

MAC TR-143

A Digitalis Therapy Advisor

Howard Silverman

December 1974

This research was supported by the Department of Health, Education, and Welfare (Public Health Service) under Grant number 1 R01 MB 00107-01.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

CAMBRIDGE

MASSACHUSETTS 02139

A DIGITALIS THERAPY ADVISOR

by

Howard Silverman

Submitted to the Department of Electrical Engineering on December 25, 1974 in partial fulfillment of the requirements for the Degree of Master of Science.

ABSTRACT

The physician administering digitalis makes use of the full richness of the clinical setting to form his/her impressions and decide on a therapeutic program. The weakness of existing programs which formulate digitalis dosage regimens lies in their inability to use all of the clinical data available - both quantitative and qualitative. This report describes the construction of a computer system which formulates digitalis dosage regimens and which adjusts this regimen by interpreting the patient's response to the original dosage regimen.

THESIS SUPERVISOR: Professor G.A. Gorry
TITLE: Associate Professor of Computer Science and Engineering

ACKNOWLEDGEMENTS

I would like to thank all of those who cared - for the long hours of discussions, patience and encouragement. A special thanks to Professor Tony Gorry and Dr. Stephen Pauker, who taught me by example as well as by words.

Table of Contents

1. Introduction.....	6
1.1 Digitalis - an Overview.....	6
1.2 Clinical Use of Digitalis.....	7
1.2.1 The Pharmacokinetics of Digitalis.....	9
1.2.2 Regimen Formulation.....	10
1.2.3 Factors Affecting Regimen Formulation.....	10
1.2.4 Digoxin vs Other Digitalis Preparations.....	13
1.3 Review of Previous Work.....	14
1.4 Capabilities of an Improved Digitalis Therapy Advisor.....	16
2. System Demonstration.....	19
2.1 The Initial Session.....	19
2.2 An Update Session.....	31
2.3 Recommendations - Assessment of Patient Response.....	37
3. A Model of Digitalis Administration.....	48
3.1 Digitalization - A Model.....	48
3.2 Construction of the PSM and Generation of the Initial Guess.....	43
3.3 The Feedback Loop.....	46
3.3.1 Characterization of Therapeutic and Toxic Response.....	46
3.3.2 Formulation of Therapeutic Action.....	49
4. The Mechanisms of ANNA.....	52
4.1 An Overview.....	52
4.2 Representation.....	55
4.3 Data Collection.....	55
4.4 Subnode Selection.....	59
4.5 Daemons.....	61
4.6 Bookkeeping.....	63
4.7 Evaluation of Patient Response.....	65
4.7.1 Building the PSM.....	65
4.7.2 The Role of the PSM in Assessing Patient Response.....	68
4.7.3 Formulation of Advice and Recommendations.....	69
4.8 Summary.....	69
5. Refining the Refinements.....	71
5.1 The Interface.....	71
5.1.1 Accepting Information from the User.....	72
5.1.2 Presentation of recommendations.....	73
5.2 Medical Critical Mass.....	74
5.2.1 How Much Data is Enough?.....	74
5.2.2 Power and Scope of the Model.....	75
5.3 Error Recovery.....	76
5.4 Temporal References.....	78
5.5 Explanation.....	79
5.6 Efficiency Considerations and Compilation.....	80
5.7 Implementation Difficulties.....	83
5.8 Conclusion.....	84
Bibliography.....	86
Appendix A.....	88

Appendix B.....91
Appendix C.....94

1. Introduction

"The use of the Foxglove (digitalis) is getting abroad, and it is better the world should derive some instruction, however imperfect, from my experience, than that the lives of men should be hazarded by its unguarded exhibition, or that a medicine of so much efficacy should be condemned and rejected as dangerous and unmanageable." - - - William Withering, 1785

The purpose of this research was to construct a computer program that can advise physicians regarding the administration of digitalis in a qualitative as well as quantitative fashion. These efforts have yielded a computer system, named ANNA, which gives such advice. In addition, many of the considerations involved in the use of digitalis have been elucidated.

I will begin with a brief overview of what digitalis is and how it is used.

1.1 Digitalis - an Overview

"The Foxglove (digitalis) when given in very large and quickly-repeated doses, occasions sickness, vomiting, purging, giddiness, confused vision, objects appearing green or yellow; increased secretion of urine, with frequent motions to part with it, and sometimes inability to retain it; slow pulse, even as slow as 35 in a minute, cold sweats, convulsions, syncope, death." - - - Withering

The term "digitalis" refers to a group of drugs known as cardiac glycosides, among which are digoxin, digitoxin, ouabain, cedalanid and digitalis leaf. The publication of "An Account of the Foxglove" by William Withering in 1785 marked the first effort to understand the effects of digitalis and to establish guidelines for its use. Withering noticed the drug caused increased urine flow and he used it to treat the abnormal accumulation of fluid known as dropsy (commonly due to weakening or failure of the heart).

In fact, the increased urine flow is a side effect of the drug's principle actions: strengthening and stabilizing of the heartbeat.

Because of its positive effects on the heart, the drug is quite useful in the management of congestive heart failure as well as rhythm disturbances and is commonly prescribed by doctors. In fact, it is estimated that one out of every five patients admitted to a hospital receives digitalis sometime during his stay (1). In 1971 it was fifth on the list of drugs most frequently prescribed by physicians through pharmacies in the United States (2).

1.2 Clinical Use of Digitalis

"Let the medicine therefore be given in the doses, and at the intervals mentioned above; let it be continued until it either acts on the kidneys, the stomach, the pulse, or the bowels; let it be stopped upon the first appearance of any one of these effects." - - - Withering

Like many drugs, digitalis can be a poison. When given in proper amounts, however, it can provide the therapeutic effects mentioned above. The physician attempts to give enough of the drug to achieve these therapeutic results but not so much as to cause toxicity. This is often quite difficult for several reasons: 1) a patient can become toxic before an adequate therapeutic effect has been achieved; 2) the difference between therapeutic and toxic levels is small, so a small increase in the amount of digitalis administered may precipitate a toxic reaction; 3) there is a great deal of overlap between therapeutic and toxic manifestations of the therapy and thus it is often difficult to tell whether or not the patient is really toxic; and finally, 4) patients exhibit a variety of individual reactions to the drug. A dosage regimen providing therapeutic results in one patient may lead to toxicity in a second patient.

In determining dosage regimens doctors have traditionally relied upon "intuition", often with poor results. Several studies indicate that as many as 20% of patients receiving the drug demonstrate toxic reactions and that the mortality rate among such toxic patients may be as high as 38% (3). It is this danger of overdose of such a widely used drug that has prompted people to seek better ways to achieve therapeutic results while preventing toxic effects.

There is no single indicator that can be used to judge the degree of toxicity in a patient. Signs of toxicity will often go unnoticed, being incorrectly interpreted as unrelated to the presence of digitalis or, even worse, as being therapeutic effects. The following are generally considered to be indicative of digitalis toxicity:

1. Gastro-intestinal symptoms such as anorexia, nausea or vomiting.
2. The appearance of premature ventricular contractions (PVCs), resulting from increased automaticity (irritability) of myocardial tissues caused by high digitalis levels.
3. Cardiac rhythms such as paroxysmal junctional tachycardia (PAT) with block or non paroxysmal junctional tachycardia following atrial fibrillation.
4. Development of heart block

Because each of the above may have some other cause than digitalis, the physician must exercise a considerable amount of clinical judgment in evaluating the degree of toxicity. For example, many hospitalized patients are very sick and commonly experience nausea and vomiting; patients with congestive heart failure may experience premature ventricular contractions due to stretching of the conduction system of their heart. Caution should therefore be exercised when assessing the meaning of possible signs of toxicity.

Treatment of Digitalis toxicity usually amounts to allowing the patient to lose the digitalis in his/her system through normal excretory pathways. This may not happen quickly enough if the patient is very toxic. In such cases, the toxic episode may be fatal unless other measures are taken (potassium administration, anti-arrhythmic drugs).

1.2.1 The Pharmacokinetics of Digitalis

The advent of radioactive tracing techniques in the early 1960's prompted attempts to better understand the pharmacokinetics of digitalis (2,4,5). A general model, best summarized by Doherty (2), was slowly pieced together. A more detailed mathematical model of digitalis kinetics can be found in Appendix A.

The following model of digoxin kinetics is drawn from articles by Doherty (2) and Jelliffe (6):

"Digoxin is 75% - 85% absorbed when taken orally and is excreted largely unchanged in the urine. Total digoxin losses from the body are proportional to the total amount of digoxin present- the greater the amount of drug that is in the body, the more that is lost or excreted per day. Because of this, single doses of digoxin disappear from the body in a logarithmic fashion. The average measured half-life of digoxin ranges from 1.6 days for patients with normal renal function to about 4.4 days in patients with no renal function. Digoxin is mainly absorbed by the tissues with approximately 7% being recycled in the liver by absorption from the digestive tract followed by excretion in the bile back into the digestive tract where it may be reabsorbed, etc. This recycling is not thought to affect the overall half-life of the drug in patients with normal liver function. About 3% of the drug is excreted in the stool."

These figures represent average values. Patients demonstrate a wide range of individual responses to the drug and care must be exercised in recognizing and dealing with these variations.

1.2.2 Regimen Formulation

The drug may be administered in a variety of ways (oral tablets, oral elixir, intravenous) with oral tablets being the most commonly used. Typically the drug regimen consists of a loading dose given to produce an initial effect followed by regular (smaller) doses of a fixed size which are referred to as maintenance doses. The maintenance dose is taken each day and serves to replace digitalis lost (via the kidneys, the bowel, and through metabolic routes) to keep the total body stores at a constant level. The goal of the physician is to keep this level high enough to provide therapeutic results but not so high as to result in toxicity. This may also be accomplished without the use of a loading dose by keeping the patient on a fixed maintenance dose and allowing enough time for him to reach equilibrium (usually about six days with digoxin - see Appendix A).

In order to formulate a proper maintenance dose the physician may do one of two things. The first (and until recently more common) is simply to guess at a proper dosage based on past experience and then to closely watch the patient's condition. If he/she becomes toxic, then the maintenance dose should be reduced - assuming, of course, that the patient recovers from the toxic episode. The second method is to assign a daily maintenance dose exactly equal to the amount of the drug lost each day, as the definition of maintenance dose would imply. In this manner steady state is achieved.

1.2.3 Factors Affecting Regimen Formulation

"Independent of the degree of the disease, or of the strength or age of the patient, I have had occasion to remark, that there are certain constitutions favourable, and others unfavourable to the success of the Digitalis." - - - Withering

Patients receiving digitalis may have a number of complicating conditions, each of which must be taken into account when formulating a regimen. These conditions can be grouped into three categories: absorption abnormalities, metabolic factors and excretion abnormalities.

Absorption Abnormalities

Abnormal absorption can be difficult to detect before administering digitalis, but little or no response to digitalis therapy may be attributed to reduced absorption of the drug (assuming it is given orally). In such cases, more digitalis should be given or it should be given intravenously (7). Intravenous doses should be less than oral doses, since the IV route circumvents the malabsorption. Caution should be exercised when administering larger oral doses, however, since reversion of the absorption abnormality would expose the patient to unusually high digitalis levels.

Metabolic Factors

There are a number of factors which affect the metabolism and effect of digitalis including hypo- and hyperthyroidism, hypo- and hyperkalemia (potassium imbalance), hypercalcemia (excess calcium), and certain conditions of the heart itself (acute infarction).

HYPO- and HYPERTHYROIDISM

In studies done with hypo- and hyperthyroid patients, it was found that "regardless of the route of administration, hyperthyroid patients exhibited lower serum levels of digoxin and the hypothyroid patients higher levels than the normal group (8)". The conclusion reached is that hypothyroid patients

should receive smaller doses of digoxin while hyperthyroid patients should receive larger doses in order to achieve therapeutic results.

HYPO- and HYPERKALEMIA

Low potassium levels can increase the overall irritability of the heart, making it more sensitive to the toxic effects of digitalis. The physician must be careful to watch potassium levels, as they may fluctuate as a result of various conditions (respiration rate, pH changes, increased fluid volume following surgery, diuretic therapy, etc.). If hypokalemia is present, the administration of digitalis should be avoided until the potassium imbalance is corrected. If this is not possible, the physician should give smaller doses and watch the patient's condition carefully.

HYPERCALCEMIA and MYOCARDOPATHY

Patients who are hypercalcemic or who suffer from a variety of primary diseases of the heart muscle known as myocardopathies tend to be more sensitive to digitalis, and care should be exercised to minimize their digitalis doses and to monitor their condition carefully.

HYPOXEMIA

The physician should exercise caution when administering digitalis to hypoxemic patients. Increased automaticity induced by digitalis causes an increased oxygen demand in myocardial tissues. In the presence of hypoxemia, this demand may not be able to be met.

Excretion Abnormalities

In addition to the metabolic factors mentioned above, deviations in the patient's ability to lose the drug through the various excretory pathways affect the construction of a dosage regimen. Digoxin is eliminated from the

body primarily by urinary excretion and recycling in the liver, with subsequent fecal losses (see Section 1.2.4).

In a normal patient, 35% of the amount of digoxin present is lost in the urine each day. It should be apparent that renal (kidney) insufficiency will affect digoxin excretion and thereby lengthen its half-life from 1.6 days in normal patients to as high as 4.4 days in patients with no renal function whatsoever. Quantitative measures of renal function such as creatinine clearance or blood urea nitrogen (BUN) can be used to compute the amount of digoxin being lost and a proper maintenance dose may be assigned. This is also possible in patients whose renal function is actively changing.

Little is known quantitatively about the effects of liver disease on the liver's ability to recycle digoxin. At present, most physicians disregard the effect of liver or gastro-intestinal dysfunction when computing digoxin losses. It is best to administer normal doses but to watch the patient's condition closely.

1.2.4 Digoxin vs Other Digitalis Preparations

It should be noted that the kinetics of digoxin differ slightly from the other digitalis preparations. Specifically, it has been found that digoxin's half-life is about one quarter that of digitoxin (see Appendix A). Digitoxin is essentially 100% absorbed when taken orally as opposed to 85% absorption of digoxin. The recycling which takes place in the liver is believed to be about fourfold (approximately 26%) that observed with digoxin, with only 16% being excreted in the urine daily. Thus liver disturbances play a greater role in digitoxin therapy and renal insufficiency a lesser role. It is also believed that about 8% of the digitoxin in the body is metabolized into digoxin - - -

an effect which should be taken into consideration.

1.3 Review of Previous Work

Computer programs to advise physicians concerning digitalis dosage regimens have been constructed. Most of the work in this area has been done by Jelliffe and his associates [6,9,10] using "conventional" programming techniques. His efforts focused primarily on programs which formulate an initial guess at a proper drug regimen, but do not have the capability of adjusting the regimen based on the patient's response. Another approach using statistical analysis and feedback (eg, serum digitalis levels) to account for individual reactions to the drug was presented by Sheiner et al [11].

Jelliffe's Work

Taking advantage of the quantitative aspects of what is known about digitalis kinetics, Jelliffe constructed a program which adjusts dosage regimens of digitalis to the patient's weight, renal function, route of administration and present computed (or measured) concentrations of digitalis. The program is intended for use "in patients with normal thyroid and hepatic (liver) function and normal electrolyte balance (potassium, sodium, etc.) who are not receiving drugs that alter the absorption or metabolism of digitalis glycosides and who have no gross clinical evidence of gastrointestinal malabsorption [10]."

