LABORATORY FOR COMPUTER SCIENCE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

MIT/LCS/TM-319

MURMUR CLINIC: AN AUSCULTATION EXPERT SYSTEM

TZE-YUN LEONG

JANUARY 1987

545 TECHNOLOGY SQUARE. CAMBRIDGE. MASSACHUSETTS 02139

Murmur Clinic:
A Cardiac Auscultation Expert System

Tze-Yun Leong

Draft of December 20, 1986

Abstract

Auscultation is a technique used in cardiac physical examination to detect irregularities by analyzing heart sounds. This paper reports on the development of *Murmur Clinic*, a cardiac auscultation expert system which is able to interpret and analyze auscultatory findings, and performs a tentative diagnosis based on a formalized diagnostic reasoning process. Descriptions of the scope addressed, the design, the diagnostic algorithm used and implementation of the system, as well as a sample session, and a discussion of limitations and possible improvements are presented.

Key Words: Artificial Intelligence in Medicine, expert system, diagnosis.

Contents

1 Introduction **1**

2 The Problem **3**

3 The Design **7**

3.1 Scope and Resources 7

3.2 Knowledge Representations and Applications 9

3.3 Decision History and User-Interface 10

4 The Algorithm **12**

5 The Implementation **17**

5.1 Data Structures 17

5.1.1 The Murmurs 17

5.1.2 The Score-Tables 18

5.2 The Rules 18

5.3 Decision History Records 20

5.3.1 Evidence-Records 20

5.3.2 Score-Templates 21

5.4 The User-Interface 21

5.4.1 Display Framework 21

5.4.2 Supporting Structures 22

5.4.3 IO-Handlers 23

5.5 Top Level Command Drivers 23

6 A Sample Session **24**

7 Discussion **35**

7.1 An Evaluation of Murmur Clinic 35

7.2 Future Directions 36

8 Conclusions **38**

A THE ALGORITHM **39**

B THE SYSTEM LAYOUT **40**

C THE REFERENCE MURMURS **41**

List of Figures

2.1 Sound Pattern of a Normal Cardiac Cycle. 4
2.2 Graphic Representation of Discrete Extra Sounds. 4
2.3 Early-diastolic-decrescendo Pattern 5
2.4 Holosystolic Pattern 5

5.1 A Murmur 18
5.2 Construction Of A Rule. 19
5.3 An Evidence-Record. 20
5.4 A Score-Template. 21
5.5 The Input Panel With Standard Configuration. 22

6.1 Initiation Of The DIAGNOSE operation. 25
6.2 First Set Of Input Findings. 25
6.3 Pop-Up Menu For Input Sets Selection. 25
6.4 Pop-Up Menu For Input Modification. 26
6.5 Second Set Of Input Findings. 26
6.6 A Message Indicating Inconsistency. 27
6.7 Third Set Of Input Findings. 27
6.8 Display Of The Diagnostic Decisions. 28

List of Tables

4.1 Interpretation of Discriminatory Significance. 14
4.2 Interpretation of Occurrence Frequency. 14

Acknowledgement

I would like to express my thanks to the following:

Bill Long, my project supervisor, for guiding and encouraging me throughout the project, and for patiently reading through drafts of this document and providing valuable comments and suggestions;

Ramesh Patil, my academic advisor, for introducing me to the Clinical Decision Making Group, and providing support, encouragement and valuable advice in all aspects of my work;

Bob Jayes, for his helpful comments and suggestions on the system;

all members in the Clinical Decision Making Group, for providing a pleasant working environment;

and finally, the MIT Undergraduate Research Opportunities Program, for making my participation in the project possible in the first place.

Chapter 1

Introduction

Auscultation is a technique used in cardiac physical examination to detect irregularities by analyzing heart sounds. This paper reports on the development of **Murmur Clinic**, a cardiac auscultation expert system which is able to interpret and analyze auscultatory findings, and then decide on the possible abnormalities that might have caused the findings.

Development of the **Murmur Clinic** is mainly motivated by our desire, within the problem domain, to:

1. understand and formalize the reasoning process of expert physicians;
2. study and apply techniques of Artificial Intelligence; and
3. ultimately, provide medical expertise serving as a decision-making aid to communities where such expertise is inadequate or unavailable.

These goals, incidentally, are in accordance with the general research objectives of the field of Artifitial Intelligence in Medicine [1].

Initiation of the **Murmur Clinic** project was borne out of the need for an auscultation component that can eventually be integrated into the Heart Failure Program, which is being developed at MIT, in collaboration with Tufts-New England Medical Center. However, the **Murmur Clinic** system, as it now stands, is an entirely independent unit. We chose to develop the system independently not only because auscultation is a self-contained subject involving sufficiently complex expertise, but also because most of our goals mentioned earlier can be achieved more easily with a specific problem domain. We believe that an interface for incorporating the system into the Heart Failure Program would not be too difficult to build. But in the mean time, construction of such an interface is of secondary importance.

As development progresses, much insight has been gained into the application of Artificial Intelligence techniques, as well as the diagnostic reasoning process of expert physicians. Development of the system is currently at the stage of performance evaluation and refinement. Although we are yet to come up with a complete system with the competence of a real expert, we believe we have established most of the components that would enable the system to behave "expertly" eventually. The issues we have examined are indeed limited, but we hope these would nonetheless put us a step forward toward a full understanding of the way in which a physician makes a diagnostic decision.

Chapter 2

The Problem

In this chapter, we give a brief description of our problem domain, i.e. we give a brief answer to the question:

- *What is Auscultation?*

To begin with, *auscultation* is a component of the cardiac physical examination in addition to inspection, palpation and percussion. An auscultatory examination involves systemic listening to different heart sounds on different locations of the body. Interpretation and analysis of the auscultatory findings can be done based on the fact that different heart sounds are produced by different forms of cardiac valve movements. Deviations from the regular sound patterns indicate abnormal valve movements, and the source(s) causing these abnormalities can then be deduced from the dynamics of the valve(s) in question.

In a normal person, usually only the "normal discrete heart sounds" can be heard when blood moves through the heart during a cardiac cycle. These sounds, which include the *first heart sound,* S1 and the *second heart sound,* S2, are associated with the movements of the heart valves in the systole (contraction) and diastole (relaxation) of the ventricles. When the heart rate is less than 100 beats per minute [1], systole is much shorter than diastole and the sound pattern can be best described as, "Lub...Dub......Lub...Dub......" In this example, the "Lub" would represent S1 and the "Dub" S2 [2]. A graphic representation of the sound pattern denoting a normal cardiac cycle (systole and diastole) used in auscultation is shown in Figure 2.1.

Irregular flow pattern of the blood through the heart, caused by valvular abnormalities or otherwise, might lead to, besides irregular first and second heart sounds, the presence of extra sounds such as the *opening snap, systolic click, third* and *fourth heart sounds*, and the *heart murmurs.* For the discrete extra sounds, the opening snap [2] and the systolic click [3] are related to valvular abnormalities, while the third and fourth

[1] A resting adult normally has a heart rate of about 70 to 80 beats per minute.

[2] Opening sound of atrioventricular valves.

[3] Opening sound of aortic or pulmonic valves.

Figure 2.1: Sound Pattern of a Normal Cardiac Cycle.

heart sounds, S3 and S4, which can be normal under certain circumstances, are usually caused by valvular abnormalities and/or other heart diseases. A graphic representation of these extra sounds is shown in Figure 2.2 [4].

Figure 2.2: Graphic Representation of Discrete Extra Sounds.

