < x \leq 2 \cdot h_{CTF} \\ 1 & \text{if } 2 \cdot h_{CTF} < x, \end{cases}$$

$$CT_C(y) = \begin{cases} 1 & \text{if } y \leq l_{CTC}/2 \\ (l_{CTC} - y)/(l_{CTC} - (l_{CTC}/2)) & \text{if } l_{CTC}/2 < y \leq l_{CTC} \\ 0 & \text{if } l_{CTC} < y \leq h_{CTC} \\ (y - h_{CTC})/(2 \cdot h_{CTC} - h_{CTC}) & \text{if } h_{CTC} < y \leq 2 \cdot h_{CTC} \\ 1 & \text{if } 2 \cdot h_{CTC} < y, \end{cases}$$

where  $x$  is the flow rate of CT,  $y$  is the cumulative volume of CT,  $l_{CTC}$  is a (time-varying) parameter,  $h_{CTC}$  is a (time-varying) parameter,  $l_{CTF}$  is a (time-varying) parameter, and  $h_{CTF}$  is a (time-varying) parameter. The parameter  $l_{CTF}$  represents the *minimum* expected chest tube flow; e.g., in a 70 kg patient, 40 ml in the first hour, 30 ml in each subsequent hour. The parameter  $h_{CTF}$  represents the *maximum* expected chest tube flow (250 ml in the first hour, 100 ml in each of the next seven hours, 50 ml/hour thereafter). These parameters are adjusted with respect to the patient's weight (the typical patient is considered to weigh 70 kg). The parameters  $l_{CTC}$  and  $h_{CTC}$  represent the minimum and maximum expected *cumulative* chest tube drainage, respectively (e.g., 250 max in the first hour, 350 max in the second hour, etc.); these parameters are adjusted with respect to the patient's weight. Note that each of these time-varying parameters is, of course, a temporal curve.

$$\mu_{UO}(x) = \begin{cases} 1 & \text{if } x \leq l_{UO}/2 \\ (l_{UO} - x)/(l_{UO} - (l_{UO}/2)) & \text{if } l_{UO}/2 < x \leq l_{UO} \\ 0 & \text{if } l_{UO} < x, \end{cases}$$

where  $x$  is the flow rate of UO and  $l_{UO}$  is a (time-varying) parameter. The parameter  $l_{UO}$  represents the *minimum* expected urine flow (80 ml in each of the first three hours, 60 ml in each of the next eight hours, 40 ml/hour thereafter); this parameter is adjusted with respect to the patient's weight.

$$\mu_{CI}(x) = \begin{cases} 1 & \text{if } x \leq .8 \\ 1 - (.1)(x - .8)/(1.6 - .8) & \text{if } .8 < x \leq 1.6 \\ .9 - (.9)(x - 1.6)/(2.4 - 1.6) & \text{if } 1.6 < x \leq 2.4 \\ 0 & \text{if } 2.4 < x, \end{cases}$$

where  $x$  is the value of CI.

$$\mu_{MAP}(x) = \begin{cases} 1 & \text{if } x \leq 40 \\ 1 - (.1)(x - 40)/(54 - 40) & \text{if } 40 < x \leq 54 \\ .9 - (.9)(x - 54)/(l_{MAP} - 54) & \text{if } 54 < x \leq l_{MAP} \\ 0 & \text{if } l_{MAP} < x, \end{cases}$$

where  $x$  is the value of MAP and  $l_{MAP}$  is a parameter.

$$\mu_{LAP-RAP}(x,y) = LR_1(x)(1 + \frac{3}{2}LR_2(y)) \text{ where}$$

$$LR_1(x) = \begin{cases} .4 & \text{if } x \leq 1 \\ (.4)(4 - x)/(4 - 1) & \text{if } 1 < x \leq 4 \\ 0 & \text{if } 4 < x, \end{cases}$$

$$LR_2(y) = \begin{cases} 0 & \text{if } y \leq h_{LAP} \\ 1 - (25 - y)/(25 - h_{LAP}) & \text{if } h_{LAP} < y \leq 25 \\ 1 & \text{if } 25 < y, \end{cases}$$

and where  $x = |LAP - RAP|$ ,  $y$  is the value of LAP, and  $h_{LAP}$  is a parameter.

The composite scoring function used for testing was:  $WS = \alpha_{CT}\mu_{CT} + \alpha_{UO}\mu_{UO} + \alpha_{CI}\mu_{CI} + \alpha_{MAP}\mu_{MAP} + \alpha_{LAP-RAP}\mu_{LAP-RAP}$  where  $\alpha_{CT} = .22$ ,  $\alpha_{UO} = .22$ ,  $\alpha_{CI} = .22$ ,  $\alpha_{MAP} = .12$ , and  $\alpha_{LAP-RAP} = .22$ .

### 5.2.1. The Tamponade Case

### **5.2.1.1. Initialization**

The following is a summary of the relevant information gathered by the questioning module: Patient # 5 was a 67 year old white male. His weight was 178 pounds. His past medical history was significant for coronary artery disease, preoperative cardiac dysfunction, and left sided congestive heart failure. The patient underwent a CABG procedure: four grafts were performed, and the LIMA was used for the grafts. The surgical procedure was uneventful.

The actions performed by the deductive module were as follows (for the sake of brevity not all of the initialization actions are described here; e.g., non-silent actions are not described): Based on the patient's above average weight, the temporal curves for chest tube drainage and urine output were shifted to higher expected ranges (1.15 times baseline).

The user was notified that "WARNING: Because the patient has had several CABG's, heavy bleeding may occur," that "WARNING: Because the patient had a LIMA cabg, may be predisposed to heavy bleeding," and, therefore, that "Chest tube drainage should be measured more frequently than typical patients during the first few hours."

The user was notified that "WARNING: Patient may be susceptible to renal problems because of significant cardiac dysfunction," and, therefore, that "urine output should be measured more frequently than typical patient during the first few hours."

Similarly, the user was notified that "WARNING: Because of significant preoperative cardiac dysfunction the patient may be predisposed to low CI,"

and, therefore, that "Cardiac output should be measured more frequently than typical patient during the first few hours."

The user was notified that "WARNING: map should be kept fairly low because patient may be predisposed to bleeding problems;" accordingly, the "lower limit" of MAP ( $l_{MAP}$ ) was set to 60 (the program's suggested limit of 55, and the user's suggested limit of 60 were reasonably close; therefore, the user's value was selected).

The user was notified that "EXPECT: lap to be high because of left sided chf," accordingly, the "higher limit" of LAP ( $h_{LAP}$ ) was set to 16 (in this case the program and the users' suggested limits were identical).

Overall, the deductive module performed well: each of the warnings was clearly appropriate, and the "limits" suggested by the computer were in close agreement with the values suggested by the user.

#### 5.2.1.2. Monitoring

During the first postoperative hour the patient's course was unremarkable: chest tube drainage was 80 ml, urine output was 270 ml, and the MAP, and LAP were "within normal limits." LAP and RAP were a bit close ( $|LAP-RAP| = 3$ ), accordingly  $\mu_{LAP-RAP}$  equaled .13. At 0100 WS was 0.02.

The second postoperative hour was also unremarkable with the exception of a slightly low CT (flow = 32 ml,  $\mu_{CT} = .14$ ). At 0200 WS was 0.06.

At time 0210 a cardiac index of 2.09 was measured; this (moderately) low value was reflected by  $\mu_{CI} = .34$ . At time 0300 WS was 0.16.

The fourth postoperative hour was unremarkable; WS remained steady

at 0.16.

During the fifth hour the patient's course began to deteriorate: the chest tube flow was high (320 ml) and the MAP was low at 58 ml. By 0500 *WS* had risen to 0.42.

During the sixth hour the chest tube drainage decreased to 60 ml. Urine output during this hour was low at 30 ml. At time 0600 *WS* equaled 0.49.

During the seventh hour the patient's course declined markedly: at 0623 *CI* was extremely low at 1.15, and *RAP* and *LAP* were equal at 13 ml. At 0700 the chest tube drainage was low (10 ml/hour), the urine output was low (10 ml/hour), and *WS* had risen to 0.73.

At 0730 the *MAP* was 42, and *WS* had risen to 0.85. At this time the patient was returned to surgery.

Two aspects of this case make the recognition of tamponade challenging. First, the pattern of chest tube drainage was unusual, ranging from extremely low (10 ml/hour) to extremely high (320 ml/hour). In retrospect, this pattern was observed because of a clog in the chest tube. Second, although the *RAP* and *LAP* were equal, the readings were not elevated (7 mm early in the postoperative course, 13 mm later). This unusual presentation was a result of the fact that the patient was suffering from hypovolemia as well as tamponade.

The scoring function, *WS*, was initially very low at 0.02, by the fifth hour had risen to 0.49, and by 0730 was quite high at 0.85. This behavior closely paralleled the clinical presentation of the patient and seems quite reasonable. This behavior was especially encouraging in view of the challenging nature of the case.

It is worth noting that the management of this patient may have been less than optimal. In particular, the chest tube drainage and urine output should probably have been measured more frequently, and the cardiac index certainly should have been measured more frequently. In fact, the program made exactly these suggestions during the initialization phase. Also note that the program could have prompted for additional measurements during the monitoring phase (although this has been only partially implemented); presumably, additional measurements could have hastened the recognition of tamponade.

At time 0730 (when monitoring was terminated)  $\int \mu_{CT} = 138$ ,  $\int \mu_{CI} = 307$ ,  $\int \mu_{MAP} = 113$ , and  $\int \mu_{LAP-RAP} = 215$ . These values certainly do *not* reflect a prototypical sequence. Note, however, that due to the confounding factors noted above (i.e., hypovolemia and the clogged chest tube) the patient's course was, in fact, not typical for tamponade (e.g., RAP and LAP were *not* elevated).

## 5.2.2. The Non-Tamponade Case

### 5.2.2.1. Initialization

The following is a summary of the relevant information gathered by the questioning module: Patient # 6 was a 60 year old white female. Her weight was 146 pounds. Her past medical history was significant for coronary artery disease and left sided congestive heart failure. The patient underwent a CABG procedure: four grafts were performed, and the LIMA

was used for the grafts. The surgical procedure was uneventful.

The actions performed by the deductive module were as follows (for the sake of brevity not all of the initialization actions are described here; e.g., non-silent actions are not described): Based on the patient's below average weight, the temporal curves for chest tube drainage and urine output were shifted to lower expected ranges (0.95 times baseline).

The user was notified that "WARNING: Because the patient has had several CABG's, heavy bleeding may occur," that "WARNING: Because the patient had a LIMA cabg, may be predisposed to heavy bleeding," and, therefore, that "Chest tube drainage should be measured more frequently than typical patients during the first few hours."

The user was notified that "WARNING: map should be kept fairly low because patient may be predisposed to bleeding problems;" accordingly, the "lower limit" of MAP ( $l_{MAP}$ ) was set to 60 (the the program's suggested limit of 55, and the user's suggested limit of 60 were reasonably close; therefore, the user's value was selected).

The user was notified that "EXPECT: lap to be high because of left sided chf," accordingly, the "higher limit" of LAP ( $h_{LAP}$ ) was set to 16 (in this case the program and the users' suggested limits were identical).

Overall, the deductive module performed well: each of the warnings was clearly appropriate, and the "limits" suggested by the computer were in close agreement with the values suggested by the user.

### 5.2.2.2. Monitoring

The patient's postoperative course was uneventful. Several of the chest tube readings were somewhat abnormal (both low and high flow rates were measured), and in one instance MAP was somewhat low; however, the vast majority of measurements were "within normal limits." WS ranged between 0.0 and 0.13 (mainly 0.0); this behavior reflects the stable nature of the postoperative course and seems quite reasonable.

At time 1130 (when monitoring was discontinued)  $\int \mu_{CT} = 132$ ,  $\int \mu_{MAP} = 5$ , and the other  $\int \mu_s$  equaled 0. These values do seem to reflect the patient's course; i.e., relatively uneventful except for some abnormalities in CT.

More complete details of both the tamponade and non-tamponade test cases can be found in the section of this thesis entitled *Details of Results*.

These performance tests were by no means rigorous, hence, it is not possible to state conclusively that our approach is valid. Overall, however, the program certainly displayed reasonable behavior; further testing seems warranted.

### 5.3. Experiment II: Test of the Initialization Interface

If a program based on our approach is to be suitable for use in the CSICU, then it must mesh with the existing clinical protocol [Draz80, Shep79]. The main consideration in regards to our system is that it cannot require a significant increase in manual input. The clinical staff already

enters much data manually (e.g., CT, UO, CI), hence, the monitoring module imposes no additional burden. In order to initialize the system, however, input is required for the questioning and deductive modules (the *initialization interface*); thus, it is essential that this interface be acceptable to the clinical staff members.

In order to test the initialization interface a staff nurse from the CSICU (Ms. Barbara E. de Vries) was selected. Ms. de Vries's computer background was quite limited (a course in high school approximately ten years previously, and experience entering and recalling data from the existing CSICU system).

The purpose of the program was briefly explained to Ms. de Vries. With no further instruction she then proceeded to interact with the initialization interface: she was able to respond to all questions, and to examine the deductions. The entire process was completed in less than fifteen minutes. Ms. de Vries reported that she felt comfortable with the interface, and would have no hesitation in using it in the CSICU.

**CHAPTER 6****DISCUSSION****6.1. Comparison With Previous Work**

The type of problem that we have selected differs in several respects from the problems addressed by most other AIM systems. In particular, our system must run continuously and on-line, whereas a diagnostic system, for example, needs to run only once and can be off-line. In addition, our system must acquire the majority of its data directly from transducers, and can expect only a limited amount of input from, and output to, the user; a diagnostic system is not similarly constrained.

Our approach has similarities to both first and second generation AIM. In particular, most of the knowledge is "shallow" (e.g., there is no causal reasoning); this is first generation in nature. The idea of having separate modules for each complication is, however, more second generation in nature (e.g., MDX uses a similar architecture).

The AIM system most closely related to ours is VM. There are some substantial resemblances: both systems are intended to run continuously, and both were designed to be superimposed on Osborn's program. The systems differ, however, in several respects. In particular, our system is concerned with the generation of alarms, whereas VM attempts to give advice. In addition, the approaches differ significantly: VM is frame based, and it considers all values to be either normal or abnormal.

There has been some previous work using fuzzy sets in medicine [Adla82], but most of this work differs substantially from ours in terms of objectives and the methods used. Esogbue and Elder [Esog79] do propose using characteristic functions to represent the abnormality of a measured value, however, these authors do not provide an explicit method (i.e., parameters) for customizing these functions (although they do briefly allude to using "more than one aspect" to determine degrees of abnormality).

## **6.2. Evaluation of the Approach**

### **6.2.1. The Fundamental Strategy**

Our fundamental strategy is to select the "important" clinical complications, and to disregard the others (i.e., build a set of "recognizers" rather than a diagnostic system). An immediate and powerful benefit of this strategy is that the problem space is greatly simplified. A potential drawback is the loss in generality since the system is ignorant of many complications. However, in environments where only a small number of complications are likely to occur, this drawback is less significant; the CSICU is such an environment.

A diagnostic system can be viewed as a *consultant* (i.e., *replacement* for the clinician). Our system is better viewed, however, as an *assistant* (i.e., *supplement*): the clinician still performs diagnosis, makes decisions, administers therapy, and so forth. What our system can provide is, in effect, an extra layer of protection—it can watch for those complications that might be

missed due to human error, ignorance, or fatigue.

There are several advantages in designing an assistant rather than a consultant. One advantage is that the acceptable level of performance can be somewhat lower. For example, even if a supplement detects only 10% of the clinical complications then it still has the potential to be of value (assuming that the false alarm rate is kept low, and assuming that the clinician might have missed some of the complications). On the other hand, a consultant that performs correctly only 10% of the time is of no value.

An assistant program is also more likely to be accepted by potential users than is a consultant program: Consultants carry the implied threat that they might replace the clinician; assistants do not.

### **6.2.2. Characteristic Functions as a Knowledge Representation**

Characteristic functions play an important role in our approach: they serve as a fundamental knowledge representation. Several benefits are derived by using these functions to represent knowledge:

- It is common to convey clinical knowledge in terms of distinguished points; e.g., "the LAP is expected to range between 4 and 14." Characteristic functions can capture this type of knowledge.
- Characteristic functions provide a natural representation of the degree of pathology of a measured value.
- The functions do not give rise to discontinuous behavior at the distinguished points.

- The functions can be customized and adjusted by means of parameters.
- The functions allow for a finer granularity when combining information (e.g., raw data is not converted into Boolean values).
- The functions act as a normalization. This normalization is useful when combining information that is measured on different numerical scales or in different units, and also allows for the construction of the sequencing technique.
- The functions are straightforward to implement for use in a computer system.

Although characteristic functions can eliminate discontinuous behavior at the distinguished points, the functions do not entirely eliminate the need for cutoffs. For example, if  $WS$  is to be used to trigger a binary alarm (e.g., *on* or *off*) then it is necessary to choose a number,  $C$ , such that the alarm is *on* if  $WS > C$ , and the alarm is *off* if  $WS \leq C$ . It is apparent that the number  $C$  is itself a cutoff point.

Note, however, that because of our approach it is necessary to choose only one cutoff point rather than many (i.e., cutoffs are not needed for each single variable and multivariable combination). In addition, note that it is possible to keep the adjustment of sensitivity and specificity (which can be adjusted by modifying  $C$ ) separate from the notion of the degree of pathology of the variables (as reflected by the characteristic functions).

Cutoffs can also arise in the sequencing technique. For example, if  $\int \mu_A \approx \int \mu_B$  then it may be necessary to select a  $\delta$  such that if  $\int \mu_A = \int \mu_B + \epsilon$ , then it is determined that event  $A$  precedes event  $B$ , if and only if  $\epsilon \geq \delta$ . Similarly, a threshold may be necessary when matching an actual event

sequence against a prototypical sequence.

The need to choose a cutoff inevitably arises when a continuous valued quantity (e.g., *WS*) is mapped into a binary quantity (e.g., alarm *on* or *off*). It is possible to eliminate this need by making the alarm signal continuous. For example, rather than a buzzer being either *on* or *off*, it could range along a spectrum from inaudible to loud. Although this might be viable in certain environments, it is probably not suitable for clinical alarms (although a three valued alarm, e.g., *low*, *medium*, and *high*, might be useful).

Our sequencing technique has several advantages over the conventional technique; however, it is still subject to certain limitations. For example, the technique, as currently conceived, is applicable only for variables which make a one-way transition; e.g., "good" -> "bad." In the domain that we have chosen (medical monitoring) this is not a major difficulty since many clinical problems can be described in terms of such transitions; this limitation might be more significant in other domains, however.

The technique breaks down when applied to bizarre patterns; see Fig. 6.1., for example.




Fig. 6.1. Bizarre patterns.

Note, however, that these are exactly the types of situations in which physicians find it difficult to order events. It might be desirable to augment the technique such that it checks for, and eliminates bizarre patterns.

The technique increases in accuracy over time; for example, in Fig. 6.2. the spike in  $A$  is not recognized as such until significantly later than time  $t_1$  ( $\int \mu_A > \int \mu_B$  at time  $t_1$ , but  $\int \mu_B > \int \mu_A$  at time  $t_2$ ).




Fig. 6.2. Hindsight is 20-20.

Of course, this type of behavior is certainly not unique to our technique: hindsight is indeed 20-20!

Although the technique can order events, it cannot determine whether events have actually occurred. For example, in Fig. 6.3.  $\int \mu_A > \int \mu_B$ , but it is not clear whether events  $A$  or  $B$  have actually occurred. It is probably not possible to overcome this limitation in a completely general way, however, certain heuristics might be useful; for example, a sliding window of  $\int \mu$  over some time interval.

Although characteristic functions are borrowed from fuzzy set theory, we have not used the standard fuzzy set operators. Fuzzy set theory has




Fig. 6.3. Have events *A* or *B* occurred?

been criticized on the grounds that its standard operators are inappropriate for solving real-life problems, and, conversely, that those operators which are useful in practice are ad hoc (Appendix I). We have demonstrated that one of the fundamental concepts from fuzzy set theory—the notion of a characteristic function—can be useful even if fuzzy set operators are not utilized. In particular, we have demonstrated that characteristic functions can serve as a fundamental knowledge representation for knowledge based systems.

A potential criticism of characteristic functions is that they are in some sense "arbitrary." It seems to us, however, that these functions are less arbitrary than cutoff points, since cutoff points inevitably lead to counter-intuitive behavior.

### 6.2.3. The Program

It proved to be straightforward to construct the program using our approach. For example, characteristic functions were easily defined in con-

sultation with a physician from the CSICU, and production rules and temporal curves were easily acquired by discussion with the clinical personnel and examination of medical texts. The ease with which the necessary knowledge was acquired and the ease with which the program was constructed is encouraging—it suggests that our knowledge representations are appropriate, and that our approach is sound.

The approach is quite modular and therefore the program is modular as well. As a result of this modularity we found that it was a simple matter to adjust the program for the purpose of tuning its performance. This is also encouraging.

#### 6.2.4. The Results

The performance tests run on the program were by no means rigorous, hence, it is not possible to state conclusively that our approach is valid. However, the test results were certainly favorable; this *suggests* that the approach is sound. Further testing seems warranted.

The approach requires only a slight increase in the amount of manual input (specifically, the initialization information must be entered); this input could be performed before the patient returns from surgery. The approach does not require any substantial modifications to the established clinical protocol. The test of the initialization interface yielded favorable results. In view of the above, we feel that the approach, if properly implemented, would be suitable for use in the CSICU.

### **6.2.5. Generality of the Approach**

Certain aspects of our approach are domain dependent. For example, expected physiologic values may differ from hospital to hospital depending upon the specific clinical protocols in use (e.g., which drugs are administered). More generally, certain aspects of the approach may be invalid in other environments; for example, it may not be possible to describe temporal curves for certain variables. However, the basic framework that we have proposed—in particular, the use of characteristic functions—would seem to have broad applicability in many domains.

**CHAPTER 7****COMMENTS AND FURTHER WORK****7.1. Adapting the Program For Use in the CSICU**

Our program was designed as a test of the feasibility of our approach, not as a fully functional alarm system; however, it was also constructed to be compatible with the existing clinical protocol in the Mount Sinai CSICU. In order to make the program suitable for use in the CSICU some additional work will be necessary:

Certain functions in the program are currently simulated, for example, data entry; these functions will need to be recoded, transducers will need to be interfaced<sup>1</sup>, and so on. There are several minor bugs in the program; these bugs are noted in the source code and should be corrected. In addition, it is probably desirable to rewrite certain sections of the program into a language more efficient than LISP.

Further programming, tuning, and testing will be necessary. In particular, *WS* may need to be adjusted, a value *C* will need to be selected if *WS* is to be mapped into a two-valued alarm, the actual and prototypical event sequences will need to be matched, and *WS* and the sequencing information will need to be integrated. Careful tuning and testing will be necessary in order to achieve a good balance between sensitivity and specificity.

In any alarm system there is an inevitable tradeoff between sensitivity and specificity: As the system begins to detect more true alarms (increased

sensitivity) it also begins to exhibit more false alarms (decreased specificity). Conversely, as the false alarm rate is lowered (increased specificity) more true alarms are missed (decreased sensitivity).

Physicians and nurses working in the CSICU at the Mount Sinai Hospital have indicated that although they are concerned with both sensitivity and specificity, the latter needs to be emphasized. Their rationale is the following: It is crucial to keep the false alarm rate as low as possible because any system with an excessive false alarm rate will be disconnected by aggravated staff members. Moreover, due to the high level of staffing in the unit, an alarm system is viewed as being a *supplement* to the clinician (i.e., a fail-safe that looks for things that the clinical staff is also watching for, but might miss due to error, ignorance or fatigue), and therefore it is acceptable that it miss a few true alarms.

Parenthetically, in a different clinical environment, it might be advisable to emphasize sensitivity. For example, in a unit that is less heavily staffed than the CSICU, an alarm system might be viewed as a *replacement* for the clinician; hence, missing true alarms would not be acceptable. This is, in fact, the choice that is made in coronary care units: alarm systems relieve nurses of the burden of continually following EKG tracings, but such systems have notoriously high false alarm rates [Long83Front]. Note that should the CSICU staff begin to rely on the alarm system to such an extent that it becomes more than a supplement, then it might be necessary to readjust the sensitivity and specificity.

One potential source of difficulty for the program is that the  $\mu$ -value for each variable is computed from the last measured value, *no matter how old*;

this can be quite misleading (e.g., a CI from several hours ago). Note that this generally will be a problem only for variables that are measured by hand (since automatically acquired variables should be measured with sufficient frequency unless there is a transducer error or the like). Clinical protocols (and judgement of the clinician) dictate the frequency at which the variables *should* be measured by hand (for example, *UO* should be measured at least every hour, and more frequently if the patient is not doing well); these protocols are not always followed, however.

A potential solution is to have the computer prompt for new measurements when values are outdated (or when otherwise appropriate). The rudiments of such a facility have, in fact, been incorporated into the program; the necessary code could be easily inserted. This is at best a partial solution, however, since the requests cannot be enforced. A better solution would be to have the system automatically measure all values.

Another potential source of difficulty is errors in the input data. The program was tested on relatively clean data (i.e., it contained few erroneous values). In actual practice, a significant portion of the data may be incorrect; the program must be able to handle this. The program does have a facility for screening out erroneous values; however, this facility is quite rudimentary at present (e.g., simple range checks). More sophisticated screening might be desirable; for example, values could be screened in terms of their deviation from previous values (e.g., using exponential weighting), or values might need to be out of range for a minimum period of time. In addition, values could be screened in a multivariate fashion; for example, if the CI is normal, then the MAP cannot be extremely low. The danger in screening

too aggressively, however, is that a genuinely abnormal value might be labeled as erroneous and discarded.

A method for minimizing data errors that is specific to this domain is to prevent data acquisition at times when momentary physiologic perturbations are likely to occur. For example, drug or fluid administration, suctioning, blood withdrawal, and movement of the patient, can all temporarily affect the hemodynamics; hence, the LAP, RAP, and MAP should not be measured when the above actions are taking place.

## **7.2. Further Research**

One area for further research is to extend the program to handle additional clinical complications; for example, hypovolemia and myocardial failure. Our approach is quite modular (and the program is as well) so this should be relatively straightforward. It also would be interesting to apply the approach in other domains; for example, non-clinical environments.

Another area for further research is the sequencing technique. For example, it might be possible to expand the technique to handle multistep transitions. In addition, heuristics could be developed for identifying bizarre patterns, and for determining whether events have actually occurred.

## **7.3. A Clinical Observation**

Most patients who are suffering from early post-surgical tamponade are receiving mechanical ventilatory support, and therefore would be expected to

exhibit reversed pulsus paradoxus. However, the literature generally reports that tamponade patients exhibit pulsus paradoxus. Since it is easy to confuse pulsus paradoxus and reversed pulsus paradoxus, this suggests that many of the clinical observations and reports in the literature are in error—further investigation may be warranted.

#### Footnotes

- 1 Highly accurate (and well calibrated) transducers will be especially essential for LAP and RAP because even slight imprecision in measurements could lead to a large perturbation in  $\mu\text{LAP}-\text{RAP}$ .

## SAMPLE CODE

### 1. The Questioning Module

#### 1.1. Sample Questions

```

(setq initq
 '(
  ("What is the patients name?(enclose in double quotes) "
   "Name: "
   ()
   (stringp response)
   name )

  ("Age? "
   "Age: "
   (
    "This system is intended for patients age 21-75 only")
    (cond ( (not(fixp response))
            (terpri)
            (patom "Should be an integer - e.g., 56")
            nil)
          ( (not(and(greaterp response 20)(lessp response 76)))
            (terpri)
            (patom "Please repeat the age: ")
            (setq repeat (read))
            (cond ((eq response repeat)
                  (terpri)
                  (patom "This system handles only adults(age 21 ")
                  (patom " - 75).")
                  (terpri)
                  (patom "System will not accept your patient. ")
                  (terpri)
                  (throw 'ques-prob tag1) )
              ( t nil) ) )
          (t))
    age )

  . . .

  ("Did the patient have a CABG? "
   "CABG? "
   ()
   (memq response '(yes no y n))
   CABG )

```

("Did the patient have any other type of cardiovascular surgery? "

"Other cv surgery? "

("

    Please note that this system has been designed to monitor only  
    "" patients who have had aortic and/or mitral and/or CABG surgery")

(cond ((not (memq response '(yes no y n))) nil)

    ((memq response '(yes y))

        (terpri)

        (patom "Please repeat your answer: ")

        (setq repeat (read))

        (cond ((memq repeat '(yes y))

            (terpri)

            (patom "This system handles only patients with mitral ")

            (patom "and/or aortic and/or CABG surgery")

            (terpri)

            (patom "System will not accept your patient")

            (terpri)

            (throw 'ques-prob tag1))

        (t nil)))

    (t)

othercardsurg )

("Has the patient had cardiac surgery before? "

"Previous cardiac surgery? "

("

    Repeat cardiac surgical cases often have heavy bleeding postop")

(memq response '(yes no y n))

reop )

("How many CABGs did the patient have in this procedure? "

"# of CABGs: "

("

    There is often a correlation between the number of CABGs  
    "" and the postoperative bleeding rate")

(memq response '(1 2 3 4 5 6))

num-of-CABG

(or (eq CABG 'yes) (eq CABG 'y)) )

("Were any of the CABGs a LIMA? "

"LIMA? "

("

    Often patients with a LIMA CABG have heavy bleeding postop")

(memq response '(yes no y n))

lima

(or (eq CABG 'yes) (eq CABG 'y)) )

("Did the patient have excessive bleeding in the OR? "

"Excessive OR bleeding? "

("

    Patients with excessive bleeding in the OR often have heavy  
    "" bleeding postoperatively")

```
(memq response '(yes no y n))
exces-bleed-OR )
```

```
("Was the patient on cardiopulmonary bypass for more than three hours? "
"Bypass more than three hours? "
("
  Patients who were on bypass for long periods of time often have
  " " heavy postoperative bleeding")
(memq response '(yes y no n))
long-bypass)
```

```
("Does the patient have a known coagulopathy? "
"Coagulopathy? "
("
  Patients with coagulation problems preoperatively are often
  " " subject to high postoperative bleeding rates. Especially
  " " significant are the PT,PTT,platelets,fibrinogen and factors I-XII")
(memq response '(yes no y n))
coagulopathy )
```

...