The main strongpoint of Jelliffe's approach is that it works - his programs can compute initial digitalis dosage regimens. He asserts that the "use of this program for the past two years has reduced adverse reactions to glycoside therapy from 31% to 12% [9]."

In spite of the improvements offered by Jelliffe's approach, it

nevertheless suffers from two major weaknesses. First, it doesn't take into account all of the factors affecting digitalis pharmacokinetics. For this reason it is ineffective when such conditions are present. Second, it provides only an initial approximation to a proper dosage regimen, leaving it to the doctor to monitor the patient's response and to adjust the dosage regimen accordingly.

Sheiner's Work

Sheiner employs statistical methodology to "provide the basis for a clinically useful computer program to suggest optimal dosage regimens for a number of drugs for individual patients [11]." After the patient is put on a dosage regimen, the blood level of digitalis is determined. This level is then used to improve the "pharmacokinetic parameters" for the patient and a new dosage regimen is computed. The feedback loop is then entered again until the patient's condition stabilizes.

This approach is attractive because it provides a framework in which feedback information can be used effectively. Each patient is accurately modeled by his individual pharmacokinetic parameters. Changes in the patient's clinical condition can be represented by changing these parameters.

Sheiner's program performs better than the one proposed by Jelliffe, but it fails in two respects. First, its goal is represented as a desired blood level of the drug. In practice it may be difficult to specify what the proper blood level for a particular patient is, especially if he/she is sensitive to digitalis for some reason (potassium depletion, hypothyroid, etc.). One is really interested in the overall effect of the dose administered more than the change in absolute blood level. In this respect, Sheiner's approach rests upon the weak assumption that a given blood level will produce some known effect.

Since the inner workings of Sheiner's program involve a great deal of complex statistical "machinery", it may be quite difficult for the user to understand how the program reaches conclusions. This may lead to skepticism on the part of the user concerning the program's answers and a corresponding decrease in its clinical effectiveness.

1.4 Capabilities of an Improved Digitalis Advisor

In a study done by Carl Peck et al [3] comparing computer-assisted therapy to that of unaided physician judgment, the computer-aided group only slightly outperformed the unaided physicians. Despite the advances of the Jelliffe and Sheiner programs it is becoming increasingly clear that they are lacking in some respects.

The physician administering digitalis makes use of the full richness of the clinical setting to form his impressions and decide on a therapeutic program. The weakness of existing programs lies in their inability to use all of the clinical data available - both quantitative and qualitative. The goal of this research was to construct a computer program which could begin to cope with the full complexity of a clinical setting, formulating its recommendations in the same way a cardiologist would.

The first step in realizing this goal was the formulation of a more complete model of digitalis administration than that used by Jelliffe and Sheiner. Although a great deal is known about the pharmacokinetics of digitalis, little work had been done identifying what components of the clinical setting are the important ones and how they are used by physicians in the formulation of digitalis dosage regimens. For example, most physicians realize that low potassium increases digitalis sensitivity but generally find

it difficult to specify precisely when and how this piece of information is used.

The initial research activity was the formulation of a better model of digitalis administration, specifying what information is necessary and how it is used (see Chapter 3). Work was then begun on a program which would be able to make use of this model. A list of necessary constituents of an effective Digitalis Advisor was formulated:

1. Computation facilities to deal with that information which is adequately described in quantitative terms (renal function, daily losses, etc.).
2. "Model-tailoring" facilities. By asking various questions about a patient, the system should be able to tailor make a patient-specific model and use this model to formulate recommendations for the patient. The system must know what questions are relevant. It must integrate incoming information into patient-specific model, realizing the worth and significance of this new information. In addition, it must have be able to change the patient-specific model when necessary and know when this model is no longer accurate.
3. Explanation capabilities. In order to test the appropriateness of the conclusions reached by the system, particularly when dealing in an area such as digitalis administration, it is important to be able to look at the reasoning behind decisions.
4. Extensibility. By using this program, inaccurate and inadequate portions of the model will be identified and corrected. If the system is to be able to incorporate this updated model, it is essential that the initial design configuration be one which allows the system to be extended and changed in an orderly way. This applies to data base maintenance as well as future programming tasks.

A computer system with rudimentary capabilities in each of the above areas that produces recommendations for digitalis dosage regimens was constructed, using an improved model of digitalis administration. In the next chapter, discussions of sample sessions with the current version of the system are presented. Chapter 3 contains a detailed description of the model of digitalis administration used by the system to produce the behavior demonstrated in Chapter 2. In Chapter 4 the structure of the computer system

that uses this model will be discussed. Chapter 5 is devoted to a discussion of further refinements which will enhance the worth of this computer system.

Readers interested in the technical details of how the system operates should consult Appendix C. For those readers more interested in the medical arguments, Chapters 1 through 3 and a brief survey of Chapter 4 and Chapter 5 are recommended.

2. System Demonstration

In the preceding chapter, a rough idea of the capabilities necessary in a Digitalis Advisor system was formulated. In this chapter, some sample sessions with the system are presented along with some commentary as to why the system is acting in this manner.

2.1 The Initial Session

Figure 2.1 contains a listing of a sample session with ANNA (the user's responses are underlined). In the following paragraphs, the reasons for asking each question will be considered and the use made of the user's responses will be discussed. The mechanisms underlying the generation of these questions will be treated in Chapter 4.

Questions 1 through 3

In Questions 1 through 3, the patient's name is requested as well as his/her identification number, which is assumed to be unique. Patient records are stored on disk and are indexed by the patient identification number (this could be a social security number or a hospital identification number). After receiving this information, the system checks its files for a patient record with the specified identification number. In this case, no such patient is found and the system correctly assumes that a new case is to be presented.

Questions 4 through 7

Figure 2.1
An Initial Session

- 11) ...MMM, WHAT IS THE PATIENT'S LAST NAME?
(ENTER VALUE): MONET
- 12) ...PLEASE TELL ME, WHAT IS THE PATIENT'S FIRST NAME?
(ENTER VALUE): CLAUDE
- 13) ...LET ME ASK YOU, PLEASE ENTER PATIENT IDENTIFICATION NUMBER
(ENTER VALUE): 386589127
...CONSULTING RECORDS - PLEASE STAND BY...
...WELL...NO RECORD FOR CLAUDE MONET
CAN BE FOUND, SO IT IS ASSUMED THAT THIS IS A NEW PATIENT
- 14) ...TELL ME, HOW OLD IS THE PATIENT?
(ENTER VALUE): 38
- 15) ...PLEASE TELL ME, WHAT IS THE PATIENT'S WEIGHT (IN POUNDS)?
(ENTER VALUE): 166
- 16) ...LET'S SEE, WHAT IS THE PATIENT'S SEX?
1 FEMALE
2 MALE
===> 2
- 17) ...TELL ME, IS THE PATIENT CURRENTLY RECEIVING DIGITALIS?
1 YES
2 NO
===> 2
- 18) ...MMM, WHAT IS THE UNDERLYING RHYTHM?
1 SINUS RHYTHM
2 PAT
3 ATRIAL FIBRILLATION
4 ATRIAL FLUTTER
5 NODAL RHYTHM
6 VENTRICULAR RHYTHM
===> 3
- 19) ...LET ME ASK YOU, WHAT IS THE CURRENT VENTRICULAR RATE?
(ENTER VALUE): 152

110) ...LET'S SEE, IS THE PATIENT HAVING PVC'S?

- 1 YES
- 2 NO

----> 2

111) ...MMM, FOR WHAT ARE YOU DIGITALIZING THIS PATIENT?

- 1 CHF
- 2 ARRHYTHMIA
- 3 PROPHYLACTIC

----> 1 2

112) ...LET ME ASK YOU, DOES THE PATIENT HAVE ORTHOPNEA?

- 1 YES
- 2 NO

----> 2

113) ...LET ME ASK YOU, DOES THE PATIENT HAVE DYSPNEA ON EXERTION?

- 1 YES
- 2 NO

----> 1

114) ...LET ME ASK YOU, ON CLIMBING STAIRS, AFTER HOW MANY STAIRS DOES DYSPNEA OCCUR?

- 1 3 STEPS
- 2 6 STEPS
- 3 1 FLIGHT
- 4 2 FLIGHTS
- 5 3 FLIGHTS
- 6 4 OR MORE FLIGHTS
- 7 UNKNOWN

----> 3

115) ...LET'S SEE, DOES THE PATIENT HAVE PAROXYSMAL NOCTURNAL DYSPNEA (PND)?

- 1 YES
- 2 NO

----> 2

116) ...PLEASE TELL ME, DOES THE PATIENT HAVE ANKLE EDEMA?

- 1 YES
- 2 NO

----> 1

117) ...LET ME ASK YOU, HOW BAD IS THE ANKLE EDEMA?

- 1 0
- 2 1+
- 3 2+
- 4 3+
- 5 4+
- 6 UNKNOWN

----> 3

1181 ...LET'S SEE, WHAT IS YOUR ESTIMATE OF THE JUGULAR VENOUS PRESSURE MEASURED IN CM ABOVE THE RIGHT ATRIUM (HEIGHT ABOVE STERNAL ANGLE OF LOUIS PLUS 6 CM)?

- 1 NORMAL (3 TO 10 CM)
- 2 MODERATE ELEVATION (10 TO 15 CM)
- 3 MARKED ELEVATION (15 TO 20 CM)
- 4 VERY HIGH (>20 CM)
- 5 UNKNOWN

====> 3

1191 ...HMMN, DOES THE PATIENT HAVE AN S3-GALLOP?

- 1 YES
- 2 NO
- 3 UNKNOWN

====> 1

1201 ...LET ME ASK YOU, WHAT IS YOUR ESTIMATE OF THE LOUDNESS OF THE S3 GALLOP?

- 1 FAINT
- 2 SOFT
- 3 LOUD
- 4 UNKNOWN

====> 3

1211 ...LET ME ASK YOU, DOES CHEST X-RAY INDICATE CARDIOEGCALY ?

- 1 YES
- 2 NO
- 3 UNKNOWN

====> 2

1221 ...HMMN, WOULD YOU PLEASE ENTER ANY SIGNS OF PULMONARY VENOUS CONGESTION?

- 1 REDISTRIBUTION
- 2 KERLEY B LINES
- 3 INTERSTITIAL EDEMA
- 4 ALVEOLAR EDEMA
- 5 FRANK PULMONARY EDEMA

====> -

1231 ...PLEASE TELL ME, WHAT IS THE PATIENT'S SERUM K?

(ENTER VALUE): 4.0

1241 ...HMMN, IS THE PATIENT RECEIVING ANY DIURETICS?

- 1 YES
- 2 NO

====> 2

125) ...HMMM, DO YOU SUSPECT THE PATIENT IS HYPOXEMIC?

- 1 YES
- 2 NO

====> 1

126) ...LET ME ASK YOU, WHAT IS THE PATIENT'S PO2?

(ENTER VALUE): 66

127) ...LET ME ASK YOU, DOES THE PATIENT HAVE A CARDIOMYOPATHY OR AN ACUTE MI?

- 1 YES
- 2 NO

====> 2

128) ...HMMM, IS THE PATIENT CLINICALLY MYXEDEMIC?

- 1 YES
- 2 NO
- 3 UNKNOWN

====> 2

129) ...TELL ME, HOW FAST DO YOU WISH TO DIGITALIZE THE PATIENT?

- 1 INSTANTANEOUS RATE
- 2 RAPID RATE
- 3 MODERATE RATE
- 4 SLOW RATE
- 5 NO PREFERENCE

====> 2

130) ...LET ME ASK YOU, IS THE PATIENT'S URINE OUTPUT GREATER THAN 400CC PER DAY?

- 1 YES
- 2 NO
- 3 UNKNOWN

====> 1

131) ...TELL ME, WHAT IS THE CONDITION OF RENAL FUNCTION?

- 1 STABLE
- 2 UNSTABLE

====> 1

132) ...HMMM, WHAT IS THE PATIENT'S CREATININE CLEARANCE? (IN ML/MIN)

(ENTER VALUE): UNKNOWN

133) ...LET'S SEE, WHAT IS THE PATIENT'S SERUM CREATININE? (IN MG %)

(ENTER VALUE): 32

...32 IS NOT A LEGAL VALUE FOR SERUM-CREATININE

...PLEASE ENTER ANOTHER

(ENTER VALUE): 3.2

134) ...HMM, WHAT ROUTE WOULD YOU LIKE TO USE?

- 1 ORAL
- 2 IV
- 3 NO PREFERENCE

==> 1

...SINCE THE PATIENT'S RATE IS GREATER THAN 140, PLEASE CONSIDER WHETHER THE IV ROUTE WOULD BE PREFERABLE.

134a) ...PLEASE TELL ME, WHAT ROUTE WOULD YOU LIKE TO USE?

- 1 ORAL
- 2 IV
- 3 NO PREFERENCE

==> 1

135) ...DO YOU WISH TO SEE THE THERAPY RECOMMENDATIONS NOW?

- 1 YES
- 2 NO

==> 1

... PLEASE TRY TO CORRECT THE PATIENT'S HYPOXEMIA AS QUICKLY AS POSSIBLE.

PRESCRIPTION FOR: CLAUDE MONET (74 11 8)

WEIGHT: 166 LBS.

AGE: 30

I. LOADING PROGRAM:

...IN ORDER TO ACHIEVE A DESIRABLE BODY STORES OF

0.448 MG FOR THIS PATIENT,

...IT IS ADVISABLE TO ADMINISTER THE FOLLOWING:

DOSE 1: 0.25 MG REPORT PATIENT RESPONSE BEFORE
ADMINISTERING NEXT DOSE (IN 2 TO 4 HOURS
OR IF CHANGE OCCURS).

DOSE 2: 0.25 MG

II. MAINTENANCE PROGRAM:

DAILY MAINTENANCE DOSE = 0.099 MG PER DAY

RENAL FUNCTION = 30%

136} ...DO YOU WANT TO SAVE THIS DATA?

1 YES

2 NO

====> 1

These questions continue to gather background information about the patient. Weight is an important consideration when deciding on an initial digitalis dosage regimen. Many patients are already taking digitalis or have taken it in the past. In such situations, it is necessary to formulate a projection of how much digitalis the patient has "on board" based on the previous dosage regimen followed and on the patient's renal function.

Questions 8 and 9

Questions 8 and 9 establish the current type and rate of the cardiac rhythm; information which is essential in subsequent evaluations of the extent of toxic reaction and of the degree of therapeutic response. This information is also used when making decisions concerning the rate of digitalization and the route of administration. For example, the patient considered here is said to be in atrial fibrillation with a ventricular response of 152. It is therefore best to digitalize him/her quickly in order to get the rate down to a more reasonable level. Furthermore, the system would suggest digitalization to be intravenous (IV) with digoxin as the preparation of choice (see Questions 35 and 35a below). The type of rhythm also serves in part as an indicator of what signs of toxicity should be expected to develop.

Question 10

The appearance of premature ventricular contractions (PVCs) in a patient receiving digitalis is often an early indicator of toxicity. However, this is not necessarily the case if PVCs were present before the patient was digitalized. Question 10 is asked in order to reduce any doubt later as to

whether the PVCs are being caused by too much digitalis or if they are related to factors present before digitalis therapy was instituted.