The continuous extra sounds—heart murmurs, are sounds that result from vibrations set up by vortices near the mural interfaces of the blood stream after it passes an obstruction or dilation (The Vortex Shedding Theory) [3]. These sounds can be attributed, according to Leatham [4], to three main factors:

1. High flow rate through normal or abnormal valves;
2. Forward flow through a constricted or irregular valve or into a dilated vessel or chamber; and
3. Backward or regurgitant flow through an incompetent valve, or septal defect.

Frequently, a combination of these factors is operative [5].

When describing a heart murmur, the following characteristics should be noted:

1. What is its *intensity*, that is, its *grade*? This is measured on a six point scale with grade 1 being the softest and grade 6 the loudest. A murmur associated with a thrill is classified as at least grade 4.
2. What is the *pattern* of the murmur? This can be further divided into three questions:
   - Is it systolic or diastolic?

[4] It is not necessary for all the extra sounds to be present at the same time.

- Is it crescendo, decrescendo, crescendo-decrescendo, decrescendo-crescendo or is it flat, with the same intensity throughout its duration?
- What is its duration and timing in the cardiac cycle?

Two examples of the murmur pattern representation are shown in Figure 2.3 and Figure 2.4.

Figure 2.3: Early-diastolic-decrescendo Pattern

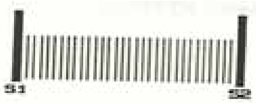

Figure 2.4: Holosystolic Pattern

3. What is the *quality* of the murmur? For example, a murmur may be described as of high, medium or low-pitched, rumbling, blowing or soft.
4. Where is the *location* of the murmur, that is, where is it best heard?
5. What areas, if any, does the murmur *radiate* to?
6. What changes, if any, are appreciated with respiration and other *physiologic* or *pharmacologic maneuvers*? These changes usually provide strong evidence to the cause(s) of the murmur. By noting the changes of the murmur in response to physiologic alterations produced by various maneuvers such as inspiration, expiration, handgrip, squatting, sudden standing, performance of the valsalva maneuver and administration of pharmacologic agents such as amyl nitrite, the origin of the vibration can usually be traced by deduction from the physiology involved. [2]

There are five major areas for auscultation–the *aortic* area, the *pulmonic* area, *3rd-left-interspace*, the *tricuspid* area and the *mitral* area (*apex*). In auscultation, each type of sound–the first and second heart sounds, extra discrete heart sounds and heart murmurs, is listened to systematically and selectively in each of the five areas. A diagnosis of the causal abnormalities can then be made from the auscultatory findings, which, in the presence of abnormalities, would involve one or more of the following:

- Presence of one or more heart murmurs [5];
- Variations of the first and/or second heart sounds; and
- Presence of one or more discrete extra heart sounds.

One major difficulty in the diagnostic process lies in the interpretation of the findings. It is a complicated task which requires considerable expertise because of the wide range of applicable information. In the interpretation of a heart murmur, for example, readily available information includes the intensity, duration in the cardiac cycle and location. These data can be supplemented by information on the quality, timing, pattern, direction(s) of radiation, and change(s) with respiration and physiologic maneuvers, any or a combination of which could provide additional evidence of the possible causes. Presence of multiple murmurs, extra heart sounds and/or variations in the first and second heart sounds further complicate the analysis procedure. These complications, however, are indeed the major factors that make auscultation an interesting and challenging problem domain for a knowledge-based system.

[5] Usually not more than three murmurs would be present at the same time.

Chapter 3

The Design

An incremental approach was adopted in the development of the Murmur Clinic system, and modularity is strongly emphasized in the design. As additional features and knowledge can be progressively incorporated, great flexibility is allowed in the modification and expansion of the system.

Even though both the structure and performance of the system are still being refined as development progresses, initial design of the system can be illustrated by answering the following questions:

- *How much knowledge should the system have?*
- *How should the knowledge be organized and applied?*
- *How should the system interact with its user?*

We will discuss each answer in turn in the following sections.

3.1 Scope and Resources

Answering the first question above addresses the scope and resources of the knowledge base. The system must have enough knowledge to make a diagnosis on auscultatory findings. In order to derive how much knowledge is involved in the diagnostic procedure, we must first understand the procedure itself; in other words, we must know the answer to the question:

- *How does a physician make a medical diagnosis?*

According to Ledley and Lusted [6], a physician's reply to the above question might be as follows. *"First, I obtain the case tests. Second, I evaluate the relative* *importance of the different signs and symptoms. Some of the data may be of first-order importance and other data of less importance. Third, to make a differential diagnosis, I list all the diseases which the specific case can reasonably resemble. Then I exclude one disease after another from the list until it becomes apparent that the case can be fitted into a definite disease category, or that it may be one of several diseases, or else that its exact nature cannot be determined."* Even though this might be a greatly simplified illustration of the actual procedure, it is a fairly accurate account of what we need to incorporate into our system.

Since the interpretation and analysis of auscultatory findings involve mostly symptomatic comparisons between the findings and the manifestations of the causal abnormalities, we decided that the corresponding associational knowledge should be a basic element of the knowledge base. An illustration of such associational knowledge is the degree of *similarity* between a finding and a manifestation of the same category. For example, we would want to establish that a finding of grade 1 in the *intensity* category is similar to a manifestation of grade 3 in the same category to a degree of 7 (in a scale from 1 to 10, with 1 being the least and 10 the most similar) ; while it is incomparable to all manifestations in any other category. The system, at the moment, does not have to know about the pathophysiology underlying the manifestations.

But the associational knowledge alone is not sufficient for the system to make accurate diagnostic judgement. Discriminatory knowledge is also needed to differentiate the degree of importance of the findings and the abnormality manifestations. This requirement arises as the manifestations of a particular causal abnormality can be varying in the degree of importance for discrimination purpose. For example, in deciding about the abnormality *mitral stenosis* (MS), presence of the opening snap is considered to be pathognomonic, while the evidence provided by a grade 1 *intensity* is less important. Moreover, the significance of a particular manifestation can be measured by the frequency with which it appears in patients with the corresponding abnormality.

In addition, judgemental knowledge for retaining only the most likely causes as diagnostic results should also be embodied into the knowledge base. This kind of knowledge is required in situations involving the exercise of medical facts and/or experience. For instance, in the case when only a systolic murmur is observed, all abnormalities which would result in one or more diastolic murmurs should be excluded from consideration.

Now that we have defined the needed amount of knowledge, we have to know how to get it. The associational knowledge, as it turns out, can be extracted mainly from various medical textbooks and publications. Extraction of the discriminatory and judgemental knowledge, however, requires work done on interrogating human experts.

3.2 Knowledge Representations and Applications

Answering the second question posed at the beginning of this chapter requires representation of the knowledge involved and its application in the diagnostic algorithm of the system. As indicated in the previous section, three different categories of knowledge are included in the knowledge base. Besides being extracted from different sources, the three categories of knowledge emphasize different epistemological structures: the associational knowledge of symptomatic relations, the discriminatory knowledge about the degree of importance of symptomatic entities, and the judgemental knowledge of medical facts. Our choice for knowledge representation is made with two basic considerations: *modularity* and *explicitness.* These two factors are essential since we wish to be able to maintain and update the three categories of knowledge independently, and operations on them should be easy to define in order for them to serve their respective purposes effectively.

We chose to represent the associational and discriminatory knowledge in discrete conceptual structures with accompanying operations. These include structures that represent individual causal abnormalities, each with a set of manifestations with their significance and frequency in the corresponding abnormality, and "tables" of similarity between auscultatory findings and manifestations of the causal abnormalities [1].