## 1.2. The Main Question Asking Function

```
; Getinfo is a function that is used to pose a series of questions
; to the user. Getinfo works as follows:
; It moves sequentially down the question list. If the question
; has a sixth component, and the component evaluates to nil, then the
; question is skipped. Otherwise, the query is posed to the user by the
; Query function.
; The queries and responses are then displayed back to the user for
; final confirmation. If the user confirms, then Getinfo returns; if the
; user disconfirms (anything besides yes or y), then Getinfo calls
; Getinfochange and a query in the question list can be posed to the user
; again. The list qsasked keeps track of which questions from the list
; are posed to the user. It is needed by the confirmation phase because
; not all questions will have been asked.
(defun Getinfo (qlist)
  ; Ask the questions
  (prog (temp count)
    (setq temp qlist) ; temp moves thru the question list
    (setq count 0) ; # of the question currently being asked
    (setq qsasked nil)
    (terpri)
    (terpri)
    continue
    (cond ((null temp) (return))))
```

```

(setq count (add1 count))
(cond ( (and (neq (cddddar temp) nil)(eq(eval(cddddar temp))nil))
      (setq temp (cdr temp))
      (go continue))
      (t (Query (car temp))
         (setq qsasked (cons count qsasked))
         (setq temp (cdr temp))
         (go continue))))
; Final confirmation
(prog ()
 (terpri)
 (terpri)
 (patom "-----")
 (cond ((eq experienced 'no)
        (terpri)
        (terpri)
        (patom "All queries and responses will now be summarized"))))
label
(terpri)
(prog (temp count)
 (setq temp qlist) ; temp moves thru qlist
 (setq count 1)    ; # of the question
 continue
 (cond ((null temp)(return)))
 (cond ( (not(memq count qsasked))
         (setq temp (cdr temp))
         (setq count (add1 count))
         (go continue) ) )

 (terpri)
 (patom "(" )
 (patom count)
 (patom ") ")
 (patom (cddddar temp))
 (patom ": ")
 (patom (eval (cddddar temp)))
 (setq temp (cdr temp))
 (setq count (add1 count))
 (go continue))
(terpri)
(terpri)
(cond
 ((eq experienced 'no)
  (patom "Are you happy with the responses?(yes or no) "))
 ((eq experienced 'yes)
  (patom "Happy? ")))
(setq response(read))
(cond ( (not (memq response '(yes y)))
      (Getinfochange qlist) (go label) )
      (t (return)))))

```

## 2. The Deductive Module

### 2.1 Sample Production Rules

```
(setq sinita-CT
```

```
  (
    (setq read-CT '((-1 0.0)) )
    (setq time-CT -1)
    (setq val-CT 0.0)
    (setq mu-CT 0.0)
    (setq cum-mu-CT 0.0)
    (setq msgs-CT nil)
    (setq W-CT-high 0)
  )
)
```

```
(setq inita-CT
```

```
  (
    "The following actions relate to chest tube drainage"

    (I-CT-1
      (if (memq reop '(yes y))
          then (setq W-CT-high (add1 W-CT-high)))
      ("WARNING: Because the patient is a reop, may have heavy bleeding")
      ())
    (I-CT-2
      (if (and(memq CABG '(yes y)) (greaterp num-of-CABG 2))
          then (setq W-CT-high (add1 W-CT-high)))
      ("WARNING: Because the patient has had several CABG's, heavy
       " " bleeding may occur")
      ())
    (I-CT-3
      (if(and(boundp 'lima)(memq lima '(yes y)))
          then (setq W-CT-high (add1 W-CT-high)))
      ("WARNING: Because the patient had a LIMA cabg, may be
       " " predisposed to heavy bleeding")
      ())
    (I-CT-4
      (if (memq exces-bleed-OR '(yes y))
          then (setq W-CT-high (add1 W-CT-high)))
      ("WARNING: Because patient had excessive bleeding in the OR, may
       " " be predisposed to heavy postop bleeding")
      ())
    (I-CT-5
      (if (memq long-bypass '(yes y))
          then (setq W-CT-high (add1 W-CT-high)))
      ("WARNING: Because patient was on cardiopulmonary bypass for an
       " " extended period of time, may also be predisposed to heavy
       " " postoperative bleeding")
    )
  )
)
```

```

    ())
(I-CT-6
  (if (memq coagulopathy '(yes y))
      then (setq W-CT-high (add1 W-CT-high)))
  ("WARNING: Because patient has a coagulopathy, may be predisposed
   " " to heavy postop bleeding")
    ())
(I-CT-7
  (if (greaterp W-CT-high 0)
      then t)
  ("Chest tube should be measured more frequently than typical patient
   " " during the first few hours")
  (" Since there seems to be an above normal chance of abnormal bleeding
   " " rate it should be watched carefully")
  ))

```

## 2.2. The Rule Interpreter

```

; Perform is a function that processes the initialization actions
; from a given action list. As noted, the actions are in the form of
; production rules, therefore, Perform is essentially a production
; system interpreter. Some observations are in order about the
; production system:
; 1) The rules seem to be commutative in the sense of Nilsson (pg 35)
;    (intuitively, with certain constraints, the order of rule
;    firings should not make any difference to the final outcome).
; 2) Rules need fire only once (if at all).
; 3) There will be only a few multistep reasoning chains (there will
;    certainly be some, however; for example, determining whether a
;    users suggested alarm limits seem to be reasonable).
; In view of these observations, it is sufficient to have the
; production system execute as follows: simply move down the action
; list a step at a time - if a rule antecedent is non nil then fire the
; consequents, otherwise don't; continue until the end of the list.
; Note that the action lists have been ordered so as to make the
; sequence of deductions seem logical to the user (and also to make
; sure that all necessary rules can fire). Of course, since the
; order of the rules is important, this means that new rules must
; be added with care.
; Note that in our implementation the ordering of the
; action lists themselves is important.
; For example, I-MAP-2 needs
; information about W-CT-high, and therefore the inita-CT list must
; be executed before the inita-MAP list. Similarly, I-RAP-7 needs
; information about E-LAP and therefore the inita-LAP list must be
; executed before the inita-RAP list.
; These dependencies could be eliminated, but at the cost of an
; increase in the length of the code.

```

```

; In addition to firing rules, Perform will also print out an English
; description of each action if it is fired, and keeps track of which
; actions have been fired. Finally, after the entire
; action list has been exhausted, the user may request explanations
; as to why particular deductions were made, and so forth.
; Note that one weakness of this phase is as follows: Some of the
; initialization actions evaluate user suggested limits. If the limits
; are deemed inappropriate by the computer, the user has a choice of
; 1) using the computers suggestion or 2) using his original suggestion.
; Ideally, there should be (but is not) a third option - the user entering
; a new suggestion. While this is certainly possible, it would entail either
; major restructuring of the code, or some inelegant "patches", and hence
; it has not been implemented.
; Note that one potential problem (although unlikely) is as follows:
; if the user types more than one character after the readc's then
; an error will be generated.
(defun Perform (alist)
  ; if rule antecedents evaluate to non nil then fire consequents,
  ; print out the description and save the rule # in fired
  (setq fired nil) ; fired is a list of the #'s of actions which have fired
  (prog (temp count)
    (setq temp (cdr alist)) ; temp move through the action list
    (setq count 0) ; the # of the action which is currently being examined
    (terpri)
    (patom "-----")
    (terpri)
    (patom (car alist))
    (terpri)
    (terpri)
    continue
    (setq count(add1 count))
    (cond ((null temp) (return)))
    (cond ((eval (cadar temp))
           (mapc (function patom)(mapcar (function eval)(caddar temp)))
           (terpri)
           (patom ' )
           (patom (caar temp))
           (patom ' )
           (readc)
           (readc)
           (setq fired (cons count fired))
           (terpri) ))
          (setq temp (cdr temp))
          (go continue))
  ; explanations for deductions
  (prog(temp)
    continue
    (terpri)
    (patom "More information? (rule# or no) ")
    (setq response(read))
    (terpri)

```

```

(cond ((memq response '(no n) ) (return))
      ((not (numberp response))
        (terpri)
        (patom "Illegal response")
        (terpri)
        (go continue))
      ((not (memq response fired))
        (terpri)
        (patom "No such rule was fired")
        (terpri)
        (go continue))
      (t (setq temp (nthcdr response alist))
        (cond ((null (caddr temp))
          (patom "Sorry - no additional information available")
          (terpri)
          (go continue))
              (t (mapc (function patom)(caddr temp))
                (terpri)
                (go continue))))))

```

### 3. The Monitoring Module

#### 3.1. Sample Functions

```

; Enter input data for CT: According to clinical
; protocol, the clinician will be entering observations periodically.
; It is assumed here, that CT is entered as the cumulative volume since
; time 0. Since the the clinician may make an observation of CT at one
; time, but not enter that volume into the system until much later, both
; the volume, and the time of observation must be entered.
; The volume is stored in memory location 2, and the time of observation
; is stored in memory location 1.
; [Enter a value and time of observation from the input file.
; Note that "float" is used to guarantee that the values are real
; (otherwise might be manipulated as fixed and this can lead to
; truncations).]

```

```

(defun Ent-CT ()
  (cond ( ( and (neq point-CT nil) ( <= (Tconvert (caar point-CT)) CLK) )
        (setq mem1-CT (caar point-CT))
          (setq mem2-CT (float (cadar point-CT)))
          (setq point-CT (cdr point-CT) ) ) )

```

```

; The updating function for CT is structured as
; follows:
; Determine if a new reading has been entered

```

```

; IF a new reading has been entered
; THEN screen the reading
;   IF reading screened successfully
;   THEN compute the current mu for the variable
;         update the cumulative mu for the variable
;         add the new reading to the list of readings
; ELSE Determine whether a new reading is overdue
(defun Upd-CT ()
  (setq reading (Rd-CT))
  (setq time (Tconvert(car reading)))
  (setq val (cadr reading))
  (setq new (not (equal time time-CT)))
  (cond ((eq new t)
        (setq screen (Scrn-CT (list time val)))
        (cond ((eq screen nil)
              (setq prev-mu mu-CT)
              (setq mu-CT (Mu-CT time val))
              (setq cum-mu-CT
                    (Comp-cum time-CT prev-mu cum-mu-CT time mu-CT))
              (setq read-CT (cons (list time val) read-CT))
              (setq val-CT val)
              (setq time-CT time)
              (terpri)
              (patom (list 'CT val-CT 'time (Ttconvert time-CT)
                          'mu (trunc mu-CT)
                          'cum-mu (Ttconvert (fix cum-mu-CT))))))
        (t (Newrd-CT) ) ) )

; Compute how pathologic CT is
; Mu's are computed for 1) cumulative volume
;   2) flow rate
; The "worst" of these is returned (i.e., max)
(defun Mu-CT (time val)
  ; compute mu for cumulative volume
  (setq cum-max (Cum-max-CT time))
  (setq cum-min (Cum-min-CT time))
  (setq mu-cum
    (cond ( (eq time 0) 0.0)
          ( (<= val (times 0.5 cum-min)) 1.0 )
          ( (<= val cum-min)
            (quotient (diff cum-min val)
                      (diff cum-min (times 0.5 cum-min))))
          ( (<= val cum-max) 0.0 )
          ( (<= val (times 2.0 cum-max))
            (quotient (diff val cum-max)
                      (diff (times 2.0 cum-max) cum-max)))
          (t 1.0)))
  ; compute mu for flow rate
  (setq flw (diff val (cadr read-CT)))
  (setq flw-max (Flw-max-CT (caar read-CT) time))
  (setq flw-min (Flw-min-CT (caar read-CT) time))

```

```

(setq mu-flw
  (cond ( (eq time 0) 0.0)
        ( (<= flw (times 0.5 flw-min)) 1.0 )
        ( (<= flw flw-min)
          (quotient (diff flw-min flw)
                    (diff flw-min (times 0.5 flw-min))))
        ( (<= flw flw-max) 0.0 )
        ( (<= flw (times 2.0 flw-max))
          (quotient (diff flw flw-max)
                    (diff (times 2.0 flw-max) flw-max)))
        (t 1.0)))
; return maximum
(setq result (max mu-cum mu-flw))
(setq result (sum grad-score elev-LAP))

; Return (normalized) maximum expected cumulative chest tube drainage.
; Note that, logically, the sum of the maximum flow rates over time
; should be => maximum cumulative volume.
; This is not the case for CT, so some adjustments should probably be made.
(defun Cum-max-CT (time)
  (times normfact
    (cond
      ( (<= time 60) (diff 250 (times 250 (quotient (diff 60 time)60.0))))
      ( (<= time 120) (diff 350 (times 100 (quotient (diff 120 time)60.0))))
      ( (<= time 180) (diff 450 (times 100 (quotient (diff 180 time)60.0))))
      ( (<= time 240) (diff 550 (times 100 (quotient (diff 240 time)60.0))))
      ( (<= time 300) (diff 650 (times 100 (quotient (diff 300 time)60.0))))
      ( (<= time 360) (diff 750 (times 100 (quotient (diff 360 time)60.0))))
      ( (<= time 420) (diff 850 (times 100 (quotient (diff 420 time)60.0))))
      ( (<= time 480) (diff 950 (times 100 (quotient (diff 480 time)60.0))))
      ( (<= time 540) (diff 1050 (times 100 (quotient (diff 540 time)60.0))))
      ( (<= time 600) (diff 1150 (times 100 (quotient (diff 600 time)60.0))))
      ( (<= time 660) (diff 1250 (times 100 (quotient (diff 660 time)60.0))))
      ( (<= time 720) (diff 1350 (times 100 (quotient (diff 720 time)60.0))))
      ( (<= time 780) (diff 1400 (times 50 (quotient (diff 780 time)60.0))))
      ( (<= time 840) (diff 1450 (times 50 (quotient (diff 840 time)60.0))))
      ( (<= time 900) (diff 1500 (times 50 (quotient (diff 900 time)60.0))))
      ( (<= time 960) (diff 1550 (times 50 (quotient (diff 960 time)60.0))))
      ( (<= time 1020) (diff 1600 (times 50 (quotient (diff 1020 time)60.0))))
      ( (<= time 1080) (diff 1650 (times 50 (quotient (diff 1080 time)60.0))))
      ( (<= time 1140) (diff 1700 (times 50 (quotient (diff 1140 time)60.0))))
      ( (<= time 1200) (diff 1750 (times 50 (quotient (diff 1200 time)60.0))))
      ( (<= time 1260) (diff 1800 (times 50 (quotient (diff 1260 time)60.0))))
      ( (<= time 1320) (diff 1850 (times 50 (quotient (diff 1320 time)60.0))))
      ( (<= time 1380) (diff 1900 (times 50 (quotient (diff 1380 time)60.0))))
      ( (<= time 1440) (diff 1950 (times 50 (quotient (diff 1440 time)60.0)))) ) ) )

; Return (normalized) maximum expected chest tube drainage between
; time T1 and time T2.

```

; 250 ml max is expected in first hour and  
 ; 100 ml max is expected in next seven hours  
 ; 50 ml max is expected in each subsequent hour.

```
(defun Flw-max-CT (T1 T2)
  (times normfact
    (diff
      (cond
        (( = T2 0) 0.0)
        (( <= T2 60) (times 250 (quotient T2 60.0)))
        (( <= T2 480) (sum 250
          (times 100 (quotient (diff T2 60) 60.0))))
        (( <= T2 1440) (sum 950
          (times 50 (quotient (diff T2 480) 60.0))))
      )
      (cond
        (( = T1 0) 0.0)
        (( <= T1 60) (times 250 (quotient T1 60.0)))
        (( <= T1 480) (sum 250
          (times 100 (quotient (diff T1 60) 60.0)))
          (( <= T1 1440) (sum 950
            (times 50 (quotient (diff T1 480) 60.0)))))))))
```

; Return (normalized) minimum expected cumulative chest tube drainage.  
 ; Note that, logically, the sum of the minimum flow rates over time  
 ; should be <= the minimum cumulative volume.  
 ; This is not the case so some adjustment should probably be made here.

```
(defun Cum-min-CT (time)
  (times normfact
    (cond
      (( <= time 60) (diff 40 (times 40 (quotient (diff 60 time)60.0))))
      (( <= time 120) (diff 60 (times 20 (quotient (diff 120 time)60.0))))
      (( <= time 180) (diff 80 (times 20 (quotient (diff 180 time)60.0))))
      (( <= time 240) (diff 100 (times 20 (quotient (diff 240 time)60.0))))
      (( <= time 300) (diff 120 (times 20 (quotient (diff 300 time)60.0))))
      (( <= time 360) (diff 140 (times 20 (quotient (diff 360 time)60.0))))
      (( <= time 420) (diff 160 (times 20 (quotient (diff 420 time)60.0))))
      (( <= time 480) (diff 180 (times 20 (quotient (diff 480 time)60.0))))
      (( <= time 540) (diff 200 (times 20 (quotient (diff 540 time)60.0))))
      (( <= time 600) (diff 220 (times 20 (quotient (diff 600 time)60.0))))
      (( <= time 660) (diff 240 (times 20 (quotient (diff 660 time)60.0))))
      (( <= time 720) (diff 260 (times 20 (quotient (diff 720 time)60.0))))
      (( <= time 780) (diff 280 (times 20 (quotient (diff 780 time)60.0))))
      (( <= time 840) (diff 300 (times 20 (quotient (diff 840 time)60.0))))
      (( <= time 900) (diff 320 (times 20 (quotient (diff 900 time)60.0))))
      (( <= time 960) (diff 340 (times 20 (quotient (diff 960 time)60.0))))
      (( <= time 1020) (diff 360 (times 20 (quotient (diff 1020 time)60.0))))
      (( <= time 1080) (diff 380 (times 20 (quotient (diff 1080 time)60.0))))
      (( <= time 1140) (diff 400 (times 20 (quotient (diff 1140 time)60.0))))
      (( <= time 1200) (diff 420 (times 20 (quotient (diff 1200 time)60.0))))
      (( <= time 1260) (diff 440 (times 20 (quotient (diff 1260 time)60.0))))
```

```
( ( <= time 1320) (diff 460 (times 20 (quotient (diff 1320 time)60.0))))
( ( <= time 1380) (diff 480 (times 20 (quotient (diff 1380 time)60.0))))
( ( <= time 1440) (diff 500 (times 20 (quotient (diff 1440 time)60.0)))) ) ) )
```

**; Return (normalized) minimum expected chest tube drainage between  
; time T1 and time T2.**

**; 40 ml min is expected in first hour and**

**; 30 ml min is expected in each subsequent hour.**

**(defun Flw-min-CT (T1 T2)**

**(times normfact**

**(diff**

**(cond**

**(( = T2 0) 0.0)**

**(( <= T2 60) (times 40 (quotient T2 60.0)))**

**(( <= T2 1440) (sum 40  
                  (times 30 (quotient (diff T2 60) 60.0))))))**

**(cond**

**(( = T1 0) 0.0)**

**(( <= T1 60) (times 40 (quotient T1 60.0)))**

**(( <= T1 1440) (sum 40  
                  (times 30 (quotient (diff T1 60) 60.0))))))**

**DETAILS OF RESULTS****PATIENT 5 (TAMPONADE)**

The following is the initialization information gathered by the *questioning module*:

```
; Initialization questions
(setq name "Mr. X")
(setq age 67)
(setq idnumber "05")
(setq bednum 9)
(setq sex 'male)
(setq height 67)
(setq weight 178)
(setq aorticsurg 'no)
(setq mitralsurg 'no)
(setq CABG 'yes)
(setq othercardsurg 'no)
(setq reop 'no)
(setq num-of-CABG 4)
(setq lima 'yes)
(setq exces-bleed-OR 'no)
(setq long-bypass 'no)
(setq coagulopathy 'no)
(setq signif-card-dysfun 'yes)
(setq signif-renal-dysfun 'no)
(setq signif-pulm-hyper 'no)
(setq low-CO 'no)
(setq fragil-aort 'no)
(setq bp 'normal)
(setq set-MAP 'yes)
(setq user-lowlim-MAP 60)
(setq user-highlim-MAP 95)
(setq lventric-hyper 'no)
(setq lventric-dilat 'no)
(setq aorvalv-prob 'no)
(setq mitral-prob 'no)
(setq left-CHF 'yes)
(setq set-LAP 'yes)
(setq user-lowlim-LAP 8)
(setq user-highlim-LAP 16)
(setq rventric-hyper 'no)
(setq rventric-dilat 'no)
(setq tricusp-prob 'no)
(setq right-CHF 'no)
(setq set-RAP 'yes)
```

```
(setq user-lowlim-RAP 5)
(setq user-highlim-RAP 15)
```

The following are the initialization actions performed by the *deductive module* (for the sake of space the silent initialization actions are not shown):

```
; General initialization
(setq normfact 1.15)
; Chest tube initialization
(setq W-CT-high 2)
; Urine output initialization
(setq W-UO-low 1)
; Cardiac index initialization
(setq W-CI-low 1)
(setq bsa 19375)
; Mean arterial pressure initialization
(setq W-MAP-keeplow 1)
(setq E-MAP-4-10 'high)
(setq E-MAP 'normal)
(setq W-E-MAP-conf 1)
(setq suggest-lowlim-MAP 55)
(setq suggest-highlim-MAP 75)
(setq lowlim-MAP 60)
(setq highlim-MAP 95)
; Left atrial pressure initialization
(setq E-LAP 1)
(setq suggest-lowlim-LAP 10)
(setq suggest-highlim-LAP 16)
(setq lowlim-LAP 8)
(setq highlim-LAP 16)
; Right atrial pressure initialization
(setq E-RAP 0)
(setq suggest-lowlim-RAP 4)
(setq suggest-highlim-RAP 14)
(setq lowlim-RAP 5)
(setq highlim-RAP 15)
; LAP/RAP initialization
; TAMP initialization
(setq alphaCT .22)
(setq alphaUO .22)
(setq alphaCI .22)
(setq alphaMA .12)
(setq alphaLR .22)
```

The following is the data that was *input* to the monitoring module. Note that the first item in each pair is the time where, for example, 0210 means two hours and ten minutes after the conclusion of surgery.

The chest tube drainage in ml (cumulative volume).

```
(setq data-CT
(0000 0)
(0100 80)
(0200 112)
(0300 160)
(0400 240)
(0500 560)
(0600 620)
(0700 630)
))
```

The urine output in ml (cumulative volume).

```
(setq data-UO
(0000 0)
(0100 270)
(0200 390)
(0300 500)
(0400 610)
(0500 680)
(0600 710)
(0700 720)
))
```

The cardiac index.

Note that the cardiac index at time 0000 is assumed to be normal (2.4).

```
(setq data-CI
(0000 2.4)
(0210 2.09)
(0623 1.15)
))
```

The mean arterial pressure.

```
(setq data-MAP
( 0000 80 ) ( 0045 86 )
( 0100 75 ) ( 0110 85 ) ( 0120 82 ) ( 0130 68 ) ( 0140 74 ) ( 0150 73 )
( 0200 76 ) ( 0210 79 ) ( 0220 80 ) ( 0230 85 ) ( 0240 85 ) ( 0250 82 )
( 0300 85 ) ( 0310 96 ) ( 0320 83 ) ( 0330 75 ) ( 0340 72 ) ( 0350 81 )
( 0400 75 ) ( 0410 72 ) ( 0420 58 ) ( 0430 71 ) ( 0440 65 ) ( 0450 55 )
( 0500 58 ) ( 0510 52 ) ( 0520 69 ) ( 0530 55 ) ( 0540 58 ) ( 0550 56 )
( 0600 52 ) ( 0610 51 ) ( 0623 61 ) ( 0640 61 ) ( 0650 48 )
( 0700 64 ) ( 0710 93 ) ( 0720 75 ) ( 0730 42 ) ) )
```

The left atrial pressure.

```
(setq data-LAP
( 0000 8 )
( 0210 7 )
( 0623 13)
```

))

The right atrial pressure.

```
(setq data-RAP
( 0000 5 )
( 0210 7 )
( 0623 13 )
))
```

The following is the output from running the *monitoring module* on the above data (Tamp = WS):

```
(Time 0)
(CT 0.0 time 0 mu 0.0 cum-mu 0)
(UO 0.0 time 0 mu 0.0 cum-mu 0)
(CI 2.4 time 0 mu 0.0 cum-mu 0)
(MAP 80.0 time 0 mu 0.0 cum-mu 0)
(LAP 8.0 time 0)
(RAP 5.0 time 0)
(LR 3.0 time 0 mu 0.13 cum-mu 0)
(Tamp 0.02)
```

```
(Time 10)
(Tamp 0.02)
```

```
(Time 20)
(Tamp 0.02)
```

```
(Time 30)
(Tamp 0.02)
```

```
(Time 40)
(Tamp 0.02)
```

```
(Time 50)
(MAP 86.0 time 45 mu 0.0 cum-mu 0)
(Tamp 0.02)
```

```
(Time 100)
(CT 80.0 time 100 mu 0.0 cum-mu 0)
(UO 270.0 time 100 mu 0.0 cum-mu 0)
(MAP 75.0 time 100 mu 0.0 cum-mu 0)
(Tamp 0.02)
```

```
(Time 110)
(MAP 85.0 time 110 mu 0.0 cum-mu 0)
(Tamp 0.02)
```

```
(Time 120)
```

(MAP 82.0 time 120 mu 0.0 cum-mu 0)  
(Tamp 0.02)

(Time 130)  
(MAP 68.0 time 130 mu 0.0 cum-mu 0)  
(Tamp 0.02)

(Time 140)  
(MAP 74.0 time 140 mu 0.0 cum-mu 0)  
(Tamp 0.02)

(Time 150)  
(MAP 73.0 time 150 mu 0.0 cum-mu 0)  
(Tamp 0.02)

(Time 200)  
(CT 112.0 time 200 mu 0.14 cum-mu 4)  
(UO 390.0 time 200 mu 0.0 cum-mu 0)  
(MAP 76.0 time 200 mu 0.0 cum-mu 0)  
(Tamp 0.06)

(Time 210)  
(CI 2.09 time 210 mu 0.34 cum-mu 22)  
(MAP 79.0 time 210 mu 0.0 cum-mu 0)  
(LAP 7.0 time 210)  
(RAP 7.0 time 210)  
(LR 0.0 time 210 mu 0.4 cum-mu 34)  
(Tamp 0.19)

(Time 220)  
(MAP 80.0 time 220 mu 0.0 cum-mu 0)  
(Tamp 0.19)

(Time 230)  
(MAP 85.0 time 230 mu 0.0 cum-mu 0)  
(Tamp 0.19)

(Time 240)  
(MAP 85.0 time 240 mu 0.0 cum-mu 0)  
(Tamp 0.19)

(Time 250)  
(MAP 82.0 time 250 mu 0.0 cum-mu 0)  
(Tamp 0.19)

(Time 300)  
(CT 160.0 time 300 mu 0.0 cum-mu 8)  
(UO 500.0 time 300 mu 0.0 cum-mu 0)  
(MAP 85.0 time 300 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 310)  
(MAP 96.0 time 310 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 320)  
(MAP 83.0 time 320 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 330)  
(MAP 75.0 time 330 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 340)  
(MAP 72.0 time 340 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 350)  
(MAP 81.0 time 350 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 400)  
(CT 240.0 time 400 mu 0.0 cum-mu 8)  
(UO 610.0 time 400 mu 0.0 cum-mu 0)  
(MAP 75.0 time 400 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 410)  
(MAP 72.0 time 410 mu 0.0 cum-mu 0)  
(Tamp 0.16)

(Time 420)  
(MAP 58.0 time 420 mu 0.29 cum-mu 1)  
(Tamp 0.2)

(Time 430)  
(MAP 71.0 time 430 mu 0.0 cum-mu 2)  
(Tamp 0.16)

(Time 440)  
(MAP 65.0 time 440 mu 0.0 cum-mu 2)  
(Tamp 0.16)

(Time 450)  
(MAP 55.0 time 450 mu 0.75 cum-mu 6)  
(Tamp 0.25)

(Time 500)  
(CT 560.0 time 500 mu 1.0 cum-mu 38)  
(UO 680.0 time 500 mu 0.0 cum-mu 0)  
(MAP 58.0 time 500 mu 0.29 cum-mu 12)  
(Tamp 0.42)

(Time 510)  
(MAP 52.0 time 510 mu 0.91 cum-mu 18)  
(Tamp 0.49)

(Time 520)  
(MAP 69.0 time 520 mu 0.0 cum-mu 22)  
(Tamp 0.38)

(Time 530)  
(MAP 55.0 time 530 mu 0.75 cum-mu 26)  
(Tamp 0.47)

(Time 540)  
(MAP 58.0 time 540 mu 0.29 cum-mu 31)  
(Tamp 0.42)

(Time 550)  
(MAP 56.0 time 550 mu 0.59 cum-mu 36)  
(Tamp 0.45)

(Time 600)  
(CT 620.0 time 600 mu 0.0 cum-mu 108)  
(UO 710.0 time 600 mu 1.0 cum-mu 30)  
(MAP 52.0 time 600 mu 0.91 cum-mu 43)  
(Tamp 0.49)

(Time 610)  
(MAP 51.0 time 610 mu 0.92 cum-mu 52)  
(Tamp 0.49)

(Time 620)  
(Tamp 0.49)

(Time 630)  
(CI 1.15 time 623 mu 0.95 cum-mu 307)  
(MAP 61.0 time 623 mu 0.0 cum-mu 58)  
(LAP 13.0 time 623)  
(RAP 13.0 time 623)  
(LR 0.0 time 623 mu 0.4 cum-mu 215)  
(Tamp 0.51)

(Time 640)  
(MAP 61.0 time 640 mu 0.0 cum-mu 58)  
(Tamp 0.51)

(Time 650)  
(MAP 48.0 time 650 mu 0.94 cum-mu 103)  
(Tamp 0.63)

(Time 700)  
(CT 630.0 time 700 mu 1.0 cum-mu 138)

(UO 720.0 time 700 mu 1.0 cum-mu 130)  
(MAP 64.0 time 700 mu 0.0 cum-mu 108)  
(Tamp 0.73)

(Time 710)  
(MAP 93.0 time 710 mu 0.0 cum-mu 108)  
(Tamp 0.73)

(Time 720)  
(MAP 75.0 time 720 mu 0.0 cum-mu 108)  
(Tamp 0.73)

(Time 730)  
(MAP 42.0 time 730 mu 0.98 cum-mu 113)  
(Tamp 0.85)

**PATIENT 6 (NON-TAMPONADE)**

The following is the initialization information gathered by the *questioning module*:

```
; Initialization questions
(setq name "Ms. Y")
(setq age 60)
(setq idnumber "6")
(setq bednum 7)
(setq sex 'female)
(setq height 67)
(setq weight 146)
(setq aorticsurg 'no)
(setq mitralsurg 'no)
(setq CABG 'yes)
(setq othercardsurg 'no)
(setq reop 'no)
(setq num-of-CABG 4)
(setq lima 'yes)
(setq exces-bleed-OR 'no)
(setq long-bypass 'no)
(setq coagulopathy 'no)
(setq signif-card-dysfun 'no)
(setq signif-renal-dysfun 'no)
(setq signif-pulm-hyper 'no)
(setq low-CO 'no)
(setq fragil-aort 'no)
(setq bp 'normal)
(setq set-MAP 'yes)
(setq user-lowlim-MAP 60)
(setq user-highlim-MAP 95)
(setq lventric-hyper 'no)
(setq lventric-dilat 'no)
(setq aorvalv-prob 'no)
(setq mitral-prob 'no)
(setq left-CHF 'yes)
(setq set-LAP 'yes)
(setq user-lowlim-LAP 8)
(setq user-highlim-LAP 16)
(setq rventric-hyper 'no)
(setq rventric-dilat 'no)
(setq tricusp-prob 'no)
(setq right-CHF 'no)
(setq set-RAP 'yes)
(setq user-lowlim-RAP 5)
(setq user-highlim-RAP 15)
```

The following are the initialization actions performed by the *deductive module* (for the sake of space the silent initialization actions are not shown):

```

; General initialization
(setq normfact 0.95)
; Chest tube initialization
(setq W-CT-high 2)
; Urine output initialization
(setq W-UO-low 0)
; Cardiac index initialization
(setq W-CI-low 0)
(setq bsa 17810)
; Mean arterial pressure initialization
(setq W-MAP-keeplow 1)
(setq E-MAP-4-10 'high)
(setq E-MAP 'normal)
(setq W-E-MAP-conf 1)
(setq suggest-lowlim-MAP 55)
(setq suggest-highlim-MAP 75)
(setq lowlim-MAP 60)
(setq highlim-MAP 95)
; Left atrial pressure initialization
(setq E-LAP 1)
(setq suggest-lowlim-LAP 10)
(setq suggest-highlim-LAP 16)
(setq lowlim-LAP 8)
(setq highlim-LAP 16)
; Right atrial pressure initialization
(setq E-RAP 0)
(setq suggest-lowlim-RAP 4)
(setq suggest-highlim-RAP 14)
(setq lowlim-RAP 5)
(setq highlim-RAP 15)
; LAP/RAP initialization
; TAMP initialization
(setq alphaCT .22)
(setq alphaUO .22)
(setq alphaCI .22)
(setq alphaMA .12)
(setq alphaLR .22)

```

The following is the data that was *input* to the monitoring module. Note that the first item in each pair is the time where, for example, 0210 means one hour and twenty minutes after the conclusion of surgery.

```

The chest tube drainage in ml (cumulative volume).
(setq data-CT
(0000 0)

```

```

(0020 40)
(0120 80)
(0220 100)
(0320 130)
(0420 160)
(0520 210)
(0620 250)
(0720 300)
(0820 320)
(0920 370)
(1020 410)
(1120 430)
))

```

The urine output in ml (cumulative volume).

```

(setq data-UO
(0000 0)
(0020 306)
(0120 426)
(0220 566)
(0320 796)
(0420 1056)
(0520 1136)
(0620 1406)
(0720 1656)
(0820 1756)
(0920 2006)
(1020 2106)
(1120 2196)
))

```

The cardiac index.

```

(setq data-CI
(0000 2.4)
(0121 2.89)
(0829 3.58)
))

```

The mean arterial pressure.

```

(setq data-MAP
( 0000 84 ) ( 0010 67 ) ( 0020 75 ) ( 0030 79 ) ( 0040 77 ) ( 0050 78 )
( 0100 80 ) ( 0110 80 ) ( 0120 80 ) ( 0130 81 ) ( 0140 82 ) ( 0150 82 )
( 0200 73 ) ( 0210 76 ) ( 0220 79 ) ( 0230 81 ) ( 0240 82 ) ( 0250 72 )
( 0300 79 ) ( 0310 70 ) ( 0320 76 ) ( 0330 98 ) ( 0340 80 ) ( 0350 71 )
( 0400 75 ) ( 0410 75 ) ( 0420 80 ) ( 0430 81 ) ( 0440 80 ) ( 0450 88 )
( 0500 72 ) ( 0510 56 ) ( 0520 94 ) ( 0530 65 ) ( 0540 72 ) ( 0550 74 )
( 0600 75 ) ( 0610 85 ) ( 0620 81 ) ( 0630 80 ) ( 0640 77 ) ( 0650 78 )

```

( 0700 75 ) ( 0710 75 ) ( 0720 79 ) ( 0730 89 ) ( 0740 80 ) ( 0750 82 )  
 ( 0800 72 ) ( 0810 77 ) ( 0820 72 ) ( 0830 71 ) ( 0840 77 ) ( 0850 74 )  
 ( 0900 75 ) ( 0910 70 ) ( 0920 72 ) ( 0930 71 ) ( 0940 74 ) ( 0950 84 )  
 ( 1000 90 ) ( 1010 80 ) ( 1020 78 ) ( 1030 74 ) ( 1040 70 ) ( 1050 88 )  
 ( 1100 73 ) ( 1110 82 ) ( 1120 73 ) )

The left atrial pressure.