Question 11

The system expects the physician to have a specific reason for giving digitalis to a patient. The only legitimate reasons considered are congestive heart failure, arrhythmia, prophylactic use or some combination of the three.

The reason for digitalization is a strong determinate of what type of therapeutic response should be sought. For example, since this patient is being digitalized for both congestive heart failure and for an arrhythmia (atrial fibrillation), the system will consider a reduction in the ventricular rate with a corresponding decrease in the symptoms and signs of congestive heart failure to be the primary therapeutic goal.

Questions 12 through 22

One of the reasons for digitalizing this patient is congestive heart failure, so the system asks questions about the manifestations of the failure. This information will be used later to determine whether or not the patient is improving as a result of the therapy. Each question concerns a symptom or sign of congestive heart failure, asking for a severity estimate when appropriate.

An alternative method for the characterization of failure would be to note which symptoms and signs were present and later ask if they had improved or not. It was felt, however, that some objective assessment of the severity of relevant symptoms and signs at the time of their appearance would be better

than asking the physician to "think back" and to decide if the manifestations of failure had improved.

Question 23 and 24

In formulating an estimate of how much digitalis to give a patient, it is essential to consider possible increased sensitivity caused by hypokalemia (low potassium). Moreover, it is often the case that digitalis is given to patients also receiving diuretics. This is particularly true with elderly patients in congestive heart failure. Chronic diuretic therapy can result in a potassium deficiency and recently initiated diuretic therapy or acute administration of diuretics can lead to unexpected (and dangerous) potassium shifts.

Question 23 determines the patient's serum potassium level which is used as an indicator of the patient's total potassium. Once this value has been obtained, the system asks about concurrent diuretic therapy (Question 24).

Questions 25 through 28

Hypoxemia, myxedema (hypothyroidism), cardiomyopathies, myocardial infarction and/or myocarditis may lead to increased digitalis sensitivity. A reduction in the digitalis dosage estimate is made for each of the above conditions suspected of being present. This reduction reflects the "better safe than sorry" maxim which has been built into the system. Whereas underdigitalization of the patient can be simply corrected by administering more digitalis, overdigitalization unnecessarily exposes the patient to the dangers of a toxic episode.

Since it was indicated that the patient might be hypoxemic (Question 25), ANNA requests the patient's arterial oxygen concentration (pO₂) in order to better estimate how much to adjust the dosage regimen in the face of possible hypoxemia.

Question 29

During the initial estimation phase, the system needs to determine how fast to administer the drug as well as how much to give. In general it is best to digitalize the patient as slowly as possible, allowing more time to detect and correct toxic responses. This goal must, however, be weighed against the urgency of therapeutic intervention.

In this instance, it is advisable to digitalize the patient relatively quickly in order to get his/her heart rate down. The user may indicate his choice for the rate of digitalization, but the system will object if it is felt to be too fast or too slow for the particular situation at hand. For example, the system always raises an objection when "instantaneous rate" is specified, since this is considered to be allowable only in those cases where the need for digitalis is most pressing, such as emergency cases in which extremely rapid rate control is desired.

Questions 30 through 33

Although it is not necessary to know a patient's excretion losses in order to choose a loading dose (assuming digitalis is to be administered rapidly), this information is essential in computing an appropriate maintenance schedule. This set of questions is designed to determine the best

available measure of renal function.

Once it has been determined that the patient has some (non-zero) renal function (Question 30) and that it is stable (Question 31), the system asks for a creatinine clearance (Question 32), considering this the most accurate indicator of renal function. This value is, however, not available, so the system requests the next best measure, serum creatinine (Question 33). Error checking is performed to make certain the values being entered are reasonable. An objection is voiced to the impossibly high value of 32 for serum creatinine and the user is asked to enter another value.

The system also uses the answers to this series of questions when determining how to compute renal function (see Section 4.4). For example, if renal function was said to be unstable, the system would select a least squares projection as the best technique for computing renal function. In this instance, it selects renal function equations which use of the available serum creatinine value.

Questions 34 and 34a

The system normally leaves the choice of route of administration to the user, but objects in this case because of the high rate. Indication that the drug is to be administered orally (Question 34a) causes the system's objection to be overridden.

Question 35

ANNA asks if the user wishes to see its therapy recommendations, having all the information it needs to formulate an initial therapy program for this

patient. First, it is suggested that the hypoxemia (as indicated by the low PO₂ value) be corrected, if possible. Following a brief summary of the state of the patient, the system indicates a total body store projection of 0.448 mg of digoxin. This figure represents an initial estimate of 0.625 mg, adjusted for sensitivities and body weight. A maintenance dose of 0.10 mg per day is indicated, given the current renal function of about 30%. Since this patient was being digitalized for an arrhythmia, it requests that the first loading dose indicated be given and any changes be reported before giving more. This step-by-step digitalization promotes a careful watch of the patient's early responses to digitalization and will help avoid "overshooting" the proper digitalis level. The feedback process starts with the first dose given, not after the patient is fully loaded.

Question 36

The final question asks if this patient record is to be filed away for future reference. As an update is to be performed on this patient in the next section, "yes" is entered.

2.2 An Update Session

In the previous section, a hypothetical patient was presented to the system and some initial therapy suggestions were generated. The following is a discussion of the user-system dialogue some four hours later (see Figure 2.2), assuming the program's advice had not been carried out and the patient was given 0.5 mg digoxin by mistake.

Figure 2.2
An Update Session

11) ...HMM, WHAT IS THE PATIENT'S LAST NAME?

(ENTER VALUE): MONET

12) ...PLEASE TELL ME, WHAT IS THE PATIENT'S FIRST NAME?

(ENTER VALUE): CLAUDE

13) ...LET ME ASK YOU, PLEASE ENTER PATIENT IDENTIFICATION NUMBER

(ENTER VALUE): 306569127

...CONSULTING RECORDS - PLEASE STAND BY...

...AH YES OF COURSE! ACCORDING TO THE RECORDS, THIS PATIENT HAS
ALREADY BEEN PRESENTED. PLEASE INDICATE WHETHER THE FOLLOWING
DESCRIBES THE PATIENT YOU HAVE IN MIND:

...PATIENT'S NAME IS CLAUDE MONET AND WAS PRESENTED ON TUESDAY
SEPTEMBER 24, 1974 TO BE DIGITALIZED FOR CHF AND ARRHYTHIA

...SEX: MALE

...AGE: 30

...WEIGHT: 165

14) ...IS THIS THE ONE?

- 1 YES
- 2 NO

==> 1

...LOADING PATIENT RECORD

15) ...HMM, WAS THE RECOMMENDED DOSE ADMINISTERED?

- 1 YES
- 2 NO MORE WAS GIVEN
- 3 NO LESS WAS GIVEN

==> 2

16) ...PLEASE TELL ME, WHAT IS THE PATIENT'S WEIGHT (IN POUNDS)?

(ENTER VALUE): 166

17) ...LET'S SEE, WHAT IS THE PATIENT'S SERUM K?

(ENTER VALUE): 4.1

181 ...TELL ME, DO YOU SUSPECT THE PATIENT IS HYPOXENIC?

- 1 YES
- 2 NO

====> 1

191 ...HMMM, WHAT IS THE PATIENT'S PO2?

(ENTER): 78

1101 ...LET ME ASK YOU, HAS THE PATIENT'S RHYTHM BECOME REGULAR?

- 1 YES
- 2 NO

====> 2

1111 ...LET'S SEE, WHAT IS THE CURRENT VENTRICULAR RATE?

(ENTER VALUE): 92

...THERE SEEMS TO BE A SIGNIFICANT DECREASE IN THE VALUE OF VENTRICULAR RATE.

1121 ...SHOULD THIS BE ASSUMED SIGNIFICANT?

- 1 YES
- 2 NO

====> 1

1131 ...HMMM, WOULD YOU PLEASE ENTER ANY OF THE FOLLOWING WHICH HAVE APPEARED?

- 1 NAUSEA
- 2 VOMITING
- 3 ANOREXIA
- 4 VISUAL DISTURBANCES (YELLOW OR GREEN TINGED VISION)

====> -

1141 ...LET ME ASK YOU, IS THE PATIENT HAVING PVC'S?

- 1 YES
- 2 NO

====> 1

1151 ...ARE THESE REALLY PVC'S OR ARE THEY ABBERANTLY CONDUCTED SUPRAVENTRICULAR BEATS (ASHMAN BEATS)?

- 1 ASHMAN BEATS
- 2 REALLY PVC'S

====> 2

1161 ...LET ME ASK YOU, HOW MANY PVC'S ARE OCCURRING?

- 1 0-5 PER MINUTE
- 2 5-15 PER MINUTE
- 3 MORE THAN 15 PER MINUTE

====> 2

117: ...LET ME ASK YOU, WHAT TYPE OF PVCS ARE THEY?

- 1 UNIFOCAL
- 2 MULTIFOCAL
- 3 UNKNOWN

===> 2

118: ...LET'S SEE, ARE THERE SALVOS?

- 1 YES
- 2 NO
- 3 UNKNOWN

===> 1

Questions 1 through 3

As in the initial session, these questions establish the identity of the patient via name and hospital number. Upon receipt of this information, the system proceeds to examine its records for knowledge of this patient. This time a patient record is found (the one which was stored at the conclusion of the last session).

Question 4

A brief description of the patient is displayed and the user is asked for verification that this is the patient to be considered. Since the description fits the patient, "yes" is entered. The system then loads the patient's record from its files, notifying the user it is ready to proceed.

Question 5

In order to interpret the patient's response to the initial regimen, it is necessary to know if the suggested regimen had been followed. In this case it is indicated that more digitalis was given than the program had suggested at the conclusion of the initial session. The administration of more digitalis than recommended may be interpreted later as a possible cause for a toxic reaction (see Section 2.3).

Questions 6 through 9

These questions explore the possibility of shifts in weight, serum potassium (hypokalemia) and pO₂ (hypoxemia), as discussed in previous sections. Here it is indicated that the patient is suspected as being hypoxemic, with a pO₂ value of 78.

Questions 10 through 12

It was originally stated this patient was to be digitalized both for congestive heart failure and for an arrhythmia (atrial fibrillation). The system regards controlling the patient's arrhythmia as its first priority, with management of the failure as its ultimate goal. Accordingly, it asks about the development of a regular rhythm (Question 10) and asks about the patient's current ventricular rate (Question 11). Although the rhythm is not yet regular, there is a reduction in rate from 152 (see Question 9 of Figure 2.1) to 92. The system notices this change and asks the user about its significance as related to the digitalis therapy (Question 12). This is important because the reduction in rate could have been caused by some condition other than digitalis response, such as abatement of a fever. We indicate that the decrease in this case is to be considered significant.

Questions 13 through 18

In these questions the system is looking for general signs of toxicity, such as nausea, vomiting, anorexia, visual disturbances or the development of PVCs. The patient is not experiencing nausea, vomiting or visual disturbances, but PVCs are indeed beginning to appear. Since this patient was in atrial fibrillation, it is possible that the PVCs are actually aberrantly conducted

supraventricular beats (Ashman Beats). In order to avoid this error, the system asks for verification that the aberrant beats are PVCs (Question 15). Having received this verification, the system attempts to further characterize these PVCs with queries concerning the amount (Question 16), type (Question 17) and whether or not the PVCs are appearing in salvos (Question 18).

2.3 Recommendations - Assessment of Patient Condition

The system now has the information necessary to evaluate the current status of the patient and to formulate the next step in the therapy program. A summarization of its conclusions regarding changes in the patient followed by specific recommendations appears in Figure 2.3. It regards the decrease in ventricular rate as a sign of increased therapeutic effect and considers the patient to be properly digitalized, since the rate has fallen below 100 (therapeutic endpoint reached). This level cannot, however, be tolerated due to the presence of a toxic reaction (toxic endpoint reached), which may have been caused by the increased amount of digitalis administered (see Question 5). Specifically, it suggests that steps be taken to correct the hypoxemia, since this may have played a large part in precipitating the toxic episode. In addition, digitalis administration should be stopped until the signs of toxicity subside. This is essential, since the patient is already toxic and more digitalis would expose him/her unnecessarily to further dangers of toxicity.

The system now has an "upper bound" on how much digitalis this patient will tolerate (this may however be influenced by the existence of hypoxemia or other factors). After the toxic reaction subsides, the system will work the patient up slowly to just below this level. The final request of the system

Figure 2.3

Presentation of Recommendations

...SUMMARY OF CLAUDE MONET :

...INDICATIONS OF AN ENDPOINT OF THERAPEUTIC-EFFECT ARE:

... VALUE VENTRICULAR-RATE LESS THAN 100

...INDICATIONS OF AN INCREASE OF THERAPEUTIC-EFFECT ARE:

... VENTRICULAR-RATE-DECREASING

...INDICATIONS OF AN ENDPOINT OF TOXICITY-PROBABILITY ARE:

... VALUE TYPE PVCs MULTIFOCAL

...INDICATIONS OF AN INCREASE OF TOXICITY-PROBABILITY ARE:

... CHANGE TYPE PVCs WORSER

... CHANGE BALANCE PVCs INCREASE

... CHANGE AMOUNT PVCs INCREASE

...POSSIBLE CAUSES FOR AN INCREASE OF TOXICITY-PROBABILITY ARE

... VALUE VALUE P02 LOW

... TOO MUCH DIG ADMINISTERED

...AS INDICATED EARLIER, THIS PATIENT IS SHOWING DEFINITE SIGNS OF TOXICITY. THE BEST ACTION TO TAKE IS TO:

1. TAKE STEPS TO CORRECT THE FOLLOWING CONDITIONS WHICH MAY CORRECT THE TOXIC TREND:
LOW P02
2. STOP DIGITALIS UNTIL SYMPTOMS OF TOXICITY SUBSIDE
3. OBTAIN A SERUM DIGOXIN LEVEL ON THIS PATIENT
4. REPORT ANY FURTHER CHANGES IN THE PATIENT'S CONDITION

is for the user to obtain a serum digoxin level on the patient, if possible. This will be used in the next update, when the system will once again evaluate the patient's response.

3. A Model of Digitalis Administration

Withering's original advice concerning the use of digitalis was to give the drug until the desired effect was noticed (usually diuresis) or until the patient got sick. At this point, either the patient would die or he/she would get better. Until recently, physicians still used Withering's coarse algorithm when digitalizing their patients. This has, however, changed in recent years. EKG's now aid in the early recognition of toxicity.

Radioactive tracing techniques have permitted a more exhaustive examination of the pharmacokinetics of digitalis and have led to the development of a mathematical model of digitalis kinetics. Some workers have constructed computer programs that use this mathematical model to formulate digitalis dosage regimens (see Section 1.4). The model used is simple and is used in one step: compute the appropriate dosage and administer it. In this section, a closer look will be taken at the nature of the problem and a more robust method for the formulation of digitalis dosage regimens will be presented.

3.1 Digitalization - a Model

"The central task of a natural science is to make the wonderful commonplace; to show that complexity, correctly viewed, is only a mask for simplicity; to find pattern hidden in apparent chaos."