Representation of the judgemental knowledge is more complex. Since this type of knowledge consists mainly of facts that can easily be expressed in a "IF...THEN...." form, a natural representation would be the production rule formalism. Nevertheless, we also realized that the judgemental knowledge in auscultation can be classified into several subcategories, each performing a well-defined subtask such as checking the consistency of the reported auscultatory findings, recognizing further differentiability of tentative conclusions of possible causes, ruling out certain possibilities, etc. In some of these subtasks, a single fact can override the rest, while in others, a few or all of the facts should be considered before a conclusion is made.

Consequently, we decided to implement the rules as structures that can be grouped together to perform different subtasks. Order of execution of the rules within a group is unimportant, but the rules are characterized by either "terminating" or "non-terminating". Execution of any rule of the former type inhibits activation of all the other rules in the same group, while execution of any rule of the latter type simply performs its actions as side effects [2].

As an illustration, consider the following cases: For a rule responsible for detecting further differentiability among the top few choices in a tentative conclusion of possible causes, the execution conditions mainly require matching the names of the top choices

[1] A more detailed description and examples of these representations will be presented in Chapter 5, where we discuss the implementation of the system.

[2] Here the word *execution* is used to indicate that the *actions* part of the rule is carried out when all the *conditions* are satisfied; while the word *activation* simply implies that the *conditions* part of the rule is being evaluated.

with a set of differentiable abnormalities. If such a match is successful, the corresponding actions are carried out, and no other rules in the same category (i.e. responsible for the same task) need to be activated. It is unnecessary to check if another match is present, since it cannot be. Hence all the rules in this category should be classified as terminating.

On the other hand, execution of a rule responsible for excluding or including certain abnormalities from being considered as possible causes should not hinder the activation of another rule in the same category that would result in the exclusion or inclusion of some other abnormalities for consideration. As a result, all the rules in this category should be classified as non-terminating.

This choice of rules representation, though lacking the conceptual simplicity of conventional production rules, preserves the modularity and explicitness emphasized in the overall design of the system. Moreover, it has proven to be very useful as a supporting structure in other parts of the system such as the user-interface.

As suggested, our choice of knowledge representation is greatly influenced by the way in which we believe a medical diagnosis is made by a physician. By using the more complicated but explicit representations in lieu of simpler but implicit structures such as the conventional production rules, we have been able to develop a diagnostic algorithm which resembles the one mentioned in a straightforward manner. We believe by modelling his behavior in this way and observing the results, the reasoning process of the physician can be studied more closely and understood more efficiently.

3.3 Decision History and User-Interface

The way the system interacts with the user is an important criterion for evaluating the performance of the system. A good user-interface should not only allow the system to be accessed easily and conveniently, but also make the system behavior transparent to the user. The former factor aims to reduce the user's reluctance to use the system, while the latter is needed to induce the user's confidence in the reliability of the system. The desire for transparency implies the need for incorporating structures to keep track of the decision history of the diagnostic process.

Two kinds of record-keeping structures for the diagnostic procedure have been developed. The first keeps track of whether a particular causal abnormality in the database is being considered for diagnosis, and the evidences which dictate this status. This kind of "evidence record" is needed because of situations such as the following: Suppose someone hears only one murmur, a holosystolic murmur, and hears an opening snap. What we really have is *mitral regurgitation* (MR), and *mitral stenosis* (MS), but the MS murmur was so faint that the examiner did not notice it. In such a situation, two sets of evidences are applicable: Since only one systolic murmur is heard, MS would not be taken into consideration initially as the MS murmur is diastolic; however, since the opening snap is present, and it is pathognomonic for MS, MS should be considered as one of the topmost decisions. Because of the existence of conflicting explanations for evidence, it is crucial to explain to the user why and how a final decision is made. This kind of explanation cannot be done without a record of the evidence applied in the decision process.

The second type of decision history record is used to keep track of the similarity score between the auscultatory findings and the manifestations of individual causal abnormalities that are being considered for diagnosis. This provides a means to explain to the user why a certain abnormality is ranked over others.

Incorporation of the decision history feature enables the system to justify and explain its reasoning to the user. This feature has been very helpful in testing and debugging the system, as it allows inconsistency in the reasoning process to be detected easily.

In addition to the decision history, we have also developed several features to make the user-interface easy and convenient to use. One of these features is the heavy use of graphics. Since descriptions of auscultatory findings involve entities such as heart sounds and murmur patterns, locations and areas of radiation, we believe that using pictorial representations would be much more accurate and convenient than asking the user to provide a word-description of the findings. Thus we have provided an input-panel with a list of the findings in pictorial representations, and allow the user to choose from these findings by clicking the mouse over the corresponding representations. Moreover,we also try to avoid asking the user to give keyboard input during interaction in general. Menu-form prompting is adopted and the user is usually asked to choose from a list of possible answers for a particular question. In these ways, we hope to provide as "user-friendly" an interface as possible.

So far we have described the problem domain we address and the motivations behind the design of our system, as well as the approach we have chosen. In the next two chapters, we will describe how our system actually works and how it is implemented.

Chapter 4

The Algorithm

In this chapter we present the reasoning algorithm of the diagnostic process of our system. The diagnostic procedure decides on the possible causal abnormalities from the input auscultatory findings, verifying the existence of single or multiple abnormalities [1]. The results are presented as a ranked list of possible causes for each abnormality determined to be present in the case. Mechanisms for explaining and justifying the reasoning involved are also provided (see **Section 3.3**).

Development of the diagnostic algorithm was greatly influenced by the way in which we believe a physician makes a diagnostic decision (see **Chapter 3**). As a result, the three major events encoded in the algorithm are:

- Get input on case;
- Weigh degree of importance of each input parameter; and
- Generate hypotheses by eliminating unlikely candidates, and then produce decision results.

Due to implementation constraints, however, these events do not necessarily appear in that order in our algorithm. In addition, an important practice in auscultatory diagnosis is being captured in the process: In the presence of multiple heart murmurs, which in turn indicate multiple abnormalities, diagnosis of the diastolic murmur, if any, is carried out first as its conclusions can provide further evidence to the diagnosis of the other, presumably systolic, murmurs [2]. Also, we produce a tentative diagnosis, a ranked list of tentative causes for each input murmur, as opposed to a definitive diagnosis, declaring particular causes as most consistent with the evidence. We expect to incorporate definitive diagnoses into the algorithm when we have established the completeness of the knowledge base, and resort to tentative diagnoses only when the evidence is inconclusive.

[1] Up to three heart murmurs, indicating the corresponding number of abnormalities, can be present at the same time.

[2] Co-occurrences of diastolic murmurs are extremely rare.

The diagnostic algorithm proceeds as follows:

Step 1

Initial auscultatory findings are entered by the user. These findings are described in terms of heart murmurs, each with characteristics of *pattern, location, intensity, quality* and optional *area(s) of radiation*, *change(s) with physiologic maneuvers*, and variations of and/or presence of extra *heart sounds*. If more than one murmur is present in the input information, the following steps are carried out successively once for each input murmur, starting with the diastolic component, if any.

Step 2

The input findings [3] are then put through a consistency test. If an inconsistent or incomplete input is detected, e.g. only the intensity of a heart murmur is given, the user will be warned and asked if he wishes to change or add more information. If the input is consistent and reasonable, or the user does not wish to change or add information after being warned of inconsistency, the procedure enters the next step. Otherwise new input is obtained and goes through the same consistency test again.