(setq data-LAP

( 0000 21 ) ( 0010 15 ) ( 0020 13 ) ( 0030 12 ) ( 0040 11 ) ( 0050 11 )  
 ( 0100 11 ) ( 0110 11 ) ( 0120 11 ) ( 0130 10 ) ( 0140 11 ) ( 0150 10 )  
 ( 0200 10 ) ( 0210 10 ) ( 0220 11 ) ( 0230 10 ) ( 0240 10 ) ( 0250 10 )  
 ( 0300 10 ) ( 0310 11 ) ( 0320 11 ) ( 0330 14 ) ( 0340 12 ) ( 0350 11 )  
 ( 0400 12 ) ( 0410 12 ) ( 0420 12 ) ( 0430 12 ) ( 0440 12 ) ( 0450 13 )  
 ( 0500 11 ) ( 0510 11 ) ( 0520 13 ) ( 0530 12 ) ( 0540 12 ) ( 0550 12 )  
 ( 0600 12 ) ( 0610 12 ) ( 0620 13 ) ( 0630 13 ) ( 0640 13 ) ( 0650 14 )  
 ( 0700 13 ) ( 0710 13 ) ( 0720 12 ) ( 0730 12 ) ( 0740 12 ) ( 0750 14 )  
 ( 0800 15 ) ( 0810 15 ) ( 0820 12 ) ( 0830 13 ) ( 0840 12 ) ( 0850 11 )  
 ( 0900 11 ) ( 0910 12 ) ( 0920 12 ) ( 0930 12 ) ( 0940 12 ) ( 0950 14 )  
 ( 1000 14 ) ( 1010 12 ) ( 1020 14 ) ( 1030 11 ) ( 1040 12 ) ( 1050 15 )  
 ( 1100 14 ) ( 1110 14 ) ( 1120 14 ) )

The right atrial pressure.

(setq data-RAP

( 0000 7 ) ( 0010 5 ) ( 0020 4 ) ( 0030 4 ) ( 0040 4 ) ( 0050 4 )  
 ( 0100 4 ) ( 0110 3 ) ( 0120 4 ) ( 0130 4 ) ( 0140 3 ) ( 0150 5 )  
 ( 0200 2 ) ( 0210 2 ) ( 0220 1 ) ( 0230 2 ) ( 0240 2 ) ( 0250 2 )  
 ( 0520 4 ) ( 0530 5 ) ( 0540 4 ) ( 0550 4 )  
 ( 0600 4 ) ( 0610 4 ) ( 0620 4 ) ( 0630 4 ) ( 0640 4 ) ( 0650 5 )  
 ( 0700 5 ) ( 0710 5 ) ( 0720 4 ) ( 0730 5 ) ( 0740 3 ) ( 0750 4 )  
 ( 0800 4 ) ( 0810 5 ) ( 0820 5 ) ( 0830 5 ) ( 0840 5 ) ( 0850 4 )  
 ( 0900 4 ) ( 0910 4 ) ( 0920 5 ) ( 0930 5 ) ( 0940 4 ) ( 0950 5 )  
 ( 1000 6 ) ( 1010 6 ) ( 1020 7 ) ( 1030 4 ) ( 1040 4 ) ( 1050 6 )  
 ( 1100 5 ) ( 1110 4 ) ( 1120 3 ) )

The following is the output from running the *monitoring module* on the above data:

(Time 0)

(CT 0.0 time 0 mu 0.0 cum-mu 0)

(UO 0.0 time 0 mu 0.0 cum-mu 0)

(CI 2.4 time 0 mu 0.0 cum-mu 0)

(MAP 84.0 time 0 mu 0.0 cum-mu 0)

(LAP 21.0 time 0)  
(RAP 7.0 time 0)  
(LR 14.0 time 0 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 10)  
(MAP 67.0 time 10 mu 0.0 cum-mu 0)  
(LAP 15.0 time 10)  
(RAP 5.0 time 10)  
(LR 10.0 time 10 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 20)  
(CT 40.0 time 20 mu 0.0 cum-mu 0)  
(UO 306.0 time 20 mu 0.0 cum-mu 0)  
(MAP 75.0 time 20 mu 0.0 cum-mu 0)  
(LAP 13.0 time 20)  
(RAP 4.0 time 20)  
(LR 9.0 time 20 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 30)  
(MAP 79.0 time 30 mu 0.0 cum-mu 0)  
(LAP 12.0 time 30)  
(RAP 4.0 time 30)  
(LR 8.0 time 30 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 40)  
(MAP 77.0 time 40 mu 0.0 cum-mu 0)  
(LAP 11.0 time 40)  
(RAP 4.0 time 40)  
(LR 7.0 time 40 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 50)  
(MAP 78.0 time 50 mu 0.0 cum-mu 0)  
(LAP 11.0 time 50)  
(RAP 4.0 time 50)  
(LR 7.0 time 50 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 100)  
(MAP 80.0 time 100 mu 0.0 cum-mu 0)  
(LAP 11.0 time 100)  
(RAP 4.0 time 100)  
(LR 7.0 time 100 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 110)  
(MAP 80.0 time 110 mu 0.0 cum-mu 0)

(LAP 11.0 time 110)  
(RAP 3.0 time 110)  
(LR 8.0 time 110 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 120)  
(CT 80.0 time 120 mu 0.0 cum-mu 0)  
(UO 426.0 time 120 mu 0.0 cum-mu 0)  
(MAP 80.0 time 120 mu 0.0 cum-mu 0)  
(LAP 11.0 time 120)  
(RAP 4.0 time 120)  
(LR 7.0 time 120 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 130)  
(CI 2.89 time 121 mu 0.0 cum-mu 0)  
(MAP 81.0 time 130 mu 0.0 cum-mu 0)  
(LAP 10.0 time 130)  
(RAP 4.0 time 130)  
(LR 6.0 time 130 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 140)  
(MAP 82.0 time 140 mu 0.0 cum-mu 0)  
(LAP 11.0 time 140)  
(RAP 3.0 time 140)  
(LR 8.0 time 140 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 150)  
(MAP 82.0 time 150 mu 0.0 cum-mu 0)  
(LAP 10.0 time 150)  
(RAP 5.0 time 150)  
(LR 5.0 time 150 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 200)  
(MAP 73.0 time 200 mu 0.0 cum-mu 0)  
(LAP 10.0 time 200)  
(RAP 2.0 time 200)  
(LR 8.0 time 200 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 210)  
(MAP 76.0 time 210 mu 0.0 cum-mu 0)  
(LAP 10.0 time 210)  
(RAP 2.0 time 210)  
(LR 8.0 time 210 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 220)

(CT 100.0 time 220 mu 0.59 cum-mu 17)  
(UO 566.0 time 220 mu 0.0 cum-mu 0)  
(MAP 79.0 time 220 mu 0.0 cum-mu 0)  
(LAP 11.0 time 220)  
(RAP 1.0 time 220)  
(LR 10.0 time 220 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 230)  
(MAP 81.0 time 230 mu 0.0 cum-mu 0)  
(LAP 10.0 time 230)  
(RAP 2.0 time 230)  
(LR 8.0 time 230 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 240)  
(MAP 82.0 time 240 mu 0.0 cum-mu 0)  
(LAP 10.0 time 240)  
(RAP 2.0 time 240)  
(LR 8.0 time 240 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 250)  
(MAP 72.0 time 250 mu 0.0 cum-mu 0)  
(LAP 10.0 time 250)  
(RAP 2.0 time 250)  
(LR 8.0 time 250 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 300)  
(MAP 79.0 time 300 mu 0.0 cum-mu 0)  
(LAP 10.0 time 300)  
(LR 8.0 time 300 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 310)  
(MAP 70.0 time 310 mu 0.0 cum-mu 0)  
(LAP 11.0 time 310)  
(LR 9.0 time 310 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 320)  
(CT 130.0 time 320 mu 0.0 cum-mu 35)  
(UO 796.0 time 320 mu 0.0 cum-mu 0)  
(MAP 76.0 time 320 mu 0.0 cum-mu 0)  
(LAP 11.0 time 320)  
(Tamp 0.0)

(Time 330)  
(MAP 98.0 time 330 mu 0.0 cum-mu 0)  
(LAP 14.0 time 330)

(Tamp 0.0)

(Time 340)

(MAP 80.0 time 340 mu 0.0 cum-mu 0)

(LAP 12.0 time 340)

(Tamp 0.0)

(Time 350)

(MAP 71.0 time 350 mu 0.0 cum-mu 0)

(LAP 11.0 time 350)

(Tamp 0.0)

(Time 400)

(MAP 75.0 time 400 mu 0.0 cum-mu 0)

(LAP 12.0 time 400)

(Tamp 0.0)

(Time 410)

(MAP 75.0 time 410 mu 0.0 cum-mu 0)

(LAP 12.0 time 410)

(Tamp 0.0)

(Time 420)

(CT 160.0 time 420 mu 0.0 cum-mu 35)

(UO 1056.0 time 420 mu 0.0 cum-mu 0)

(MAP 80.0 time 420 mu 0.0 cum-mu 0)

(LAP 12.0 time 420)

(Tamp 0.0)

(Time 430)

(MAP 81.0 time 430 mu 0.0 cum-mu 0)

(LAP 12.0 time 430)

(Tamp 0.0)

(Time 440)

(MAP 80.0 time 440 mu 0.0 cum-mu 0)

(LAP 12.0 time 440)

(Tamp 0.0)

(Time 450)

(MAP 88.0 time 450 mu 0.0 cum-mu 0)

(LAP 13.0 time 450)

(Tamp 0.0)

(Time 500)

(MAP 72.0 time 500 mu 0.0 cum-mu 0)

(LAP 11.0 time 500)

(Tamp 0.0)

(Time 510)

(MAP 56.0 time 510 mu 0.59 cum-mu 2)

(LAP 11.0 time 510)  
(Tamp 0.07)

(Time 520)  
(CT 210.0 time 520 mu 0.0 cum-mu 35)  
(UO 1136.0 time 520 mu 0.0 cum-mu 0)  
(MAP 94.0 time 520 mu 0.0 cum-mu 5)  
(LAP 13.0 time 520)  
(RAP 4.0 time 520)  
(LR 9.0 time 520 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 530)  
(MAP 65.0 time 530 mu 0.0 cum-mu 5)  
(LAP 12.0 time 530)  
(RAP 5.0 time 530)  
(LR 7.0 time 530 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 540)  
(MAP 72.0 time 540 mu 0.0 cum-mu 5)  
(LAP 12.0 time 540)  
(RAP 4.0 time 540)  
(LR 8.0 time 540 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 550)  
(MAP 74.0 time 550 mu 0.0 cum-mu 5)  
(LAP 12.0 time 550)  
(RAP 4.0 time 550)  
(LR 8.0 time 550 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 600)  
(MAP 75.0 time 600 mu 0.0 cum-mu 5)  
(LAP 12.0 time 600)  
(RAP 4.0 time 600)  
(LR 8.0 time 600 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 610)  
(MAP 85.0 time 610 mu 0.0 cum-mu 5)  
(LAP 12.0 time 610)  
(RAP 4.0 time 610)  
(LR 8.0 time 610 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 620)  
(CT 250.0 time 620 mu 0.0 cum-mu 35)  
(UO 1406.0 time 620 mu 0.0 cum-mu 0)  
(MAP 81.0 time 620 mu 0.0 cum-mu 5)

(LAP 13.0 time 620)  
(RAP 4.0 time 620)  
(LR 9.0 time 620 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 630)  
(MAP 80.0 time 630 mu 0.0 cum-mu 5)  
(LAP 13.0 time 630)  
(RAP 4.0 time 630)  
(LR 9.0 time 630 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 640)  
(MAP 77.0 time 640 mu 0.0 cum-mu 5)  
(LAP 13.0 time 640)  
(RAP 4.0 time 640)  
(LR 9.0 time 640 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 650)  
(MAP 78.0 time 650 mu 0.0 cum-mu 5)  
(LAP 14.0 time 650)  
(RAP 5.0 time 650)  
(LR 9.0 time 650 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 700)  
(MAP 75.0 time 700 mu 0.0 cum-mu 5)  
(LAP 13.0 time 700)  
(RAP 5.0 time 700)  
(LR 8.0 time 700 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 710)  
(MAP 75.0 time 710 mu 0.0 cum-mu 5)  
(LAP 13.0 time 710)  
(RAP 5.0 time 710)  
(LR 8.0 time 710 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 720)  
(CT 300.0 time 720 mu 0.0 cum-mu 35)  
(UO 1656.0 time 720 mu 0.0 cum-mu 0)  
(MAP 79.0 time 720 mu 0.0 cum-mu 5)  
(LAP 12.0 time 720)  
(RAP 4.0 time 720)  
(LR 8.0 time 720 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 730)  
(MAP 89.0 time 730 mu 0.0 cum-mu 5)

(LAP 12.0 time 730)  
(RAP 5.0 time 730)  
(LR 7.0 time 730 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 740)  
(MAP 80.0 time 740 mu 0.0 cum-mu 5)  
(LAP 12.0 time 740)  
(RAP 3.0 time 740)  
(LR 9.0 time 740 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 750)  
(MAP 82.0 time 750 mu 0.0 cum-mu 5)  
(LAP 14.0 time 750)  
(RAP 4.0 time 750)  
(LR 10.0 time 750 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 800)  
(MAP 72.0 time 800 mu 0.0 cum-mu 5)  
(LAP 15.0 time 800)  
(RAP 4.0 time 800)  
(LR 11.0 time 800 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 810)  
(MAP 77.0 time 810 mu 0.0 cum-mu 5)  
(LAP 15.0 time 810)  
(RAP 5.0 time 810)  
(LR 10.0 time 810 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 820)  
(CT 320.0 time 820 mu 0.59 cum-mu 53)  
(UO 1756.0 time 820 mu 0.0 cum-mu 0)  
(MAP 72.0 time 820 mu 0.0 cum-mu 5)  
(LAP 12.0 time 820)  
(RAP 5.0 time 820)  
(LR 7.0 time 820 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 830)  
(CI 3.58 time 829 mu 0.0 cum-mu 0)  
(MAP 71.0 time 830 mu 0.0 cum-mu 5)  
(LAP 13.0 time 830)  
(RAP 5.0 time 830)  
(LR 8.0 time 830 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 840)

(MAP 77.0 time 840 mu 0.0 cum-mu 5)  
(LAP 12.0 time 840)  
(RAP 5.0 time 840)  
(LR 7.0 time 840 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 850)  
(MAP 74.0 time 850 mu 0.0 cum-mu 5)  
(LAP 11.0 time 850)  
(RAP 4.0 time 850)  
(LR 7.0 time 850 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 900)  
(MAP 75.0 time 900 mu 0.0 cum-mu 5)  
(LAP 11.0 time 900)  
(RAP 4.0 time 900)  
(LR 7.0 time 900 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 910)  
(MAP 70.0 time 910 mu 0.0 cum-mu 5)  
(LAP 12.0 time 910)  
(RAP 4.0 time 910)  
(LR 8.0 time 910 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 920)  
(CT 370.0 time 920 mu 0.05 cum-mu 113)  
(UO 2006.0 time 920 mu 0.0 cum-mu 0)  
(MAP 72.0 time 920 mu 0.0 cum-mu 5)  
(LAP 12.0 time 920)  
(RAP 5.0 time 920)  
(LR 7.0 time 920 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 930)  
(MAP 71.0 time 930 mu 0.0 cum-mu 5)  
(LAP 12.0 time 930)  
(RAP 5.0 time 930)  
(LR 7.0 time 930 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 940)  
(MAP 74.0 time 940 mu 0.0 cum-mu 5)  
(LAP 12.0 time 940)  
(RAP 4.0 time 940)  
(LR 8.0 time 940 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 950)

(MAP 84.0 time 950 mu 0.0 cum-mu 5)  
(LAP 14.0 time 950)  
(RAP 5.0 time 950)  
(LR 9.0 time 950 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 1000)  
(MAP 90.0 time 1000 mu 0.0 cum-mu 5)  
(LAP 14.0 time 1000)  
(RAP 6.0 time 1000)  
(LR 8.0 time 1000 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 1010)  
(MAP 80.0 time 1010 mu 0.0 cum-mu 5)  
(LAP 12.0 time 1010)  
(RAP 6.0 time 1010)  
(LR 6.0 time 1010 mu 0.0 cum-mu 0)  
(Tamp 0.01)

(Time 1020)  
(CT 410.0 time 1020 mu 0.0 cum-mu 114)  
(UO 2106.0 time 1020 mu 0.0 cum-mu 0)  
(MAP 78.0 time 1020 mu 0.0 cum-mu 5)  
(LAP 14.0 time 1020)  
(RAP 7.0 time 1020)  
(LR 7.0 time 1020 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1030)  
(MAP 74.0 time 1030 mu 0.0 cum-mu 5)  
(LAP 11.0 time 1030)  
(RAP 4.0 time 1030)  
(LR 7.0 time 1030 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1040)  
(MAP 70.0 time 1040 mu 0.0 cum-mu 5)  
(LAP 12.0 time 1040)  
(RAP 4.0 time 1040)  
(LR 8.0 time 1040 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1050)  
(MAP 88.0 time 1050 mu 0.0 cum-mu 5)  
(LAP 15.0 time 1050)  
(RAP 6.0 time 1050)  
(LR 9.0 time 1050 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1100)

(MAP 73.0 time 1100 mu 0.0 cum-mu 5)  
(LAP 14.0 time 1100)  
(RAP 5.0 time 1100)  
(LR 9.0 time 1100 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1110)  
(MAP 82.0 time 1110 mu 0.0 cum-mu 5)  
(LAP 14.0 time 1110)  
(RAP 4.0 time 1110)  
(LR 10.0 time 1110 mu 0.0 cum-mu 0)  
(Tamp 0.0)

(Time 1120)  
(CT 430.0 time 1120 mu 0.59 cum-mu 132)  
(UO 2196.0 time 1120 mu 0.0 cum-mu 0)  
(MAP 73.0 time 1120 mu 0.0 cum-mu 5)  
(LAP 14.0 time 1120)  
(RAP 3.0 time 1120)  
(LR 11.0 time 1120 mu 0.0 cum-mu 0)  
(Tamp 0.13)

(Time 1130)  
(Tamp 0.13)

## APPENDIX I

## AN ANALYSIS OF FOUR UNCERTAINTY CALCULI

**Note**

This appendix has been published in slightly edited form in [Henkind SJ, Harrison MC. An analysis of four uncertainty calculi. IEEE-SMC 1988;18:700-14]. Reprinted by permission of the publisher.

**AI.1. Introduction**

An important issue faced by contemporary AI workers is how to deal with uncertain information. Early AI systems were generally designed to function in highly restricted environments; e.g., "Blocks World." The information in such environments is usually completely certain; e.g., "block *A* is on top of block *B*."

In recent years, however, there has been much effort expended on developing AI systems that can solve real-life problems. These *expert* (or *knowledge-based*) systems must deal with information that can be uncertain; e.g., "John *may* have disease *X*."

There are multiple sources and types of uncertainty in knowledge-based systems. For example, uncertainty may occur due to less than perfect *reliability* of measurements; e.g., "the reading is accurate to within plus or minus 5 percent." As another example, uncertainty may occur because of *proba-*

*bilistic* information; e.g., "John has a one in a million chance of winning the lottery."

Uncertainty can also arise due to inferential mechanisms within a knowledge-based system. For example, in a *weak implication*, premise and conclusion are only partially correlated; e.g., given that John is driving a Rolls Royce, one may infer that John is rich, but the conclusion is not completely certain because John may have stolen the car, or John may be the chauffeur, etc.

Another source of uncertainty is *lexical imprecision* [Henk88Impr]. Lexical imprecision is a lack of uniform meanings for the semantic primitives used to represent knowledge. For example, the expert who constructs a knowledge-based system may consider *fever* to be a temperature of 37.5 (Celsius) or above, whereas the user of the system, or another expert, may consider *fever* to be a temperature of 38.0 or above.

Bonissone and Tong [Boni85] provide a more thorough review of the sources and types of uncertainty to be found in knowledge-based systems.

Various numerical calculi have been proposed as methods to represent and propagate uncertainty in expert systems [Lee87].<sup>1</sup> Among the more prominent calculi are *probabilistic* (in particular, *Bayesian*) methods, the *Dempster-Shafer theory*, *fuzzy set theory*, and the *MYCIN* and *EMYGIN* calculi. Since the *MYCIN* and *EMYCIN* calculi are so similar, they will be referred to collectively here as *MYCIN/EMYCIN*, unless we need to distinguish between them. It should be noted that only *MYCIN/EMYCIN* was specifically developed in the context of knowledge-based systems; the other calculi were developed independently, and have been recently adapted and

applied by AI workers.

In the Bayesian calculus, uncertainty is viewed as probability, where probability can be interpreted as a relative frequency, as a degree of belief, or in some other manner. The probabilities are numerical values in the unit interval  $[0,1]$  where (using the relative frequency interpretation) 1 means that an event always occurs, 0 means that an event never occurs, and an intermediate value means that an event sometimes occurs.

In the Dempster-Shafer theory, uncertainty is viewed as a degree of belief. These degrees of belief are numerical values in the interval  $[0,1]$  where 1 means total belief, 0 means a lack of belief (not to be confused with disbelief), and an intermediate value means partial belief.

In fuzzy set theory, uncertainty is viewed as a degree of set membership (e.g., the degree of presence of a symptom). The degrees of membership are numerical values in the interval  $[0,1]$  where 1 means that an object is a member, 0 means that an object is not a member, and an intermediate value means that an object is a partial member.

In MYCIN/EMYCIN, uncertainty is viewed as a degree of confirmation. The degrees of confirmation are numerical values in the interval  $[-1,1]$  where 1 means that the evidence confirms the hypothesis, -1 means that the evidence disconfirms the hypothesis, 0 means that the confirming evidence balances the disconfirming evidence (or that there is no relevant evidence), and an intermediate value means that the hypothesis is partially confirmed (or partially disconfirmed) by the evidence.

Since each of the calculi views uncertainty from a different perspective, comparisons can be problematic. In order to provide a common framework

for our analyses, and in order to allow for meaningful comparisons, we examine how each calculus handles the following problem: given some premise,  $e$  (which has some degree of certainty), compute the certainty of the conclusion,  $d$ , where  $e \rightarrow d$  (this implication may itself have some degree of certainty). This should not be strictly interpreted as an application of modus ponens, since several of the calculi do not have such a notion; all the calculi, however, have some mechanism for computing the certainty of a conclusion given some premise. This problem has been selected since it is representative of the fundamental operations carried out on nearly all knowledge representations; e.g., the firing of a rule in a production system, the propagation of an inference across two adjacent nodes in a semantic net, etc.

We will examine this problem in the context of medical diagnosis. Medical diagnosis often takes place in several steps: starting with a set of signs and symptoms (the evidence),  $e$ , progressing through several intermediate hypotheses,  $h_1, h_2, \dots$ , to a final diagnosis,  $d$ . For the sake of simplicity, we assume diagnosis to be only one step:  $e$  to  $d$ . This assumption is made without loss of generality, since  $e$  and  $d$  can be viewed as intermediate hypotheses,  $h_i$  and  $h_{i+1}$ , in a multistep diagnosis.

The set of all signs and symptoms will be denoted by  $S$ ,  $|S| = n$  (where  $|S|$  is the cardinality of  $S$ ),  $s_1, \dots, s_n$  will be an enumeration of  $S$ , and  $s$  will denote an arbitrary element of  $S$ .

The set of all diseases will be denoted by  $D$ ,  $|D| = m$ ,  $d_1, \dots, d_m$  will be an enumeration of  $D$ , and  $d$  will denote an arbitrary element of  $D$ .

An individual sign or symptom will also be called a *piece of evidence*.

The conjunction of zero or more pieces of evidence will be called an *evidence pattern*; e.g.,  $s_2$  &  $s_7$  is an evidence pattern. Alternatively, an evidence pattern may be represented in set notation; e.g.,  $\{s_2, s_7\}$ , where conjunction is understood. If  $|S| = n$ , then there are  $2^n$  possible evidence patterns;  $e_1, \dots, e_{2^n}$  will be an enumeration of these patterns, and  $e$  will denote an arbitrary evidence pattern. The number of signs and symptoms that a given  $e$  is composed from will be denoted by either  $|e|$  or  $p$ ; e.g., if  $e_j = \{s_2, s_7\}$ , then  $|e_j| = 2$ , and  $p = 2$ .

In the Bayesian and Dempster-Shafer calculi, each  $s$  (sign or symptom) is an attribute which is considered to be either *present* or *not present (absent)* in a patient. For example, a patient either does, or does not, have a fever. Therefore, an evidence pattern such as  $\{s_2, s_7\}$  is taken to mean that  $s_2$  and  $s_7$  are present in the patient and that all other signs and symptoms are absent.<sup>2</sup>

There are techniques, some of which will be discussed later, that can extend the Bayesian and Dempster-Shafer calculi to handle  $s$ 's which are *partially present*. These techniques, however, lead to a significant increase in computational complexity, and/or depend on the imposition of strong assumptions. Therefore, unless otherwise noted, we will view signs and symptoms as being either present or absent in the Bayesian and Dempster-Shafer calculi.

In fuzzy set theory and MYCIN/EMYCIN, signs and symptoms are attributes which are supplied along with numeric weights. These weights can be viewed as measurements of the degree of certainty (or, equivalently, uncertainty) of the individual pieces of evidence. Therefore, an overall cer-

tainty for each evidence pattern,  $e$ , must be computed before the certainty of  $d$  can be computed. The certainty of  $e$  is computed by combining, in some fashion, the certainties of the individual pieces of evidence; e.g., in fuzzy set theory if  $e_j = \{s_2, s_7\}$ ,  $s_2$  has degree of presence .3, and  $s_7$  has degree of presence .8, then (depending on the how the operators are defined) the degree of presence of  $e_j$  would be given by  $\min(.3, .8) = .3$ .

Thus, in fuzzy set theory and MYCIN/EMYCIN, the computation of the certainty of  $d$ , given  $e$ , is conceptually a two step process: (1) compute the certainty of  $e$ ; (2) compute the certainty of  $d$ . In the Bayesian and Dempster-Shafer calculi, the first step is not needed since, implicitly,  $e$  has a certainty of one (but see [Shaf76, pg. 66]).

In the pages that follow we review and analyze the aforementioned uncertainty calculi. Particular attention is paid to the computational complexity of these calculi, as exponential complexity can lead to pragmatic difficulties. The complexity analyses do not assume any particular knowledge representation; nor do we assume any particular structure for the expert system. Unless otherwise stated, the complexity analyses are *worst case*.

One complexity measure of particular concern is *information complexity*. Information complexity is a measure of the amount of auxiliary information that is needed, in addition to the implications, in order to be able to compute the certainty of an *arbitrary* disease  $d$ , given *arbitrary* evidence pattern  $e$ . The auxiliary information is either: (1) precomputed and *stored* by the system (e.g., prior probabilities in the Bayesian approach), or (2) *supplied* along with pieces of evidence (e.g., certainty factors in the EMYCIN calculus). For example (as will be shown later), in the Bayesian framework,

the system needs to store  $O(m \cdot 2^n)$  probabilities in order to be able to compute the probability of *arbitrary* disease,  $d$ , given *arbitrary* evidence pattern  $e$ . As another example, in the EMYCIN calculus, a certainty factor needs to be supplied for each individual piece of evidence. Note that information complexity does not serve as a measure of how much effort is necessary to *gather* the required auxiliary information.

The *time complexity* is also of concern. It is a measure of how much computational time is required to compute the certainty of a *particular* disease  $d_i$ , given the *particular* evidence pattern  $e_j$ . Note that, in our formulation, information complexity is a measure of the information required to solve an *arbitrary* diagnostic problem, whereas time complexity is a measure of the time required to solve a *particular* diagnostic problem.