The Sciences of the Artificial by Herbert A. Simon

Jelliffe's program made use of a quantitative model that is based on the mathematical relationships between maintenance dose, renal function, weight, etc. (see Appendix A). This model is implemented in the form of a procedure that accepts various parameters (renal function, weight, etc) and yields the appropriate maintenance dose (see Figure 3.1). Different patients are modeled by changing these input parameters, but the (internal) procedural model

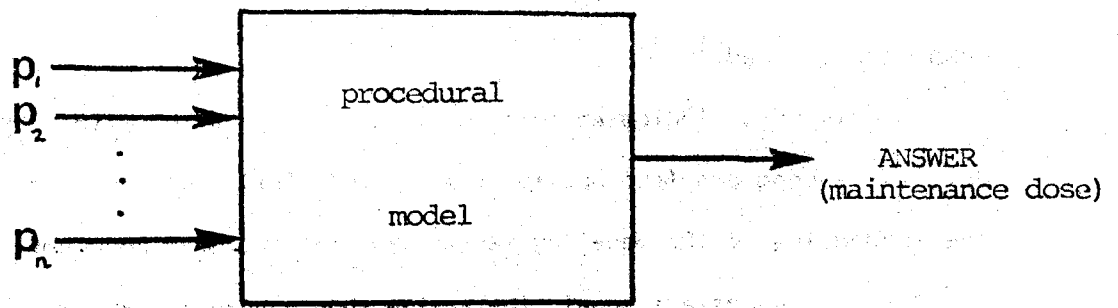


Figure 3.1
Procedural Model of Jelliffe

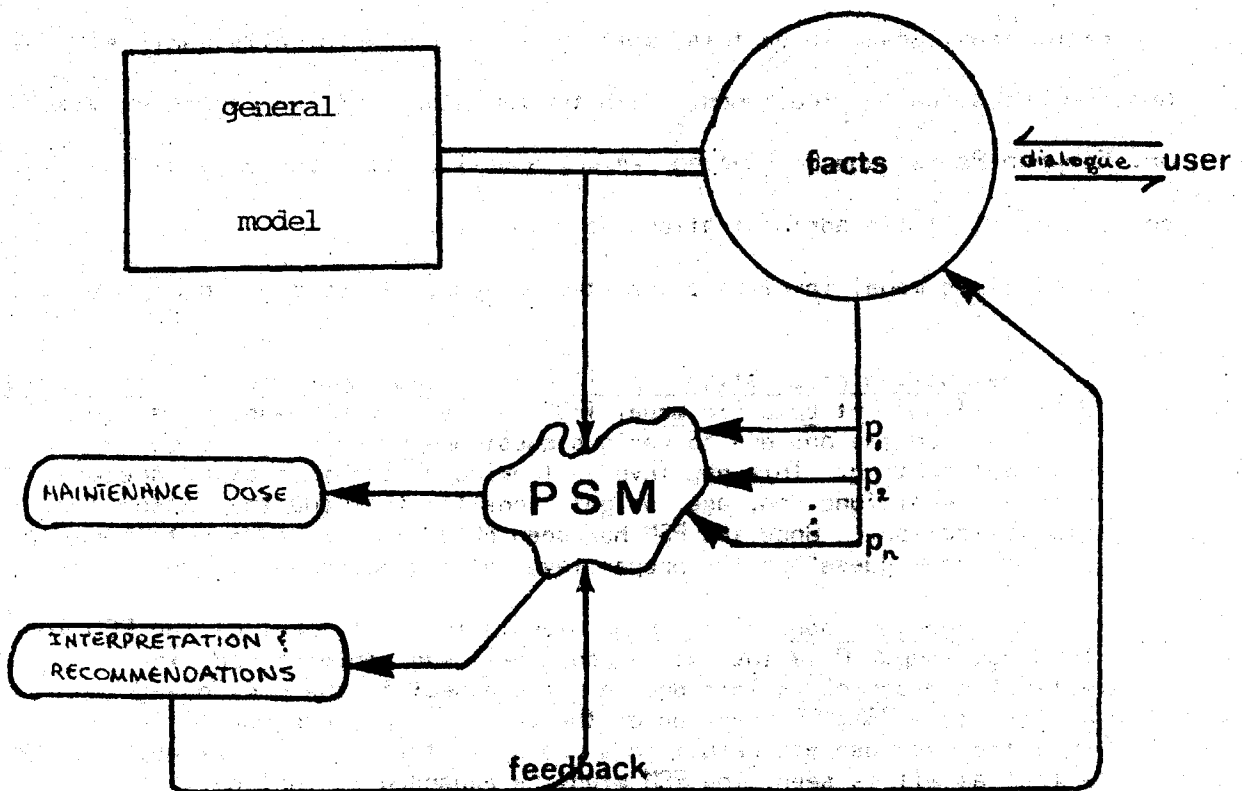


Figure 3.2 Model Tailoring

remains unaltered.

Studies have indicated that there is a limit to the effectiveness of Jelliffe's programs (see Section 1.4). It was felt that this limit arose from the inadequacy of the modeling facilities used by these programs. Instead of using a single patient model, the system should have the ability to "tailor-make" a model for each patient it considers, basing its recommendations on this patient specific model.

Construction of a computer program with such improved modeling capabilities requires 1) a clear understanding of what the necessary constituents of a patient specific model are; 2) how a patient specific model can be produced and 3) how it can be used to produce therapeutic suggestions and recommendations. To do this, what is known about digitalis administration must be methodically structured. With the assistance of two cardiologists, Dr. Stephen Pauker and Dr. Harvey Zarren, a well structured model of the process of digitalis administration was developed.

Use of this model involves a two step process, as depicted in Figure 3.2:

1. Construction of a Patient Specific Model and Generation of the Initial Guess. A patient specific model (PSM) is formulated based on a general model and on the answers to various questions concerning the current clinical setting. This may involve looking at a number of parameters, such as renal function, age, weight, sensitivities and the reason for digitalization. Once the PSM has been obtained, it is used to formulate an "educated guess" at the proper level of digitalis for the patient.

2. The feedback loop. The second step is the refinement of the PSM in a feedback loop. Once the initial dose has been administered, the patient's response is interpreted with respect to the previously constructed PSM. Comparison of the demonstrated response to the expected response may result in a change in the PSM. For example, if no effect at all is seen, the PSM might be expanded to include the possibility of malabsorption. Once the PSM has been updated, it is used to formulate the next step in therapy.

In the following sections PSM construction and use of the PSM in the feedback loop will be discussed.

3.2 Constructing the PSM and Generation of the Initial Guess

The PSM contains two types of information. First, the values of clinical variables such as weight, creatinine clearance, age, sex, etc. for the individual patient under consideration are recorded. These values are used as input to a Regimen Formulation Procedure similar to that of Jelliffe when computing an initial digitalis dose regimen. Second, the PSM contains assertions describing the type of therapeutic and toxic symptoms and signs to be expected in the patient. These are used in the feedback loop to characterize the degree of therapeutic effect and of toxic reaction demonstrated in response to therapy (see Section 3.3.1).

Construction of the PSM involves a series of subtasks, each having specific information requirements and information outputs. Moreover, these subtasks must be executed in a particular sequence if their information requirements are to be met. This sequence can be deduced by matching information inputs and outputs of each subtask. It turns out that very few legitimate sequences exist⁶. The following sequence is the one used by ANNA:

1. Characterization of Cardiac Rhythm

It is often useful to examine changes in the cardiac rhythm when interpreting the patient's response to the drug. In addition, the type of rhythm affects the reason for digitalization. For example, a patient in sinus rhythm should not be digitalized unless he/she is in failure or is being digitalized prophylactically. This

⁶This bears a strong analogy to much of the work done in computer recognition of visual scenes. Early attempts experienced limited success, primarily because of "oversimplification" of the domain. Later work which made use of the rich assortment of clues in a visual scene such as shadows and cracks produced much better results.

is commonly done to prevent possible failure or to control ventricular rate should certain arrhythmias emerge under various stresses, such as surgery.

2. Characterization of Purpose

The physician should specify why the patient is being digitalized. Without a clear idea of what is to be gained by digitalization, it will be difficult to determine whether or not it has been gained. Furthermore, this decision determines the initial dosage estimate (IDE) which is adjusted by subsequent subtasks as described below. The IDE used by the system for failure, arrhythmias and prophylactic use is 0.625 mg (digoxin).

3. Characterization of possible Metabolic Effects

The "educated guess" should take into account any metabolic factors that might increase or decrease digitalis sensitivity in a patient (see Section 1.2.3). For each condition causing increased sensitivity, an appropriate adjustment should be made according to the following:

<u>Sensitivity</u>	<u>Reduction Factor</u>
pO ₂ (65 - 80)	3/4
(50 - 75)	2/3
(< 50)	1/2
Hypothyroid	2/3
Hypokalemia	2/3
Cardiomyopathies	2/3

Decreases in sensitivity are of less concern, since they may be corrected in the feedback loop.

4. Decision Concerning Rate of Digitalization

The physician should consider the patient's overall clinical

status and the effect which is desired in determining the rate at which to digitalize the patient. Once this decision has been made, the appropriate dosages can be computed using the desired rate and the IDE obtained from (2) and (3).

5. Characterization of Renal Function

It is not necessary to characterize renal function for the purposes of loading the patient (if loading is done rapidly), but it is essential to know renal function when computing a maintenance dose, since urinary excretion forms the primary losses of digoxin from the body (this is not necessarily true with digitoxin...see Section 1.2.4). For patients with stable renal function, this information can be gathered now and used later in computing daily losses.

6. Decision Concerning Route

The choice of route is primarily a function of the rate of digitalization and of clinical convenience. If the patient is to be digitalized rapidly, intravenous administration is appropriate. This also true of patients unable to take digitalis orally (eg, comatose patients).

7. Decision Concerning Preparation

In general, digoxin is the preparation of choice for several reasons. First, its quick onset and short half life allow closer control over its effect. Second, more pharmacological research has been done with this preparation and its kinetics are better understood than those of the other digitalis preparations.

This completes the tasks necessary for construction of the PSM. The system has formulated the IDE (subtasks 2 and 3), has obtained values for weight and renal function and is now able to generate an initial dose regimen by using these values as input to the Regimen Formulation Procedure mentioned earlier. The next step to consider is the feedback loop.

3.3 The Feedback Loop

Once digitalis has been administered, the patient's response to it is evaluated and the next step in the therapy program is formulated, based on this assessment. The feedback loop consists of two subtasks: characterization of therapeutic and toxic response and formulation (and execution) of therapeutic action.

3.3.1 Characterization of Therapeutic and Toxic Response

As the amount of digitalis in the body is increased, a corresponding increase in both therapeutic and toxic manifestations is seen, each of which can be viewed as increasing at a different rate, hopefully with therapeutic gain increasing faster than toxic response. In the feedback loop, changes in the patient's clinical state are interpreted in light of the information contained in the PSM and each is considered an independent characterization of the degree of therapeutic gain and degree of toxic reaction. Before discussing this process in more detail, a descriptive framework for the expression of these characterizations will be presented.

A Descriptive Framework

A straightforward method for describing the degree of therapeutic or

toxic response would be to assign the patient a rating on some continuum, say from zero to 100. A patient could then be described as being 43% "toxic" or 22% "therapeutic". The formulation of such an exact assessment, however, would prove to be difficult. Furthermore, such a fine characterization is not really necessary.

A more attractive method would be to make this assessment discrete, using terms such as "none", "some" or "fully" to describe the location of the patient on it. In this way a patient might be described as being "no therapeutic but some toxic" or "fully therapeutic with some toxicity." This approach has the advantage that only three possibilities must be considered instead of the hundred possible characterizations of the previous scheme. This reduction makes it considerably easier to decide which category is appropriate. The drawback of this type of description is that it has a low "resolution", but more resolution is not needed, since this assessment will be re-evaluated the next time through the feedback loop.

The following discussion details the construction and appearance of the portion of the PSM containing assertions that describe the type of therapeutic and toxic symptoms and signs expected in the patient.

Characterization of Therapeutic Gain

The original reason for digitalization determines the criteria contained in the PSM for the measurement of therapeutic response. The possible reasons are:

1. Digitalization for Arrhythmia

If the reason for digitalization was atrial fibrillation, a decrease in ventricular rate would be considered as a sign of increased therapeutic effect. The extent of this increased therapeutic response

(eg. "none", "some" or "fully") is dependent on the size of the decrease in rate as well as the magnitude of the current rate. For example, if the rate falls under 100, the patient is considered to be "fully" therapeutic.

2. Digitalization for Congestive Heart Failure

If the patient was digitalized for congestive heart failure, the therapeutic response can be measured by changes in relevant symptoms such as orthopnea, dyspnea, paroxysmal nocturnal dyspnea (pnd), etc. and signs such as ankle edema, neck vein distension, etc. Each of the symptoms and signs is assigned a severity scale. For example, ankle edema is expressed as 1+, 2+, 3+, etc. If the extent of edema changes from 3+ to 1+, a net change of (plus) two units results. Similar computations can be carried out for the other symptoms of failure. We can then define "fully therapeutic" as a change of Y units or more, "some therapeutic" as a net change between X and Y units and "none" therapeutic as a net change of less than X units, where X and Y are fixed threshold values. Adjustment of the threshold values will make the system more or less demanding when judging the therapeutic response of patients in failure.

3. Arrhythmia and Failure

In cases where digitalis is given to remedy both an arrhythmia and failure, one should deal with the arrhythmia first in the manner described above. Digitalis treatment of the rhythm disturbance generally results in the improvement of failure, since it is often a manifestation of the arrhythmia.

4. Prophylactic Use

Patients being digitalized prophylactically are generally given smaller amounts of digitalis and hence are less likely to experience digitalis toxicity. Because of this, characterizations of therapeutic gain and toxic response are generally not relevant. It is nevertheless important to watch for possible signs of toxicity, as described above.

Characterization of Toxicity

The characterization of toxicity is largely independent of the reason for digitalization. There are a number of conditions which, if present, immediately result in a classification of "fully" toxic. These include:

1. PVCs; appearance of multifocal PVCs, PVCs in salvos or PVCs in excess of 15 per minute.
2. Development of paroxysmal atrial tachycardia (PAT) with block.
3. Development of second or third degree heart block .

In addition to above conditions, there are a number of developments suggestive of toxicity. The presence of each contributes a certain amount to a "score", the value of which is used in determining the extent of toxicity (similar to judging therapeutic gain for digitalization of failure). These include:

1. Nausea, vomiting, anorexia, certain visual disturbances
2. Small increases in the amount of PVCs occurring per minute
3. Excessive slowing (ventricular rate dropping below 60 beats per minute).

3.3.2 Formulation of Therapeutic Action

Once the evaluation of the patient response has been carried out, it is a relatively simple matter to decide on the next step in the therapy program.

Since the descriptive framework permits three possibilities for both therapeutic effect and for toxic effect, only nine descriptive states are possible. Each descriptive state has a well defined course of action associated with it. The states and course of action for each are:

1. Fully therapeutic - no toxic

The patient has achieved the therapeutic goal and shows no signs of toxicity. Place patient on maintenance dose which fixes the digitalis load at this level. Digitalization is completed unless patient's condition changes.

2. Fully therapeutic - some toxic

The patient has met the therapeutic goal but is experiencing some toxicity. Correct any sensitivities which may be causing increased sensitivity. Hold digitalis until toxic manifestations subside and increase slowly to just below that level. Treat toxic manifestations if necessary (administration of potassium or anti-arrhythmia therapy). Obtain a serum digoxin level if possible.