Step 3

The input findings are then checked for evidence to rule out the abnormalities in the database as possible causes. Each causal abnormality is encoded as a reference murmur with seven attributes (six characteristics and accompanying heart sounds) in the database. Based on the evidences from the input, each abnormality is classified into either the active or the dormant state, indicating whether it is suitable for being considered in the diagnosis. In the mean time, an evidence-record is generated for each of these abnormalities, keeping track of the evidence which induced the classification. When the evidence serves to reinforce the presence of a particular abnormality, this evidence is recorded together with the *evidence-score* in the evidence-record. The higher the evidence-score, the stronger the indication of the presence of the abnormality.

Step 4

For each of the abnormalities in the list of active candidates for diagnostic considerations, a score-template is generated. The score-template is a structure with slots to hold the *similarity scores* between the input findings and the manifestations of each attribute of the respective abnormality, together with its evidence-score from the evidence-record. Filling in the score-template for a particular abnormality involves the following steps:

[3] Corresponding to a single murmur from here on.

1. The similarity score of each manifestation with respect to the corresponding input is determined from the score-table. There are seven score-tables for the seven attributes of a heart murmur. The similarity scores are encoded in a scale from 0 to 10, with 10 indicating an exact match, and 0 inconsistency.

   Each of the similarity scores are scaled by a *decision-factor* on a scale from 0 to 1. The decision-factor is determined by the *discriminatory significance* (70 percent) and the *occurrence frequency* (30 percent), of the particular manifestation. The discriminatory significance indicates the importance of the manifestation in distinguishing the corresponding abnormality from the others, while the occurrence frequency indicates how often a patient with the abnormality has the particular manifestation. Both of these factors can be interpreted on a scale from 0 to 10. Table 4.1 and Table 4.2 show the interpretations of the discriminatory significance (DS) and occurrence frequency (FS) values.

DS	Value	Interpretation
1	10	Almost pathognomonic.
2	7	Important--Strongly weighted evidence for the corresponding abnormality.
3	3	Negligible--Characteristic but not a discriminating finding.

Table 4.1: Interpretation of Discriminatory Significance.

OF	Value	Interpretation
A	10	Always present.
U	7	Usually present.
S	3	Sometimes present (Rare).

Table 4.2: Interpretation of Occurrence Frequency.

2. The weighted similarity scores of a particular attribute is summed and entered into the corresponding slot in the template. Only those attributes in which input information is provided are evaluated; otherwise a default value is provided.

3. The *overall similarity score* is generated by summing up the similarity scores of all its attributes and the evidence-score from the corresponding evidence-record for the abnormality.

Step 5

The abnormalities are ranked according to their overall similarity scores. The abnormality with the highest score is the most likely cause of the murmur.

Step 6

The ranked list is checked for further differentiability as follows: If the first two or three abnormalities are further differentiable by some physiologic maneuver(s), and input information in this regard is not provided, the user is prompted for additional information. If additional maneuver information is provided, the process returns to **Step 2**. If the user does not provide more information, or the differentiability test failed, the process returns to **Step 2** and continues with the next set of murmur findings, if any. Notice that each tentative diagnosis can provide further evidence for the diagnosis of the murmur(s) that follows. This is most true of diastolic murmur evidence followed by systolic murmur(s). The order in which co-existing systolic murmurs, if any, are diagnosed is unimportant. The cycle terminates when the list of input murmur(s) is exhausted.

Step 7

At this point we have a ranked list of possible causal abnormalities with respect to each input heart murmur. The evidence-records for these abnormalities are then checked to see if those with pathognomonic evidence are among the highest ranked choices. If not, each of the non-highest ranked abnormalities with pathognomonic evidence is separated into a new ranked list by itself. This last step is necessary because of the possibility of cases like the unobvious co-existence of *mitral regurgitation* and *mitral stenosis* mentioned in **Section 3.3**

Step 8

All the ranked lists go through a final test for discrepancies, and are displayed for the user as the diagnostic results, together with a brief description on how the decisions were made.

A summary of the algorithm presented can be found in **Appendix A**.

Four types of questions can be answered by the system about the reasoning in the diagnostic procedure. The questions, pertaining to each list of possible causes produced (i.e. corresponding to each input murmur), are as follows:

- WHY is [abnormality] ranked first?
- WHY is [abnormality] NOT ranked first?
- WHY is [abnormality] ranked as it is?
- WHY is [abnormality-1] ranked OVER [abnormality-2]?

The answers are generated by tracing the decision history. Each answer is a list of similarity scores from the score-template, and any relevant evidence in the evidence-record (see algorithm above). Detailed examples will be provided in the sample session in Chapter 6.

Chapter 5

The Implementation

The **Murmur Clinic** system is developed on, and makes extensive use of the window system and graphics facilities of the Symbolics 3600 Lisp Machines. The system can be divided into five basic units: *data structures, rules, decision history records, user-interface* and *top-level command drivers.* In this chapter we shall discuss the implementation of these units individually. A figure summarizing the interaction of the units can be found in **Appendix B**.

5.1 Data Structures

The data structures represent conceptual entities such as the murmurs and the score-tables.

5.1.1 The Murmurs

Both the characteristics of possible causal abnormalities and the auscultatory findings entered by the user are encoded in structures called murmurs. A murmur is an object with *name, intensity, quality, pattern, location,* accompanying *heart-sounds, maneuver* and *radiation.* Each of these attributes takes one or more manifestation values. Except in murmurs that represent input findings, each manifestation is accompanied by its *discriminatory significance* and *occurrence frequency* (see **Chapter 4**). The murmur representation of a causal abnormality is shown in Figure 5.1.

Each causal abnormality in the database is represented as a murmur structure. There are 16 reference murmurs in the database, as listed in **Appendix C**.

Each input murmur is encoded similarly except there are no *discriminatory significance* and *occurrence frequency* accompanying the reported manifestations.

```
AORTIC-STENOSIS
INTENSITY:     Grade 1                        with SIG:3 and FRE:U
               Grade 2                        with SIG:3 and FRE:U
               Grade 3                        with SIG:3 and FRE:U
               Grade 4                        with SIG:3 and FRE:U
               Grade 5                        with SIG:3 and FRE:S
               Grade 6                        with SIG:3 and FRE:S
QUALITY:       HARSH                          with SIG:2 and FRE:U
               SOFT                           with SIG:3 and FRE:U
MANEUVER:      DECREASED with VALSALVA        with SIG:2 and FRE:U
               UNCHANGED with INSPIRATION     with SIG:3 and FRE:U
               DECREASED with INSPIRATION     with SIG:3 and FRE:U
PATTERN:       EARLY-PEAK-SYSTOLIC            with SIG:3 and FRE:U
               MID-PEAK-SYSTOLIC              with SIG:3 and FRE:U
               LATE-PEAK-SYSTOLIC             with SIG:3 and FRE:U
HEART-SOUNDS:  A2-DECREASED                   with SIG:2 and FRE:U
               SYSTOLIC-CLICK                 with SIG:3 and FRE:U
LOCATION:      AORTIC                         with SIG:3 and FRE:A
               MITRAL                         with SIG:3 and FRE:U
RADIATION:     From AORTIC to LEFT-NECK       with SIG:2 and FRE:U
               From AORTIC to RIGHT-NECK      with SIG:2 and FRE:U
               From MITRAL to LEFT-NECK       with SIG:2 and FRE:U
               From MITRAL to RIGHT-NECK      with SIG:2 and FRE:U
```

Figure 5.1: A Murmur

5.1.2 The Score-Tables

A score-table holds the pair-wise *similarity scores* of all possible manifestation values of a particular murmur attribute. Mechanisms are provided for looking up, changing and adding the similarity score of two entries to the table. There is a score-table for each murmur attribute.