The specific details of how an uncertainty calculus is integrated into an expert system will generally affect the overall complexity of that system. For example, in the MYCIN program a certainty factor threshold of .2 was empirically selected [Shor84]; this threshold is used to limit search, and thus affects the overall time complexity of the program.

The specific details of the expert system will also interact with the complexity of the chosen calculus. For example, if an expert system uses the Bayesian approach, and contains only one rule, " $e_7 \rightarrow d_4$ ," then there is a need only for  $P(d_4 | e_7)$ ; i.e., the number of required probabilities is *not* exponential.

These sorts of interactions between expert systems and uncertainty calculi are quite important; however, they are beyond the scope of this paper.

In order to make the presentation of this paper consistent, some liberties have been taken with the notation and terminology used in the source references. In addition, certain aspects of the calculi have been simplified for ease of presentation.

## AI.2. Probabilistic Methods (Bayes' Rule)

Modern probability theory began in the 1600s when Blaise Pascal and other mathematicians set out to rigorously analyze games of chance [Olki80]. Since then, probability theory has been extensively developed, and applied to many fields. In recent years, artificial intelligence researchers have used probabilistic methods to handle uncertainty in knowledge-based systems; e.g., [Duda76].

### *Fundamentals of the Bayesian approach:*

In order to utilize a probabilistic method for handling uncertainty, it is necessary to gather a set of *prior probabilities* which describe the population. Evidence pertaining to a given individual is interpreted in the light of the prior probabilities. The result of this analysis is a set of *posterior probabilities*.

Prior probabilities may be determined by means of a frequency or other statistical analysis. Such a statistical analysis presupposes that adequate data is available to describe the population; if such data is not available then prior probabilities may be acquired by more subjective means (e.g., experts' estimates).

There has been much philosophical discussion concerning the definition of probability. It has been defined as a relative frequency ("objective probability"), as a measure of belief ("subjective probability"), and in other ways as well. Regardless of the exact definition chosen, expositions of probability theory are based on a three part axiomatization. The axioms are generally of the form:

- (1)  $P(\text{impossible-}t) = 0$ ;
- (2)  $P(\text{sure-}t) = 1$ ;
- (3) If  $t_1$  and  $t_2$  are mutually exclusive, then  
 $P(t_1 \cup t_2) = P(t_1) + P(t_2)$ ,

where  $P$  denotes probability,  $t$  is an event, and  $0 \leq P(t) \leq 1$ .

Computational methods can be constructed using the axioms as a foundation, and these methods can then be applied to problem solving in various domains. For example, there has been much work on devising probabilistic methods for medical diagnosis; these techniques are well reviewed in [Crof72], [Roge79], and [Ward78].

Many computational methods are based on the theorem known as *Bayes' rule*. One formulation of Bayes' rule is the following:

$$P(d_i | e_j) = \frac{P(d_i \cap e_j)}{P(e_j)} = \frac{P(d_i)P(e_j | d_i)}{\sum_{k=1}^m P(d_k)P(e_j | d_k)}$$

where  $d_i$  is one of  $m$  mutually exclusive (disjoint) diseases,  $e_j$  is an evidence pattern,  $P(d_i)$  is the (a priori) probability of disease  $d_i$ ,  $P(e_j | d_i)$  is the (conditional) probability of evidence pattern  $e_j$  given disease  $d_i$ , and  $P(d_i | e_j)$  is the (conditional) probability of disease  $d_i$  given evidence pattern  $e_j$ . The  $d$ 's must partition the sample space (i.e., must be exhaustive);

therefore, depending on the particular domain, one of the  $d$  might be *health* (the absence of disease). It should be noted that Bayes' rule is a direct consequence of the three probability axioms.

*Computational complexity:*

In the Bayesian<sup>3</sup> framework, the computation of the certainty of  $d_i$ , given  $e_j$ , is simply the calculation of  $P(d_i|e_j)$ . In order to analyze the complexity we begin with the following model: Assume that there exists only one disease,  $d_1$ . Since Bayes' rule requires partitioning of the sample space,  $d_2$  will denote the absence of disease  $d_1$  (i.e.,  $D = \{d_1, d_2\}$ ).

Given this model, Bayes' rule becomes

$$P(d_i|e_j) = \frac{P(d_i)P(e_j|d_i)}{\sum_{k=1}^2 P(d_k)P(e_j|d_k)} \quad (1)$$

In this model, the problem of interest is to compute the probability of  $d_i$  given  $e_j$  (i.e.,  $P(d_i|e_j)$ ). Given a *particular*  $e_j$ , the problem is solved by looking up values for  $P(d_1)$ ,  $P(d_2)$ ,  $P(e_j|d_1)$ , and  $P(e_j|d_2)$  (these probabilities having been determined during the initial analysis), and then applying (1).

Since there are  $2^n$   $e_j$ 's, there are  $2^n$  probabilities  $P(e_j|d_1)$ , and  $2^n$  probabilities  $P(e_j|d_2)$ . Therefore, in order to be able to compute  $P(d_i|e_j)$  given *arbitrary*  $e$ , it is necessary to determine  $1+1+2^n+2^n = 2(1+2^n)$  probabilities during the initial analysis. Even in this simple model, the information complexity is  $O(2^n)$ . The time complexity, assuming fixed time to look up the required probabilities, is a constant.

In a more general model, many diseases are possible. If there are

$m - 1$  possible disjoint diseases (with  $d_m$  denoting the absence of disease), then in order to compute  $P(d_i|e_j)$  it is necessary to look up values for  $P(d_1), \dots, P(d_m), P(e_j|d_1), \dots, P(e_j|d_m)$ ; therefore,  $m(1 + 2^n)$  probabilities are required for arbitrary  $e$ , and the information complexity is  $O(m \cdot 2^n)$ . In this model, assuming fixed time to look up the required probabilities, the time complexity is  $O(m)$ .

In an even more general model, diseases need not be disjoint (this requires a slight reformulation of Bayes' rule). In this model, the problem of interest is to compute  $P(d'|e_j)$  where  $d'$  is a disease complex (conjunction of zero or more diseases). If there are  $m$  diseases, then there are  $2^m$  disease complexes; therefore, the number of required probabilities =  $2^m(1+2^n) = O(2^{m+n})$ . In this model, assuming fixed time to look up the required probabilities, the time complexity is  $O(2^m)$ .

It is apparent that the probabilities  $P(d_i|e_j)$  can be precomputed for all  $i$  and  $j$ . This does not affect the order of the information complexity (it remains exponential), but it does reduce the time complexity to a constant. The reason that this is not done is that, in practice, Bayes' rule is always used with one or more assumptions (e.g., independence); these assumptions greatly reduce the information complexity.

The above models assume that a given  $s$  (e.g., *fever*) is either present or absent. For certain signs and symptoms merely stating the absence or presence may not be adequate (e.g., how severe was the fever?). In such cases it is necessary to be more fine-grained. This can be done by subdividing  $s$ , hence expanding the set  $S$ ; e.g.,  $S = \{\text{fever, chills}\}$  is expanded to  $S = \{\text{low-fever, moderate-fever, high-fever, chills}\}$ . Note that by per-

forming this expansion each  $s$  is still considered to be present or absent; however, there are now more  $s$ 's to choose from. The price of such an expansion is an increase in complexity.

It is possible to adapt the Bayesian approach to allow the  $s$ 's to be continuous valued [Crof72]; however, this requires strong assumptions, (e.g., modeling the  $s$  as continuous probability distributions), and a reformulation of Bayes' rule.

*Reducing the complexity:*

Due to the exponential number of required prior probabilities, Bayes' rule is intractable for all but the simplest of problems. As a result, systems that use Bayes' rule inevitably make one or more assumptions which reduce the number of required probabilities. Assumptions can be made either about the possible diseases (e.g., the *disjoint disease* assumption), or about the evidence (e.g., the *independent evidence*, *restricted evidence*, *probability formula*, and *maximum entropy* assumptions).

The *disjoint disease* assumption is that a patient can have only one disease at a time. This assumption reduces the  $2^m$  term to  $m$ , e.g.,  $2^m(1 + 2^n)$  is reduced to  $m(1 + 2^n)$ . In some cases, this is a reasonable assumption; for example, young patients generally only have one disease. In other cases, however, this is not a reasonable assumption; for example, elderly patients tend to have several concurrent diseases.

The *independent evidence* assumption is that the conditional probability of an evidence pattern equals the product of the conditional probabilities of each of its constituent elements [Crof72] (remember that each evi-

dence pattern,  $e$ , is a conjunction of zero or more signs and symptoms). For example, if  $e_j = s_2 \ \& \ s_7 \ \& \ s_9$  (where  $s_2, s_7, s_9 \in S$ ) then

$$P(e_j | d_i) = P(s_2 | d_i) \cdot P(s_7 | d_i) \cdot P(s_9 | d_i) .$$

This assumption reduces the  $2^n$  term to  $n$ , e.g.,  $m(1 + 2^n)$  is reduced to  $m(1 + n) = O(mn)$ . Other formulations of independence are possible [Char83], [Schw87], but the reduction in complexity is the same.

There is much discussion in the literature concerning the validity of the independent evidence assumption. It is clear that the signs and symptoms seen in a given disease are generally *not* independent, since they are manifestations of the same underlying pathology. Nevertheless, Charniak [Char83] argues that the independent evidence assumption is reasonable in practice.

On the other hand, Norusis [Noru75] presents results which demonstrate that, in cases of slight dependence, the assumption of independence can lead to large rates of error. This has been observed, in fact, in a system for the diagnosis of heart disease [Szo178]. In addition, [Lesm85] and [Pedn81] discuss some difficulties with the Prospector model [Duda76] (Bayesian with independence assumptions); but see also [Char83] and [Pear82].

The *restricted evidence* assumption is that only certain evidence patterns will be observed in practice, and that, therefore, only the probabilities for those evidence patterns need to be acquired. Since the probabilities  $P(e_j | d_i)$  must sum to 1 over all  $j$ , for a given  $d_i$ , this assumption may require a normalization of the acquired probabilities (e.g., compute the sum,  $k$ , of each of the acquired probabilities, and then divide each proba-

bility by  $k$ ). The more restricted the evidence (the less the number of probabilities to be acquired), the more the  $2^n$  term is reduced. The validity of the restricted evidence assumption increases as  $k$  approaches 1.

The *probability formula* assumption is that it is possible to supply all of the prior probabilities  $P(e_j|d_i)$ , for each  $d_i$ , by means of a formula,  $f_i(j)$ , where  $f_i(j)$  defines some probability mass function (e.g.,  $P(e_j|d_i) = f_i(e_j) = \binom{w}{j} z^j (1-z)^{(w-j)}$  where  $w$  and  $z$  are constants, and  $j$  is a variable). Given such a formula, any of the  $2^n$  probabilities for a given  $d_i$  can be computed as needed. Therefore, since only the formulas need to be stored, the  $2^n$  term is reduced to 1; e.g.,  $m(1 + 2^n)$  is reduced to  $m(1 + 1) = 2m$ .

In certain cases, it may be reasonable to assume that prior probabilities can be given by a formula—for example, if the  $e$  are conjunctions of only a few (e.g., one or two)  $s$ 's. However, when the  $e$  are conjunctions of many  $s$ 's (such as in the models developed above) it is generally not feasible to provide a probability formula. (Recall that  $e_1, \dots, e_{2^n}$  is an enumeration of all conjunctions of signs and symptoms; it is extremely unlikely that the  $e$ 's could be ordered such that their probabilities are a function of their position in the enumeration.)

A variation on the probability formula assumption is the *maximum entropy* assumption [Chee83], [Lemm82]. This assumption distributes prior probabilities in such a way as to maximize the entropy function,  $H$ , where

$$H = -\sum_{ij} P(d_i \cap e_j) \text{Log}(P(d_i \cap e_j)).$$

The basic idea is to acquire all available prior probabilities (called

*constraints* in this approach) and then distribute the remaining probability to maximize  $H$ . The less the number of constraints supplied, the more the information complexity is reduced.

The rationale for the maximum entropy principle is that unless there is knowledge to the contrary, probability should be distributed as "evenly" as possible. The validity of the maximum entropy assumption is related to the nature of the underlying domain.

The disjoint disease assumption is used in nearly all Bayesian systems. Most Bayesian systems also assume independent evidence. The restricted evidence assumption is sometimes made in systems which use estimated probabilities (to be described). The probability formula assumption is only infrequently used in Bayesian systems, although it is in wide use as the basis for most statistical procedures.

Variations of the assumptions are also possible. For example, Pearl investigates the case where the probabilistic dependencies may be modeled as a tree or graph structure [Pear87]. As another example, Hummel and Manevitz [Humm87] develop a notion of  $\alpha$ -independence, where  $\alpha$  is a parameter,  $\alpha(s_1, s_2) = 1$  means complete independence, and  $\alpha(s_1, s_2) = 0$  means complete dependence. Other assumptions can also be made; e.g., "all diseases are equally likely."

*Advantages of the Bayesian approach:*

- (1) Computational methods based on Bayes' rule have an axiomatic foundation, and well understood mathematical properties.
- (2) Bayesian methods generally require only a modest amount of compu-

tation time, since the time complexity is not a function of  $n$ .<sup>4</sup>

*Disadvantages of the Bayesian approach:*

- (1) The major difficulty with the Bayesian approach is the exponential number of prior probabilities that are required. This combinatorial explosion inevitably leads to the imposition of simplifying assumptions which may or may not be valid for the given domain.
- (2) Another difficulty is that the statistical analyses to determine the prior probabilities tend to require a massive amount of data. Even after having reduced the number of required probabilities by making assumptions, the volume of data needed may still be extremely large.

Ideally, prospective studies on the population are used to gather the needed data. However, population studies are expensive and not always possible. Another alternative is exhaustive review of the literature—but this, too, is not always possible.

If the required data is not available (or is too massive) then the prior probabilities may be acquired by a more subjective method. One approach that is used is to ask "experts" to supply estimates of the probabilities. Empirical studies on the effectiveness of this approach have given mixed results. For example, Tversky and Kahneman [Tver74] report that people are not very good at estimating probabilities. In addition, Leaper et al [Leap72] reports that a system for the diagnosis of acute abdominal pain (the system uses "a variant of Bayes' theorem" [Horr72]) was 83% accurate when using subjective estimates, but 91% accurate when using survey data from the litera-

ture. On the other hand, Gustafson et al [Gust71] reports that a Bayesian system for the diagnosis of thyroid disease performed nearly as well using estimated probabilities, as it did using probabilities derived from actuarial data. Gustafson et al [Gust73] also reports that the incorporation of subjective estimates of conditional dependencies improved the performance of a system for the diagnosis of thyroid dysfunction.

In [Humm87] Hummel and Manevitz extend the notion of experts' estimates. They suggest that by polling collections of experts, "subjective" opinions can be mapped into "objective" statistics.

- (3) Yet another difficulty with the Bayesian approach is that prior probabilities are highly dependent on context. For example, the probability of a given disease depends on geographic location (e.g., malaria is much more common in the tropics). Disease probabilities also change over time (e.g., fifty years ago polio was common; now it is extremely rare).

### **AI.3. The Dempster-Shafer Theory**

The Dempster-Shafer theory is based on Dempster's work from the 1960s [Demp67], and later extensions by Shafer [Shaf76]. In the past few years, the theory has begun to receive increasing attention from the AI community as a potentially useful tool for manipulating uncertain information [Garv81].

*Fundamentals of the Dempster-Shafer theory:*

Let  $\Theta$  be a set of  $m$  diseases which are mutually exclusive and exhaustive.<sup>5</sup> In the Dempster-Shafer theory the set  $\Theta$  is called a *frame of discernment*, and each disease in  $\Theta$  is called a *singleton*. The power set of  $\Theta$  (the set of all subsets of  $\Theta$ ) will be denoted by  $2^\Theta$ . Any subset of  $\Theta$  (including  $\Theta$  itself, the empty set,  $\phi$ , and each of the singletons) is called a *hypothesis*.

**Example 1:**

Suppose that there are five infectious disease organisms,  $V_1, V_2, V_3, B_1$ , and  $B_2$ . Suppose, in addition, that  $V_1$  and  $V_2$  are the two members of the virus family  $V_{1,2}$ , that  $V_3$  is also a virus, and that  $B_1$  and  $B_2$  are bacteria. Then the set  $\{V_1, V_2, V_3, B_1, B_2\}$  is the frame of discernment, and the singletons are  $V_1, V_2, V_3, B_1$ , and  $B_2$ . Any subset of  $\{V_1, V_2, V_3, B_1, B_2\}$  is a hypothesis; e.g.,  $\{V_1, V_2, V_3\}$  is the hypothesis that the organism is a virus.

Recall that each piece of evidence (sign or symptom) is considered to be either present or absent. Associated with each piece of evidence is a *basic probability assignment* (abbreviated *bpa*). A bpa is a function which assigns to each hypothesis a *basic probability number* (abbreviated *bpn*). The basic probability numbers are in the range  $[0,1]$  and must sum to 1 over all hypotheses. If the sum of the bpn's for all of  $\Theta$ 's proper subsets equals  $k$ , then the bpn for  $\Theta$  is defined to be  $1 - k$ . Furthermore, the bpn for  $\phi$  is defined to be zero. If  $g$  is a bpa, and  $A \in 2^\Theta$ , then  $g(A)$  denotes the bpn for  $A$ .

**Example 2:**

Given a patient who is known to have an infectious disease, and given a piece of evidence (e.g., result of a blood test) which suggests organism family  $V_{1,2}$  to the degree .6, and organism  $V_3$  to the degree .2. Then  $\text{bpa}(\{V_1, V_2\}) = .6$ , and  $\text{bpa}(\{V_3\}) = .2$ . The bpn's must sum to 1; therefore,  $\text{bpa}(\{V_1, V_2, V_3, B_1, B_2\}) = 1 - (.6 + .2) = 1 - .8 = .2$ . The bpn for all other subsets is zero (e.g.,  $\text{bpa}(\{V_1\}) = 0$  and  $\text{bpa}(\{V_2\}) = 0$ ). Intuitively,  $\text{bpa}(\{V_1, V_2\}) = .6$  because the piece of evidence suggests organism family  $V_{1,2}$  to the degree .6; however,  $\text{bpa}(\{V_1\}) = 0$  because the evidence does not specifically suggest  $V_1$  (likewise for  $V_2$ ).

The *belief* in  $A$  (sometimes also called the *lower probability* of  $A$ ) is the sum of the bpn of  $A$  plus the bpn's of all of  $A$ 's proper subsets. A function which specifies the belief for all  $A$ ,  $A \in 2^\Theta$ , is known as a *belief function* (abbreviated *Bel*). If  $G$  is a belief function, and  $A \in 2^\Theta$ , then  $G(A)$  denotes the belief in  $A$ . It should be clear that each bpa has one and only one corresponding Bel, and vice versa.

**Example 3:**

The belief that the organism is a virus (i.e.,  $\text{Bel}(\{V_1, V_2, V_3\}) = \text{bpa}(\{V_1, V_2, V_3\}) + \text{bpa}(\{V_1, V_2\}) + \text{bpa}(\{V_1, V_3\}) + \text{bpa}(\{V_2, V_3\}) + \text{bpa}(\{V_1\}) + \text{bpa}(\{V_2\}) + \text{bpa}(\{V_3\}) = 0 + .6 + 0 + 0 + 0 + 0 + .2 = .8$ ).

Note that if  $A$  is a hypothesis, then  $0 \leq \text{bpa}(A) \leq 1$ , and  $0 \leq \text{Bel}(A)$

$\leq 1$ . Intuitively, 1 means total belief in a hypothesis, and 0 means lack of belief (not to be confused with disbelief). The distinction between  $\text{bpa}(A)$  and  $\text{Bel}(A)$  is crucial:  $\text{bpa}(A)$  is the amount of belief committed *exactly* to  $A$ , whereas  $\text{Bel}(A)$  is the *total* amount of belief committed to  $A$ . The two numbers are not necessarily the same ( $\text{bpa}(A) \leq \text{Bel}(A)$ ).

**Example 4:**

$\text{Bel}(\{V_1, V_2, V_3\}) = .8$ , but  $\text{bpa}(\{V_1, V_2, V_3\}) = 0$ , since the piece of evidence does not commit any belief to *exactly* that hypothesis.

Each piece of evidence will lead to a distribution of belief (alternatively,  $\text{bpa}$ 's) among the hypotheses (subsets). Intuitively, the more conclusive a piece of evidence, the more belief is distributed into smaller subsets.

**Example 5:**

A fever suggests infection (either viral or bacterial), but a low white blood count (*WBC*) suggests viral infection. Therefore, the piece of evidence, *fever*, might induce  $\text{bpa}(\{V_1, V_2, V_3, B_1, B_2\}) = 1$ ; whereas the piece of evidence, *low-WBC* might induce  $\text{bpa}(\{V_1, V_2, V_3\}) = .8$ , and  $\text{bpa}(\{V_1, V_2, V_3, B_1, B_2\}) = .2$ .

As noted above, a belief of 0 for a hypothesis does not mean disbelief; rather it means a lack of belief. The disbelief in a hypothesis,  $A$ , is given by  $\text{Bel}(A^c)$ , where  $A^c$  is used to denote the set-theoretic complement of  $A$ . The *plausibility* of  $A$  (sometimes also called the *upper probability* of  $A$ ) is defined to be  $1 - \text{Bel}(A^c)$ . Plausibility can be thought of as the extent

to which the evidence does not support the negation of a hypothesis.

**Example 6:**

(Same situation as in examples (1), (2), (3), and (4).) The disbelief that the organism is a virus =  $\text{Bel}(\{V_1, V_2, V_3\}^c) = \text{Bel}(\{B_1, B_2\}) = 0$ . The plausibility that the organism is a virus =  $1 - \text{Bel}(\{V_1, V_2, V_3\}^c) = 1 - \text{Bel}(\{B_1, B_2\}) = 1 - 0 = 1$ . Intuitively, the plausibility that the organism is a virus equals 1, because there is no evidence to the contrary.

Note that  $\text{Bel}(A)$  must be  $\leq$  plausibility( $A$ ) for all  $A$ . The interval between  $\text{Bel}(A)$  and plausibility( $A$ ) is known as the *belief interval*. Intuitively, the smaller the magnitude of the belief interval, the more conclusive the piece of evidence is with respect to the hypothesis.

If  $s_i$  is a piece of evidence which induces bpa  $g_i$ , and  $s_j$  is a piece of evidence which induces bpa  $g_j$  (where  $s_i, s_j \in S$ ), then the bpa induced by the conjunction of evidence ( $s_i$  &  $s_j$ ) is denoted by  $g_i \oplus g_j$ . The Dempster-Shafer theory defines the following combination rule (*Dempster's rule*) for determining  $g_i \oplus g_j$ :

$$(1) \quad g_i \oplus g_j(A) = \frac{\sum_{X \cap Y = A} g_i(X)g_j(Y)}{1-k}, \text{ when } A \neq \phi;$$

$$(2) \quad g_i \oplus g_j(A) = 0, \text{ when } A = \phi;$$

$$\text{where } k = \sum_{X \cap Y = \phi} g_i(X)g_j(Y), \text{ and where } X, Y \in 2^\Theta.$$

Intuitively, the combination rule narrows the hypothesis set by distributing the bpn's (hence, belief) into smaller and smaller subsets. Note that, by

definition,  $\phi$  must be assigned bpn zero. The division by  $1 - k$  is a *normalization* to force the bpn's to sum to 1.

One methodology for using the Dempster-Shafer theory is as follows: Each piece of evidence induces a bpa. (Evidence against a hypothesis is treated as evidence in favor of the complement of the hypothesis). The bpa's are then combined using the combination rule. The resulting bpa is then used to compute a Bel, and from Bel the plausibilities can be computed. For a more detailed introduction to the Dempster-Shafer theory, the reader is referred to the excellent presentation by Gordon and Shortliffe [Gord84].

*Computational complexity:*

In the Dempster-Shafer framework, the computation of the certainty of  $d_i$ , given  $e_j$ , is the calculation of  $\text{bpa}(d_i)$ . Recall that  $e_j$  is a conjunction of  $p$  pieces of evidence. Each individual piece of evidence induces a bpa, hence  $\text{bpa}(d_i)$  is computed by  $p - 1$  applications of Dempster's rule.

If  $|\Theta| = m$ , then in the worst case each bpa will need to specify  $(2^m - 2)$  bpn's (remember that the bpn for  $\phi$  is defined to be zero, and that any "leftover" belief is given to the frame of discernment). Therefore, since there are  $n$  possible pieces of evidence, the worst case information complexity is  $(n \cdot (2^m - 2)) = O(n \cdot 2^m)$ .

The Dempster combination rule computes products over all subsets of  $\Theta$ . Given that  $|\Theta| = m$ , then there are  $2^m$  subsets. Therefore, the (worst case) time complexity to combine two pieces of evidence is  $O(2^m \cdot 2^m) = O(2^{2m})$ . By induction, the worst case time complexity to combine  $p$  pieces

of evidence is  $O(2^m)$ .

The complexities are significantly reduced if the bpa's assign non-zero bpn's to only a few hypotheses (implicitly assigning zero to the other hypotheses). For example, it is possible to reduce the time complexity to linear by assuming that "each piece of evidence either confirms or denies a single proposition rather than a disjunction" [Barn81] (this assumption is roughly analogous to the restricted evidence assumption in the Bayesian approach). As another example, each bpa may specify as few as one bpn (giving bpn of 1 to a single hypothesis); in this case the information complexity is  $(n-1) = O(n)$ . In actual practice the complexities will range from linear to exponential depending on the particular bpa's (the authors would be grateful to receive any empirical data on this!).

The Dempster-Shafer theory assumes that all hypotheses are mutually exclusive. In many domains (e.g., medicine) this is not a tenable assumption. By defining the singletons within the frame of discernment to be disease complexes, this limitation can be overcome; however, the cost is a significant increase in complexity (e.g.,  $2^m$  hypotheses).

The Dempster-Shafer theory assumes that each  $s$  is either present or absent. As in the Bayesian approach, each  $s$  can be further subdivided if necessary (e.g., *fever* is expanded to *low-fever*, *moderate-fever*, *high-fever*); the cost is an increase in complexity. The  $s$  can also be modeled as continuous valued [Stra84], but this requires strong assumptions.

*Advantages of the Dempster-Shafer theory:*

(1) The information and time complexities *can* be quite low if certain con-

ditions are met.

- (2) Several aspects of the theory are intuitively pleasing. For example, the notion that evidence can be relevant to subsets, rather than just singletons, nicely captures the idea of narrowing in on a diagnosis; e.g., cardiovascular disease -> heart disease -> left ventricular disease -> specific diagnosis.

*Disadvantages of the Dempster-Shafer theory:*

- (1) The theory assumes that pieces of evidence are independent [Barn81], [Shaf76, pg. 57], [Zade84]. As was discussed in the section on Bayesian methods, it is not always reasonable to assume independent evidence.
- (2) There is no theoretical justification for the combination rule. According to Shafer [Shaf76, pg. 57], there is "no conclusive *a priori* argument for Dempster's rule ... the rule does *seem* (italics added) to reflect the pooling of evidence."
- (3) The potentially exponential complexities can be a source of difficulty.

#### **AI.4. Fuzzy Set Theory**

In 1965 Zadeh [Zade65] first presented the idea of fuzzy sets. Since then, there has been a great deal of work exploring the mathematical properties of these sets; there has also been work on applying fuzzy sets in various problem domains [Dubo80]. In recent years researchers have begun to investigate fuzzy sets as a potential tool for handling uncertainty

in artificial intelligence systems [Yage84].

*Fundamentals of fuzzy set theory:*

One motivation for fuzzy set theory is to avoid certain types of paradoxes which arise in the context of classical set theory. The following is an example of one such paradox:

A heap containing one stone is small.  
 If you add one stone to a small heap it remains small.  
 Therefore (by induction), every heap is small.

Before presenting the fundamentals of fuzzy set theory we first review some concepts from classical set theory. Let  $X$  denote a universe of objects, and let  $x$  denote an individual element from that universe. Let  $A$  be a subset of  $X$ . Then we say that  $\mu_A$  is the *characteristic function* (or *membership function*) of  $A$  if

$$\mu_A(x) = \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise.} \end{cases}$$

In addition, fuzzy theorists call the set  $\{0,1\}$  the *valuation set*, and call the value assigned by the function for a given  $x$  the *degree of membership*. Note that in classical set theory a degree of membership is either 0 or 1.

**Example:**

Let  $X$  be the set of integers and let  $E$  be the set of even integers. Choose  $x$  to be the integer 3. Then  $\mu_E(x) = 0$ . Now choose  $x$  to be the integer 8. Then  $\mu_E(x) = 1$ .

In *fuzzy set theory* [Dubo79], [Gain76] the valuation set is expanded to be the real interval  $[0,1]$ . That is, the degree of membership can be 0,1, or any value in between. Intuitively, the larger the value of  $\mu_A(x)$ , the more  $x \in A$ .

**Example:**

Let  $X$  be the set of all people and let  $T$  be the set of tall people. Choose  $x$  to be Tom, where Tom is 7 feet tall. Then  $\mu_T(x) = 1$ . Intuitively, Tom is tall. Now choose  $x$  to be Dick, where Dick is 6 feet tall. Then  $\mu_T(x) = .5$ . Intuitively, Dick is "somewhat tall." Finally, choose  $x$  to be Harry, where Harry is 5 feet tall. Then  $\mu_T(x) = 0$ . Intuitively, Harry is not tall.

In order to use fuzzy set techniques it is necessary to supply a characteristic function for each set of interest. For example, a formula for the tallness of a person (where  $h$  is a person's height) might be

$$\mu_T(h) = \begin{cases} 0 & h < 5 \\ (h-5)/2 & 5 \leq h \leq 7 \\ 1 & h > 7. \end{cases}$$

It should be noted that the values assigned by a characteristic function are chosen by the person who constructs the function. For instance, in the preceding example, we felt that .5 is a reasonable number for a six foot individual (the first author's height is 6 feet 4 inches!); other authors might have chosen .46 or .65.

Also note that characteristic functions depend upon the particular context (i.e., universe): If  $X$  were the set of basketball players then  $\mu_T(x)$  might equal .9 for Tom, rather than 1.

Values assigned by a characteristic function should not be interpreted as probabilities. For example, there is *not* a 50% chance that Dick is tall; rather, Dick is 50% tall.

It should now be clear how fuzzy set theory avoids the paradox

presented before: if  $X$  is the set of heaps, and  $S$  is the set of small heaps, then  $\mu_S$  for a heap with eight stones might equal .999, whereas  $\mu_S$  for a heap with nine stones might equal .998.

Various operations have been defined on fuzzy sets, the two most fundamental being  $\cap$  (set intersection) and  $\cup$  (set union). In his seminal paper Zadeh [Zade65] defined the operations as follows:

$$\mu_{A \cap B}(x) = \min(\mu_A(x), \mu_B(x));$$

$$\mu_{A \cup B}(x) = \max(\mu_A(x), \mu_B(x)).$$

Many other definitions have been proposed for intersection and union; see [Yage80], for some examples. Most proposed definitions for the operators seek to satisfy certain algebraic properties; e.g, monotonicity, commutativity, and associativity. In particular, intersection is usually, but not always, formulated as a triangular norm ( $T$ -norm),<sup>6</sup> and union is usually, but not always, formulated as a triangular conorm ( $T$ -conorm).<sup>7</sup> Most, but not all, work in fuzzy set theory uses the min and max definitions for intersection and union.