3. Fully therapeutic - fully toxic

The patient has met the therapeutic goal but is experiencing dangerous level of toxicity. Perform same actions as in (2) above.

4. Some therapeutic - no toxic

The patient is showing some therapeutic response to the drug but not enough; no toxic reaction as yet. Increase the amount of digitalis slightly.

5. Some therapeutic - some toxic

The patient is showing some therapeutic response to the drug but not

enough; some toxicity is beginning to appear. Perform same actions as in (2) above.

6. Some therapeutic - fully toxic

The patient is showing some therapeutic response to the drug but not enough; experiencing dangerous level of toxicity. Correct any sensitivities, perform same actions as in (2) above. Think about using another type of treatment.

7. No therapeutic - no toxic

The patient is not responding to treatment. Make sure drug is being taken and investigate the possibility of malabsorption. Obtain a serum digoxin level. Correct any conditions which are decreasing sensitivity. Administer more digitalis, exercising caution due to increased digitalis load.

8. No therapeutic - some toxic

The patient shows no therapeutic response but some toxic effect is evidenced. Perform same actions as in (2) above. Think about using another type of treatment.

9. No therapeutic - fully toxic

The patient is showing dangerous level of toxicity and no therapeutic response. Hold digitalis and correct sensitivities if present. If no sensitivities present, discontinue use of digitalis.

The feedback loop is continued until one of the "halting states" is encountered (eg, states 1, 6 and 9 above), at which time digitalization is completed and the patient can be continued on the current maintenance schedule.

4. The Mechanisms of ANNA

4.1 An Overview

In the previous chapter it was shown that an ordering of the subtasks involved in digitalis administration exists. The system's internal representation of this ordering is the Therapy Transition Network (TTNET) appearing in Figure 4.1. The TTNET consists of a number of nodes linked by arcs. For the purposes of the present discussion, one can think of each node as a procedure, although this is implemented somewhat differently in the actual system (see Appendix C). Each of these procedures may be linked to sub-procedures by three types of links: non-selective (solid arcs), semi-selective (not shown in Figure 4.1) and selective (dotted arcs). When a node in the TTNET completes its execution (called a Network Procedure Execution or NPE), one or more of the subprocedures linked to it are executed depending on the type of linkage. Selective linkages imply that at most one of the subprocedures may be selected and executed, whereas non-selective linkages result in the serial (top to bottom) execution of the subnode procedures. Semi-selective linkages result in the execution of one or more of the subnodes and are considered later in Section 4.5.

For example, upon completion of execution of the procedure corresponding to the node RATE-OF-DIGITALIZATION, at most one of the subnodes, INSTANTANEOUS-RATE, RAPID-RATE, MODERATE-RATE or SLOW-RATE will be chosen and executed (see Section 4.4 for a discussion of the mechanism for making this choice). On the other hand, execution of the node BEGIN-THERAPY eventually results in the (serial) execution of the subnodes:

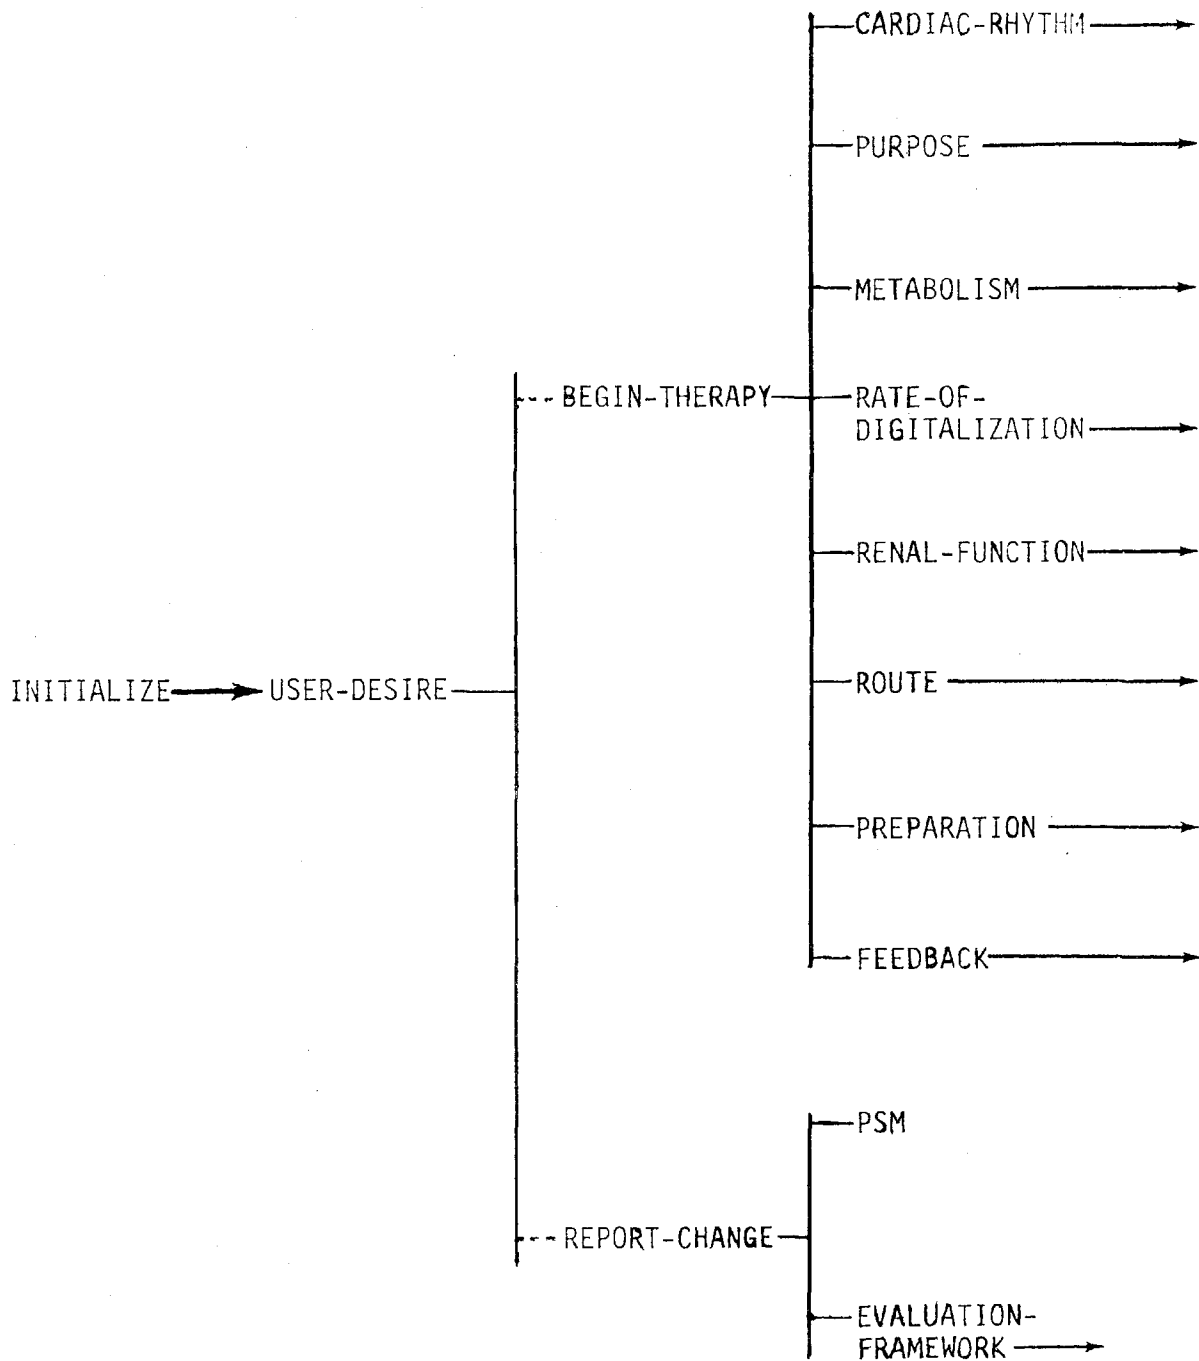


Figure 4.1 The TTNET (continued on next page)




Figure 4.1 (continued)

CARDIAC-RHYTHM
PURPOSE
METABOLISM
RATE-OF-DIGITALIZATION
RENAL-FUNCTION
ROUTE
PREPARATION
FEEDBACK

The system begins with execution of the node, INITIALIZE. Execution of this node results in the execution of a number of other nodes, as described above. The effect of these executions is to produce the system behavior demonstrated in Chapter 2. The rest of this chapter will be devoted to a more detailed discussion of the processes which occur during the execution of nodes of the TTNET. Before presenting such a discussion, it is appropriate to make some comments concerning the representation used by the system.

4.2 Representation

ANNA uses a set of LISP programs collectively known as GOBBLE that were developed by members of the Clinical Decision Making Group at Project MAC. GOBBLE facilitates the representation of medical knowledge in a declarative fashion and allows pattern matching. In order not to burden the reader, a full discussion of this facility is presented elsewhere (see Appendix B) and a stylized english version, appearing in italics, will be used instead of the actual GOBBLE assertions.

4.3 Data Collection

The collection of data about a patient is a formidable task in itself, particularly when the program is to deal with busy physicians who have little

time or patience for unnecessary questions or questions phrased in unfamiliar terms. A great deal of time was spent developing a question asking module that would know what facts were needed, in what order they should be asked and what logical dependencies exist between facts. For example, when confronted with a patient demonstrating a low serum potassium level, it is advisable to ask about concurrent diuretic therapy, since the existence of such may account for the low potassium level. If diuretic therapy is being administered, further inquiries as to the type of therapy (acute, chronic, episodic) and the existence of potassium supplement use should be made. The question asking module has the ability to accept such logical dependencies and to use them when determining what questions need to be asked.

In order to maintain a consistent data base, it is essential to detect erroneous user replies. After a more detailed discussion of the question asking module, the problem of data verification will be considered.

Question Asking

In the course of a NPE, a request may be issued to the question asking module to gather data about some facet of the patient's clinical state. This is implemented by associating a procedural fragment called a Question Directing Subroutine (QDS) with each of the nodes of the TTNET.

The system is equipped with about fifty QDS's which serve to guide its information collection. Each of these QDS's is a simple procedure fragment such as those exhibited above. These represented pre-compiled "depth first" procedures for data collection, such as those imparted to medical students learning to perform physical examinations. They are not to be confused with the overall control structure of the system that invokes them.

A simple QDS would be:

(ASK (VALUE WEIGHT))

which would result in a question similar to Question 5 of Figure 2.1. The answer is checked by lower level routines for validity and entered into the data base (see Validity Checking below).

In order to reflect logical dependencies between facts, more complex QDS's involving conditionals can be composed. For example, the following QDS would ask about the existence of orthopnea, requesting further information, only if the user indicates that orthopnea is present:

(ASK (STATUS ORTHOPNEA))
 (IF (STATUS ORTHOPNEA PRESENT)
 (ASK (AMOUNT ORTHOPNEA)))

A further extension of this mechanism is to allow QDS's to call for the execution of other QDS's. For example, the QDS for CONGESTIVE-HEART-FAILURE might look like:

(ASK ORTHOPNEA
 DYSPNEA
 PND
 ANKLE/ EDEMA
 NECK-VEIN-ELEVATION
 S3-GALLOP
 CARDIOMEGALY
 PULMONARY-VEINUS-HYPERTENSION)

Interpreting this QDS would result in the generation of calls to the question asking module to interpret the QDS's for orthopnea, dyspnea, pnd, etc.

Because of the modularity and simplicity of the QDS's, it would be feasible to construct a special interpreter to explain what each one does. In this way, the program could explain what information it is trying to gather as well as how it thinks it should go about doing so. Furthermore, this simplicity makes them easy to change. The data collection behavior of the

program may be adjusted in a simple and straightforward manner*.

Validity Checking

Multiple choice format for question responses was chosen in order to simplify the task of error checking, since the alternatives displayed are considered to be the only legitimate responses. Some of the questions, however, ask for values (eg, weight, serum potassium, ventricular rate, etc.). A response to this type of question is checked by comparing it to the range of admissible values specified in the system's dictionary. If the value entered by the user falls outside this range, an error message is printed and the system accepts another value from the user.

The system also checks two or more related facts for validity. For example, the verification module would complain about a weight change from 165 pounds to 330 pounds in one day, even though both of these values lie within the admissible range of weights. Similar checks are performed to detect unlikely changes in serum potassium levels, ventricular rate, etc.

The above strategy works well when considering responses to individual facts or to simply related groups of facts ("syntactic error checking"), but some higher level error checking is needed ("semantic error checking"). For example, if the user specifies slow digitalization for a patient in atrial fibrillation, the system should interpret this as a semantic error, indicating that "rapid rate" would be more appropriate. Semantic error checking of this type is implemented via a special mechanism called daemons, which will be discussed at length in Section 4.5.

*An interesting extension of the question asking module would be another sub-module which would accept some higher level description (eg, English or otherwise) and generate (or modify) QDS's automatically.

4.4 Subnode Selection

Until now, little has been said concerning how the data being collected are used by the system. In this section, one way in which the system uses the data it collects is presented: subnode selection.

As mentioned earlier, the last thing done in an NPE of a node containing selective links to a number of subnodes is the selection of one (or none) of those subnodes to be executed. This is done in the following manner. The supervisor issues a request to a pattern matching module (PMM) to choose the "best" subnode to be executed next. The PMM carries out this activity in two phases. First, it retrieves pattern assertions associated with all subnodes which are (selectively) linked to the node finishing execution. On the basis of a comparison of these pattern assertions and the contents of the data base, each subnode is classified as being either qualified or unqualified for selection (see below). Second, the PMM chooses the "best" subnode from among the qualified subnodes from the first phase. In this manner, at most one of the subnodes will be selected. The following discussion indicates what these pattern assertions look like and contain an example of PMM operation.

Phase I - Examination of Subnodes

Each node of the TNET has associated with it a series of pattern assertions composed of a pattern type (PREREQUISITE, PRECLUDES, SUFFICIENT), and the pattern itself. The rules for interpreting these pattern assertions are the following:

1. A PREREQUISITE pattern assertion must be satisfied in order for the match to be successful (see 3 below).
2. A PRECLUDES assertion which is found to be true invalidates the match.

3. A SUFFICIENT assertion which is found to be true immediately results in a successful match, even if the pattern contains unprovable PREREQUISITE assertions or valid PRECLUDES assertions.

A successful match results in the classification of the subnode as qualified and it will be considered in the second phase of selection. Otherwise it is noted as being unqualified for selection. For example, the subnode SLOW-RATE of the node RATE-OF-DIGITALIZATION has two pattern assertions associated with it:

"Selection of SLOW-RATE is precluded if the reason for digitalization is an arrhythmia, if pulmonary edema is present, or if the user specifies some other rate of digitalization."

"Specification by the user that slow digitalization is desired is sufficient to qualify SLOW-RATE."

For example, the PMM would find SLOW-RATE to be a successful match if the data base contained:

"The user prefers to digitalize the patient slowly."

If this assertion was not found, the match would still be successful if none of the following is true:

"The patient is being digitalized for an arrhythmia"

"Pulmonary edema is present"

"The user's preference for the rate of digitalization is other than slow rate."