5.2 The Rules

Rules are structures with six fields: *number, name, nature, conds, acts* and *doc*. The *number* and *name* fields classify the rule. The *conds* is a list of conditions that have to be satisfied before the list of actions in the *acts* can be carried out. The *nature* field indicates whether the rule is of "terminating" or "non-terminating" nature (see **Section 3.2**). The *doc*, finally, holds a explanatory description of the rule.

Each of the *conds, acts* and *doc* fields of a rule can be defined using the rulepar constructor. This constructor takes in two arguments: Identification name of the field, e.g. m2-conds, m2-acts, etc.; and a list of events to be carried out. In the *conds* field, this list represents the list of conditions to be satisfied, while in the *acts* and *doc* fields, a list of actions to be carried out. The rule itself, in turn, can be constructed using the defrule operator. Figure 5.2 shows the construction of a rule used in the test for further differentiability of the abnormalities (see **Chapter 4**).

As we have mentioned in **Section 3.2**, rules are structures used to capture judgemental knowledge in the system. Since this kind of knowledge can be differentiated into various categories, structures called rule-tables are devised to hold all the rules serving a

```
;;; MAN2
;;; MR-IHSS

(rulepar man2-conds
        '((case-m 'idiopathic-hypertrophic-subaortic-stenosis
                  'mitral-regurgitation
                  'mitral-regurgitation-with-prolapsed-mitral-valve)
          (> 5 (score-diff (get-tscore 1 (interpret-outposition 'current))
                           (get-tscore 2 (interpret-outposition 'current))))
          (> 5 (score-diff (get-tscore 2 (interpret-outposition 'current))
                           (get-tscore 3 (interpret-outposition 'current))))
          (not (or (included? 'squatting (inmurmur-maneuver *current-inmurmur*))
                   (included? 'valsalva (inmurmur-maneuver *current-inmurmur*))))))

(rulepar man-squval
        '((format terminal-io"~&
  The diagnoses of the ~S set of findings can be further differentiated
  by conducting squatting and-or valsalva maneuver(s)."
                    (get-inposition 'current))
           (add-man-handler)))

(rulepar man2-doc
        "IF the first three ranked murmurs include MR, MRPMV and IHSS,
    AND the respective score-difference between
        the first and second, and the second and third
        ranked murmurs are less than 5,
    AND the input does not include information on SQUATTING
        or VALSALVA
   THEN ask if additional input information on SQUATTING
      or VALSALVA is available")

(defrule man2
        mr-ihss
        term                                  ; Terminating rule
        man2-conds
        man-squval
        man2-doc)
```

Figure 5.2: Construction Of A Rule.

particular purpose. Each rule-table has a dispatcher, which, when invoked, executes all rules in the table successively. The dispatcher can recognize the nature of the executed rules. If a terminating rule is executed, i.e. its *acts* field is invoked, the dispatcher terminates itself, preventing any other rules in the same rule-table from being invoked. This is true, as mentioned in **Section 3.2**, when the execution of a particular rule, whose *conds* field usually calls for a test or checking, renders any other test or checking within the same task category unnecessary. Otherwise, execution of a rule is simply allowed as side effects.

Currently four sets of rules, i.e. four rule-tables are used in the system. These include rules that check for input consistency, evidences for classification and confirmation of the causal abnormalities, further differentiability of the highest ranked abnormalities, and those used in the implementation of the query system in the user-interface. All of these rules are terminating except for those used to check for classification and confirmation evidences.

5.3 Decision History Records

Evidence-records and score-templates are implemented to keep track of the decision-history of the diagnostic procedure. Information contained in these records is used to provide explanations and justifications of the reasoning process.

5.3.1 Evidence-Records

An evidence-record is a structure which keeps track of the *status* of a particular abnormality, i.e. whether it is being considered for diagnosis, confirmed, or not being considered for diagnosis with respect to a certain set of input information. The *evidences* that brought about this status, together with the *evidence-scores* which represent the extent of confirmation for the abnormality provided by these evidences, if any, are also registered in the record. Figure 5.3 shows a representation of the evidence-record.

```
MURMUR-NAME:          AORTIC-STENOSIS
STATUS:               Dormant
POSITIVE-EVIDENCE:    E1
NEGATIVE-EVIDENCE:
EVIDENCE-SCORE:       10
```

Figure 5.3: An Evidence-Record.

```
MURMUR-NAME:          AORTIC-STENOSIS
INTENSITY-SCORE:      4.20
QUALITY-SCORE:        10.00
MANEUVER-SCORE:       0.00
PATTERN-SCORE:        7.60
SOUNDS-SCORE:         0.00
LOCATION-SCORE:       10.00
RADIATION-SCORE:      4.20
EVIDENCE-SCORE:       10.00
TOTAL-SCORE:          46.00
```

Figure 5.4: A Score-Template.

5.3.2 Score-Templates

A score-template is a structure with slots to hold the *similarity scores* for each attribute of an abnormality, together with its total *evidence-score*, if any. Filling in a score-template involves looking up and weighing the similarity scores of manifestations of each murmur attribute with respect to its input counterpart from the corresponding score-table, and the overall evidence-score from the evidence-record of the corresponding abnormality (see **Chapter 4**). Mechanisms are also provided for sorting the list of score-templates into descending overall score order, and extracting specific information from the templates. Figure 5.4 shows a representation of the score-template.

5.4 The User-Interface

Implementation of the user-interface of the system relies heavily on the window system facilities of the Symbolics Lisp Machine. The interface is a combination of three parts: display framework, supporting structures and io-handlers.

5.4.1 Display Framework

When the system is invoked, the screen display as shown in Figure 5.5 is set up. The display is built on underlying window system facilities of the hardware and is divided into the following parts: a command-menu-pane where a top-level operation can be selected and invoked, a picture-pane where manifestation descriptions for each murmur attribute are displayed and chosen from, an interaction-pane where communication between the system and the user takes place, and, which are not shown in the figure, a diagnose-menu-pane where operations facilitating the input process can be selected and invoked, and an information-pane where useful information is displayed during interaction. Pop-up menus and displays prompting for further input information appear at appropriate times. The interaction-pane is extended during query sessions to allow larger display