There has been a good deal of work relating multivalent logics to fuzzy set theory. Of particular concern are logics in which the truth of a proposition is a number in the real interval  $[0,1]$ ; i.e., if  $Q$  is a proposition, and  $v(Q)$  is the truth of  $Q$ , then  $0 \leq v(Q) \leq 1$ . These logics are sometimes called *fuzzy logics*, although Zadeh generally uses the term to refer to the logic underlying the theory of approximate reasoning (discussed below) [Dubo80]; there are other uses for the term *fuzzy logic*, as well [Gain76].

Several definitions for the logical connectives have been proposed.

These definitions are based on simple arithmetic functions and max and min [Dubo80]. For example, in one fuzzy logic,  $v(\neg Q) = 1 - v(Q)$ ;  $v(Q \rightarrow R) = \max(1 - v(Q), v(R))$ ; and so forth (where  $Q$  and  $R$  are propositions).

An extension of fuzzy set theory is the *theory of approximate reasoning*. One of the fundamental notions of approximate reasoning is that of a *possibility distribution* [Zade78]. Intuitively, a possibility is the degree to which something is feasible. For example, if Tom is a tall person, then it is quite possible that Tom is 6.5 feet high (e.g., the possibility that Tom is 6.5 feet high is .75). It is less possible that Tom is 5.5 feet high (e.g., the possibility that Tom is 5.5 feet high is .25). Note that possibility is not the same as probability; e.g., the probability that Tom is 6.5 feet high is certainly not .75. Also note that whereas probabilities must sum to 1, the same is not true for possibilities. Another notion from the theory of approximate reasoning is *necessity*; intuitively, necessity is the degree to which something must be true.

A fundamental tool in the theory of approximate reasoning is the *rule of compositional inference*. In a classical inference system, given the premises

$x$  is  $A$  implies  $y$  is  $B$ ,

$x$  is  $A$

the conclusion,  $y$  is  $B$ , is inferred by modus ponens. The rule of compositional inference extends the above to the following situation: from the premises

$x$  is  $A$  implies  $y$  is  $B$ ,

$x$  is  $A'$

the conclusion,  $y$  is  $B'$ , is inferred, where  $A$  and  $A'$ , and  $B$  and  $B'$ , are "similar" fuzzy sets. There are actually several formulations of the inference rule [Fuka80]; these formulations are based on compositions of max and min.

Note that fuzzy logics deal with propositions of the form " $x \in A$ ," where  $A$  is a fuzzy set and the truth of the proposition is given by  $\mu_A(x)$ . Approximate reasoning, on the other hand, deals with propositions of the form " $x$  is  $A$ ," where  $A$  is a fuzzy set. Thus, fuzzy logics manipulate numerical values (acquired from a fuzzy membership function), whereas approximate reasoning manipulates fuzzy sets.

*Computational complexity:*

Since there are many variations and extensions of fuzzy set theory, there are multiple ways of computing the certainty of  $d_i$ , given  $e_j$ . Adlassnig [Adla82] reviews some of the methods that have been proposed in the context of medical diagnosis.

One simple method is the following: Each  $s$  is viewed as a fuzzy set. Each observation about the patient is mapped into a fuzzy degree of membership for a particular sign or symptom (e.g., temperature=38.0  $\rightarrow$   $\mu_{fever} = .81$ ). This mapping is done by a characteristic (membership) function. There are  $n$  possible signs and symptoms, hence,  $n$  characteristic functions are needed.

Each disease is viewed as having a prototypical pattern of signs and symptoms (expressed as a sequence of conjunctions and disjunctions; the maximum length of the sequence is  $n$ ). The computation of the certainty

of  $d_i$ , then, is just the computation of the sequence of conjunctions and disjunctions using the degrees of membership. Since  $n$  characteristic functions are needed, the information complexity is  $O(n)$ . If the usual min and max definitions are chosen for intersection and union, the time complexity is also  $O(n)$ .

More sophisticated formulations are possible. For example, the pieces of evidence may themselves be fuzzy sets (rather than degrees of membership), and the certainty of  $d_i$  would then be computed by using the rule of compositional inference [Sanc79]. We have not analyzed the complexity of this approach in detail; however, it appears to be polynomial in information and time.

*Advantages of fuzzy set theory:*

- (1) Fuzzy set techniques have low information and time complexity (depending on the definitions of the operators and the particular methods used).
- (2) Fuzzy set techniques seem to be a good solution for some of the problems that arise due to lexical imprecision [Henk88Char].
- (3) Since there are many possible definitions for the operators, and many ways to formulate a problem, fuzzy set techniques are quite flexible.
- (4) Several aspects of the theory are intuitively pleasing. For example, defining the valuation set to be  $[0,1]$  rather than  $\{0,1\}$  allows one to easily represent the relative severity of an entity (e.g.,  $\mu_{\text{Fever}}(37.5) = .77$ , but  $\mu_{\text{Fever}}(38.5) = .99$ ).

*Disadvantages of fuzzy set theory:*

- (1) It is not always clear how to construct reasonable membership functions. Various methods have been proposed [Norw84] including the use of statistical data, and the composition of simpler functions, but no completely general approach seems to exist yet.
- (2) The choice of appropriate definitions for the operators can be problematic. As Zadeh himself has acknowledged, different definitions are needed in different situations [Zade76]; however, it is not always clear which definitions should be used. Unfortunately, those definitions with nice mathematical properties often perform poorly when applied to real-world problem solving [Zimm80]. On the other hand, those definitions which perform well in a given problem may be ad-hoc and lack mathematical rigor.
- (3) The inherent flexibility of fuzzy set theory can also be a disadvantage, since there is little guidance as to which methods to use to solve a given problem.

#### **AI.5. The MYCIN and EMYCIN Calculi**

The MYCIN calculus [Shor75Amod] was developed during the course of Shortliffe's doctoral work on the MYCIN system. This calculus was modified during work on EMYCIN [Buch84, pp. 215-216], and the modified version was incorporated into all of the EMYCIN systems (including MYCIN itself).<sup>8</sup> The EMYCIN calculus has also been incorporated into several commercially available expert systems shells. There is some con-

fusion regarding the calculus: most papers in the literature discuss the original (MYCIN) version, whereas it is the modified (EMYCIN) version that is (and was) actually used in practice.

*Fundamentals of the MYCIN calculus:*

Given a parameter  $P$  which can take on one value at a time from the set  $V$ . A *hypothesis* is a statement of the form "the value of  $P$  is  $v$ ," where  $v \in V$ .

The fundamental idea behind the the MYCIN calculus is that for each hypothesis, certain pieces of evidence tend to confirm it, while others tend to disconfirm it.

More formally, if  $h$  is a hypothesis, and  $s$  is a piece of evidence that tends to confirm  $h$ , then  $MB[h,s]$  is the belief in  $h$  based on  $s$ .<sup>9</sup>  $MB$  is called the *measure of belief*, and for each  $h$  and  $s$ ,  $0 \leq MB[h,s] \leq 1$ . Intuitively, if  $s$  absolutely confirms  $h$ , then  $MB[h,s] = 1$ ; if  $s$  has "nothing to do" with  $h$  then  $MB[h,s] = 0$ ; and if  $s$  tends to confirm  $h$ , but is not conclusive, then  $MB[h,s] = x$ , where  $0 < x < 1$ .

An analogous measure,  $MD$ , is defined for evidence which tends to disconfirm a hypothesis.  $MD$  is called the *measure of disbelief*. For a given  $h$  and  $s$ , only one of  $MB[h,s]$  and  $MD[h,s]$  can be non-zero, since a single piece of evidence cannot both confirm and disconfirm a given hypothesis.

In general, more than one piece of evidence may tend to confirm a hypothesis. The MYCIN calculus defines the following combining function:

$$MB[h, s_i \& s_j] = MB[h, s_i] + MB[h, s_j] - MB[h, s_i]MB[h, s_j] \quad (2)$$

where  $MB[h, s_i \& s_j]$  is the measure of belief in  $h$  due to both pieces of evidence ( $s_i$  and  $s_j$ ). An analogous function is defined for  $MD[h, s_i \& s_j]$ . As before,  $e$  is used to denote the conjunction of zero or more pieces of evidence; e.g.,  $e$  could replace  $s_i$  and  $s_j$  in the above formula.

In order to compute the measure of belief of the conjunction of two hypotheses, the following combining function is defined:

$$MB[h_i \& h_j, e] = \min(MB[h_i, e], MB[h_j, e]) . \quad (3)$$

The measure of belief in the disjunction of two hypotheses is defined to be:

$$MB[h_i \text{ or } h_j, e] = \max(MB[h_i, e], MB[h_j, e]) . \quad (4)$$

Analogous formulas are defined for MD (but with max replaced by min and vice versa).

The *certainty factor* of a given hypothesis is defined to be the *net* belief in that hypothesis. It is given by:

$$CF[h, e] = MB[h, e] - MD[h, e] . \quad (5)$$

Note that  $-1 \leq CF \leq 1$ . Intuitively,  $CF[h, e] = 1$  if, based on the evidence, the hypothesis is known with certainty to be true;  $CF[h, e] = -1$  if, based on the evidence,  $h$  is known with certainty to *not* be true; and  $CF[h, e] = 0$  if there is no evidence regarding the the hypothesis, or the confirming evidence is balanced by the disconfirming evidence ( $MB = MD$ ).

One difficulty with the calculus described above is the computational expense of storing separate MB's and MD's for every hypothesis (i.e., MB and MD for every possible value of every parameter). A more serious difficulty is that one piece of disconfirming evidence can overwhelm many pieces of confirming evidence (and vice versa).

**Example:**

If ten pieces of evidence support a hypothesis with  $MB[h, s_i] = .8, i = 1..10$ , then the cumulative  $MB[h, e] \approx .999$  by (2). Suppose that the piece of evidence  $s_{11}$  disconfirms the hypothesis with  $MD[h, s_{11}] = .8$ . By (5),  $CF = MB - MD \approx .999 - .8 \approx .2$ . Note that the computed CF would be essentially the same even if there were a million pieces of evidence confirming the hypothesis with  $MB = .8$ ; this is not in accord with intuition.

**The EMYCIN calculus:**

The EMYCIN calculus was developed to overcome the difficulties of the MYCIN calculus. One modification was to keep track of only the CF for each hypothesis, rather than separate MB's and MD's. This change, in turn, required the redefinition of the combining functions. Two of the functions used in EMYCIN are:

$$CF[h, s_i \& s_j] = \begin{cases} X + Y - XY & X, Y \text{ both } > 0 \\ (X + Y) / (1 - \min(|X|, |Y|)) & \text{one of } X, Y < 0 \\ X + Y + XY & X, Y \text{ both } < 0 \end{cases} \quad (6)$$

where  $X = CF[h, s_i]$  and  $Y = CF[h, s_j]$ ;

$$CF[h_1 \& \dots \& h_n, e] = \min(CF[h_1, e], \dots, CF[h_n, e]). \quad (7)$$

In both calculi, a CF is associated with each implication (i.e., rule):  $h_i \rightarrow h_j (C)$ . For example,  $h_i \rightarrow h_j (.7)$ . Intuitively, if  $h_i$  is known with complete certainty ( $CF[h_i, e] = 1$ ), then  $h_j$  is known with certainty .7. Since the premise of an implication is not always known with certainty, another combining function is needed. EMYCIN defines

$$CF[h_j, h_i] = C \cdot (CF[h_i, e]) \quad (8)$$

where  $C$  is the certainty of  $h_i \rightarrow h_j$ . (Similar formulas are defined in the MYCIN calculus to compute MB and MD for implications).

In practice, the EMYCIN calculus initializes each hypothesis to  $CF = 0$ . The CF's may be assigned directly by the user in response to a query, or may be changed by means of the combining functions as evidence is combined and deductions are made.

*Computational complexity:*

In order to utilize a piece of evidence, the MYCIN calculus needs 2 additional items of information (i.e., an MB and MD must be provided with each piece of evidence). The EMYCIN calculus needs one additional piece of information (i.e., a CF with each piece of evidence). Therefore, both MYCIN and EMYCIN have linear information complexity.

By examination of the combining functions (2), (3), and (4), it is clear that, in the MYCIN calculus, the time complexity for the combination of  $p$  hypotheses and/or pieces of evidence is  $O(p)$ . Likewise, examination of functions (6), (7), and (8), reveals that, in the EMYCIN calculus, the time complexity for the combination of  $p$  hypotheses and/or pieces of evidence is  $O(p)$ .

*Advantages of the MYCIN and EMYCIN calculi:*

- (1) Both calculi have linear information complexity and linear time complexity.

*Disadvantages of the calculi in a rule-based context:*

Certain limitations of the calculi became apparent within the context of the MYCIN program:

- (1) One limitation relates to the construction of rules. For example, the three rules

$$A \rightarrow D, B \rightarrow D, C \rightarrow D \quad (9)$$

are *logically* equivalent to the single rule

$$(\$AND(\$OR A B C)) \rightarrow D \quad (10)$$

(\$AND is similar to Boolean AND—see [Shor84] for a definition; likewise for \$OR). However, invoking the three rules (9) can lead to a cumulative CF for *D* which differs from that computed by invocation of the single rule (10) [Buch84, pg. 212].

- (2) The order of rule invocation can also cause difficulties. For example, in a rule-based system certain rules can be self referential (recursive). It can be shown that certainty factors are *not* commutative with respect to the order of invocation of these rules; i.e., given the same evidence it is possible to get different CF's depending on the order in which the rules are invoked [Cend84, pp. 265-266], [Shor84].

*General disadvantages of the calculi:*

- (1) The combining functions are ad-hoc [Gord84].
- (2) Because of the nature of the combining functions, the MB's and MD's tend to converge quickly to the asymptote 1, while the CF's stay near 0. Therefore, the calculi do not perform well if there is a great deal of evidence to be combined. A possible solution is to use adjustable weights to damp the combining functions [Buch84, pp. 213-214], but

this has the disadvantage of being highly dependent on the particular problem to be solved.

- (3) The MYCIN and EMYCIN calculi were developed in the context of a very limited domain, i.e., the diagnosis and therapy of bacteremia. Certain characteristics of this context led to (implicit) assumptions in the construction of the combining functions.<sup>10</sup> For example, one assumption is that the prior probabilities for all competing hypotheses are small [Buch84, pg. 211]. Another assumption is that the various pieces of evidence are independent of each other [Shor75Amod]. Because of these assumptions, as well as other limitations, the calculi are simply not appropriate for many types of problems [Buch84, pg. 214].

#### AI.6. Comparison of the Four Calculi

Each of the four calculi views uncertainty from a different perspective: in the Bayesian calculus uncertainty is probability, in the Dempster-Shafer theory uncertainty is the degree of belief, in fuzzy set theory uncertainty is the degree of set membership, and in MYCIN/EMYCIN uncertainty is the degree of confirmation. In the Bayesian, Dempster-Shafer and fuzzy set calculi the uncertainty of an entity is specified by a numerical value in the range  $[0,1]$ ; in MYCIN/EMYCIN the range is  $[-1,1]$ .

Recall our diagnostic formulation: an uncertainty value is acquired for each piece of evidence,  $s$ , and these values are manipulated in some fashion to yield the uncertainty value of the diagnosis,  $d$ . The calculi differ

from each other in how and when they acquire uncertainty values, and in how these uncertainty values are manipulated.

In the Bayesian calculus uncertainty values (prior probabilities) may be acquired by means of a statistical analysis (e.g., frequency analysis). Given a large enough sample size, all observers would compute essentially the same uncertainty value for a piece of evidence; hence, these uncertainty values can be considered objective. Subjective methods can also be used to acquire probabilities; for example, expert's estimates may be used. The other three calculi do not have well defined objective methods for acquiring uncertainty values, hence, the values are determined by subjective methods. (One might argue that fuzzy set theory does have an objective method, i.e., characteristic functions; note, however, that there is no uniformly accepted construction technique for these functions, hence, they are ultimately subjective.)

Uncertainty values may either be acquired in advance and stored for each piece of evidence, or they may be supplied to the system as needed. In the Bayesian calculus the uncertainty values (prior probabilities) are generally acquired in advance and stored by the system. In the other three calculi the uncertainty values are supplied along with the pieces of evidence: in the Dempster-Shafer calculus a bpa is provided along with each piece of evidence, in the fuzzy set calculus a degree of membership is provided along with each piece of evidence, and in MYCIN/EMYCIN a CF (or MB and MD) is supplied along with each piece of evidence. Note that the times at which the uncertainty values are acquired is not a fundamental aspect of any of the calculi; what we described above are the conventions

generally followed in practice.

In the Bayesian and Dempster-Shafer calculi each  $s$  (sign or symptom) is an attribute which is considered to be either present or absent in a patient. As a result, in these calculi the computation of the certainty of  $d_i$  given  $e_j$  is conceptually a one step process ( $e_j \rightarrow d_i$ ). In fuzzy set theory and MYCIN/EMYCIN, however, each sign and symptom has an associated uncertainty value. Hence, the computation of the certainty of  $d_i$  is conceptually a two step process: (1) compute the certainty of  $e_j$  by combining in some fashion the uncertainty value of each of the constituent signs and symptoms, (2) compute the certainty of  $d_i$ . It is interesting to note, in this regard, that even though fuzzy set theory and MYCIN/EMYCIN attempt to "do more" than the other two calculi (i.e., allow  $e_j$  itself to be uncertain), they have lower complexities (the Bayesian calculus has exponential information complexity and the Dempster-Shafer calculus is potentially exponential in both information and time, whereas fuzzy set theory and MYCIN/EMYCIN are linear in information and time). Of course, the Bayesian and Dempster-Shafer calculi can be adapted to handle uncertain  $e_j$ , but this requires an increase in complexity (or strong assumptions).

It is also worth noting that in the Bayesian and Dempster-Shafer calculi the weights of inferences; e.g.,

$$e_j \xrightarrow{w} d_i$$

are implicitly contained by the prior probabilities (or bpa's), whereas in MYCIN/EMYCIN and fuzzy set theory (depending upon the formulation), these weights are made explicit; e.g.,  $e_j \rightarrow d_i(.7)$ .

In a sense, the Bayesian and Dempster-Shafer calculi contain more

information than the fuzzy set and MYCIN/EMYCIN calculi: In the Bayesian calculus, given the necessary prior probabilities,  $P(d_i|e_j)$  can be immediately computed. Similarly, in the Dempster-Shafer theory, given the bpa's for each piece of evidence, the bpn of  $d_i$  can be immediately computed. In MYCIN/EMYCIN, however, the CF of  $e_j$  is not, by itself, enough to compute the CF of  $d_i$ ; additional knowledge, e.g., a production rule, is also necessary (fuzzy set theory is similar, depending on the particular formulation used). In other words, one might say that the fuzzy set and MYCIN/EMYCIN calculi are more tightly coupled to the underlying knowledge of the expert system, than are the Bayesian and Dempster-Shafer calculi.

It is instructive to examine how the four calculi distinguish between ignorance (e.g., "I don't know my chances of winning the lottery"), and uncertainty (e.g., "I have a .001% chance of winning the lottery").

In the Bayesian calculus it is difficult to represent ignorance since each entity must be assigned a probability. One potential method that has been suggested is to represent ignorance by assigning equal prior probabilities to all events, but this method is fraught with difficulties [Shaf76, pg. 23]. Another possible method is to use whatever knowledge is available to make a "reasonable" assignment. For example, suppose that  $d_1$  has prior probability .6,  $d_2$  has prior probability .4,  $P(d_1 | s)$  is unknown, and  $P(d_2 | s)$  is unknown. Then it might seem reasonable to assign  $P(d_1 | s) = .6$  and  $P(d_2 | s) = .4$ , but there is no way of distinguishing this arbitrary assignment from the case in which the conditional probabilities actually are .6 and .4.

In the Dempster-Shafer calculus, ignorance is represented by assigning belief to large subsets (i.e, given more knowledge, the belief would be assigned to a smaller subset, or even a singleton). In addition, the magnitude of the belief interval (plausibility-Bel) also reflects the degree of ignorance.

In fuzzy set theory, the magnitude of the possibility-necessity interval indicates the degree of ignorance.

In MYCIN/EMYCIN, a CF equal to zero can be used to represent ignorance, but it is not possible to distinguish this from the case in which the confirming evidence is balanced by the disconfirming evidence.

#### **AI.7. Connections Between the Four Calculi**

There are some connections between the four calculi that are worth discussing:

##### *The Bayesian and Dempster-Shafer calculi:*

Both the Bayesian and Dempster-Shafer calculi assume that diseases are mutually exclusive and exhaustive. Both calculi can be extended to handle concurrent diseases; however, this leads to a significant increase in complexity. In the Dempster-Shafer theory, evidence is assumed to be independent; this assumption is almost always made in Bayesian systems, although it is not a fundamental aspect of the Bayesian calculus.

The Bayesian calculus is based on three axioms:

- (1)  $P(\text{impossible-}t) = 0;$

(2)  $P(\text{sure-}t) = 1$ ;

(3) If  $t_1$  and  $t_2$  are mutually exclusive, then

$$P(t_1 \cup t_2) = P(t_1) + P(t_2).$$

The Dempster-Shafer theory retains axioms (1) and (2) (substituting *Bel* for *P*), but not axiom (3) (actually, it has a third, but quite different, axiom [Shaf76, pg. 5]). This leads to several consequences. First of all, in the Bayesian calculus if  $P(A) = x$ , then  $P(A^c) = 1 - x$ ; the analogous identity does not hold in the Dempster-Shafer theory.

More importantly, in the Dempster-Shafer theory it is possible to assign belief to a subset without assigning belief to any of the individual elements of that subset; this is not possible in the Bayesian calculus.

**Example:**

In the Dempster-Shafer calculus it is possible to specify that:  $\text{Bel}(\{V_1, V_2\}) = x, x > 0$ ;  $\text{Bel}(\{V_1\}) = 0$ ;  $\text{Bel}(\{V_2\}) = 0$ . In the Bayesian calculus,  $P(\{V_1, V_2\})$  can only be non-zero if  $P(\{V_1, V_2\}) = P(\{V_1\}) + P(\{V_2\}) > 0$ ; i.e.,  $P(\{V_1\}) > 0$  and/or  $P(\{V_2\}) > 0$ .

As might be expected, when belief is assigned only to singletons, the Dempster-Shafer calculus reduces to the Bayesian calculus with independence. See Hummel and Landy [Humm88] and Kyburg [Kybu87] for additional discussion of the theoretical relationships between the Dempster-Shafer and Bayesian calculi.

*The Bayesian and MYCIN calculi:*

Adams [Adam76] has shown that given independent evidence, and small prior probabilities, the Bayesian and MYCIN calculi produce similar (but not identical) results. It is not true, however, (as is sometimes stated) that the MYCIN calculus is simply "Bayes' theorem with the independence assumption."

*The Dempster-Shafer and MYCIN calculi:*

There are also connections between the Dempster-Shafer theory and the MYCIN calculus: for example, under certain restrictions, Dempster's rule reduces to one of the MYCIN combination formulas [Gord84], [Shaf76, pp. 75-77]. See Groszoff [Gros86a] for further discussion.

*The fuzzy set and MYCIN/EMYCIN calculi:*

There are some interesting connections between fuzzy set theory and MYCIN/EMYCIN. For example, both use max and min to combine information. In addition, MB and MD in MYCIN are analogous, respectively, to necessity and complement of possibility from fuzzy set theory [Prad83]. Scheffe [Sche80] explores the relationship between fuzzy set theory and MYCIN/EMYCIN in a more formal manner.

*Additional Connections:*

There are several papers in the literature which discuss the intuitive, philosophical, and theoretical connections between the four calculi. For example, Groszoff [Gros86b] discusses some of the relationships of the Bayesian, Dempster-Shafer and Mycin calculi. Prade [Prad85] presents a uniform framework for the Bayesian, Dempster-Shafer and fuzzy set cal-

culi, while Thompson [Thom85] discusses the structural similarities of these calculi. See also [Prad83] and [Sche80].

### **AI.8. Empirical Studies**

There are some empirical studies in the literature that report on the success of a given calculus in a specific program; e.g., [DeDo74], [DeDo72], [Duda76], [Horr72] for the Bayesian approach, [Fies82] for fuzzy set theory, and [Buch84] for the MYCIN calculus. There are few reports in the literature of failures (perhaps, not surprisingly).

Several empirical studies have claimed to show that the exact choice of numbers is not crucial. For example, [Buch84, pp. 218-219] reports that the MYCIN calculus is relatively insensitive to small changes in the CF's. (In addition, [Szol78] reports that small changes in INTERNIST's numbers, make little difference to the final outcome; Rogers et al [Roge79] reports that using *fewer* factors in a statistical diagnostic system leads to better performance.)

Tong and Shapiro [Tong85] have compared the performance of different calculi on the same problem. They report, in regards to an information retrieval system, that the choice of an uncertainty calculus becomes less crucial as the complexity of the queries increase.

### **AI.9. A Simple Example/Future Research**

In [Zade84] Zadeh presents the following example of an application

of the Dempster-Shafer theory:

**Example:**

Suppose that piece of evidence,  $X$ , induces  $\text{bpa}(d_1) = .99$ ,  $\text{bpa}(d_2) = .01$ ,  $\text{bpa}(d_3) = 0$ , while the piece of evidence,  $Y$ , induces  $\text{bpa}(d_1) = 0$ ,  $\text{bpa}(d_2) = .01$ ,  $\text{bpa}(d_3) = .99$ . Then, by Dempster's rule,  $\text{bpa}(d_1) = 0$ ,  $\text{bpa}(d_2) = 1$ ,  $\text{bpa}(d_3) = 0$ .

Note that the disbelief in  $d_2$  induced by  $X$  is quite high ( $\text{disbelief}(\{d_2\}) = \text{Bel}(\{d_2\}^c) = \text{Bel}(\{d_1, d_3\}) = .99$ ). Similarly, the disbelief in  $d_2$  induced by  $Y$  is quite high (.99). Yet, according to the combination rule, the belief in  $d_2$  equals 1. Zadeh suggests that this is counter-intuitive since extremely high disbeliefs are combined to yield total belief. An analogy may help to clarify Zadeh's argument.<sup>11</sup>

**Analogy 1:**

Suppose that your doctor thinks that you have either lymphoma, the plague, or leukemia. Not being sure of the diagnosis, he or she refers you to two specialists. Specialist  $X$  tells you that you probably have lymphoma, probably don't have the plague, and definitely don't have leukemia. Specialist  $Y$  tells you that you probably have leukemia, probably don't have the plague, and definitely don't have lymphoma. Most people would *not* conclude that they definitely have the plague; therefore, Dempster's rule does not seem to work in this interpretation.

There is another interpretation, however: Note that the disbelief in  $d_1$  induced by  $Y$  equals 1 (disbelief( $\{d_1\}$ ) = Bel( $\{d_1\}^c$ ) = Bel( $\{d_2, d_3\}$ ) = 1). Similarly, the disbelief in  $d_3$  induced by  $X$  equals 1. If one accepts that "When you have eliminated the impossible, whatever remains, however improbable, must be the truth," then the belief in  $d_2$  should equal 1. An analogy is also useful here:

**Analogy 2:**

Suppose that someone was murdered in Chicago. There are three suspects: Tom, Dick and Harry. Detective  $X$  states that Tom probably was in Chicago, Dick probably wasn't in Chicago, and Harry definitely was not in Chicago. Detective  $Y$  states that Harry probably was in Chicago, Dick probably wasn't in Chicago, and Tom definitely was not in Chicago. Most people would conclude that Dick is the murderer. In this interpretation Dempster's rule appears to work.

In the Bayesian calculus (replacing  $bpa$  with  $P$ , etc.) the probabilities  $P(d_1 | X \& Y)$ ,  $P(d_2 | X \& Y)$ ,  $P(d_3 | X \& Y)$  cannot be calculated, but rather are prior probabilities. Given independence assumptions, however, the probabilities can be computed and the resulting numbers are identical to those computed by Dempster's rule.

In fuzzy set theory (replacing  $bpa$  with  $\mu$ , etc.) there are many ways to combine pieces of evidence. One possible result would be  $\mu(d_1) = 0$ ,  $\mu(d_2) = .01$ , and  $\mu(d_3) = 0$  (computed by using the min operator). A different choice of operators would give a different result.

Recall that in the MYCIN calculus only one of MB and MD can be non-zero for a given piece of evidence. In the example, the belief in  $d_1$  induced by  $X$  equals .99, and the disbelief in  $d_1$  induced by  $X$  equals .01. Since both of these are non-zero it is not clear what result the MYCIN calculus would compute.

It is apparent from the example and the two analogies that people combine evidence in different ways depending on the particular situation (e.g., "doctor's opinions" vs. "detective's alibis"). The Bayesian, Dempster-Shafer and MYCIN/EMYCIN calculi do not reflect this since they always combine evidence in the same way. Fuzzy set theory has a bit more flexibility in this regard, but, as noted before, it is often not clear which operators or methods are appropriate in a given situation.

A fundamental characteristic of the example above is that the two pieces of evidence reflect a wide disparity (if not contradiction) in opinion. The four calculi do not capture this fact—they compute aggregate numbers, but keep no record of the divergence in opinion. From a statistical perspective, it is as if the calculi compute the mean, but not the standard deviation. There does seem to be an advantage (such as in the example above) to keeping track of higher order statistics. The only work in this direction that we are aware of is by Hummel and Landy [Humm88]; there would seem to be much grounds for future research in this direction

#### AI.10. Summary/Conclusion

There has been an unfortunate tendency to blindly choose an uncer-

tainty calculus for incorporation into a knowledge-based system. In particular, some of the commercial expert systems shells offer only one calculus (and those with more flexibility give little guidance as to which calculus to choose). We offer the following suggestions:

The Bayesian calculus is well suited for applications where probabilities are known (or can be acquired with a reasonable effort). The calculus is especially attractive because of its strong theoretical foundation (note, however, that the calculus is generally unusable in its purest form due to exponential information complexity).

The Dempster-Shafer calculus is a good choice for applications where uncertainty is best thought of as being distributed in sets rather than just single items. Depending upon the particular domain, it can also have acceptable information and time complexity.

The fuzzy set calculus is well suited for applications where the evidence is itself fuzzy in nature. Fuzzy set techniques also have the advantages of great flexibility and low information and time complexities.

The main appeal of MYCIN/EMYCIN is its low information and time complexity, and its current wide availability.

Each of the four calculi has a different perspective on uncertainty, and each manipulates uncertain information in a different way. Despite what some authors have claimed [Chee85] there does not seem to be one calculus that is "the best" for all situations. Each of the calculi has its strongpoints; the main disadvantage that we see in all of the calculi is that they compute aggregate numbers, but keep no record of divergence in opinions.