Hence the pattern associated with SLOW-RATE can be paraphrased in the following manner:

"SLOW-RATE is to be considered a viable candidate for the RATE-OF-DIGITALIZATION unless the patient is being digitalized for rapid atrial fibrillation, or there is pulmonary edema present, or the user has specified some other rate."

Phase II - Selection of Best Subnode

After the PMI has established which subnodes are qualified, it issues a request to a special submodule to choose the best among them. In order to do this, the relative merits of the qualified subnodes must be weighed. Although some heuristics could be used to accomplish this, a rather simple method has proven effective. The relative selection submodule has access to fixed priority lists for each group of subnodes. For example, the priority list for the subnodes linked to RATE-OF-DIGITALIZATION is:

(SLOW-RATE MODERATE-RATE RAPID-RATE INSTANTANEOUS-RATE)

If MODERATE-RATE and RAPID-RATE were the qualified subnodes from Phase I, the relative selection submodule would choose the former, based on the ordering specified in the priority list. This particular priority list reflects the general maxim: "digitalize patients as slowly as possible", since SLOW-RATE has priority over MODERATE-RATE, etc.

4.5 Daemons

The daemon mechanism performs higher level semantic error checking. A daemon node is (optionally) associated with each node of the TINET, being treated exactly as any other node, with one difference: it is executed immediately after the node to which it is attached (eg, before any other subnode selections or executions). This is indicated in Figure 4.2 by the wavy line connecting the daemon node, RATE-DAEMONS, to the node RATE-OF-DIGITALIZATION.

Since daemon nodes are essentially treated like any other node, it is possible to have a daemon node linked to a number of daemon subnodes. Such linkages are denoted as being semi-selective, appearing in Figure 4.2 as

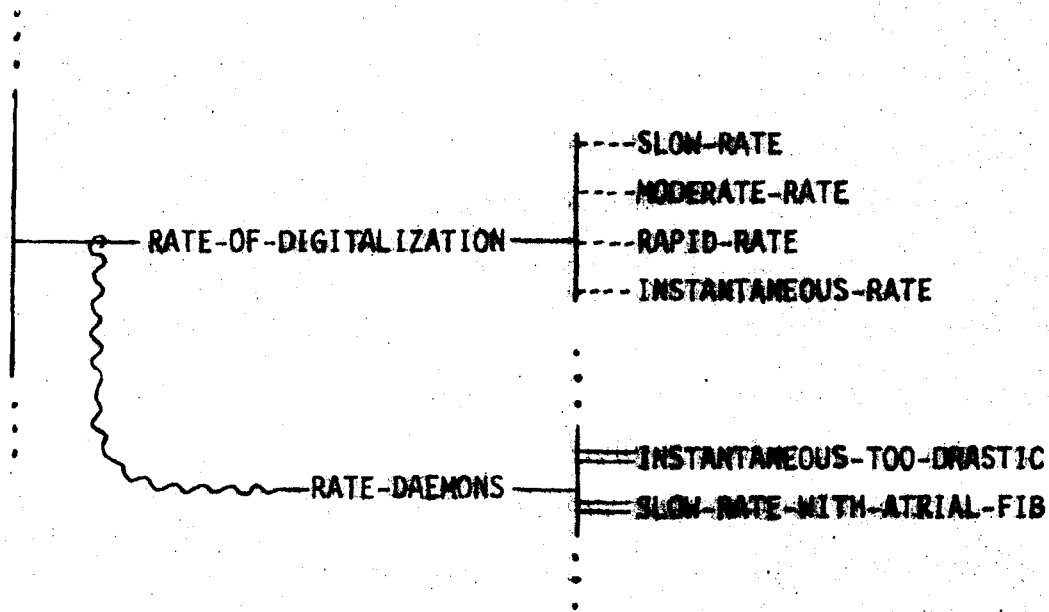


Figure 4.2 Daemons and Semi-selective Linkages

double lines. Semi-selective linkages are similar to selective linkages, however instead of performing the second phase of the selection process, the PMM causes all qualified subnodes to be executed, in the order in which they were examined. Each of the subnodes contains a message to be displayed which informs the user what type of objection the system has to the incoming response.

The session appearing in Figure 2.1 contains an example of daemon operation. The system notices that the route preferred by the user was "oral" in spite of the fact that the patient's rate was quite high. The daemon subnode which recognizes this error is **ORAL-WITH-QUICK-RATE** and contains the pattern assertions:

"The user prefers oral administration and the ventricular rate is currently greater than 140."

In practice, this allows construction of daemon subnodes with patterns which detect semantic inconsistencies in the incoming data. As should be

apparent from Chapter 1, many "medical common sense" checks need to be made to insure the quality and consistency of the data contained in the data base.

4.6 Bookkeeping

What is done during a NPE other than information gathering and validity checking? Clearly something else must go on if the system is to formulate useful recommendations and/or conclusions.

One useful side effect of a NPE, particularly when formulating the initial regimen, is to assign values to internal variables. These values are later referenced by the computation routines. For example, execution of the node NO-RENAL-FUNCTION will result in the following assignment being made:

"Set the value of renal function to zero"

Subsequent requests to the computation module for the value of renal function will result in a value of zero being returned.

It is possible to specify the value of an internal variable as being constant (age or weight), variable (maintenance dose) or procedural (in the face of changing renal function, a least squares procedure is executed in order to obtain a projected value). The example above indicates how constants are assigned. A slightly more complex situation in which variable value assignments are made is discussed below.

If the patient's renal function was not zero, a node other than NO-RENAL-FUNCTION would have been chosen and executed. Assume that STABLE-RENAL-FUNCTION had been chosen instead. This time a measure of renal function is needed along with a method for computing the patient's renal function given this measure. To do this, RENAL-FUNCTION-MEASURE is executed and asks questions about the various types of renal function measures available such as

creatinine clearance, serum creatinine, etc. If a serum creatinine value is available, the node SERUM-CREATININE will be chosen and executed immediately after RENAL-FUNCTION-MEASURE finishes and will carry out the assignments:

"Set the value of serum creatinine to <some value>"

and

"Set the value of renal function to 0.9 times the value of weight divided by 154.0 times the value of serum creatinine"

Whereas the first assignment has the effect of binding the variable SERUM-CREATININE to some constant value, the second assignment binds RENAL-FUNCTION to an expression, whose evaluation will return the appropriate value. Such expressions may contain arithmetic operators (eg, sums, differences, exponentials, etc.) and variables which have been assigned some value (either constant, variable or procedural). When the computation module receives a request for the value of RENAL-FUNCTION, it determines that its value can be obtained from the values for SERUM-CREATININE and WEIGHT, each of which is obtained by a recursive call to the computation module. Eventually, the evaluation is completed and the appropriate value returned.

Consider now the case in which renal function is found to be unstable. A similar series of node executions will be carried out. The system will realize that an estimation of renal function will have to be made, and will ask the user which measure is available. Depending on the reply, either ESTIMATE-CREATININE-CLEARANCE or ESTIMATE-SERUM-CREATININE will be executed. Assuming that serum creatinine values are to be used, execution of ESTIMATE-SERUM-CREATININE will carry out the following assignments:

"Set the value of serum creatinine to the procedure APPROXIMATE-RF"

and

"Set the value of renal function to 0.9 times the value of weight

divided by 154.0 times the value of serum creatinine"

where APPROXIMATE-RF is a procedure for performing some sort of projection, given past serum creatinine values. In the last case, a similar series of assignments was made, the difference being that SERUM-CREATININE now has a value of procedural type. When the computation module receives a request for the value of RENAL-FUNCTION, it will evaluate the defining equation for renal function. As before, this evaluation requires recursive calls to obtain values for SERUM-CREATININE and WEIGHT. This time, the procedure APPROXIMATE-RF will be executed in order to arrive at a value for SERUM-CREATININE and WEIGHT remains a constant. Finally, RENAL FUNCTION is computed by substituting these values into its defining expression and evaluating it.

4.7 Evaluation of Patient Status

One of the most important activities of the system is the construction of the PSM (see Section 3.2) and the assessment of the current status of the patient during the feedback loop using the PSM. In this section the following topics will be considered: 1) construction of the PSM, 2) its role in evaluating the patient response and 3) the formulation of advice and recommendations.

4.7.1 Building the PSM

The PSM is constructed during the execution of the node FEEDBACK, occurring upon completion of the question asking activities during the initial session (see the TINET in Figure 4.1). The PSM is built by fitting together appropriate pieces of a number of patient prototypes in its data

base. The effect of the execution of the node, FEEDBACK, is to select one of the subnodes:

ATRIAL-FIB-PROTOTYPE
 ATRIAL-FLUTTER-PROTOTYPE
 PAT-PROTOTYPE
 CHF-PROTOTYPE

The selection process is primarily dependent on the reason for digitalization recorded in the data base. For example, associated with the sub-node ATRIAL-FIB-PROTOTYPE is the pattern assertion:

"A prerequisite for ATRIAL-FIB-PROTOTYPE is that the type of arrhythmia is atrial fibrillation and that the patient is being digitalized for an arrhythmia"

Linked to each of these patient prototype nodes (via a non-selective link) is the node STANDARD-PROTOTYPE. Hence two patient prototype nodes may be executed: STANDARD-PROTOTYPE and one of the nodes linked to FEEDBACK.

Each patient prototype node contains two types of information: prototype clauses and relevant daemons.

Prototype Clauses

Execution of a patient prototype causes each of the prototype clauses associated with that prototype to be examined. A prototype clause consists of a pattern, a QDS (optional) and an interpretation clause. The interpretation clause in turn consists of an interpretation pattern and a summarization statement.

The prototype clause is processed in the following fashion. If a successful match of its pattern and the data base results, the QDS and the interpretation clause are inserted into the PSM (interpretation clauses are discussed later in Section 4.7.2). For example, one of the prototype clauses of the CHF-PROTOTYPE is:

"If orthopnea is present, ask about changes in orthopnea and interpret any improvements in the patient's orthopnea as a sign of increasing therapeutic gain."

Breaking this prototype clause apart, the pattern is: *"Orthopnea is present"* and the QDS involves asking questions about the existence and severity of orthopnea. The interpretation pattern is *"If the orthopnea is less severe than it was earlier"* and the summarization statement is "A sign of an increase in therapeutic gain is the reduction of the patient's orthopnea." Similarly, a prototype clause appearing in the patient prototype, ATRIAL-FIB-PROTOTYPE is:

"Always ask about ventricular rate changes; if the rate is below 100, interpret this as having reached the therapeutic goal."

In this case the pattern is always true, the QDS will ask about ventricular rate changes. The interpretation pattern is *"If the ventricular rate is less than 100"* and the summarization statement is *"A therapeutic goal has been reached."*

Relevant Daemons

The second type of information contained in the patient prototype is a list of relevant daemons. These daemons are similar to those in the patient prototype nodes, the difference being that the daemons contain prototype clauses which apply to signs of toxicity as opposed to signs of therapeutic gain. This distinction reflects a subtle difference in the way therapeutic gain and toxic response are judged. Whereas the system will actively search for expected signs of therapeutic gain, toxic symptoms are interpreted only when the system is told of their appearance (this is not altogether true at present, since the system steers the information gathering process - see Section 5.1).

When the system has finished examining the prototype clauses contained in

the patient prototype, it executes whatever daemons are specified as being relevant within that patient prototype. For example, the STANDARD-PROTOTYPE specifies the relevant daemons, EXCESSIVE-SLOWING, VENTRICULAR-IRRITABILITY and NON-CARDIAC-SIGNS, each of which interprets some general sign of toxicity (see Section 1.2). On the other hand, the patient prototype ATRIAL-FIB-PROTOTYPE lists NON-PAROXYSMAL-JUNCTIONAL-TACH as a relevant daemon, reflecting the fact that the development of non paroxysmal junctional tachycardia in a patient with atrial fibrillation should be interpreted as a sign of toxicity.

4.7.2 The Role of the PSM in Assessing Patient Response

When the user instructs the system that he/she wishes to discuss a previously presented patient, the node REPORT-CHANGE is executed. This node is linked (via non-selective linkages) to two nodes: PSM and EVALUATION-FRAMEWORK. Execution of the former results in the formation of an assessment of the patient's response to treatment as described below.

The system uses the previously constructed PSM to assess the patient's response to therapy in two steps. First, it passes the QDS's contained in the PSM to the question asking module in order to assemble the necessary information. The questions appearing in Figure 2.2 were generated in this fashion. Second, it matches each of the interpretation patterns against the data base (note that the QDS's have just finished adding new information to the data base). For each of these patterns which is successfully matched, the corresponding summarization statement is inserted into the data base. This summary represents the system's assessment of the patient's current response to therapy.

4.7.3 Formulation of Advice and Recommendations

Following the execution of the node, PSM, the node EVALUATION-FRAMEWORK will be executed. This node is connected (via selective links) to a series of subnodes, each corresponding to one of the nine possible descriptive states mentioned in Chapter 3. The patterns associated with these states refer to the various summarization statements inserted into the data base during execution of the PSM node. For example, the pattern assertion of the subnode THERAPEUTIC-ENDPOINT (eg, descriptive state 1 of Section 3.3.2) is:

"A prerequisite for THERAPEUTIC-ENDPOINT is that a therapeutic goal has been reached."

Execution of the node corresponding to particular descriptive states has two effects. First, it results in the system's recommendations being displayed as described in Section 3.3.2. Second, it may cause changes to be made in the PSM. For example, if descriptive state 6 (some therapeutic - fully toxic) was selected, the system would suggest that the user obtain a serum digoxin or (digitoxin) level. A QDS which asks about the results of this test is added to the list of QDS's contained in the PSM and appropriate interpretive clauses are inserted into the PSM. During the next update session, the system will ask for a serum digoxin (digitoxin) level and use it in assessing the patient response.

4.8 Summary

The discussion so far has centered around a description of the problems involved in digitalis administration and the manner in which the current version of ANNA functions in dealing with these problems. The system was

built so that it would be relatively simple to extend. The next chapter contains a discussion of what extensions are advisable in the next version of the system and some thoughts on how these extensions might be implemented.

5. Refining the Refinements

The process of developing a Digitalis Therapy Advisor is not unlike the method of administering digitalis described in the previous chapters. Each is best achieved by constructing an "initial guess" and improving this starting effort based on its performance. In the same way that Jelliffe's early efforts gave rise to this research, it is to be expected that further versions of ANNA will be constructed, based on experiences with the existing system. Initial work with ANNA has pointed to a number of areas where refinements are needed. The remainder of this chapter will be devoted to a consideration of these refinements.

5.1 The Interface

The worth of a Digitalis Therapy Advisor depends on the quality of the conclusions and recommendations it formulates and on the efficiency and ease of communication between the program and the user. This is particularly true when the user community consists largely of individuals having little or no previous exposure to computers (eg, doctors or nurses). The interface should provide a means for fluid and efficient communication between the user and the system, 1) allowing the user to transfer what he/she knows about a patient to the system and 2) allowing the system, in turn, to present its recommendations to the user. This communication should be as "comfortable" as possible for the user in order to assure effective interaction and communication with the system.

The nature of the interface is heavily dependent on the form of the "internal machinery" of the system. For this reason, work on the interface

was delayed until the construction of this "machinery" was completed. The current interface is described in detail in Chapter 4. In this section improvements to the existing interface will be discussed.