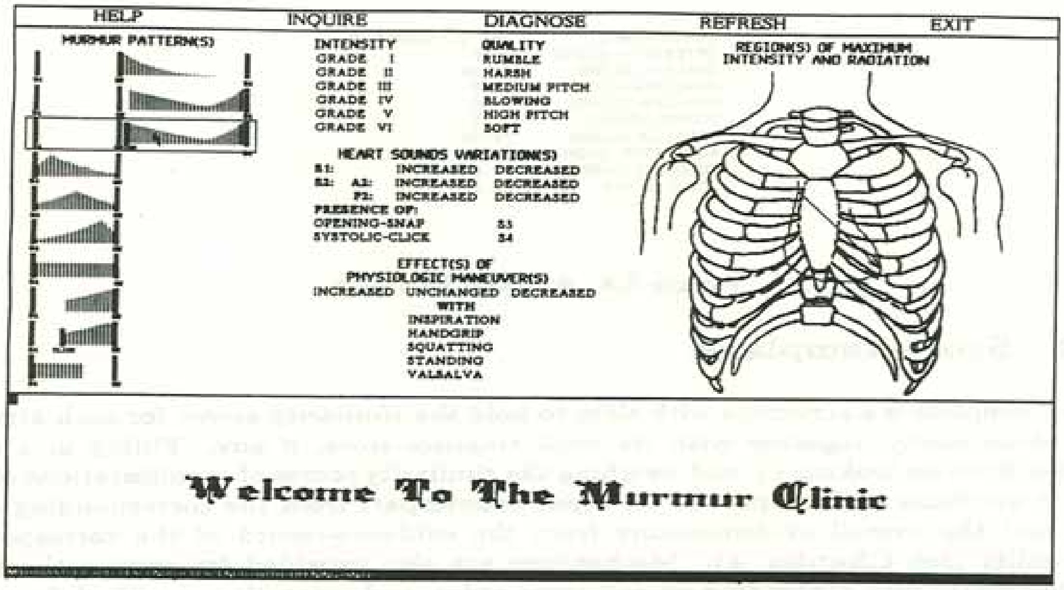

Figure 5.5: The Input Panel With Standard Configuration.

area. Examples of the different configurations and features of the input panel will be presented in the sample session in Chapter 6.

5.4.2 Supporting Structures

Regions and arrows are the two supporting structures that enable and aid the selection of manifestation values from the screen.

A region is defined as part of the screen whose *contents*, usually denoted by a name, can be extracted when selected. Each of the displayed manifestation descriptions of the murmur attributes is associated with a region, whose *contents* corresponds to the name of the option. The hollow rectangular box in Figure 5.5 is a region as seen on the user-interface. The corresponding name of the region is displayed on the black line [1] at bottom of the panel as shown.

An arrow is a structure, as shown also in Figure 5.5, which is displayed on the screen when the user inputs a radiation direction of a murmur.

[1]The *who-line*.

5.4.3 IO-Handlers

Input from the user and output from the system are handled by the io-handlers (input/output handlers). These handlers dispatch the input and output information to appropriate temporary buffers for further manipulation by the top-level command drivers and the internal scoring mechanisms. Consequently, the command drivers can be implemented with great flexibility since they do not have to keep track of all the operations being done on the input and output information. The DIAGNOSE command driver, for instance, is written in close accordance with the reasoning algorithm we adopted. Changes to the algorithm can be made on the abstract level without having to worry about the lower level details, for they will be handled by the io-handlers. Conversely, modifying the lower level operations will not affect the command driver algorithm.

5.5 Top Level Command Drivers

There are four top-level command drivers in the system, each corresponding to an operation invocable from the menu-pane that functions as follows:

HELP	Gives a brief introduction to the system.
INQUIRE	Returns the name(s), if any, of the manifestation description(s) chosen by the user from the screen.
DIAGNOSE	Invokes the diagnostic procedure.
REFRESH	Refreshes the screen display.
EXIT	Exits the Murmur Clinic.

When invoked, io-handlers and other procedures are called in the DIAGNOSE command driver to process the input and output information accordingly.

Chapter 6

A Sample Session

We now present an example of the invocation of the DIAGNOSE operation of the system.

The DIAGNOSE operation is invoked by clicking on the choice shown on the command-menu-pane. A display set-up as shown in Figure 6.1 will be initiated. The new configuration consists of the command-menu-pane and the picture-pane which are unchanged, a diagnose-menu-pane where operations can be invoke to facilitate the input process, an interaction-pane where the system communicates with the user, and an information-pane where the input given by the user will be monitored.

Input findings can now be entered by clicking on the different descriptive options shown on the panel. Each murmur is entered separately, with the diastolic and systolic components being considered as different murmurs. However, it is sufficient to enter the *heart sounds* findings only once, as they will be updated automatically for all the murmur descriptions entered during the session.

As each finding is chosen, it will be reflected on the information-pane. Moreover, when a *radiation* finding is entered, an arrow indicating the direction of the radiation will also be displayed on the picture-pane. A complete set of input findings is shown in Figure 6.2.

To begin description of an additional murmur, click on the ADD-MURMUR option on the diagnose-menu-pane and the information-pane will be refreshed. Findings are then input as before.

At any point of the input session, if the user wishes to change any information he has given so far, he can click on the diagnose menu-pane option CHANGE-MURMUR. A pop-up menu as shown in Figure 6.3 will be invoked, prompting the user to enter the set of input findings that he wishes to modify.

If the user indicates that he wishes to modify the current set of input findings, a second pop-up menu, corresponding to the findings indicated on the information-pane, will be invoked as shown in Figure 6.4. The user then indicates the particular input finding(s) that he wishes to modify by clicking left on the corresponding finding(s), and

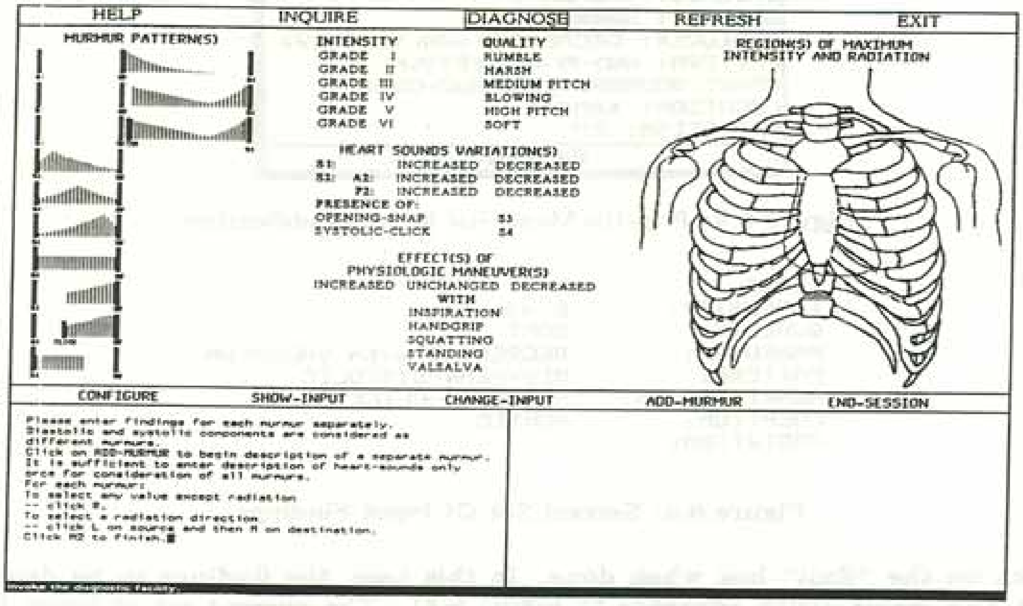

Figure 6.1: Initiation Of The DIAGNOSE operation.

```
INPUT
INTENSITY:      Grade 2
QUALITY:        BLOWING
MANEUVER:
PATTERN:        MID-DIASTOLIC-DECRESCENDO-CRESCENDO
HEART-SOUNDS:   A2-DECREASED
LOCATION:       AORTIC
RADIATION:      From AORTIC to MITRAL
```

Figure 6.2: First Set Of Input Findings.

Which set of input findings?
CURRENT
FIRST

Figure 6.3: Pop-Up Menu For Input Sets Selection.