## Footnotes

- 1 Non-numerical methods for handling uncertainty have also been proposed, e.g., [Coh83]; these methods are not discussed here.
- 2 We assume, unrealistically, but without loss of generality, that the physician examines the patient for all possible signs and symptoms. It should be noted that, in certain cases, the *order* in which information is gathered is crucial since the administration of one diagnostic test may affect the outcome of subsequent tests. For example, an upper GI series can interfere with the interpretation of a subsequent intravenous pyelogram [Szol78].
- 3 Many authors use the term *Bayesian* to refer to "subjective" (as opposed to "objective") probabilities; we use the term to refer to computational methods based on Bayes' rule.
- 4 Of course, significant computational time may be needed in the initial analysis to gather the prior probabilities. In addition, significant time may be required if  $m$  is large.
- 5 Expositions of the Dempster-Shafer theory generally denote the set of diseases by  $\Theta$ ; we use that notation ( $\Theta = D$ ). In this section we closely follow Gordon and Shortliffe [Gord84].
- 6 A triangular norm is a function from  $[0,1] \times [0,1]$  to  $[0,1]$  which is monotonic, commutative, and associative, and which has boundary conditions that satisfy the truth table of logical AND.
- 7 A triangular conorm is a function from  $[0,1] \times [0,1]$  to  $[0,1]$  which is monotonic, commutative, and associative, and which has boundary conditions that satisfy the truth of logical OR.
- 8 Recall that MYCIN is a program for the diagnosis and therapy of bacteremia, that EMYCIN is an expert systems shell based on the control structure of MYCIN, and that MYCIN was later re-implemented in EMYCIN.
- 9 In this section we follow Shortliffe and Buchanan's original notation [Shor75Amod]. In the MYCIN and EMYCIN calculi there is a duality between evidence and hypotheses; e.g., the hypothesis confirmed to some degree by one piece of evidence may itself be a piece of evidence for another hypothesis. Moreover,  $MB[h,s]$  may mean the degree of confirmation for sign  $h$ , based on observation  $s$  (e.g.,  $h = fever$ ,  $s = temperature$ ), or it may mean the degree of confirmation of disease  $h$  based on sign  $s$  (e.g.,  $h = infection$ ,  $s = fever$ ). The appropriate meaning is usually made clear by the context; in any event, this is of little consequence for the following discussion.
- 10 A more explicit influence on the construction of the combining functions was that they should satisfy certain mathematical criteria; e.g., commutativity.
- 11 We would like to thank Prof Ernest Davis for suggesting the two following analogies.

**APPENDIX II****LEXICAL IMPRECISION****Note**

A version of this paper was presented at the 1986 Conference on Uncertainty in Artificial Intelligence (August 1986). A slightly edited version of this paper has been published in [Henkind SJ. Imprecise Meanings as cause of uncertainty in medical knowledge-based systems. In: Lemmer J, Kanal LN eds. Uncertainty in artificial intelligence, vol II. New York: North-Holland, 1988:35-41]. Reprinted by permission of the publisher.

**AII.1. Introduction: Lexical Imprecision**

Specialized fields of knowledge can be viewed as having their own languages. These languages, which are subsets of natural language, are known as *sublanguages* [Kitt82]. Typically, sublanguages have a specialized vocabulary: for example, the vocabulary of medicine as found in a medical dictionary, or the vocabulary of law as found in a legal dictionary.

Since sublanguages are generally used in complex situations where there are difficult problems to be solved, and critical decisions to be made, it is desirable that the underlying vocabularies be well defined. In particular, the words need to have precise meanings. Clearly, if a given word does not

have a precise meaning, then a sentence containing that word may be imprecise as well. If the sentence is intended to convey information relevant to the solution of a problem, then the imprecision can lead to an incorrect solution.

Unfortunately, the vocabularies underlying most sublanguages are much less well defined than is commonly believed. For example, a word may be used to denote many similar, but not identical phenomena: In a recent review of the medical literature [Henk87] we found more than a dozen definitions for the important clinical phenomenon *pulsus paradoxus*. All the definitions were intended to define the same phenomenon. Yet they differed enough that two physicians, given the same patient, could reach opposite conclusions as to the presence or absence of a paradoxical pulse. We will say that words which have several similar, but not identical, meanings are *lexically imprecise*.

The lexical imprecision to be found in medical terminology is not limited to certain exceptional words, but is, in fact, extremely common. For example, given that a patient has produced only 300 ml of urine during the past day, nearly every physician would state that the patient has oliguria (low urine output). But, if the physicians were also told that the patient weighs 80 pounds, has received no fluids in the past 24 hours, and has been exercising heavily, then some would say that the patient is not oliguric, while others would continue to state that the patient has oliguria, albeit "to be expected under the circumstances." *Oliguria* is a term that is well-known to all physicians, but it is apparent that the use of the word is highly dependent upon who is using the word, under what circumstances, and so forth. A similar analysis can be performed for many other clinical entities, e.g., hyperten-

sion, etc.

It should be noted that we are not concerned here with words that are used to denote inherently imprecise entities. For example, the word *lethargic* must be somewhat imprecise because lethargy is very subjective in nature. However, we are extremely concerned with words that are used to denote ostensibly objective entities. The word *oliguria* can cause problems because although it is not well defined, nearly all physicians believe that it is. The word *lethargy*, on the other hand, is much less likely to cause difficulty because physicians recognize that it is subjective.

Lexical imprecision is more than just a theoretical curiosity: it has, in fact, caused widespread difficulties in medicine. For example, "There are varying schools of thought among specialists in interpreting petit mal seizures. Some assign the designation *petit mal* to 3 percent of all forms of epilepsy; others classify 80 percent of seizures under this rubric .... In three recent papers, the results reported by the authors on a newly introduced anti-epileptic drug for the treatment of petit mal were respectively that it was highly effective, moderately effective, and ineffective. How much of this discrepancy is to be attributed to the drug or conditions of the trials, and how much to the different conditions regarded by the experiments as being petit mal?" [Kenn73].

Lexical imprecision has also caused difficulties in other areas. For example, imprecision in word meanings is known to decrease the effectiveness of computerized information retrieval systems.

## **AII.2. Effects on Knowledge-Based Systems**

One of the fundamental choices to be made in the construction of a knowledge-based system is the selection of an appropriate knowledge representation. Although rule-based representations are currently the most common choice, there are many other options, e.g., frames, semantic nets, scripts, etc.

At some point, however, all knowledge representations require a choice of semantic primitives. These primitives are the fundamental objects which a system will manipulate. In a knowledge-based system designed to perform medical diagnosis, for example, the semantic primitives would be various signs, symptoms, lab values, and diagnoses. Note that the semantic primitives are generally either lexical items (words), or numbers.

Consider a system which, when given some collection of signs and symptoms, deduces the patient's disease state. Call the set of all signs and symptoms,  $S$ , the set of all subsets of  $S$  (power set of  $S$ ),  $P(S)$ , and the set of all diseases,  $D$ . Then diagnosis can be viewed as a map from  $P(S)$  to  $D$ ; i.e.,  $P(S) \rightarrow D$ . For the sake of argument, it is assumed here that each patient has one and only one disease. Furthermore, it is assumed that the diagnostic map is provided by a single domain expert.

Suppose, now, that a user makes a certain set of observations  $O$  about a patient. In order to use the diagnostic system, the user will need to express these observations in terms of the semantic primitives of the system. In other words,  $O$  must be mapped by the user into  $P(S)$ . Therefore, the computer-assisted diagnostic process is actually  $O \rightarrow P(S) \rightarrow D$ .

Unfortunately, due to lexical imprecision, two individuals may map the same observations differently. This can, in turn, lead to a different set of deductions by the system. For example, suppose that an expert creates the following system:

```

IF pulsus-paradoxus
THEN tamponade-likely
ELSE tamponade-not-likely.

```

Also, suppose that user U1 defines pulsus paradoxus as "an inspiratory decline in systolic arterial pressure of 10 mm Hg or more," and user U2 defines it as "a decline of 13 or more." If a patient has an inspiratory decline of 12, then U1 would map his observation into the semantic primitive pulsus-paradoxus, but U2 would not. Therefore, the system would provide U1 with the conclusion that tamponade is likely, but would provide U2 with the opposite conclusion.

Certainly, both conclusions cannot be correct. The difference in conclusions is due to the fact that the map  $O \rightarrow P(S)$  is not uniquely specified. In particular, this map varies from individual to individual, depending on the definitions that each person happens to use. It is worth noting that the developers of the INTERNIST system [Popl86] have, in fact, documented difficulties with the interface to their system due to variability in word meanings.

If a system is constructed by two or more domain experts, then lexical imprecision can lead to less than optimal performance of the diagnostic map  $P(S) \rightarrow D$ . Suppose, for example, that expert E1 provides the rule "IF A

THEN B," and expert E2 provides the rule "IF B THEN C." Given A, the system will deduce C. But this may be incorrect if E1 and E2 have different definitions for the semantic primitive B. Of course, if a given expert is able to provide correct solutions to problems, then his knowledge must be, in some sense, internally consistent. The problem here, is that the components of his knowledge may be inconsistent with the components of another expert.

### **AII.3. Lexical Imprecision Is Not Lexical Ambiguity**

Many words have several distinct meanings. For example, the word *beat* can be used as a verb to denote the act of physically abusing someone, as a verb to denote the act of sailing a boat close to the wind, as a noun to denote a policeman's patrol area, and so forth. This phenomenon of multiple distinct meanings is referred to as *lexical ambiguity*. Note that lexical ambiguity is not the same thing as lexical imprecision. A useful analogy is the following: If you open a dictionary and choose a word, it will have several distinct definitions. This is lexical ambiguity. If you open two dictionaries and choose the same word you will find sets of very similar, but not identical definitions. This is lexical imprecision.

Lexical ambiguity can lead to difficulties in the processing of natural language by computer. From a syntactic standpoint, the possibility that a given word may have multiple meanings makes it necessary to select the correct part-of-speech for the word. For example, in parsing the sentence "He is on his beat," it must be determined that *beat* is being used as a noun, and not as a verb. Since the meanings of a lexically imprecise word are all

the same part-of-speech, lexical imprecision does not lead to problems in parsing.

Lexical ambiguity can also lead to difficulties in semantic analysis. Semantic analysis requires that the correct word-sense of an item be selected. For example, in analyzing the sentence fragment "Beat until you see the buoy," it must be determined that *beat* refers to the action of sailing close to the wind, and not to the infliction of physical violence. The dictionary of a typical natural language processing system may contain multiple definitions for a given word, but these definitions are invariably distinct. Hence, lexical imprecision has not posed many problems for semantic analysis.

It is at the level of pragmatics that lexical imprecision will cause the most difficulty for systems that use natural language. While subtle distinctions in meaning are of little consequence in the syntactic and semantic decomposition of sentences, these same distinctions can have profound consequences on a system that attempts to use those sentences.

There is much discussion in the literature of techniques for handling lexical ambiguity. Birnbaum [Birn85] provides a detailed review and analysis. On the other hand, there seems to be little, if any, discussion of methods for coping with lexical imprecision. The following section describes some possible techniques.

#### **AII.4. Coping With Lexical Imprecision**

The ideal solution to the problem of lexical imprecision would be to eliminate it entirely. This would require that a precise set of definitions be

established in each specialized field of knowledge—presumably by a committee of experts. Furthermore, every individual in that field would need to agree upon and use those definitions. There has been, in fact, a great deal of effort in this direction, but with only mixed results. A notable success is the science of chemistry: once a molecule's structure has been determined, that substance has a name assigned to it that conveys the same meaning to every chemist. In the field of medicine, such efforts have been much less successful. For example, the American College of Cardiology, and the American Heart Association have published what they feel to be a standard definition for *pulsus paradoxus* [Amer71], yet this definition was not to be found in any of the more than sixty papers in the literature that we surveyed [Henk87].

Although efforts to standardize terminology are extremely important, it will never be possible to completely eliminate lexical imprecision. As a field of knowledge expands, new discoveries are made, new measurement techniques devised, and entities are viewed at finer levels of granularity. This in turn renders some previously precise words less precise. For example, to characterize a patient as being "hypertensive" was actually quite precise a hundred years ago; but in modern medicine, such a characterization is far from adequate. Lexical imprecision is, and will remain, a ubiquitous phenomenon, and high performance knowledge-based systems will need to handle it in a reasonable fashion.

As was discussed previously, lexical imprecision can lead to difficulties both at the interface to a knowledge-based system, i.e.,  $O \rightarrow P(S)$ , and within the system, i.e.,  $P(S) \rightarrow D$ . There are several ways to lessen the impact at the interface level. One way is to eliminate that level as much as possible. In

particular, the more observations that a system can make directly (rather than through a human intermediary), the less the damage that will be done due to individual differences in the mapping  $O \rightarrow P(S)$ . In some situations a system should be able to gather most of its input directly: for instance, monitoring a chemical plant. In many other situations, however, a human intermediary is essential. For example, computer vision and robotics notwithstanding, no machine is yet capable of performing a complete physical examination of a patient.

Problems at the interface level can also be decreased by insisting on quantification. For example, the question "Does the patient have oliguria?" is highly susceptible to individual differences in the mapping  $O \rightarrow P(S)$ , but a request for the volume of urine output is much less problematic [Henk86]. Unfortunately, not all phenomena can be quantified—e.g., petit mal. Even for those phenomena which can be quantified, there are still potential difficulties to be aware of. For example, the normal ranges of various biomedical tests, e.g., enzyme assays, are not standardized, but actually vary from laboratory to laboratory.

Another consideration is that it must be certain that all observers are in fact measuring the same phenomenon (e.g., is pulsus paradoxus the inspiratory decrease in pulse pressure, or the inspiratory decrease in systolic pressure?). This could be encouraged as follows: when the system requests the value of an entity, it also provides a definition. For example, "Please input the measured value of pulsus paradoxus (inspiratory decline in systolic arterial pressure)." Of course, these definitions would need to be built into the system.

The incorporation of definitions into a system is also a potential solution to problems in the construction of the diagnostic map  $P(S) \rightarrow D$ . In particular, each expert could be encouraged to record his definitions for the semantic primitives with which he is working. Thus, if expert E1 produces the rule "IF A THEN B," he would be expected to provide definitions for the primitives A and B. Before expert E2 entered the rule "IF B THEN C," he would be expected to check the definition of B, and so forth.

Unfortunately, it may be unreasonable to expect that definitions be provided for all the primitives of a large system. The demands in terms of increased development time and overhead could be enormous; nor is it clear how such definitions could be incorporated into existing systems. Furthermore, since the definitions of primitives are themselves composed from other primitives, it is impossible to enforce the complete consistency of the definitions within a system. It should be noted that there are currently a few knowledge-based systems which have (limited) facilities for the incorporation of definitions; e.g., MYCIN [Buch84 pg. 210] and AI/RHEUM [Lind86].

Fuzzy set theory [Zade65] has some applicability as a tool to handle lexical imprecision. Many words are lexically imprecise because they are based upon cutoffs. For example, the lexical imprecision of *pulsus paradoxus* stems largely from the fact that different experts use different cutoffs, e.g., 10, 13, etc. These cutoffs in turn lead to discontinuous behavior, e.g., *pulsus paradoxus* is considered to be present given a cutoff of 10, but is considered to be absent given a cutoff of 13. By modeling *pulsus paradoxus* as a fuzzy concept, the damaging effects of lexical imprecision could be greatly reduced because the presence or absence of *pulsus* would no longer be treated in a

discontinuous fashion.

Note, however, that fuzzy techniques are not a solution for all cases of lexical imprecision. For example, fuzzy techniques could not reconcile the imprecision due to one observer measuring the decrease in pulse pressure, and another another observer measuring the decrease in systolic pressure.

### **AII.5. Summary**

Uncertainty is a major source of difficulty in the construction and use of knowledge-based systems. One type of uncertainty arises from the strength of the implication operator in inferences such as "IF A THEN B"; e.g., if A then there is a 40% chance of B. Another type of uncertainty arises from the weight of evidence: e.g., there is a 70% chance that the patient has A. In this paper, we have discussed another type of uncertainty—that which is due to imprecision in the underlying primitives: e.g., two experts have a different conception of A. Lexical imprecision has been described by other researchers [Buch84 pg. 210, Boni85] but it has received much less attention than other types of uncertainty.

Lexical imprecision can degrade the performance of knowledge-based systems. Effects can surface at both the user interface and inferencing levels.

Among the techniques to handle lexical imprecision are the direct acquisition of input data, quantification, the inclusion of definitions, and fuzzy set methods. Currently we are examining ways of incorporating these techniques into medical knowledge-based systems.

## **AII.6. Postscript**

The techniques described in this paper have been incorporated into the alarm system. In particular, the system directly acquires data, quantifies phenomena, has a limited facility for definitions (in the questioning phase), and uses fuzzy set techniques.

The system has not been extensively tested yet, hence, it is not possible to determine whether the system is susceptible to problems due to lexical imprecision, and if so, whether the above techniques are a viable solution. (Note that since the system was constructed by only one "expert"—me—it seems unlikely that lexical imprecision will cause problems at the inferencing level.)

**APPENDIX III****REVIEW OF PULSUS PARADOXUS****Note**

This appendix has been published in slightly edited form in [Henkind SJ, Benis AM, Teichholz LE. The paradox of pulsus paradoxus. *Am Heart J* 1987;114:198-203]. Reprinted by permission of the publisher.

**AIII.1. Introduction**

An observation of pulsus paradoxus may have occurred as early as 1669, when Lower [Lowe1669] noted with regard to a patient with constrictive pericarditis "that intermission of the pulse followed every successive inspiration with such constancy." In 1698 Floyer [Floy1717] reported an inspiratory disappearance of the arterial pulse in patients having attacks of bronchial asthma. In 1850 Williams [Will1850] described the phenomenon, and it was also detected by German physicians in 1854 in a patient with purulent pericarditis [Vier1855]. In 1873, Kussmaul [Kuss1873] first used the term *pulsus paradoxus* to describe a patient in whom the arterial pulse seemed to vanish during inspiration. Since then, the phenomenon of a marked decrease or complete obliteration of the arterial pulse upon inspiration, has generally been labeled as pulsus paradoxus.

### **AIII.2. Occurrence of Pulsus Paradoxus**

Pulsus paradoxus (abbreviated here as *pulsus*) is not a highly specific phenomenon. It can be found relatively frequently in various circulatory abnormalities. Cohn [Cohn67] reported its presence in 30 out of 61 cases of hypovolemic shock. Sapira [Sapi83] reports it to be present in 38% of patients with constrictive pericarditis, but Spodick [Spod84Puls] states that "Pulsus paradoxus occurs with constrictive pericarditis if it is effusive-constrictive pericarditis. In pure constriction (i.e., with no accompanying pericardial fluid under pressure) pulsus paradoxus is absent, except for rare cases." It has also been observed in right ventricular infarction [Lore79], severe congestive heart failure [Cons76], cardiomegaly [Jaco78], and in patent ductus arteriosus, myocarditis, tumors of the mediastinum, paramediastinal effusion, endocardial fibrosis, fibroelastosis, myocardial amyloidosis, scleroderma, mitral stenosis, tricuspid stenosis and other cardiac pathologies [Sapi83].

Pulsus is seen in various respiratory conditions [Vais74]. Knowles [Know73] reported it in 12 out of 15 cases of exacerbated asthma, and Cohen [Cohe73] in 2 out of 9 cases of pulmonary embolism. It is also found in emphysema [Sett80] and other respiratory diseases. In addition, pulsus has been reported in extreme obesity and ascites [Lang67], pregnancy [Brau84], and high speed centrifugal rotation [Ursc67].

Pulsus paradoxus is most commonly associated with cardiac tamponade (although recent studies seem to indicate that pulsus paradoxus may not occur until relatively late in the course of this disease [Frie80, Klop85]). Reported frequencies of occurrence in tamponade vary: one source states

that pulsus is "almost invariably present, unless hypotension is profound" [Shab70]. Another source found significant pulsus in 10 out of 14 cases of tamponade (none of which were profoundly hypotensive) [Redd78]. In a study of 40 dogs induced into tamponade, pulsus was present in 70% of the animals [Gunt67].

Despite the diversity of diseases in which pulsus can be seen and the differing reported frequencies, it is nevertheless valuable clinically. The presence of a marked pulsus paradoxus strongly suggests the presence of underlying pathology. Under certain circumstances, the phenomenon is nearly always found [Hill82]: e.g., a study found pulsus to be present in 100% of patients with post cardiac-surgical tamponade [Nels69].

### **AIII.3. Conditions That May Prevent Or Mimic Pulsus Paradoxus**

There exist conditions which may prevent the occurrence of pulsus, under circumstances in which it is usually expected (e.g., tamponade). Among the conditions that have been reported are severe calcific aortic stenosis [Lang67], aortic regurgitation [Gube]81, uremia with left ventricular dysfunction [Gube81], "underlying heart disease" causing "a marked elevation of left ventricular pressure so that the two ventricles are unequally compressed" [Brau84], severe musculoskeletal abnormalities of the chest [Wine79], and atrial septal defect [Wine79]. In addition, a reversed pulsus paradoxus (to be described later) can mask underlying pulsus paradoxus. Reversed pulsus paradoxus may be found in idiopathic hypertrophic subaortic stenosis, isorhythmic ventricular rhythms, and in patients on intermittent

positive pressure ventilation [Mass73].

The importance of being aware of the above conditions is illustrated by the case of a woman with tamponade and atrial septal defect: "One patient ... died. This patient had progressive hypotension over several days; the absence of a paradoxical pulse was considered evidence that tamponade was not the cause of her hypotension" [Wine79].

In addition, it should be verified that the phenomenon being observed is really pulsus paradoxus. Salel [Sale73] reports on a case where a fortuitous synchronization of respiration with atrio-ventricular dissociation resembled pulsus. Swinton [Swin69] notes that various thoracic outlet syndromes can mimic pulsus paradoxus; e.g., a "pulsus paradoxus" in the left radial artery only, due to a fibrous anomaly of the anterior scalenius muscle.

#### **AIII.4. Effects of Respiration on the Arterial Pressure**

In a healthy individual one expects to see a drop in systolic pressure during normal inspiration. In one study, a drop of up to 5 mm Hg was detected in 20 normal volunteers [Know 73]. A cardiology text states that a drop of 2 to 6 mm Hg is normal [Cons76]. Other authors [Shab70] cite slightly different results: "It is difficult to assign a precise boundary between normal and abnormal; however 8 to 10 mm Hg is generally taken as the upper limit of normal."

It is interesting to note, that since pulsus paradoxus can be thought of as an accentuation of a normal physiologic phenomenon, many consider the word *paradoxus* to be a poor choice. Wagner [Wagn73] states that "The great

clinical observer Kussmaul gave two reasons to justify his selection of this term: 'I propose to call this pulse paradoxical first because of the obvious discrepancy in (regular) precordial activity and in (irregular) peripheral arterial pulsation, and second because of the additional paradox that the peripheral pulse in spite of the apparent irregularity decreases and disappears with repetitive regularity.'

Abnormal ventilation in the normal individual can greatly alter the peripheral arterial pressure. At least as early as 1899 it was noted that a normal individual inspiring through external resistance will have a pronounced weakening of the pulse [Sapi 83]. Other observers have noted that deep breathing will accentuate the normal drop of arterial pressure in inspiration: a drop of up to 15 mm Hg in systolic pressure is typical for very deep inspiration [Cons 76]. Another investigator claims, however, that high lung volume and external resistance together will lead to a marked accentuation of the normal drop on inspiration, but that neither factor alone is sufficient to cause such a change [Rebu73]. An investigation of five normal subjects induced into having pulsus by hyperinflation and breathing through external resistance revealed that "Changes in frequency of respiration and the duration of inspiration, singly or in combination, may influence the magnitude of pulsus paradoxus" [Mart 81]. It is also known that an individual on intermittent positive pressure respiration will actually have an increase in inspiratory systolic pressure. This phenomenon has been called *reversed pulsus paradoxus* [Mass73].

Changes in ventilation affect the character of pulsus paradoxus in various disease states: "In patients with bronchial asthma, a different pattern of breathing may alter the degree of paradox at a given degree of obstruction of

outflow" [Mart81]. Along similar lines, another author notes about patients with effusion, that "pulsus paradoxus is inconstant with tranquil breathing but is regularly induced by deep inspiration" [Lang67].

### **AIII.5. Mechanism of Pulsus Paradoxus**

In order to understand the mechanism of pulsus paradoxus, it is first desirable to understand the basis of the physiologic drop in systolic pressure that occurs in normal individuals during inspiration. The mechanism of this phenomenon has been controversial: "For more than 100 years, there have been conflicting studies analyzing the cardiovascular effects of a single spontaneous respiratory cycle" [Robo83].

One of the most common explanations for this phenomenon is that inspiration leads to a pooling of blood in the pulmonic vasculature, and an increase in venous return to the right heart. The former effect predominates, and so the net effect is a decrease in left ventricular output and hence pressure [Swar81]. In addition, it is postulated that the decreased intrathoracic pressure is transmitted to the aorta, thus also lowering systolic pressure.

Another common explanation is that inspiration leads to a reduced intrathoracic pressure, thereby augmenting right ventricular filling. This in turn leads to compression of the left ventricle, hence decreased left-sided filling, and therefore decreased output and pressure [Brau84].

The mechanism of pulsus paradoxus has been extensively investigated for the condition of cardiac tamponade. The earliest proposed mechanism for pulsus in tamponade is probably that of Lower(1669) [Lowe1669] who postu-

lated that it is due to inspiratory traction of the diaphragm on the pericardium. In 1961 Dock [Dock61] performed experiments on cadavers that lend support to this theory, but more recent investigators have discounted the theory [Shab70].

In 1924, Katz and Gauchat [Katz24] proposed that pulsus paradoxus in cardiac tamponade is due to decreased filling of both ventricles during inspiration. A 1960 study presents evidence tending to confirm this mechanism in humans [Shar60], but more recent investigations have demonstrated an inspiratory augmentation of right ventricular filling [Gunt67].

An explanation that is in favor now is that "In cardiac tamponade, inspiration probably augments systemic venous return to a greater percentage than normal and thus produces a greater than normal respiratory variation in the systemic flow. The increased inspiratory augmentation of right ventricular output is manifest in the systemic circulation 2 or 3 heartbeats later as an expiratory augmentation, in systemic flow and pressure" [Shab70].

Another current explanation is that the paradoxical pulse in tamponade is due essentially to an accentuation of normal physiologic mechanisms: reduced intrathoracic pressure leads to increased right ventricular filling; this in turn impedes left ventricular filling and hence reduces flow and pressure [Brau84, McGr79]. A recent hemodynamic study provided evidence consistent with this theory but also concludes that "other mechanisms predominate in severe tamponade" [Fuen85].

Other potential mechanisms for pulsus paradoxus in tamponade can be found in [Goli63, Spod64]. Mechanisms for pulsus paradoxus have also been proposed in the context of hypovolemic shock [Cohn67], obesity [Cons76], pregnancy [Cons76], ventilatory obstruction [Cohn67], and asthma [Lang67,

Rebu73]. The mechanism for reversed pulsus paradoxus has been examined as well [Robo83].

There are many conflicting theories regarding the cause of pulsus paradoxus. In fact, some of the older mechanisms that we have presented, e.g., [Dock61, Goli63, Katz24, Lowe69, Shar60, Spod64] have been rather convincingly ruled out by more recent evidence. Nevertheless, there still exists a lack of consensus as to which mechanisms are correct. Shabetai [Shab70] has suggested that "Failure to discriminate between tamponade and constriction, and between severe and mild compression accounts in large measure for the contradictory findings on the mechanisms of hemodynamic alterations...." We also speculate that, perhaps, failure to control adequately for ventilatory effects may have influenced some of the reported results. Undoubtedly, there are multiple mechanisms that cause pulsus paradoxus [Katz24, Shab81]. It is likely that various combinations of these mechanisms are responsible for the phenomenon under different conditions.

#### **AIII.6. Quantification of Pulsus Paradoxus**

It is vital that pulsus paradoxus be quantified. Since qualitative findings are highly subjective, the statement that a patient does or does not have pulsus paradoxus, based on purely qualitative techniques (e.g., palpation of a "marked drop in pulse strength"), is problematical.

There exists evidence that the degree of pulsus paradoxus is related to the *type* of the underlying disease. In 1924 Gauchat and Katz [Gauc24] noted that patients with pericardial disease seem to have a more pronounced varia-

tion of pulse amplitude than patients with respiratory disease. Recent investigations tend to support this observation. For example, in a study of 10 patients with chronic obstructive airway disease, the mean inspiratory drop in systolic pressure was about 18 mm Hg [Sett80], whereas in a study of 53 medical tamponade patients, the value was about 49 mm Hg [Gube81]. However, Lange [Lang66] reports a mean value of slightly under 10 mm Hg in 16 patients with constrictive pericarditis.

There also exists evidence that the degree of pulsus is related to the *severity* of the underlying disease: two investigations found a direct correlation between the degree of pulsus and the severity of asthma (as quantified by various ventilatory measurements) [Know73, Rebu73]. Similarly, a recent hemodynamic study found a correlation between the degree of pulsus, and the severity of cardiac tamponade [Curt85].

There are several features of the arterial pressure which should be quantified. It is known that as a patient with cardiac tamponade deteriorates, compensatory mechanisms (e.g., increased vascular resistance) become inadequate and systemic arterial pressure progressively falls [Brau84]. Therefore, these patients, already having very low arterial pressures, would not be expected to have a pronounced pulsus paradoxus [Spod67]. Hence, merely noting the magnitude of an inspiratory decline in systolic pressure may be misleading. Accordingly, both inspiratory and expiratory systolic arterial pressures should be recorded in order that the clinician be able to evaluate the relative significance of the observed phenomenon.

If the diastolic arterial pressure remains constant during the inspiratory cycle, then the inspiratory drop in pulse pressure will be equal to the inspiratory drop in systolic arterial pressure. In fact, however, inspiration may

affect the diastolic pressure. For example, in a patient with severe chronic bronchitis and emphysema, diastolic pressure declined to such an extent on inspiration, that even though the patient had a marked decline in systolic pressure, he maintained an essentially normal pulse pressure [Shab70]. Similarly, another author states that "the diastolic pressure in tamponade changes very little, while in bronchospasm it varies as much as does systolic pressure" [Cons76]. In view of the above, it is apparent that the levels of the inspiratory and expiratory diastolic arterial pressures are potentially important and should be recorded.

In order to acquire inspiratory and expiratory systolic arterial pressure readings, either sphygmomanometry or arterial cannulation may be used [Brau84, Cons76, Shab81]. Accurate diastolic values, however, are notoriously difficult to gather by the cuff, and therefore should ideally be acquired by invasive means. Since this is not always realistic in a clinical setting, it may be possible to obtain only approximations of the diastolic values by use of the sphygmomanometer. Note, however, that in seriously ill patients, an invasive transducer system generally will be in use.

Finally, it is emphasized that our insistence on quantification is not meant to denigrate the value of the physical examination. There is no doubt that a skilled clinician might, by palpation, be the first to detect pulsus paradoxus. Fowler [Fow180] has noted that a paradoxical pulse is sometimes palpable in the femoral or carotid arteries even when the radial pulse is imperceptible.

#### **AIII.7. Definitions of Pulsus Paradoxus From the Literature**

Many distinct definitions of pulsus paradoxus can be found in the literature: in the course of our literature review we encountered over ten distinct definitions that could lead to different conclusions as to the presence or absence of a paradoxical pulse [Amer71, Curt82, Dock61, Hurs86, Lang67, Redd78, Selz83, Spod64, Urda83, Wagn73, Wyng85]. Many of these definitions are from relatively current literature.

Since pulsus paradoxus was described by Kussmaul before auscultatory blood pressure determinations had been devised [Lang67], the earliest definitions of the phenomenon were purely qualitative. Qualitative definitions, however, are still in use today [Dorl81, Sted82, Thom81]. For example, one current cardiology text [Selz83] defines pulsus paradoxus as follows: "Reduction in pulse amplitude during inspiration (which is merely an exaggeration of the normal respiratory effect upon the arterial pulse)."