5.1.1 Accepting Information From the User

The current interface operates in "active" mode when obtaining information from the user. The essential feature of this approach is that the user is unable to take control of the dialogue, but must answer a series of questions generated by the system. This method has two advantages. First, it is easy to implement, since the program can be provided with a simple flowchart which directs the information gathering procedures. Second, after asking all of the questions specified in the flowchart, the system is guaranteed to have all of the information it needs. Were this not the case, the flowchart can be updated to include previously unasked questions.

Although this method gathers the appropriate data, it may not be altogether comfortable from the user's standpoint, since it places constraints on what the user may tell the program about a patient as well as when he may tell it. Typically, the user wishes to present some initial set of facts to the program, such as "this is a 53 year old woman to be digitalized for atrial fibrillation with a low serum potassium and a creatinine clearance of 120 ml per minute" and it may be frustrating for him/her to wait until the system gets around to asking for this information.

A better method would be to allow the user to enter an initial description of the patient, after which the system would run through a flowchart (such as that mentioned earlier) in order to "fill in the blanks." This approach is more difficult to implement than the previous one, since it

requires the development of a language for expressing the initial patient description and techniques for translating this description into the system's internal representation. Despite these difficulties, the development of such an interface would be feasible for several reasons:

1. Recent developments in natural language processing (12) make construction of crude English parsers possible within a relatively short time. A parser could be built that would translate an initial patient description expressed in English (or some subset thereof) into the system's internal representation. When the parser had completed this task, the normal question asking machinery would be called upon to gather additional information.
2. The statements of interest to the system are limited in number, each corresponding to an entry in the dictionary. By consulting the dictionary, the translation routine can determine what facts are relevant and what in particular about those facts is of interest. For example, the dictionary contains an entry for PVCs, specifying that the type, number and existence of salvos are all important properties of PVCs. If the initial patient description contains a mention of PVCs, the system would recognize the importance of this and enter the appropriate assertions into the data base. Conversely, if the user indicated that the patient "has brown hair and four fingers on each hand", the system would consider this information unimportant, since there are no dictionary entries about hair, fingers or hands.

5.1.2 Presentation of Recommendations

Concurrent with the development of an interface as described above, it is necessary to construct an additional interface module capable of presenting recommendations to the user in a neat, concise and logical manner. This module must "know" such things as how to format recommendations on display screens. Furthermore, it must determine which portions of the recommendations need be displayed in various situations. For example, if a patient was being digitalized for atrial fibrillation, it is advisable to display the first dose to be given followed by a request to "report back" before administering subsequent doses (see Section 2.2). On the other hand, for patients being digitalized prophylactically, a maintenance dose is usually of primary interest.

Extensions to this basic interface could be made. For example, simple warning messages might be suppressed when conversing with a cardiologist but would be displayed to a medical student. Each user could have a "personalized initialization file" which would automatically specify individual preferences (eg, "always use digoxin and always administer intravenously").

5.2 Medical Critical Mass

If the system produces an inaccurate recommendation, is it because the system's reasoning was faulty or because it did not know enough to arrive at the proper conclusion? This question is of primary importance and one which is difficult to answer. Such considerations can be divided into two categories:

1. Does the program consider enough of the available data?
2. Is the model used by the program of sufficient power and scope to formulate accurate recommendations?

5.2.1 How Much Information is Enough?

When determining what data to gather from the user, one runs the risk of burdening the user with an excess of questions. On the other hand, if the system doesn't have sufficient information, it will have difficulty formulating useful and accurate recommendations. In the construction of ANNA, a compromise was made between the number of questions asked and the information requirements of the system. The current version of the program reflects the minimum amount of questioning necessary to allow the production of useful recommendations. The current version of the system is considered to be about 80% complete in terms of the adequacy of its information gathering activities. Rigorous testing of the system in a clinical environment will expose any major "gaps" in its information requirements. Appropriate adjustments can be made to correct these deficiencies.

5.2.2 Power and Scope of the Model

The limitations seen in Jelliffe's programs were primarily due to the insufficient scope and power of the mathematical model he used. At this point, it is difficult to assess the ability of ANNA to produce accurate and useful recommendations in a representative sample of clinical settings. The most critical area to be evaluated is the performance of patient specific models. Questions which need to be evaluated by rigorous clinical testing are:

1. Do they model the patient accurately enough or should more information be contained in them?
2. How careful should the system be when determining if a PSM is still valid?
3. How many possible PSMs are there? If there are not many (say less than 100), should the system select one "off the shelf" instead of constructing them for each patient presented?

4. Are there patients for whom the system will be unable to construct a PSM? If so, does this indicate a fundamental inadequacy of this approach?

Initial experiences indicate the notion of PSMs and their current use by the system to interpret changes in the patient's condition are effective. Future adjustments in this area will focus primarily on extending the range and content of the PSMs rather than any fundamental readjustments in the way they are built or used. It is hoped that the strategies employed by ANNA will constitute a significant improvement over those previously available.

5.3 Error Recovery

Unreliable or inadequate data are an unwelcome but ever present problem in real world clinical situations. The following are problems which need to be addressed in this area: 1) how can the system go about detecting erroneous information, 2) what strategies can be employed to correct any decisions based on the erroneous data and 3) what are the appropriate assumptions in situations where the available data is inadequate?

Detection of Erroneous Information

The system currently has the ability to run both syntactic (via the dictionary) and semantic checks (via the daemon mechanism) on incoming data. An initial attack on the recognition of erroneous information would be to extend both of these facilities, particularly the daemon mechanism. This would involve a minimum of work, due to the ease of adding new daemons to the system. Most of the effort would involve identifying those areas where checks for erroneous information should be made. These include:

1. Impossible items. The system should check for enormous weight changes, changes in sex, large shifts in age, etc.

2. Improbable items. A 120 pound 3-year-old, renal function of more than 100% of normal, unusually large maintenance doses are all to be considered improbable.

3. Suspicious items. Suspicious items are things such as an increase in the ventricular rate following digitalis administration, large shifts in potassium, arterial oxygen, etc.

A simpler method for the detection of erroneous information is direct notification by the user that a previously entered item is not valid. The difficulty with this method is deciding what to do to correct the error.

Correcting Decisions Based on Erroneous Data

Since the system assumes all incoming data to be reliable, there are no facilities to recover from the input of invalid data. It would be possible, however, to implement an error recovery facility in the following manner. Each fact is associated with a severity class, with essential facts such as the reason for digitalization considered the most severe. The lowest severity class would include non-essential facts (such as the sex of the patient). When a datum is identified as being erroneous, the system looks up its severity class. The erroneous fact may move to a different severity class based on the following:

1. Difference between the erroneous item and the actual item. Thus an error in weight of two pounds would be put into a lower severity class than normal (eg, record the true value but don't process the error further).
2. The point at which the erroneous information was entered. If the information was just entered, no error recovery may be necessary (note this may also be true if the information is old and has already been replaced by subsequent values).

Error handlers of varying abilities are invoked, depending on the severity class of the error made. At the highest level, the system would correct for the error by completely reworking everything it did since the fact was entered. Lower level error handlers would employ tracing facilities which

can connect each fact to the places where it had been used. Assuming these situations do not affect further decisions, the decision can be "re-run". This is particularly true of errors occurring in information used by the computation routines when generating the initial guess. In such cases, the value may be corrected and the computation carried out again.

Making Appropriate Assumptions

What should the system do if the user is unable to respond to one of its questions? Clearly the system must have some provision for making reasonable assumptions when it cannot have all the data it needs. The ability to do this requires knowing when it is valid to make assumptions as well as what assumptions to make. The former represents a major difficulty in developing assumption-making capabilities, since it requires a firm definition of the minimal set of facts necessary to make meaningful assumptions. Despite this obstacle, it would be possible to equip the system with the ability to make reasonable assumptions in certain areas where data may commonly not be available, such as estimation of renal function, values of serum potassium, etc. Once again, the system would invoke assumption making routines of varying abilities, depending on the particular data under consideration. For example, it might settle for a classification of "low", "medium" or "high" if no serum potassium value is available. If the user could not specify the type of cardiac rhythm, however, the system would try to get as much information as possible (eg, asking about the patient's EKG) before making any assumptions, since knowing the cardiac rhythm is quite important.

5.4 Temporal References

Valuable information can be obtained by the correct interpretation of the sequence of events and the timing between events. A significant inadequacy of the current representational scheme is the way in which the system handles time references. This is in large part due to the fact that time references were introduced after the representation had already stabilized (see Appendix B).

At the present time, the system relies primarily on its ability to compare sequences of events to known patterns, with the sequence comprised of only two events. In order to enhance the ability of the system to make use of temporal references, several improvements are necessary. First, a representational scheme must be implemented which allows convenient representation of time expressions. Second, the system needs a time specialist capable of utilizing these time expressions. This would include detecting "trends" (eg, the serum digoxin level has been rising slightly over the past week) and establishing the affect of time intervals when interpreting the significance of changes in weight, ventricular rate, serum potassium, etc.

5.5 Explanation

In Chapter 1, explanation facilities were indicated to be an essential component of a Digitalis Advisor program. Although such facilities are as yet unimplemented, ANNA has been designed in such a way as to promote such an activity. In order to generate explanations, the system must do a considerable amount of "bookkeeping" which results in an increase in program size and a corresponding decrease in efficiency. In order to circumvent this problem, different levels of explanation can be implemented according to the following scheme:

1. No Explanations Possible (greatest efficiency)
2. Informational Questions - retrieving simple facts from the data base, such as the date of the initial session, the patient's age, etc.
3. Procedural Questions - queries concerning decisions made by the program. This includes such things as "why was that question asked?", "how was the maintenance dose computed?" or "why was the IV route suggested?"
4. Projectional Questions - requests regarding the use made of some fact. Possible questions are "what is the patient's weight used for?" or "what is the effect of my answering 'no' to this question?"

There is little practical experience to draw on regarding the construction of programs capable of generating such explanations. Although it could be done using the current representational scheme employed by ANNA, preliminary efforts indicate that a considerable time investment would be necessary in order to complete this task. In addition, it is not clear what type of explanations will be desirable from the user's point of view. The most advisable route is the development of rudimentary explanation capabilities in all areas mentioned above followed by clinical testing.

Given the development of some specialized explanation facilities, the system could perform as a powerful teaching tool. The program could explain each step in its reasoning process to a medical student, allowing him/her to learn how to administer digitalis by becoming familiar with the internal model used by the program.

5.6 Efficiency Considerations and Compilation

"I once was invited to a party and was told the address was number 64. 'How am I to remember that number?' I asked. 'Simple, remember it as being four cubed.' That evening I went to number 27 by mistake."

The above incident underlines the observation that people often find it easier to remember procedures for doing things rather than the actual things

generated by the procedures. In the construction of ANNA, a similar problem was encountered: how much of the system should be written as a procedure and how much should be represented in a strictly declarative fashion? The latter was emphasized in this work, primarily because it supports explanation generation. The benefits associated with record keeping and explanation generation must, however, be balanced against the resultant decrease in program efficiency (see Section 5.5). The optimal situation would be to have both procedural and declarative representations available and to be able to switch from one to the other, depending on what types of things the user wants done. The following discussion presents a workable means of doing this.

In order to promote efficient interaction with the system, it is important that the system's response time be kept to a minimum. Currently, most of the system's activities are carried out under the direction of specialized interpreters ("interpretive execution"). One way to achieve a significant gain in program efficiency would be writing the system as a pure procedure ("procedural execution") - an undertaking which would involve a considerable amount of effort. Furthermore, making changes in this procedure would be difficult. A more attractive alternative for generation of a procedural version of the system is the construction of a compiler which would take as input the network nodes with their associated patterns and actions and produce as output a procedure whose execution carries out the appropriate actions. For example, consider the nodes for renal function depicted in Figure 5. Looking at their pattern assertions, it can be seen that only two facts are being considered: the status and condition of renal function. Furthermore, the actions to be carried out in each of the subnodes involves either some simple action (eg, set the value of some internal variable) or results in the activation of another series of node executions. From this

Figure 5
Nodes for Renal Function

```
(NODE-IS RENAL-FUNCTION
  (SUBNODES-ARE RENAL-FUNCTION
    (NO-RENAL-FUNCTION STABLE-RENAL-FUNCTION
      CHANGING-RENAL-FUNCTION)))
```

Renal Function Sub Nodes:

```
(NODE-IS STABLE-RENAL-FUNCTION
  (PREREQUISITE STABLE-RENAL-FUNCTION
    (AND (CONDITION RENAL-FUNCTION STABLE)
      (STATUS RENAL-FUNCTION PRESENT)))
  (ACTION STABLE-RENAL-FUNCTION
    (ACTIVATE RENAL-FUNCTION-MEASURE)))
```

```
(NODE-IS CHANGING-RENAL-FUNCTION
  (PREREQUISITE CHANGING-RENAL-FUNCTION
    (AND (CONDITION RENAL-FUNCTION UNSTABLE)
      (STATUS RENAL-FUNCTION PRESENT)))
  (ACTION CHANGING-RENAL-FUNCTION
    (SET-VALUE RENAL-FUNCTION APPROXIMATE-RF)))
```

```
(NODE-IS NO-RENAL-FUNCTION
  (PREREQUISITE NO-RENAL-FUNCTION
    (STATUS RENAL-FUNCTION ABSENT))
  (ACTION NO-RENAL-FUNCTION (SET-VALUE RENAL-FUNCTION 0.0)))
```

information alone, it would be possible to (automatically) generate a procedural version of these nodes which might look like the following LISP procedure:

```
(DEFINE RENAL-FUNCTION NIL
  (COND ((IF (STATUS RENAL-FUNCTION ABSENT))
    (SET-VALUE RENAL-FUNCTION 0.0))
    ((IF (STATUS RENAL-FUNCTION PRESENT))
    (COND ((IF (CONDITION RENAL-FUNCTION UNSTABLE))
      (SET-VALUE RENAL-FUNCTION
        APPROXIMATE-RF))
      ((IF (CONDITION RENAL-FUNCTION STABLE))
        (RENAL-FUNCTION-MEASURE))))))
```

The procedure RENAL-FUNCTION-MEASURE, called inside this procedure could be generated in a similar fashion.

An improvement to the above would be the construction of an optimizing compiler capable of using "medical common sense" rules to optimize these procedures. For example, knowing that the status of renal function must be either present or absent and that the condition of renal function is either stable or unstable allows a more efficient procedure to be generated:

```
(DEFINE RENAL-FUNCTION NIL
  (COND ((IF (STATUS RENAL-FUNCTION ABSENT))
    (SET-VALUE RENAL-FUNCTION 0.0))
    ((IF (CONDITION RENAL-FUNCTION UNSTABLE))
    (SET-VALUE RENAL-FUNCTION
      APPROXIMATE-RF))
    (RENAL-FUNCTION-MEASURE)))
```

The addition of control structures to supervise switching between procedural execution and interpretive execution would result in a truly versatile and efficient system.

5.7 Implementation Difficulties

The preceding sections have touched on a number of important considerations to be included in subsequent versions of ANNA. There are a number of problems which, although of lesser theoretical interest, are nonetheless important to consider. Among these are refinement of internal representation, storing and retrieval methods for working with patient records and logistical problems as terminal location, terminal access, bookkeeping procedures (eg, maintaining a record of sessions with the system) and instruction of the user community in use of the system.