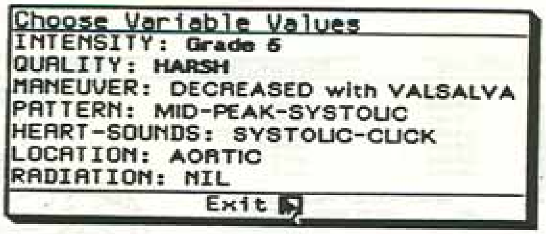

Figure 6.4: Pop-Up Menu For Input Modification.

```
INPUT
INTENSITY:        Grade 4
QUALITY:          SOFT
MANEUVER:         DECREASED with VALSALVA
PATTERN:          MID-PEAK-SYSTOLIC
HEART-SOUNDS:     SYSTOLIC-CLICK
LOCATION:         AORTIC
RADIATION:
```

Figure 6.5: Second Set Of Input Findings.

then click on the "Exit" box when done. In this case, the findings to be deleted are *Grade 5* and *Harsh* (with reference to Figure 6.4). The current set of input findings will be updated correspondingly. The user can now continue, if he wishes, to enter additional findings for the murmur.

If the user wishes to modify the description of a previous, other than the current set of input findings, the system will first check to see if the current set of findings is complete for a diagnosis to be carried out. If not, the user is asked to finish the current description before making any modifications. Otherwise, the current set of findings is saved [1], and the indicated set of findings will be displayed on the information-pane. A similar modification process to the one mentioned will be carried out.

The second set of input findings consists of parameters shown in Figure 6.5.

Throughout the input session (and also after the diagnosis is made, as can be seen later), the user can also inspect the different sets, if any, of the findings that are given so far. This is done by clicking on the SHOW-INPUT option on the diagnose-menu-pane. A pop-up menu similar to that in Figure 6.3 will be invoked, prompting the user to indicate which set of findings description he wishes to see. The corresponding set of findings chosen will then be displayed on the information-pane.

Whenever a new finding is added to the description, the whole set of findings will be checked for consistency. Should any indiscrepancy be detected, the user would be warned or prompted for a modification. For example, if we now enter a second *location* value: *3rd-left-interspace* to the current set of input, inconsistency will be detected, and

[1] Provided that no inconsistency is detected prior to saving.

the message as shown in Figure 6.6 will result. If we answer "YES" to the prompt, a new environment with *3rd-left-interspace* as the *location* entry will be set up for the description of a new murmur. Additional findings can then be entered as mentioned.

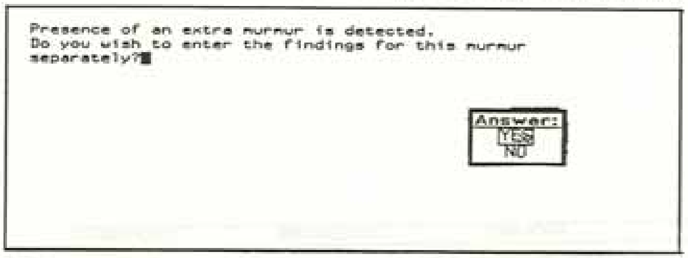

Figure 6.6: A Message Indicating Inconsistency.

The third set of input findings entered is shown in Figure 6.7.

Now if we click twice on the middle-button of the mouse, a message prompting for confirmation of completion will appear. If we answer "NO" to the prompt, nothing will happen and input modification(s) and/or addition(s) can continue to be made; otherwise, the diagnostic process will begin and the results will be displayed as in Figure 6.8. The score accompanying each decision in the ranked lists indicates how consistent the corresponding decision is to the particular set of findings as compared to the other decisions. The higher the score, the more consistent the decision.

In this case the diagnostic process detected further differentiability among the top few choices corresponding to the third set of findings. Hence the user is asked whether or not he wishes to enter additional information to enable a more accurate diagnosis. If we answer "YES", the standard configuration used in the input session is set up, and we enter the additional information on the *valsalva* maneuver.

After we declare completion of input as described earlier, the diagnostic process is again invoked, and a new set of results is displayed.

```
INPUT
INTENSITY:      Grade 3
QUALITY:        HARSH
MANEUVER:
PATTERN:        HOLOSYSTOLIC
HEART-SOUNDS:   S4
LOCATION:       3RD-LEFT-INTERSPACE
RADIATION:      From 3RD-LEFT-INTERSPACE to PRECORDIUM
```

Figure 6.7: Third Set Of Input Findings.

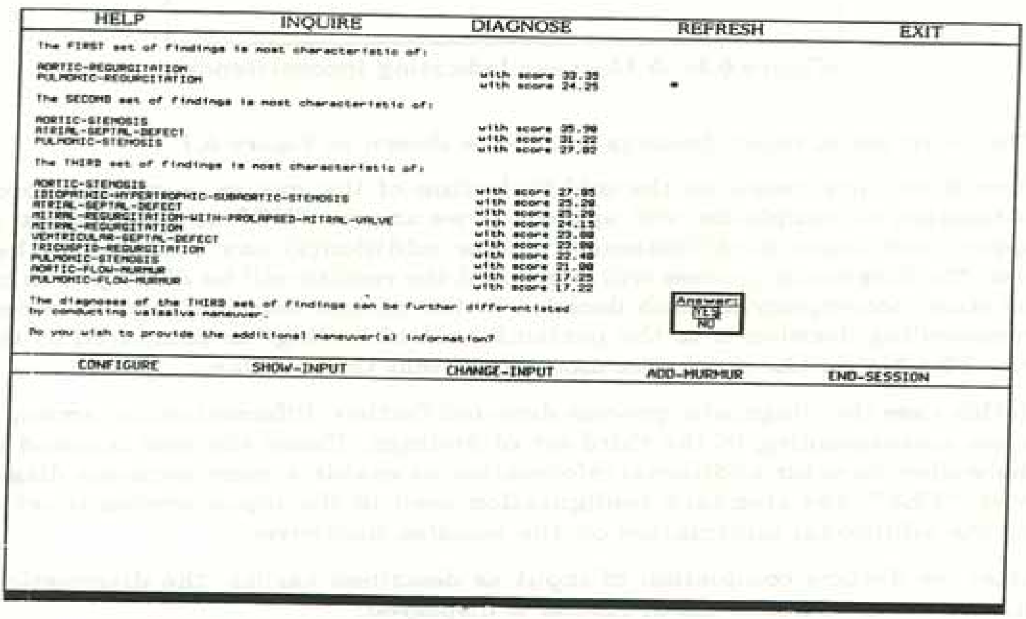

Figure 6.8: Display Of The Diagnostic Decisions.

If we choose not to provide additional information when further differentiability is detected, the query handler is invoked. During the query session, questions pertaining to the decisions made can be answered. Formats of the answerable questions can be found by typing "H" to the query-prompt , and the list will be displayed as shown below.

```
Question? ( H for options, Q to quit)
H

The following questions can be answered:

WHY
WHY <murmur-name> or <position> [<order of ranked-list>]
WHY NOT <murmur-name> or <position> [<order of ranked-list>]
WHY <murmur-name> or <position> OVER <murmur-name> or <position> [<order of ranked-list>]
DESCRIBE <murmur-name> or <rule-number> or <rule-name>

Type:
H -- To redisplay this message.
L -- To redisplay the decision list(s).
N -- To display the list of murmur names and their abbreviations.
M -- To enable diagnose-menu choose.
Q -- To end the current DIAGNOSE session.

Note: <murmur-name> can be in full as listed or abrreviated.

      <position> is the position of the murmur in the result list and an integer.

      <order of ranked-list> is an integer indicating the decisions on the correspond
                             set of input findings. Default to 1 if no value is given.

      <rule-number> is the number of a particular rule in the form of, e.g. c2.

      <rule-name> is the name of a particular rule.
```

At any point during the query session, should the user wish to invoke the operations on the diagnose-menu-pane, he can type "M" to the query-prompt and then click on the corresponding operation name. A less tedious way to invoke these operations is to be devised.