In general, recent definitions tend to be more quantitative: For example, pulsus paradoxus has been defined as an inspiratory drop in systolic pressure of greater than 5 mm Hg [Spod64], greater than 13 mm Hg [Curt83], greater than 20 mm Hg [Lang67], and greater than 10 mm Hg [Wagn73]. The latter is the most commonly accepted value. Other authors have defined pulsus in terms of the percentage decline in systolic pressure [Redd78], in terms of the pulse pressure (i.e., systolic - diastolic difference) [Urda83], and in a mixed quantitative - qualitative fashion [Hurs86]. Yet other definitions have tied pulsus paradoxus to "a rise or no significant fall in venous pressure" [Dock61] (i.e., Kussmaul's sign), or do not specify whether pulsus is an inspiratory or expiratory phenomenon - "A change of systolic blood pressure of greater than 10 mm Hg with respiration" [Wyng85].

There are several difficulties with the definitions of pulsus paradoxus that

are to be found in the literature. As was discussed at length, it is crucial to quantify pulsus. Since qualitative definitions tend to encourage lack of quantification, it is not desirable to use definitions that are qualitative in whole or part. Several definitions are based on the decline in pulse pressure, but as was discussed previously, an inspiratory fall in systolic pressure is not synonymous with an inspiratory fall in pulse pressure.

Quantitative definitions in the literature invariably attempt to define pulsus paradoxus in terms of a numerical threshold: if the inspiratory decline in systolic arterial pressure (or pulse pressure) exceeds this threshold, then the sign is considered to be present, otherwise, the sign is considered to be absent. There are several difficulties with this approach. First of all, the success of such an approach is predicated upon the ability to determine a meaningful threshold. After years of intensive investigation it seems that an upper limit somewhere in the range of 5-10 mm Hg is reasonable; however, it is still not possible to pick an exact cutoff.

Even if the data were available to establish a precise threshold, the threshold would simply not be meaningful for many patients. For example, in a healthy individual 7 mm Hg might be below the threshold, but in the decompensating patient with cardiac tamponade, an inspiratory decline of 7 mm Hg might represent complete obliteration of the arterial pulse.

*The above problems stem from viewing pulsus paradoxus as a sign which is either present or absent.* It seems to us that pulsus paradoxus should be viewed from a different perspective: it is a phenomenon that should be quantified; and it can only be interpreted in the light of other physiologic variables.

### **AIII.8. Confusion About Pulsus Paradoxus**

There have been some major difficulties in the reporting of pulsus paradoxus in both patient records, and in the literature. As has been emphasized, it is crucial to quantify pulsus paradoxus, yet often the phenomenon is not quantified. For example, in a study of medical tamponade patients, 3 out of 56 patient records noted the presence of pulsus, but gave no indication of its magnitude [Gube81].

Even more disturbing is the fact that pulsus is frequently not noted for those patients in whom it clearly should be. For example, in a study of 310 open-heart surgery cases, pulsus paradoxus was "not specifically discussed in most of the patients' records" [Enge70]. Even for patients in whom a diagnosis of tamponade has been established, pulsus is not always noted: in a literature case report of post-surgical cardiac tamponade, no mention is made of pulsus [Rufi77], and in a hemodynamic study of 13 patients with wound-induced tamponade, pulsus is not discussed [Shoe73].

A few years ago a survey was done to determine what percentage of medically trained individuals could "correctly define" pulsus paradoxus [Sapi83]. This survey considered pulsus paradoxus to be an inspiratory decrease in systolic pressure of greater than 10. The result was that 13% of the medical students, 23% of the house staff, and 57% of the attending physicians could correctly define pulsus, although "only medical students produced definitions that were clearly wrong" [Sapi83]. It is impossible to draw any absolute conclusions; nevertheless, one must wonder whether the multiplicity of definitions in the literature accounts in some degree for the confusion.

### **AIII.9. A Definition**

In order to overcome the limitations of definitions in the literature we will suggest a definition based on the following points: 1) Qualitative criteria are not adequate. 2) A slight inspiratory drop in systolic pressure is generally normal, but it is difficult to say at exactly what point this drop becomes pathologic. 3) Abnormal ventilation can influence the presence and/or magnitude of pulsus paradoxus. Therefore, it is not realistic to tie the definition of pulsus paradoxus to "normal inspiration."

The definition that we suggest is as follows: *Pulsus paradoxus is the measured value of the fall in systolic arterial pressure during inspiration.* In order to properly use this definition, certain pieces of information should be provided for each patient when noting pulsus paradoxus. In particular, systolic and diastolic pressures should be provided in inspiration and expiration if possible. In addition, the ventilatory status should be clearly described: rate, rhythm and depth, and other factors as deemed appropriate (e.g., anatomic factors, spirometry, pulmonary function tests, etc.).

Note that this proposed definition of pulsus paradoxus virtually forces quantification. In addition, for the reasons discussed above, we have chosen not to set an exact value at which to consider the phenomenon pathologic. If the clinician wishes to draw attention to a seemingly pathologic degree of pulsus, then he or she might note the presence of an "exaggerated " or "pathologic" pulsus paradoxus. Since quantification should always be provided, problems of ambiguity and misinterpretation will be avoided.

### **AIII.10. Summary**

The paradoxical pulse has been known to clinicians for many years. Unfortunately, there is no uniformity in usage of the term *pulsus paradoxus*. Many definitions have been proposed, yet none is agreed upon in the literature. The phenomenon is frequently not noted when it should be, and appropriate quantification is rarely provided. The underlying mechanisms have not been definitively established. It is the pervasive nebulosity surrounding this frequently-occurring phenomenon that we consider to be the true "paradox of pulsus paradoxus."

We have suggested that pulsus paradoxus be viewed as a continuous phenomenon, rather than a clinical sign that is either present or absent. Our suggested definition for the phenomenon is that "pulsus paradoxus is the measured value of the fall in systolic arterial pressure during inspiration." We have also emphasized the importance of quantifying the phenomenon. In particular, systolic and diastolic pressures should be noted during both inspiration and expiration, and a description of relevant ventilatory variables should be provided.

**APPENDIX IV****CARDIAC TAMPONADE IN POST-CARDIAC SURGICAL PATIENTS****Note**

This appendix has been submitted for publication.

**AIV.1. Introduction**

A serious complication of open-heart surgery is cardiac tamponade. Occurring only infrequently, tamponade is of great significance since its outcome may be fatal, unless it is promptly recognized and treated. In this report we describe the case of a patient who developed cardiac tamponade shortly after open intracardiac surgery, and who manifested many of the typical signs and symptoms. We also include a review of the literature.

**AIV.2. Case Report**

A 71 year old white female was admitted for mitral valve replacement. Her past medical history was significant for hypertension, rheumatic heart disease, congestive heart failure, and long standing atrial fibrillation. Approximately eight years prior to admission the patient had a pacemaker implant because of refractory bradycardia (secondary to sick sinus syndrome). Preoperative cardiac catheterization revealed a severely calcified

mitral annulus, 4+ mitral regurgitation, and mild aortic insufficiency. Coronary arteriography revealed modest narrowing of the left anterior descending and distal right coronary arteries; these narrowings were judged to be of a magnitude not requiring surgical intervention. Preoperative partial thromboplastin time (PTT) was 27.2/25.5, platelet count was 307,000, and prothrombin time (PT) was 13.3/12.5 (Coumadin had been stopped one week prior to surgery).

At operation the heart was approached via a median sternotomy and the patient placed on cardiopulmonary bypass. The incompetent mitral valve was found to be heavily calcified and was fused to the heavily calcified annulus. The calcifications extended into the papillary muscles and the left ventricular wall; these massive calcifications, coupled with significant ventricular hypertrophy, made the surgery inordinately difficult. The valve was excised and replaced with a 25 mm porcine bioprosthesis (this was a suboptimal size but because of the hypertrophy and calcifications it was not possible to insert a larger prosthesis; due to cracking of the calcifications sutures would not hold in the annulus, therefore the prosthesis was anchored by placing pledgeted horizontal mattress sutures adjacent to the annulus in the left atrium). The pacemaker generator (Medtronic Spectrax #8422) was then replaced. Anterior and posterior mediastinal drainage tubes were placed, and the pericardium was left open. The sternum was closed and the patient was transferred to the cardiac surgical intensive care unit (CSICU). Due to the difficulty of the surgery the patient was on bypass for a prolonged period of time (approximately five hours); the total length of surgery was approximately eight hours.

The patient's postoperative course was initially marginal: mean arterial pressure was in the 60s and 70s, and the cardiac output was 2.9 (cardiac index = 2.27), however, the mean right atrial pressure (*MRAP*) was 17, and mean left atrial pressure (*MLAP*) was 27. In addition, there was significant drainage from the chest tubes: 475cc of blood in the first 100 postoperative minutes. Over the next two hours the chest tubes drained another 600cc. The PT was measured as 13/10, and the PTT was 67/31; 2 units of fresh frozen plasma were administered. The platelet count was 198,000; 12 units of platelets were administered.

Approximately forty minutes later (four hours postoperatively) the patient's condition began to markedly deteriorate. The *MRAP* and *MLAP* were elevated and nearly equal at 23 and 25 respectively. The mean arterial pressure was reduced to the 30s to 40s and urine output had ceased. Epinephrine was administered in order to increase the blood pressure, but the patient continued to decompensate.

In view of the patient's grave condition, immediate reoperation was undertaken in the CSICU. The patient's chest was reopened revealing a large volume of blood and clots surrounding the heart. An active bleeding site was found in the atrioventricular groove underneath the left atrial appendage (the site corresponding to the location of the left circumflex coronary artery). Bleeding was controlled with finger pressure, and the patient was brought back to the operating room. She was placed on cardiopulmonary bypass, the laceration was repaired, and a saphenous vein graft from the aorta to the distal circumflex was performed in order to bypass the bleeding site. An intra-aortic balloon pump was placed and the

patient came of bypass easily.

Although the patient's cardiac status improved after reoperation she had a stormy postoperative course: One day subsequent to reoperation an embolus lodged in her right brachial artery - an embolectomy was performed; she also had a cerebrovascular accident which resulted in right sided hemiparesis and dysphasia. Several weeks later the patient developed renal failure; she recovered from this; she also developed an upper gastrointestinal bleed. The patient was eventually discharged 86 days after admission; she was sent to a skilled nursing facility because of right sided hemiparesis and dysphasia.

#### **AIV.3. Tamponade in Cardiac Surgical Patients**

Although virtually any complication is possible, there are only a few complications which are *typically* seen in patients after cardiac surgery. Among the complications which are typically seen in the early postoperative period (up to the first 48 hours) are arrhythmias, fluid and electrolyte imbalances (e.g., hypovolemia), low cardiac output, hypertension (especially in coronary artery bypass and graft patients), oliguria or frank renal failure, low output syndrome, and *early* cardiac tamponade [Behr80]. Although cardiac tamponade is not a common source of postoperative morbidity and mortality, it is by no means rare: e.g., Engelman et al [Enge70] reported an incidence of 1.3% (4 out of 310 cases), and Nelson et al [Nels69] reported an incidence of 2.0% (9 out of 446 cases).

Tamponade can also occur later in the postoperative course; it is then

known as *delayed* or *late* tamponade [Bork81,Elli74]. Typically, delayed tamponade occurs from 4 to 40 days postoperatively; however, it can also occur much later, e.g., 3 months after surgery [Fras73]. Other typical complications in the later postoperative period include pulmonary embolism, postpericardiotomy syndrome, and congestive heart failure. This paper concentrates on early tamponade, although much of the discussion is valid for late tamponade as well.

#### **AIV.4. Causes of Cardiac Tamponade**

Cardiac tamponade is the condition wherein blood or some other substance compresses the heart to the extent that the venous return is impaired, thereby compromising circulation. Cardiac tamponade occurs in patients who have undergone cardiac surgery; it also occurs in patients who have not undergone cardiac surgery.

In non-surgical patients tamponade is classically associated with a pericardial effusion. A pericardial effusion may be secondary to infection (e.g., tuberculosis, viral pericarditis), neoplasia (e.g., lung, breast, lymphoma), uremia, or irradiation [Brau87]. Sometimes no cause can be identified for the effusion.

Non-surgical patients may also develop tamponade because of blood compressing the heart (hemopericardium). Hemopericardium may be secondary to a ruptured ventricle following a myocardial infarction [Choo85]. Other sources of hemopericardium include ruptured ventricular aneurysms, a ruptured sinus of Valsalva [Kill87], aortic rupture into the pericardium

[Frie87], penetrating wounds of the heart, and blunt trauma to the chest resulting in cardiac rupture [Lear87]. There are iatrogenic causes as well; e.g., cardiac catheterization. In addition, patients with acute pericarditis who are on anticoagulants may develop hemopericardium [Brau87].

Rare instances have been reported of air or some other gas compressing the heart (pneumopericardium) [Cumm84]. Gas may be formed by an infectious organism; air may gather due to an anomalous connection to an air containing organ (e.g., lung, esophagus).

In the post-cardiac surgical patient tamponade is most frequently associated with blood compressing the heart. Blood may collect in the pericardium, or it may collect in the mediastinum; it may be in liquid form or it may be clotted. The source of blood is not always clear: in some cases reoperation may reveal an active bleeding site, however, in many cases a bleeding source is not found [Nels69]. In rare instances blood may collect because of ventricular rupture following a postoperative myocardial infarction [Boja87].

Tamponade may also occur in post-cardiac surgical patients because of pericardial effusions. These effusions may occur shortly after surgery, e.g., [Rufi77], or they may occur much later in the postoperative course, e.g. [McCa74].

Another cause of postoperative tamponade is acute dilation of the heart (e.g. due to valvular regurgitation) [Behr80]. Rare instances have also been reported of chyle compressing the heart many days after surgery [Rose82], and of postoperative pneumopericardium [Cumm84].

Although it is quite uncommon, cardiac tamponade may also occur *during* cardiac surgery. In these cases tamponade occurs when the sternum is

closed over a markedly distended or edematous heart (*tight mediastinal syndrome*).

#### **AIV.5. Physiology of Cardiac Tamponade**

The physiology of tamponade has been well described in non-(cardiac) surgical patients [Lang67,Shab70,Shar60] as well as in victims of penetrating injuries to the heart [Shoe73]; the physiologic mechanisms appear to be the same in cardiac surgical patients [Hill82].

Tamponade begins when some agent (e.g., blood, effusion, etc.) compresses the heart. This compression leads to impaired diastolic filling, hence a reduced stroke volume. Compensatory mechanisms are then activated in order to maintain adequate perfusion: the heart rate increases and peripheral vessels constrict. As tamponade progresses the compensatory mechanisms become inadequate: the cardiac output falls, the arterial pressure then falls, and, if the patient is not promptly treated, death ensues.

As would be expected, a greater degree of compression leads to a greater degree of cardiac impairment [Brau87]. In addition, the time period over which compression develops is also of significance; for example, in acute situations as little as 200cc of fluid in the pericardium can lead to tamponade [Ste173]; however, if fluid accumulates slowly then as many as 1000cc of fluid may accumulate before tamponade develops [Brau87].

Generally, the entire heart is compressed uniformly; however, rare instances have been reported in which only isolated chambers of the heart are compressed: [Bate82] reports a case of right atrial tamponade, [Hill69]

reports a case of right ventricular tamponade, [Yaco66] reports a case of left atrial tamponade, [Jone79] reports a case of left ventricular tamponade, [Mill78] and [Phil83] report right sided tamponade, and [Simp73] reports left sided tamponade. The physiology in these cases depends on the particular chambers involved; for example, the physiology of left ventricular tamponade is similar to that of left ventricular failure [Jone79].

#### **AIV.6. Signs and Symptoms of Cardiac Tamponade**

The signs and symptoms of cardiac tamponade have been well described in the literature, e.g., [Brau87,Frie87]. Physical examination typically reveals a small quiet heart and faint breath sounds. The jugular venous pressure is elevated (because of cardiac compression) resulting in jugular venous distension. The impaired circulatory status is reflected in a low arterial pressure and, if the impairment is severe, urine output is decreased. Features of congestive heart failure may also be present; e.g., dyspnea, orthopnea, and hepatic engorgement.

Catheterization typically reveals an elevated venous pressure (possibly exceeding 30 cm [Dail81]), with a prominent x descent. The right atrial, right ventricular and pulmonary artery diastolic pressures, and the pulmonary wedge pressure are typically elevated and approximately equal (usually in the mid 20s). Cardiac output is generally reduced. Chest x-ray may reveal a widened mediastinum [Dail81]; the lung fields are usually clear [Brau87] (note, however, that x-ray findings may be quite unreliable due to technical factors; e.g., it is difficult to judge mediastinal size on a portable

AP projection). EKG changes may be present; e.g., decreased QRS amplitude, and electrical alternans of the P, QRS, and T waveforms [Brau87].

Not all signs and symptoms will necessarily be seen in a given patient. The particular pattern of signs and symptoms that is present depends on several factors: e.g., the severity of the cardiac compression, and the effectiveness of the compensatory mechanisms. In addition, the rapidity of the development of cardiac compression also plays a role: for example, features of congestive heart failure are more common in tamponade that develops gradually rather than abruptly [Brau87].

Tamponade in post-cardiac surgical patients typically presents as described above. However, alternative presentations may occur. For example, heart sounds may be of normal intensity if the anterior pericardium is absent [Week76].

It is frequently stated that pulsus paradoxus is a common, if not universal, phenomenon in tamponade; e.g., [Curt82,Dail81,DeGo87]. The occurrence of pulsus paradoxus in post-cardiac surgical tamponade, however, is not clear. Some authors state that the phenomenon is frequently present; e.g., Nelson et al [Nels69] reports pulsus paradoxus in 13 out of 13 cases of post-surgical tamponade. See also [Hill82]. On the other hand, Ellison and Kirsch [Elli74] report that in 9 cases of post-surgical tamponade none had pulsus paradoxus. See also [Jone79] and [Yaco66]. This wide discrepancy in the reported rates of occurrence may be due, in part, to the fact that pulsus paradoxus is not a well defined phenomenon [Henk87]. In addition, the presence of pulsus paradoxus in post-surgical patients may be obscured since these patients are frequently on ventilators: [Mass73] reports that intermit-

tent positive pressure ventilation may cause reversed pulsus paradoxus; this reversed pulsus paradoxus can mask an underlying true pulsus paradoxus [Week76]. Whatever the occurrence of pulsus paradoxus in tamponade may be, it is crucial not to put too much weight on the presence or absence of the phenomenon: [Wine79], for example, reports on a case (non-surgical) where the diagnosis of tamponade was missed because pulsus paradoxus was not present - this patient died.

It is also unclear what the occurrence rate of Kussmaul's sign (inspiratory rise in venous pressure) is in post-surgical tamponade: certain authors, for example, state that the sign is frequently present; e.g., [Dail81]; other authors, e.g., [Lang67,Spod84Ann] state that Kussmaul's sign is not found in tamponade.

#### **AIV.7. Predisposition to and Prevention of Cardiac Tamponade**

It is well known that patients who have been bleeding heavily after surgery are more likely to develop tamponade [Dail81]. Certain patients are predisposed to heavy bleeding: for example, patients who were on cardiopulmonary bypass for long periods of time, and patients who have had previous cardiac surgery [Behr80]. Other individuals who are predisposed to heavy postoperative bleeding include patients with preexisting coagulation defects, patients who have received high levels of anticoagulants during surgery [Fras73,Hill82], and patients who have received high doses of antiplatelet agents [Boja87].

Many cardiac surgeons leave the pericardium open, believing that this

will lower the probability of tamponade, e.g., [Hill82]. Certain investigators, however, have questioned the efficacy of this approach. For example, in a series of 450 patients, Asanza et al [Asan76] found essentially no difference in the frequency of tamponade between those patients who had closed pericardiums and those who had open pericardiums. Furthermore, based on a study of 100 patients, Cunningham et al [Cunn75] conclude that closing the pericardium can prevent the occurrence of cardiac tamponade (they suggest that the main sources of bleeding that lead to tamponade are extrapericardial).

Drainage tubes are invariably placed intraoperatively in order to prevent excess blood from accumulating postoperatively. These tubes do seem to exert a protective effect: if the drainage from the tubes suddenly ceases (e.g., due to clotting), tamponade may well follow [Behr80].

If drainage from the tubes is excessive, then reoperation is indicated to prevent tamponade (as well as exsanguination). One criteria that has been suggested is to reoperate if bleeding exceeds 300ml/hr for 4 to 6 hours [Behr80].

#### **AIV.8. Recognition of Cardiac Tamponade**

It is important to recognize the presence of tamponade as quickly as possible since immediate therapeutic intervention may be lifesaving. Most of the typical complications that arise in the early postoperative period, e.g., arrhythmias, hypertension, etc., are readily recognized and not easily confused with tamponade; distinguishing between low output syndrome (myo-

cardial failure) and tamponade can, however, be difficult at times.

Unfortunately, physical examination is often not helpful in making the distinction. One diagnostic clue is that tamponade is almost invariably associated with a high bleeding rate, whereas patients with myocardial failure usually have had normal bleeding rates [Nels69]. Intracardiac pressures are extremely useful: in tamponade the intracardiac pressures are elevated and approximately equal<sup>1</sup>; in myocardial failure the pressures are usually lower and not equal [Week76]. In addition, a steep x descent in the jugular venous pulse is characteristic of tamponade, in myocardial failure there is a steep y descent [Yaco66]. Echocardiography can also provide strong evidence suggestive of tamponade.

Chest x-ray may or may not be useful in recognizing tamponade; an enhanced technique using juxtaposed epicardial and pericardial clips is described in [Meck74]. According to Malcolm and Poirier [Malc87] epicardial pacing wires may malfunction in tamponade (these authors postulate that clotted blood surrounding the heart acts as an electrical ground and prevents ventricular capture). See also [Frat75].

#### **AIV.9. Treatment of Cardiac Tamponade**

Treatment of cardiac tamponade is aimed at decompressing the heart, thereby allowing increased diastolic filling, hence improving circulatory dynamics. In acute situations, immediate decompression is mandatory. This may be accomplished by performing mediastinocentesis or pericardiocentesis using a large bore needle. More effective removal of blood and clots may be

achieved by opening the lower end of the median sternotomy incision and inserting a sucker. It should be emphasized that the removal of as little as 25cc of fluid can decompress the heart sufficiently to prevent death [Elli74,Fish85]. Another technique that can be effective in life threatening situations is to increase the mediastinal space by opening the lower end of the sternotomy incision, inserting a finger, and exerting upward traction on the sternum [Hill82].

In some cases, evacuation of fluid may be sufficient therapy, e.g., certain instances of anterior tamponade [Hill82]; however, more definitive treatment is usually needed. In cases of hemopericardium surgical exploration is generally required in order to effectively evacuate clots, and to repair bleeding sites (if they can be located). In cases of pericardial effusion radiation therapy [Rufi77] or surgical resection of the pericardium [McCa74] may be used.

Cardiac tamponade that occurs during surgery (tight mediastinal syndrome) is treated by delaying sternal closure. Alternative treatments include partial pericardiectomy [Fund86], and a method using sternal traction [Riha75].

#### **AIV. 10. Summary**

Tamponade is a potentially life threatening complication of cardiac surgery. In post-cardiac surgical patients tamponade is most frequently due to hemopericardium; pericardial effusions and other mechanisms are also possible. Tamponade generally occurs in patients who have had excessive post-

surgical bleeding. Physical signs include jugular venous distension, hypotension, and low urine output. Intracardiac pressures are elevated and approximately equal; the cardiac output is low. Echocardiography can be helpful in making the diagnosis. Pulsus paradoxus is *not* a reliable sign in post-surgical tamponade. Immediate therapy for tamponade is decompression of the heart by evacuation of fluid; more definitive treatment is usually also required (i.e., reoperation).

#### Footnotes

- 1 Extremely rare exceptions occur when tamponade affects an isolated chamber; for example, in isolated right atrial tamponade the right atrial pressure is greater than the pulmonary artery diastolic and wedge pressures [Bate82]. As another example, in isolated right ventricular tamponade the right ventricular pressure is much greater than the pulmonary artery pressure [Hill69].

**BIBLIOGRAPHY**

Adams JB. A probability model of medical reasoning and the MYCIN model. *Mathematical Biosciences* 1976;32:177-86.

Adlassnig KP. A survey on medical diagnosis and fuzzy subsets. In: Gupta MM, Sanchez E, eds. *Approximate reasoning in decision analysis*. New York: North-Holland, 1982:203-17.

Aikins JS, Kunz JC, Shortliffe EH. PUFF: an expert system for interpretation of pulmonary function data. *Computers and Biomedical Research* 1983;16:199-208.

American College of Cardiology and American Heart Association. Glossary of cardiologic terms related to physical diagnosis: Part IV. Arterial Pulses. *Am J Cardiol* 1971;27:708-9.

Asanza L, Rao G, Voleti C, Harstein ML, Wisoff BG. Should the pericardium be closed after an open-heart operation? *Ann Thorac Surg* 1976;22:532-4.

Barnett JA. Computational methods for a mathematical theory of evidence. *Proc seventh IJCAI* 1981:868-75.

Barr A, Feigenbaum EA, Cohen PR. *The handbook of artificial intelligence*. Los Altos CA: William Kaufman, 1981.

Bateman T, Gray R, Chaux A, et al. Right atrial tamponade complicating cardiac operation. *J Thorac Cardiovasc Surg* 1982;84:413-9.

Behrendt DM, Austen WG. *Patient care in cardiac surgery*. 3rd ed. Boston: Little Brown, 1980:79,101-46.

Benis AM, Fitzkee HL, Jurado RA, Litwak RS. Improved detection of adverse cardiovascular trends with the use of a two-variable computer alarm. *Crit Care Med* 1980;8:341-4.

Birnbaum L. Lexical ambiguity as a touchstone for theories of language analysis. *Proceedings ninth IJCAI* 1985:815-20.

Bojar RM, Overton JW, Madoff IM. Successful management of left ventricular rupture following myocardial revascularization. *Ann Thorac Surg* 1987;44:312-4.

Bonissone PP, Tong RM. Editorial: reasoning with uncertainty in expert systems. *Int J Man-Machine Studies* 1985;22:241-50.

**Borkon AM, Schaff HV, Gardner TJ.** Diagnosis and management of postoperative pericardial effusions and late cardiac tamponade following open-heart surgery. *Ann Thorac Surg* 1981;31:512-9.

**Braunwald E, ed.** Heart disease a textbook of cardiovascular medicine. 2nd ed. Philadelphia: W B Saunders, 1984:25:1481-2.

**Braunwald E, Isselbacher KJ, Petersdorf RG, eds.** Harrison's principles of internal medicine. 11th ed. New York: McGraw-Hill, 1987:1008-14.

**Buchanan BG, Shortliffe EH, eds.** Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project. Reading: Addison-Wesley, 1984.

**Cendrowska J, Bramer MA.** A rational reconstruction of the MYCIN consultation system. *Int J Man-Machine Studies* 1984;20:229-317.

**Charniak E.** The Bayesian basis of common sense medical diagnosis. *Proc AAAI* 1983:70-3.

**Cheeseman P.** A method for computing generalized Bayesian probability values for expert systems. *Proc eighth IJCAI* 1983:198-202.

**Cheeseman P.** In defense of probability. *Proc ninth IJCAI* 1985:1002-9.

**Choo MH, Chia BL, Chia F.** Cardiac tamponade from ventricular rupture: value of two-dimensional echocardiography in guiding acute surgical management. *Crit Care Med* 1985;13:446-7.

**Clancey WJ.** Dialogue management for rule-based tutorials. *Proc sixth IJCAI* 1979:155-61.

**Clancey WJ.** Details of the revised therapy algorithm. In: Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project. Reading: Addison-Wesley, 1984:133-46.

**Clancey WJ, Letsinger R.** NEOMYCIN: reconfiguring a rule-based expert system for application to teaching. *Proc seventh IJCAI* 1981:829-36.

**Clancey WJ, Shortliffe EH, eds.** Readings in medical artificial intelligence. Reading Mass: Addison-Wesley, 1984.

**Cohen PR, Ginsberg MR.** A framework for heuristic reasoning about uncertainty. *Proc eighth IJCAI* 1983:355-7.

Cohen S, Kupersmith J, Aroesty J, Rowe JW. Pulsus paradoxus and Kussmaul's sign in acute pulmonary embolism. *Am J Cardiol* 1973;32:271-5.

Cohn JN, Pinkerson AL, Tristani FE. Mechanism of pulsus paradoxus in clinical shock. *J Clin Invest* 1967;46:1744-55.

Connely DP, Dean DW, Hultman BK. Physician-oriented result reporting in an intensive care environment. In: Salamon R, Blum B, Jorgensen M, eds. *Proceedings of the fifth conference on medical informatics*. Amsterdam: North-Holland, 1986:810-2.

Constant J. *Bedside cardiology*. 2nd ed. Boston: Little Brown, 1976:61-4.

Croft DJ. Is computerized diagnosis possible? *Computers and Biomedical Research* 1972;5:351-67.

Cullen DJ, Teplick R. The role of computers in the future of intensive care. *Proc IEEE* 1979;67:1307-8.

Cummings RG, Wesley RLR, Adams DH, Lowe JF. Pneumopericardium resulting in cardiac tamponade. *Ann Thorac Surg* 1984; 37:511-8.

Cunningham JN Jr., Spencer FC, Zeff R, Williams CD, Cukingnan R, Mullin M. Influence of primary closure of the pericardium after open-heart surgery on the frequency of tamponade, postcardiotomy syndrome, and pulmonary complications. *J Thorac Cardiovasc Surg* 1975;70:119-25.

Curtiss, EI, Lindsey, RL, Reddy, PS. Diagnostic criteria for pulsus paradoxus and abnormal inspiratory decrease of ejection time in cardiac tamponade. In: Reddy, PS ed. *Pericardial disease*. New York: Raven Press, 1982:203-14.

Curtiss E, Reddy P, Uretsky B, Cecchetti A. Relation between cardiac tamponade severity and pulsus paradoxus. *J Am Col Cardiol* 1985;5:529.

Daily EK, Schroeder JS. *Techniques in bedside hemodynamic monitoring*. 2nd ed. St. Louis: CV Mosby, 1981:165.

Davis R. Interactive transfer of expertise: acquisition of new inference rules. *Proc fifth IJCAI* 1977:321-8.

Davis R. Meta-rules: reasoning about control. *Artificial Intelligence* 1980;15:179-222.

Davis R, Buchanan B, Shortliffe E. Production rules as a representation for a knowledge-based consultation program. *Artificial Intelligence* 1977;8:15-

45.

DeAsla RA, Benis AM, Jurado RA, Litwak RS. Management of postcardiotomy hypertension by microcomputer-controlled administration of sodium nitroprusside. *J Thorac Cardiovasc Surg* 1985;89:115-20.

DeDombal FT, Leaper DJ, Horrocks JC, Staniland JR, McCann AP. Human and computer-aided diagnosis of abdominal pain: further report with emphasis on performance of clinicians. *Brit Med J* 1974;1:376-80.

DeDombal FT, Leaper DJ, Staniland JR, McCann AP, Horrocks JC. Computer-aided diagnosis of acute abdominal pain. *Brit Med J* 1972;2:9-13.