5.8 Conclusion

"He felt a rush of pity at the mingled sight and remembrance, and, recalling the relief his mother had found from a simple preparation of foxglove, he promised Sally Oates to bring her something that would ease her. ..the fact of her having found relief from drinking Silas Marner's stuff became a matter of general discourse. When Doctor Kimble gave physic, it was natural that it should have an effect; but when a weaver, who came from nobody knew where, worked wonders with a bottle of brown waters, the occult character of the process was evident."

from Silas Marner by George Eliot

One must eventually come to grips with what is perhaps the central question involved in any research: of what value is it? The construction of ANNA did not bring forth any significant new technological devices. I view this research as a pioneering attempt to build computer programs which can perform complex problem solving tasks in a non-trivial real world domains. In fact, the most reassuring result of this work is that such domains (at least the domain of digitalis administration) are in fact susceptible to detailed and rigorous analysis. Such analysis, in my opinion, represent the first step toward the development of computer programs proficient at solving problems in complex real world domains. I believe that a better understanding of the general problem solving mechanisms used by people will follow from these views

of specific complex domains.

Pragmatically speaking, it is difficult to assess the effect systems such as ANNA will have on the practice of medicine. Certainly it would be exciting if this system turned out to be better at digitalis administration than unaided physician judgement. I feel, however, that this is a long way off. The real value of this work is, in my opinion, the effect it could have on the medical education system. Hopefully, the view of digitalis therapy presented herein is more precise and easier to assimilate than that currently available. Medical students allowed to "play" with the program would begin to understand the algorithm which it uses. It would be a short step for them to later use this same algorithm with their own patients. The net effect would be that they would be better doctors and better teachers.

BIBLIOGRAPHY

1. Ogilvie RI, Reudy J: An educational program in digitalis therapy. JAMA 222:50-55, 1972
2. Doherty JE: Digitalis Glycosides: Pharmacokinetics and their clinical implications. Ann Intern Med 79:229-238, 1973
3. Peck CC, Sheiner LB et al: Computer-assisted digoxin therapy. N Eng J Med 289:441-446, 1973
4. Doherty JE, Flanigan WJ et al: Tritiated Digoxin XIV. Enterohepatic circulation, absorption and excretion studies in human volunteers. Circulation 42:867-873, 1970
5. Doherty JE, Perkins WH, Mitchell GK: Tritiated digoxin studies in human subjects. Arch Intern Med 108:531-539, 1961
6. Jelliffe RW: An improved method of digoxin therapy. Ann Intern Med 69:703-717, 1968
7. Heizer WD, Smith TW, Goldfinger SE: Absorption of digoxin in patients with malabsorption syndromes. N Eng J Med 285:257-259, 1971
8. Doherty JE, Perkins WH: Digoxin metabolism in hypo- and hyperthyroidism: Studies with tritiated digoxin in thyroid disease. Ann Intern Med 64:489-507, 1966
9. Jelliffe RW, Buell J, Kalaba R et al: A computer program for digitalis dosage regimens. Math Biosci 9:179-193, 1970
10. Jelliffe RW, Buell J, Kalaba R: Reduction of digitalis toxicity by computer-assisted glycoside dosage regimens. Ann Intern Med 77:891-906, 1972
11. Sheiner LB, Rosenberg B, Melmon K: Modelling of individual pharmacokinetics for computer-aided drug dosage. Computers and Biomedical Research 5:441-459, 1972
12. Pratt, Vaughan R.: A Linguistics Oriented Programming Language. A. I. Memo 277, February 1973
13. McDermott, D. and Sussman, G: THE CONNIVER REFERENCE MANUAL. A. I. Memo 259a, January 1974
14. Sussman, G., Winograd, T. and Charniak E.: Micro-Planner Reference Manual. A. I. Memo 203a, December, 1971

15. Gorry, G. A., Safran C. and Kahn, K.: GOBBLE. unpublished memo

Appendix A

A Mathematical Model of Digitalis Kinetics

I. Half Life

It is known that digitalis is lost from the body through excretory pathways in an amount proportional to the amount present. This relationship can be phrased in the following manner:

$$n = n_0 e^{-\lambda T} \quad (1)$$

where n = amount left after time t
 n_0 = amount present at time t_0
 λ = excretion constant
 T = $t - t_0$

The half life, $t_{1/2}$, of the drug is by definition equal to the time it takes for half of the drug to disappear:

$$\frac{1}{2} n_0 = n_0 e^{-\lambda t_{1/2}}$$

$$\text{or } t_{1/2} = \frac{\log (2)}{\lambda} = \frac{0.69315}{\lambda} \quad (2)$$

II. Excretion Constants

Experimentally measured half lives for digoxin average around 1.6 days, which yields an excretion constant of 0.4332 day^{-1} . The excretion constant, λ , is proportional to the rate at which the drug is lost from the body. Since digitalis is lost by renal and non-renal (fecal) routes, λ can be computed in

the following manner:

$$\lambda = \lambda_{\text{renal}} + \lambda_{\text{non-renal}} \quad (3)$$

Experimental evidence indicates that in patients with no renal function ($\lambda_{\text{renal}} = 0$), the half life of digoxin increases to about 4.1 days. In this case, digoxin is being lost exclusively through non-renal routes, $\lambda_{\text{non-renal}}$ can be calculated using equation (3):

$$\lambda_{\text{non-renal}} = \frac{\log(2)}{4.1} = 0.169 \text{ day}^{-1} \quad (4)$$

Combining equations (3) and (4), a value of $.2642 \text{ day}^{-1}$ is obtained for λ_{renal} (assuming normal renal function).

Thus the excretion of digitalis can be modelled by these two parameters, λ_{renal} and $\lambda_{\text{non-renal}}$. It is generally assumed that $\lambda_{\text{non-renal}}$ is constant for all patients, but λ_{renal} is a function of the patient's renal function. Similar calculations can be done for other preparations.

III. Renal Function

If a patient with no renal function is said to have renal function = 0 and a patient with normal renal function has renal function = 1, renal function can be expressed as a number between zero and one. This is done by using various clinical measures available to the physician such as creatinine clearance or serum creatinine values in the following manner:

$$\text{renal-function} = \frac{\text{creatinine clearance}}{125.0}$$

$$= \frac{0.98 \times \text{weight}}{\text{serum creatinine} \times 154.0} \quad (6)$$

If renal function is changing, it is necessary to approximate a value by extrapolation of previous values.

IV. Maintenance Dose vs Body Stores

The daily maintenance dose, M , is by definition equal to the amount of the drug lost in a period of one day. This can be expressed as the difference between n_0 and n when $t = 1$ day:

$$M = n_0 - n$$

$$= n_0 (1 - e^{-\lambda}) \quad (7)$$

Thus if a patient is given an initial dose, body stores, at time $t = 0$, and a constant daily maintenance dose, M , (eg, taken at time = 1,2,3,...), the peak body stores at time $t = i$ is:

$$\text{body stores}_{i}^{\text{peak}} = \text{body-stores}_{i-1}^{\text{peak}} e^{-\lambda} + M \quad (8)$$

and the minimum body stores is given by:

$$\text{body-stores}_{i}^{\text{minimum}} = (\text{body-stores}_{i-1}^{\text{minimum}} + M) e^{-\lambda} \quad (9)$$

Appendix B
GOBBLE Syntax and Use

ANNA uses a data base facility written in LISP called GOBBLE. Facts may be entered into the data base as non-circular list structures (see below) and may be associated with a particular context when added. In this way, the data base consists of a number of independent contexts, each containing a series of GOBBLE assertions. This is quite similar to the data base features of PLANNER and CONNIVER (13, 14), with the following differences:

1. GOBBLE contexts are independent of each other, unlike the notion of context trees in CONNIVER.
2. GOBBLE lacks the pattern directed procedure (method) execution present in CONNIVER and PLANNER.

A more detailed discussion can be found elsewhere (15). For now, I will briefly review the syntax used by ANNA and present a few examples indicating how patient data are transformed into internal representation.

GOBBLE assertions are generally expressed as a list of three elements of the form:

(<function> <argument> <value>)

where <function> is some atomic function, <argument> is some argument of that function, and <value> is <function>(<argument>). Although <function> must be atomic, it is acceptable for <argument> or <value> to be GOBBLE assertions, in a recursive manner.

The following would be legitimate GOBBLE representations of the fact "The

patient weighs 165 pounds."

(VALUE WEIGHT 165.0) or (WEIGHT PATIENT 165.0)

Note that "weight" can serve either as a <function> or as an <argument> to the function "value". In order to maintain consistency, the system has a dictionary which unambiguously specifies legal functions, legal arguments for each function and legal values for the arguments of functions. For example, the system would object to each of the following:

(WEIGHT FREEDOM 165.0) - "FREEDOM" is not a valid argument of "WEIGHT"

(VALUE WEIGHT 13000.0) - illegal value for "WEIGHT"

However, the assertion (VALUE WEIGHT 165.0) would be found acceptable.

Remember that <argument> and <value> could also be non-atomic GOBBLE assertions. Thus the statement:

(TIME-OF (VALUE WEIGHT 165.0) DAY-1)

would be accepted by the system, providing "TIME-OF" was a legal function, etc. In fact, this is how temporal knowledge is represented by the system.

Using this format, the dictionary was compiled specifying all legal functions, arguments and values which are useful. For example, each of the following is specified as a legal assertion:

(TYPE ARRHYTHMIA ATRIAL-FIBRILLATION)

(VALUE VENTRICULAR-RATE 124)

(STATUS RENAL-FUNCTION PRESENT)

(CONDITION RENAL-FUNCTION STABLE).

By using the declarative representation described above, the knowledge of

the system is always explicitly represented by the assertions contained in the data base. This has two distinct advantages over procedural (flowchart) data acquisition, where current knowledge is implicitly represented by the current location in the procedure. First, one can generate a reasonable description of what is known about the patient simply by displaying the assertions currently in the data base. This is very difficult using procedures, since an explanation would require examination of what the procedure had been doing since execution started. Second, at each point where new information is gathered, the system has ready access to the entire body of knowledge gathered so far. This enhances the ability of the program to deal with unexpected responses or to advise a confused user what responses might be appropriate. Once again, such activities would be quite difficult while in the middle of a procedure execution.

Appendix C Implementation Details

I. Primitive Concepts and Operations

In Appendix B, the data representation mechanisms of GOBBLE were presented and the notion of GOBBLE contexts introduced. In fact, the nodes of the TTNET described in Chapter 4 each correspond to a GOBBLE context, having a unique name and containing various GOBBLE assertions. The type of assertions in a given context depend on what type of node it represents (see below) but generally include such things as pattern assertions, GOBs and an assertion indicating the type of node the context represents. There are seven legal node types: *descriptor*, *candidate*, *action*, *daemon*, *revision*, *suggestion*, and *resume*. The data base is initialized with the appropriate nodes (eg, GOBBLE contexts) and one special context, FACTS which is initially empty. This context is used to store incoming assertions (see Section 4.3)

The system consists of a data base containing the contexts, a node interpreter for each type of node, a list of nodes to be executed (called the PROCESS-LIST) and a supervisor procedure to control the execution of nodes.

The fundamental activity of the system is the execution of nodes. The supervisor procedure performs this in the following manner: the node type of the first node on the PROCESS-LIST is determined. This node is then passed to the node interpreter for that type of node. Upon completion of execution, the supervisor deletes the processed node from the PROCESS-LIST and repeats the

process until the list is empty.

II. Activation

Activation of a node is accomplished by adding the name of the node to the PROCESS-LIST. It is put into the list immediately after the current node (since nodes can only be activated while inside other nodes). In this way a network corresponding to the TNET is constructed. Notice, however, that the nodes are "tied together" via various assertions contained in the nodes. Since these assertions are stored declaratively, the system has the potential to dynamically change the TNET by altering these assertions (the current version of the system does not make use of this facility).

III. The Node Interpreters

There are eight node interpreters, each of which is an "expert" at carrying out the execution of a specific class of nodes. They are listed below along with a brief synopsis of their action:

1. DESCRIPTOR - This interpreter is responsible for the execution of nodes containing selective linkages to subnodes. Its primary activity is the selection of one of these subnodes, as described in Section 4.4. It is assumed that nodes of type descriptor contain an assertion of the form:

(CANDIDATES-ARE <name of descriptor node> <list of subnodes>)

Since this assertion is stored in a declarative fashion, it is possible to alter the <list of subnodes> portion to add or remove potential candidates from consideration.

2. CANDIDATE - This interpreter looks for an assertion of the form:

(ACTION <name of node> <action to be carried out>)

and executes it. Usually the <action to be carried out> involves setting the value of some internal variable or the activation of another node.

3. ACTION - This interpreter supervises the execution of user or system defined LISP procedures, such as those which display recommendations to the user, etc. For example, a node which generates a display of a patient summary looks like:

(NODE-IS PATIENT-SUMMARY
(NODE-TYPE PATIENT-SUMMARY ACTION)
(PRINT-PATIENT-SUMMARY))

Execution of this node will result in evaluation of the LISP function, PRINT-PATIENT-SUMMARY.

4. DAEMON - This interpreter oversees the execution of nodes corresponding to daemon nodes. Its functions are quite similar to that of the DESCRIPTOR interpreter, with the exception that all qualified candidate sub-nodes are activated (see Section 4.4).

5. REVISION - Subnodes of nodes of type DAEMON are of type REVISION or SUGGESTION (with some exceptions). The REVISION interpreter is responsible for the reconsideration of some fact received from the user which is felt by the system to be in error (note that such facts are identified by the pattern matching activities of the DAEMON interpreter - see Section 4.5). This interpreter proceeds by notifying the user of the problem and, if necessary, deletes the erroneous fact from the data base and inserts the correct one.

6. SUGGESTION - The other type of sub-node of DAEMON nodes is the SUGGESTION node. The interpreter for this node displays a suggestion contained in the node for the user. SUGGESTION nodes are responsible for messages suggesting the use of potassium supplements and correction of various disorders (hypokalemia, hypoxemia, etc.)

7. RESUME - During the course of a pattern matching activity, the system may create a new context of type RESUME which contains information about the match (eg, what facts were used, which ones were not found, the patterns used, etc.). Although the RESUME interpreter is as yet unimplemented, it is intended that its execution would result in the generation of an explanation for a choice made on the basis of some pattern matching activity.

Bookkeeping

The supervisor keeps track of node activations, recording the "parent" and "children" for each node which is executed. This information is valuable when performing backup operations or generating explanations. The system also associates each incoming fact with the node which was being executed when it was entered into the data base. In addition, each node contains a list of facts which was gathered during its execution. In this way it is possible to tell what facts were entered while in some particular node or, conversely, what node was active when a particular fact was entered.

Candidate Directory

When the descriptor interpreter chooses and activates a subnode, an assertion of the form:

(SUBNODE-FOR <descriptor> <subnode>)

is put into a special context called the SUBNODE-DIRECTORY. These statements can be used like any other fact in pattern assertions. For example, a pattern assertion which is true only if digitalis is being administered orally would look like:

(PREREQUISITE <some node> (SUBNODE-FOR ROUTE ORAL))

In addition to referencing this information in pattern assertions, it is often convenient to print out the subnode directory so that the user can see decisions have been made by the program.