DeGowin RL. DeGowin & DeGowin's bedside diagnostic examination. 5th ed. New York: Macmillan, 1987:407-8.

Dempster AP. Upper and lower probabilities induced by a multivalued mapping. *Annals Mathematical Statistics* 1967;38:325-39.

Dock W. Inspiratory traction on the pericardium. *Arch Intern Med* 1961;108:837-40.

Dorland's Illustrated Medical Dictionary. 26th ed. Philadelphia: W B Saunders, 1981:1095.

Drazen EL. Use of computerized monitoring systems in intensive care. In: *Proceedings of the IEEE conference on computers in cardiology*. 1980:323-5.

Dubois D, Prade H. Outline of fuzzy set theory: an introduction. In: Gupta MM, Ragade RK, Yager RR, eds. *Advances in fuzzy set theory and applications*. New York: North-Holland, 1979:27-48.

Dubois D, Prade H. *Fuzzy sets and systems: theory and applications*. New York: Academic Press, 1980.

Duda RO, Hart PE, Nilsson NJ. Subjective Bayesian methods for rule-based inference systems. In: *AFIPS proceedings of the 1976 National Computer Conference*, vol 45. New York, 1976:1075-82.

Duda RO, Shortliffe EH. Expert systems research. *Science* 1983;220:261-8.

Ellison LH, Kirsch MM. Delayed mediastinal tamponade after open-heart surgery. *Chest* 1974;65:64-6.

Engleman RM, Spencer FC, Reed GE, Tice DA. Cardiac tamponade following open-heart surgery. *Circulation* 1970;41 and 42(supp 2):165-71.

**Esogbue AO, Elder RC.** Fuzzy sets and the modelling of physician decision processes, part 1: the initial interview-information gathering session. *Fuzzy Sets and Systems* 1979;2:279-91.

**Estafanous FG, Tarazi RC.** Systemic arterial hypertension associated with cardiac surgery. *Am J C* 1980;46:685-94.

**Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ.** Representation of dynamic clinical knowledge: measurement interpretation in the intensive care unit. *Proc sixth IJCAI* 1979:260-2.

**Fagan LM, Kunz JC, Feigenbaum EA, Osborn JJ.** Extensions to the rule-based formalism for a monitoring task. In: Buchanan BG, Shortliffe EH, eds. *Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project*. Reading: Addison-Wesley, 1984:397-423.

**Fagan LM, Shortliffe EH, Buchanan BG.** Computer based medical decision making: from MYCIN to VM. In: Clancey WW, Shortliffe EH, eds. *Readings in medical artificial intelligence*. Reading Mass: Addison-Wesley, 1984:241-55.

**Fieschi M, Joubert M, Fieschi D, Soula G, Roux M.** Sphinx: an interactive system for medical diagnosis aids. In: Gupta MM, Sanchez E, eds. *Approximate reasoning in decision analysis*. New York: North-Holland, 1982:269-75.

**Fishman MC, Hoffman AR, Klausner RD, Thaler MS.** *Medicine*. 2nd ed. Philadelphia: Lippincott, 1985:72-3.

**Floyer J.** *A treatise of the asthma*. 2nd ed. London: R. Wilkin; 1717. [cited by Settle].

**Fowler NO, ed.** *Cardiac diagnosis and treatment*. 3rd ed. New York: Harper and Row, 1980:981-3,994.

**Fraser DG, Ulliyot DJ.** Mediastinal tamponade after open-heart surgery. *J Thorac Cardiovasc Surg* 1973;66:629-31.

**Frater RWM.** Techniques for distinguishing between pericardial tamponade and myocardial failure [Letter]. *J Thorac Cardiovasc Surg* 1975;70:541-2.

**Friedman HH.** *Problem-oriented medical diagnosis*. 4th ed. Boston: Little Brown, 1987:40.

**Friedman HS, Sakurai H, Lajam F.** Pulsus paradoxus: a manifestation of a marked reduction of left ventricular end-diastolic volume in cardiac tamponade. *J Thorac Cardiovasc Surg* 1980;79:74-82.

Fuenning C, Wise R, Brower R, Permutt S. The Mechanism of pulsus paradoxus in cardiac tamponade [Abstract]. Federation Proceedings 1985;44(5):1379.

Fukami S, Mizumoto M, Tanaka K. Some considerations on fuzzy conditional inference. Fuzzy Sets and Systems 1980;4:243-73.

Fundaro P, Santoli E, Rossi FS, Botta M, Santoli C. Partial pericardiectomy to prevent cardiac compression in open-heart surgery. Ann Thorac Surg 1986;41:581-4.

Gaines BR. Foundations of fuzzy reasoning. Int J Man-Machine Studies 1976;8:623-68.

Gallanti M, Guida G, Spampinato L, Stefanini A. Representing procedural knowledge in expert systems: an application to process control. Proc ninth IJCAI 1985:345-52.

Garvey TD, Lowrance JD, Fischler MA. An inference technique for integrating knowledge from disparate sources. Proc seventh IJCAI 1981;319-25.

Gauchat HW, Katz LN. Observations on pulsus paradoxus (with special reference to pericardial effusion) I. clinical. Arch Intern Med 1924;33:350-70.

Golinko RJ, Kaplan N, Rudolph AM. The mechanism of pulsus paradoxus during acute pericardial tamponade. J Clin Invest 1963;42:249-57.

Gomez F, Chandrasekaran B. Knowledge organization and distribution for medical diagnosis. IEEE-SMC 1981;11:34-42.

Gordon J, Shortliffe EH. The Dempster-Shafer theory of evidence. In: Buchanan BG, Shortliffe EH, eds. Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project. Reading: Addison-Wesley, 1984:133-46.

Gorry GA, Silverman H, Pauker SG. Capturing clinical expertise: a computer program that considers clinical responses to digitalis. Am J Med 1978;64:452-60.

Groszoff BN. Evidential confirmation as transformed probability: on the duality of priors and updates. In: Kanal LN, Lemmer JF, eds. Uncertainty in artificial intelligence. New York: North-Holland, 1986:153-66.

Groszoff BN. An inequality paradigm for probabilistic knowledge: the logic of conditional probability intervals. In: Kanal LN, Lemmer JF, eds.

Uncertainty in artificial intelligence. New York: North-Holland, 1986:259-75.

Guberman BA, Fowler NO, Engel PJ, Gueron M, Allen JM. Cardiac tamponade in medical patients. *Circulation* 1981;64:633-40.

Guntheroth WG, Morgan BC, Mullins G. Effect of respiration on venous return and stroke volume in cardiac tamponade. *Circulation Research* 1967;30:381-90.

Gustafson DH, Kestly JJ, Greist JH, Jensen NM. Initial evaluation of a subjective Bayesian diagnostic system. *Health Services Research* 1971;6:204-13.

Gustafson DH, Kestly JJ, Ludke RL, Larson F. Probabilistic information processing: implementation and evaluation of a semi-P.I.P. diagnostic system. *Comput and Biomed Res* 1973;6:355-70.

Halpern JL, Levine R. *Bypass*. New York: Times Books, 1985:145.

Hayes-Roth F, Waterman DA, Lenat DB, eds. *Building expert systems*. Reading Mass: Addison-Wesley, 1983.

Henkind SJ. Imprecise meanings as a cause of uncertainty in medical knowledge-based systems. In: Lemmer J, Kanal LN eds. *Uncertainty in artificial intelligence, vol II*. New York: North-Holland, 1988:35-41.

Henkind SJ. Characteristic functions as a knowledge representation. 1989: In preparation.

Henkind SJ, Benis AM, Teichholz LE. Early tamponade in a post-cardiac surgical patient: case report and literature review. 1989: Submitted for publication.

Henkind SJ, Benis AM, Teichholz LE. Quantification as a means to increase the utility of nomenclature-classification systems. In: Salamon R, Blum B, Jorgensen M, eds. *Proceedings of the fifth conference on medical informatics*. Amsterdam: North-Holland, 1986:858-61.

Henkind SJ, Benis AM, Teichholz LE. The paradox of pulsus paradoxus. *Am Heart J* 1987;114:198-203.

Henkind SJ, Harrison MC. An analysis of four uncertainty calculi. *IEEE-SMC* 1988;18:700-14.

Hill JD, Johnson DC, Miller GE, Kerth WJ, Gerbode F. Latent mediastinal tamponade after open-heart surgery. *Arch Surg* 1969;99:808-14.

Hill JD, Rodvien R, Mielke CH. Bleeding and hemorrhagic complications. In: Litwak RS, Jurado RA, eds. Care of the cardiac surgical patient. Norwalk: Appleton-Century-Crofts, 1982:367-86.

Horrocks JC, McCann AP, Staniland JR, Leaper DJ, DeDombal FT. Computer-aided diagnosis: description of an adaptable system, and operational experience with 2,034 cases. Brit Med J 1972;2:5-9.

Hummel RA, Landy MS. Evidence as opinions of experts. In: Lemmer J, Kanal LN eds. Uncertainty in artificial intelligence, vol II. New York: North-Holland, 1988:43-53.

Hummel RA, Manevitz LM. Combining bodies of dependent information. Proc tenth IJCAI 1987:1015-7.

Hurst JW, ed. The heart. 6th ed. New York: McGraw-Hill, 1986:146.

Jacobs WR, Talano JV. Cardiomegaly and paradoxical pulse. Arch Intern Med 1978;138:1125-6.

Jones MR, Vine DL, Attas M, Todd EP. Isolated left ventricular tamponade. J Thorac Cardiovasc Surg 1979;77:142-6.

Jurado RA, Fitzkee HL, De Asla RA, Lukban SB, Litwak RS, Osborn JJ. Reduction of unexpected, life-threatening events in postoperative cardiac surgical patients: the role of computerized surveillance. Circulation Sup II 1977;56:44-9.

Jurado RA, Osborn JJ. Patient Surveillance and general care. In: Litwak RS, Jurado RA, eds. Care of the cardiac surgical patient. Norwalk: Appleton-Century-Crofts, 1982:119-60.

Kahn K, Gorry GA. Mechanizing temporal knowledge. Artificial Intelligence 1977;9:87-108.

Kahn MG, Fagan LM, Shortliffe EH. Context-specific interpretation of patient records for a therapy advice system. In: Salamon R, Blum B, Jorgensen M, eds. Proceedings of the fifth conference on medical informatics. Amsterdam: North-Holland, 1986:910-2.

Katz LH, Gauchat HW. Observations on pulsus paradoxus (with special reference to pericardial effusion) II. experimental. Arch Intern Med 1924;33:371-93.

Kennedy J, Kossman C. Nomenclatures in medicine. Bulletin of the Medical Library Association 1973;61:238-52.

Killen DA, Wathanacharoen S, Pogson GW. Repair of intrapericardial rupture of left sinus of Valsalva aneurysm. *Ann Thorac Surg* 1987;44:310-1.

Kittredege R, Lehrberger J, eds. *Sublanguage: studies of language in restricted semantic domains* New York: Walter de Gruyter, 1982.

Klopfenstein H, Schuchard G, Wann L, et al. The relative merits of pulsus paradoxus and right ventricular diastolic collapse in the early detection of cardiac tamponade: an experimental echocardiographic study. *Circulation* 1985;71:829-33.

Knowles GK, Clark TJH. Pulsus paradoxus as a valuable sign indicating severity of asthma. *Lancet* Dec. 15 1973;1356-9.

Koyama T, Kaihara S, Minamikawa T, Kurokawa T. Time-oriented features for a medical consultation system. *Proc seventh IJCAI* 1981:910-2.

Kulikowski CA. Artificial intelligence methods and systems for medical consultation. *IEEE PAMI* 1980;2:464-76.

Kussmaul A. Ueber schwierige mediastino-pericarditis und den paradoxen puls. *Berlin Klin Wschft.* 1873;10:433-5. [cited by Sharp].

Kyburg HE. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence* 1987;31:271-93.

Lange RL. Compressive cardiac and circulatory disorders: clinical and laboratory correlation. *Am Heart J* 1967;74:419-30.

Lange RL, Botticelli JT, Tsagaris TJ, Walker JA, Gani M, Bustamante RA. Diagnostic signs in compressive cardiac disorders, constrictive pericarditis, pericardial effusion, and tamponade. *Circulation* 1966;33:763-77.

Leaper DJ, Horrocks JC, Staniland JR, DeDombal FT. Computer assisted diagnosis of abdominal pain using "estimates" provided by clinicians. *Brit Med J* 1972;4:350-4.

Leavitt BJ, Meyer JA, Morton JR, Clark DE, Herbert WE, Hiebert CA. Survival following nonpenetrating traumatic rupture of cardiac chambers. *Ann Thorac Surg* 1987;44:532-5.

Lee NS, Grize YL, Dehnad K. Quantitative models for reasoning under uncertainty in knowledge-based expert systems. *Int J Intelligent Systems* 1987;2:15-38.

Lemmer JF, Barth SW. Efficient minimum information updating for Bayesian inferencing in expert systems. *Proc AAAI* 1982:424-7.

Lesmo L, Saitta L, Torasso P. Evidence combination in expert systems. *Int J Man-Machine Studies* 1985;22:307-26.

Lindberg DAB, Kingsland III LC, Grant KD, Sharp GC. The AIRHEUM knowledge-based consultant system in rheumatology [Abstract]. *Proceedings of the fifth world congress on medical informatics* 1986:1144.

Lindsay RK, Buchanan BG, Feigenbaum EA, Lederberg J. *Applications of artificial intelligence to chemistry: the DENDRAL project*. New York: McGraw-Hill, 1980.

Linkens DA. On-line control in medicine. In: Salamon R, Blum B, Jorgensen M, eds. *Proceedings of the fifth conference on medical informatics*. Amsterdam: North-Holland, 1986:405-11.

Long W. Criteria for computer generated therapy advice in a clinical domain. In: *Proceedings of the IEEE conference on computers in cardiology*. 1980:285-8.

Long WJ. Reasoning about state from causation and time in a medical domain. *Proc AAAI* 1983:251-4.

Long WJ. Causal reasoning in a physiologic model as a computational paradigm. In: *Proceeding of the IEEE MEDCOMP conference*. 1983.

Long WJ, Naimi S, Criscitiello MG, Kurzrok S. Reasoning about therapy from a physiologic model. In: Salamon R, Blum B, Jorgensen M, eds. *Proceedings of the fifth conference on medical informatics*. Amsterdam: North-Holland, 1986:756-60.

Long WT, Russ TA, Locke WB. Reasoning from multiple sources in arrhythmia management. In: *Proceedings of the IEEE conference on the frontiers of engineering and computing in health care*. 1983.

Lorell B, Leimbach RC, Pohost G, et al. Right ventricular infarction clinical diagnosis and differentiation from cardiac tamponade and pericardial constriction. *Am J Cardiol* 1979;43:465-71.

Lower R. *Tractatus de corde, item de motu et colore sanguinis et chyli in eum transitu*. Amsterdam: Elsevier Press; 1669:105-6. [cited by Dock].

Malcolm I, Poirier NL. Cardiac pacing in pericardial tamponade [letter to the editor]. *Ann Thorac Surg* 1987;44:331.

Maloney JV. The trouble with patient monitoring. *Ann Surg* 1968;168:605-14.

Martin J, Jardin J, Sampson M, Engel LE. Factors influencing pulsus paradoxus in asthma. *Chest* 1981;80:543-9.

Massumi RA, Mason DT, Vera Z, Zelis R, Otero J, Amsterdam EA. Reversed pulsus paradoxus. *N Engl J Med* 1973;289:1272-5.

McCabe JC, Engle MA, Ebert PA. Chronic pericardial effusion requiring pericardiectomy in the postpericardiotomy syndrome. *J Thorac Cardiovasc Surg* 1974;67:814-7.

McDermott J. R1: A rule-based configurer of computer systems. *Artificial Intelligence* 1982;19:39-88.

Meckstroth CV, Cattaneo SM. Juxtaposed epicardial-pericardial clips. *J Thorac Cardiovasc Surg* 1974;68:447-54.

McGregor M. Pulsus paradoxus. *N Eng J Med* 1979;301:480-2.

Michalski RS, Baskin AB. Integrating multiple knowledge representations and learning capabilities in an expert system: the ADVISE system. *Proc Eighth IJCAI* 1983;256-8.

Miller DC, Oyer PE, Ricks W, et al. Localized tamponade of the right atrium and right ventricle. *Arch Surg* 1978;113:764-6.

Miller RA, Pople HE, Myers JD. INTERNIST-I, an experimental computer-based diagnostic consultant for general internal medicine. *N Engl J Med* 1982;307:468-476.

Moses J. A MACSYMA primer. Mathlab Memo no. 2, Computer Science Laboratory, MIT, 1975.

Naghdy F et al. Development of a microprocessor-based monitoring system for post-surgical cardiac patients. *J Microcomput Appl* 1984;7:41-9.

Nelson RM, Jenson CB, Smoot WM. Pericardial tamponade following open-heart surgery. *J Thor Cardiovascular Surgery* 1969;58:510-6.

Nilsson NJ. Principles of artificial intelligence. Palo Alta: Tioga, 1980.

Norusis MJ, Jacquez JA. Diagnosis I. Symptom nonindependence in mathematical models for diagnosis. *Comput and Biomed Res* 1975;8:156-72.

Norwich AM, Turksen IB. A model for the measurement of membership and the consequences of its empirical implementation. *Fuzzy Sets and Systems* 1984;12:1-25.

**Olkin I, Gleser LJ, Derman C.** Probability models and applications. New York: Macmillan, 1980.

**Osborn JJ, Beaumont JO, Raison JC, Russel J, Gerbode P.** Measurement and monitoring of acutely ill patients by digital computer. *Surgery* 1968;64:1057-70.

**Patil RS, Szolovits P, Schwartz WB.** Causal understanding of patient illness in medical diagnosis. *Proc seventh IJCAI* 1981:893-99.

**Pauker SG, Gorry GA, Kassirer JP, Schwartz WB.** Towards the simulation of clinical cognition: taking a present illness by computer. *Am J Med* 1976;60:981-96.

**Pearl J.** Reverend Bayes on inference engines: a distributed hierarchical approach. *Proc AAAI* 1982:133-6.

**Pearl J.** Fusion, propagation, and structuring in belief networks. *Artificial Intelligence* 1987;29:241-88.

**Pednault EPD, Zucker SW, Muvesan LV.** On the independence assumption underlying Bayesian updating. *Artificial Intelligence* 1981;16:213-22.

**Phillips TF, Rodriguez A, Cowley RA.** Right ventricular outflow obstruction secondary to right-sided tamponade following myocardial trauma. *Ann Thorac Surg* 1983;36:353-8.

**Pople HE.** On the mechanization of abductive logic. *Proc third IJCAI* 1973:147-52.

**Pople HE.** The formation of composite hypotheses in diagnostic problem solving: an exercise in synthetic reasoning. *Proc fifth IJCAI* 1977:1030-7.

**Pople HE.** Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In: Szolovits P, ed. *Artificial intelligence in medicine*. Boulder Co: Westview Press, 1982:119-90.

**Pople HE.** Freeing the language of discourse for medical consultation systems. *Proc fifth national AMSI congress* 1986:272-3.

**Pople HE, Myers JD, Miller RA.** Dialog: a model of diagnostic logic for internal medicine. *Proc fourth IJCAI* 1975:848-55.

**Prade H.** A synthetic view of approximate reasoning techniques. *Proc eighth IJCAI* 1983:130-6.

Prade H. A computational approach to approximate and plausible reasoning with applications to expert systems. *IEEE PAMI* 1985;7:260-83 (also corrections pp. 747-8).

Raison JC, Beaumont JO, Russel JA, Osborn JJ, Gerbode F. Alarms in the intensive care unit: an interim compromise. *Comput and Biomed Res* 1968;1:556-64.

Rebuck AS, Pengelly LD. Development of pulsus paradoxus in the presence of airways obstruction. *N Engl J Med* 1973;288:66-9.

Reddy PS, Curtiss EI, O'Toole JD, Shaver JA. Cardiac tamponade: hemodynamic observations in man. *Circulation* 1978;58:265-72.

Rihai M, Tomatis LA, Schlosser RJ et al. Cardiac compression due to closure of the median sternotomy in open-heart surgery. *Chest* 1975;67:113-4.

Robotham JL, Cherry D, Mitzner W, Rabson JL, Lixfeld W, Bromberger-Barnea B. A re-evaluation of the hemodynamic consequences of intermittent positive pressure ventilation. *Crit Care Med* 1983;11:783-93.

Rodewald LE. *BABY: An expert system for patient monitoring in a newborn intensive care unit.* M.S. thesis, Computer Science Department, University of Illinois Champaign-Urbana 1984.

Rogers W, Ryack B, Moeller G. Computer-aided medical diagnosis: literature review. *Int J Biomed Comput* 1979;10:267-89.

Rose DM, Colvin SM, Danilowicz D, Isom OW. Cardiac tamponade secondary to chylopericardium following cardiac surgery: case report and review of the literature. *Ann Thorac Surg* 1982;34:333-6.

Rufilanchas JJ, Iglesias A, Villagra F, Maronas JM, Figuera D. Pericardial effusion causing early cardiac tamponade after open-heart surgery. *Scand J Thorac Cardiovasc Surg* 1977;11:217-9.

Russ TA. A system for using time dependent data in patient management. In: Salamon R, Blum B, Jorgensen M, eds. *Proceedings of the fifth conference on medical informatics.* Amsterdam: North-Holland, 1986:165-9.

Russel R. Hemodynamic monitoring in a coronary intensive care unit. *Mount Kisco: Futura*, 1974:175.

Salel A, Amsterdam EA, Zelis R. Pseudopulsus paradoxus. *Chest*. 1973;64:671-2.

Sanchez E. Medical diagnosis and composite fuzzy relations. In: Gupta MM, Ragade RK, Yager RR, eds. *Advances in fuzzy set theory and applications*. New York: North-Holland, 1979:437-47.

Sapira JD, Kirkpatrick MB. On pulsus paradoxus. *South Med J* 1983;76:1163-4.

Sauers R, Walsh R. On the requirements of future expert systems. *Proceedings of the eighth IJCAI* 1983:110-5.

Scarl EA, Jamieson JR, Delaune CI. A fault detection and isolation method applied to liquid oxygen loading for the space shuttle. *Proc ninth IJCAI* 1985:414-6.

Schefe P. On foundations of reasoning with uncertain facts and vague concepts. *Int J Man-Machine Studies* 1980;12:35-62.

Schoolman HM, Bernstein LM. Computer use in diagnosis, prognosis, and therapy. *Science* 1978;200:926-31.

Schwartz SM, Baron J, Clarke JR. A causal Bayesian model for the diagnosis of appendicitis. *Proc second workshop on uncertainty in artificial intelligence* 1987:229-36.

Selzer A. *Principles and practice of clinical cardiology*. 2nd ed. Philadelphia: Saunders, 1983:41.

Settle HP, Engel PJ, Fowler NO et al. Echocardiographic study of the paradoxical arterial pulse in chronic obstructive lung disease. *Circulation* 1980;62:1297-307.

Shabetai R. *The pericardium*. New York: Grune and Stratton, 1981:279-324.

Shabetai R, Fowler NO, Guntheroth WG. The hemodynamics of cardiac tamponade and constrictive pericarditis. *Am J Cardiol* 1970;26:480-9.

Shafer G. *A mathematical theory of evidence*. Princeton NJ: Princeton University Press, 1976.

Sharp J, Bunnell I, Holland J, Griffith G, Greene D. Hemodynamics during induced cardiac tamponade in man. *Am J Med* 1960;29:640-6.

Sheppard LC. The computer in the care of critically ill patients. *Proc IEEE* 1979;67:1300-6.

Sheppard LC. Computer control of the infusion of vasoactive drugs. *Ann Biomed Engin* 1980;8:431-44.

Sheppard LC, Kouchoukos NT, Kurtts MA, Kirklin JW. Automated treatment of critically ill patients following operation. *Ann Surg* 1968;168:596-604.

Shoemaker WC, Carey JS, Yao ST, Mohr PA, Printen KJ, Kark AE. Hemodynamic monitoring for physiologic evaluation, diagnosis, and therapy of acute hemopericardial tamponade from penetrating wounds. *J Trauma* 1973;13:36-44.

Shortliffe EH. *Computer-based medical consultations: MYCIN*. New York: American Elsevier, 1976.

Shortliffe EH. Details of the Consultation System. In: Buchanan BG, Shortliffe EH, eds. *Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project*. Reading: Addison-Wesley, 1984:78-132.

Shortliffe EH, Axline SG, Buchanan BG, Merigan TC, Cohen SN. An artificial intelligence program to advise physicians regarding antimicrobial therapy. *Computers and Biomedical Research* 1973;6:544-60.

Shortliffe EH, Buchanan BG. A model of inexact reasoning in medicine. *Mathematical Biosciences* 1975;23:351-79.

Shortliffe EH, Buchanan BG, Feigenbaum EA. Knowledge engineering for medical decision making: a review of computer-based clinical decision aids. *Proc IEEE* 1979;67:1207-24.

Shortliffe EH, Davis R, Axline SG, Buchanan BG, Green CC, Cohen SN. Computer-based consultations in clinical therapeutics: explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research* 1975;8:303-20.

Shortliffe EH, Scott AC, Bischoff MB, Campbell AB, Van Melle W, Jacobs CD. ONCOCIN: an expert system for oncology protocol management. *Proc seventh IJCAI* 1981:876-81.

Simpkin P, Hedley-Brown A, Ersoz A, Braimbridge MV. Chronic left heart tamponade. *J Thorac Cardiovasc Surg* 1973;65:531-3.

Sochurek H. Medicine's new vision. *National Geographic* 1987;171:2-41.

Spodick DH. *Chronic and constrictive pericarditis*. New York: Grune and Stratton, 1964:139.

Spodick DH. Acute cardiac tamponade: pathologic physiology, diagnosis and management. *Prog Cardiovasc Dis* 1967;10:64-96.

Spodick DH. Kussmaul's sign [letter to the editor]. *Ann Thorac Surg* 1984;37:519.

Spodick DH. Pulsus paradoxus [Letter]. *South Med J* 1984;77:804.

Stedman's Medical Dictionary. 24th ed. Baltimore, Williams and Wilkins, 1982: 1171.

Stefik M, Aikins J, Balzer R et al. The organization of expert systems, a tutorial. *Artificial Intelligence* 1982;18:135-73.

Stein L, Shubin H, Weil MH. Recognition and management of pericardial tamponade. *JAMA* 1973;225:503-6.

Strat TM. Continuous belief functions for evidential reasoning. *Proc AAAI* 1984:308-13.

Swartout WR. A digitalis therapy advisor with explanations. *Proc fifth IJCAI* 1977:819-25.

Swartz MH, ed. An introduction to physical diagnosis. New York: Raven Press, 1981:138-9.

Swinton NW, Nelson WP, Hall RJ, Harrell JE. Paradoxical pulse a manifestation of the thoracic outlet syndromes. *Arch Intern Med* 1969;124:492-4.

Szolovits P, ed. Artificial intelligence in medicine. Boulder Co: Westview Press, 1982.

Szolovits P, Pauker SG. Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence* 1978;11:115-44.

Thomas CL, ed. Taber's cyclopedic medical dictionary. 14th ed. Philadelphia: F A Davis Co., 1981:P-168.

Thompson TR. Parallel formulation of evidential reasoning theories. *Proc ninth IJCAI* 1985:321-7.

Tong RM, Shapiro DG. Experimental investigations of uncertainty in a rule-based system for information retrieval. *Int J Man-Machine Studies* 1985;22:265-82.

Trigoboff M, Kulikowski CA. IRIS: a system for propagation of inferences in a semantic net. Proc Fifth IJCAI 1977:274-80.

Tversky A, Kahneman D. Judgement under uncertainty: heuristics and biases. Science 1974;185:1124-31.

Urdang, L ed. Mosby's medical and nursing dictionary. St Louis: CV Mosby, 1983:910.

Urschel CW. Pulsus Paradoxus. Effect of gravity and acceleration in its production. Am J Cardiol 1967;19:360-4.

Vaisrub S. Pulsus paradoxus pulmonale. JAMA 1974;228:1030-1.

van Melle W. A domain-independent production-rule system for consultation programs. Proc sixth IJCAI 1979:923-5.

Vierordt K. Die lehre vom arterienpuls in gesunden und kranken zustanden. Braunschweig F Vieweg. 1855; p 266 and figure 2. [cited by Braunwald].

Wagner HR. Paradoxical pulse: 100 years later. Am J Cardiol 1973;32:91-2.

Wardle A, Wardle L. Computer-aided diagnosis—a review of research. Meth Inform Med 1978;17:15-28.

Warner HR. Computer-based patient monitoring. In: Stacy RW, Waxman BD, eds. Computers in biomedical research vol 3. New York: Academic Press, 1969:239-51.

Weeks KR, Chatterjee K, Block S, Matloff JM, Swan HJC. Bedside hemodynamic monitoring: its value in the diagnosis of tamponade complicating cardiac surgery. J Thorac Cardiovasc Surg 1976;71:250-2.

Weil MH, Shubin H, Rand W. Experience with a digital computer for study and improved management of the critically ill. JAMA 1966;198:1011-6.

Weiss SM, Kulikowski CA. EXPERT: a system for developing consultation models. Proc sixth IJCAI 1979:942-7.

Weiss SM, Kulikowski CA, Amarel S, Safir A. A model-based method for computer-aided medical decision-making. Artificial Intelligence 1978;11:145-72.

Weiss SM, Kulikowski CA, Safir A. A model-based consultation system for the long term management of glaucoma. Proc fifth IJCAI 1977:826-32.

Weiss S, Kulikowski CA, Safir A. Glaucoma consultation by computer. *Comput Biol Med* 1978;8:25-40.

Williams CJB. *London J Med.* 1850;2:164. [cited by Knowles].

Winer HE, Kronzon I. Absence of paradoxical pulse in patients with cardiac tamponade and atrial septal defects. *Am J Cardiol* 1979;44:378-80.

Winston PH. *Artificial Intelligence.* 2nd ed. Reading Mass: Addison-Wesley, 1984.

Wyngaarden JB, Smith LH, eds. *Cecil textbook of medicine.* Philadelphia: WB Saunders, 1985:152.

Yacoub MH, Cleland WP, Deal CW. Left atrial tamponade. *Thorax* 1966;21:305-9.

Yager RR. On a general class of fuzzy connectives. *Fuzzy Sets and Systems* 1980;4:235-42.

Yager RR. Approximate reasoning as a basis for rule-based expert systems. *IEEE SMC* 1984;14:636-43.

Yu VL, Fagan LM, Bennett SW, Clancey WJ, Scott AC, Hannigan JF, Buchanan BG, Cohen SN. An evaluation of MYCIN's advice. In: Buchanan BG, Shortliffe EH, eds. *Rule-based expert systems: the MYCIN experiments of the Stanford heuristic programming project.* Reading: Addison-Wesley, 1984:589-96.

Zadeh LA. Fuzzy sets. *Information Control* 1965;8:338-53.

Zadeh LA. A fuzzy-algorithmic approach to the definition of complex or imprecise concepts. *Int J Man-Machine Studies* 1976;8:249-91.

Zadeh LA. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems* 1978;1:3-28.

Zadeh LA. A mathematical theory of evidence [Book Review]. *AI Magazine* 1984;5 #3:81-3.

Zimmermann HJ, Zysno P. Latent connectives in human decision making. *Fuzzy Sets and Systems* 1980;4:37-51.