We now give an illustration of the answers to each kind of questions that can be asked during the query session. Note that the corresponding set of input findings, if any, is actually displayed on the information-pane below the diagnose menu-pane (with reference to Figure 6.8). Note also that values of the *heart-sounds* attribute of each set of input findings have been updated with entries in other sets as specified.

A query of the WHY format is shown below. The answer shows the scores of each attribute of the first ranked abnormality with respect to the first set of findings. If the finding for a particular attribute is not entered, e.g. the *maneuver* attribute in this case, the score for the attribute defaults to 0; and it is not shown. The SIG and FRE values of each attribute value correspond to the *discriminating significance* and *occurrence frequency* as mentioned in Chapter 4.

```
Question? ( H for options, Q to quit)
WHY

[ i.e. WHY is AORTIC-REGURGITATION ranked first corresponding to the first set of findings? ]

AORTIC-REGURGITATION
is ranked first because it has the following properties:

INTENSITY:      Grade 1                                  with SIG:3 and FRE:U
                Grade 2                                  with SIG:3 and FRE:U
                Grade 3                                  with SIG:3 and FRE:U
with SCORE      4.20

QUALITY:        BLOWING                                  with SIG:2 and FRE:U
                HIGH                                     with SIG:3 and FRE:U
with SCORE      7.00

PATTERN:        EARLY-DIASTOLIC-DECRESCENDO              with SIG:2 and FRE:A
                EARLY-PEAK-SYSTOLIC                      with SIG:3 and FRE:S
with SCORE      3.95

HEART-SOUNDS:   A2-DECREASED                             with SIG:2 and FRE:U
                A2-INCREASED                             with SIG:3 and FRE:S
                S3                                       with SIG:3 and FRE:S
with SCORE      7.00

LOCATION:       AORTIC                                   with SIG:2 and FRE:U
                PULMONIC                                 with SIG:3 and FRE:S
with SCORE      7.00

RADIATION:      From AORTIC to MITRAL                    with SIG:3 and FRE:U
                From AORTIC to RIGHT-STERNAL-BORDER      with SIG:3 and FRE:S
                From PULMONIC to MITRAL                  with SIG:3 and FRE:S
with SCORE      4.20

The TOTAL-SCORE is 33.35.

INPUT
INTENSITY:      Grade 2
QUALITY:        BLOWING
MANEUVER:
PATTERN:        MID-DIASTOLIC-DECRESCENDO-CRESCENDO
HEART-SOUNDS:   A2-DECREASED
                SYSTOLIC-CLICK
                S4
LOCATION:       AORTIC
RADIATION:      From AORTIC to MITRAL
```

A query of the format: WHY [ABNORMALITY] [LIST-ORDER] is shown below. As mentioned above, if the optional [list-order] is not provided, the reference set of input findings will default to the first set.

```
Question? ( H for options, Q to quit)
WHY 2 2

[ i.e. WHY is ATRIAL-SEPTAL-DEFECT ranked as it is in the SECOND list? ]

ATRIAL-SEPTAL-DEFECT
is ranked as it is because it has the following properties:

INTENSITY:        Grade 1                           with SIG:3 and FRE:U
                  Grade 2                           with SIG:3 and FRE:U
                  Grade 3                           with SIG:3 and FRE:U
                  Grade 4                           with SIG:3 and FRE:S
with SCORE        3.78

QUALITY:          HARSH                             with SIG:3 and FRE:U
                  SOFT                              with SIG:3 and FRE:U
with SCORE        4.20

MANEUVER:         DECREASED with VALSALVA           with SIG:2 and FRE:U
                  INCREASED with INSPIRATION        with SIG:3 and FRE:U
                  INCREASED with SQUATTING          with SIG:3 and FRE:U
with SCORE        7.00

PATTERN:          EARLY-PEAK-SYSTOLIC               with SIG:3 and FRE:U
with SCORE        2.94

HEART-SOUNDS:     SYSTOLIC-CLICK                    with SIG:3 and FRE:U
                  S4                                with SIG:3 and FRE:U
with SCORE        8.40

LOCATION:         PULMONIC                          with SIG:2 and FRE:U
                  3RD-LEFT-INTERSPACE               with SIG:3 and FRE:U
with SCORE        4.90

The TOTAL-SCORE is 31.22.
Its score-difference from the first-ranked murmur is -4.68.


INPUT
INTENSITY:        Grade 4
QUALITY:          SOFT
MANEUVER:         DECREASED with VALSALVA
PATTERN:          MID-PEAK-SYSTOLIC
HEART-SOUNDS:     A2-DECREASED
                  SYSTOLIC-CLICK
                  S4
LOCATION:         AORTIC
RADIATION:
```

Answer to queries of the format WHY NOT [ABNORMALITY] [LIST-ORDER] is similar to the two formats above, except when the [ABNORMALITY] specified was not considered for diagnosis, as shown below.

```
Question? ( H for options, Q to quit)
WHY NOT PR 2

[ i.e. WHY is PULMONIC-REGURGITATION NOT ranked first in the SECOND list? ]

PULMONIC-REGURGITATION
was not considered for diagnosis because of
Rule PURE-SYSTOLIC


INPUT
INTENSITY:      Grade 4
QUALITY:        SOFT
MANEUVER:       DECREASED with VALSALVA
PATTERN:        MID-PEAK-SYSTOLIC
HEART-SOUNDS:   A2-DECREASED
                SYSTOLIC-CLICK
                S4
LOCATION:       AORTIC
RADIATION:
```

The rule PURE-SYSTOLIC above indicates that the findings imply the presence of only systolic murmur(s), but the abnormality in question is a diastolic murmur. Whenever an abnormality which is not considered for diagnosis is referenced in a query, the reason why it is excluded, usually captured in one or more rules, will be presented. Queries of the DESCRIBE [ABNORMALITY/RULE] can be made to find out the attributes of an abnormality or the contents of a rule as shown below.

- *Example 1*

```
Question? ( H for options, Q to quit)
DESCRIBE IHSS

[Interpretation of SIG:   1: Almost pathognomonic
                          2: Important, strongly weighted evidence for the correspoding
                             abnormality
                          3: Negligible, characteristic but not a discriminating finding
 Interpretation of FRE:   A: Always present
                          U: Usually present
                          S: Sometimes present [Rare] ]

IDIOPATHIC-HYPERTROPHIC-SUBAORTIC-STENOSIS
INTENSITY:       Grade 1                                with SIG:3 and FRE:U
                 Grade 2                                with SIG:3 and FRE:U
                 Grade 3                                with SIG:3 and FRE:U
                 Grade 4                                with SIG:3 and FRE:U
                 Grade 5                                with SIG:3 and FRE:S
                 Grade 6                                with SIG:3 and FRE:S
QUALITY:         HARSH                                  with SIG:3 and FRE:U
MANEUVER:        DECREASED with SQUATTING               with SIG:2 and FRE:U
                 INCREASED with STANDING                with SIG:2 and FRE:U
                 INCREASED with VALSALVA                with SIG:2 and FRE:U
PATTERN:         LATE-PEAK-SYSTOLIC                     with SIG:3 and FRE:U
                 HOLOSYSTOLIC                           with SIG:3 and FRE:U
HEART-SOUNDS:    S4                                     with SIG:3 and FRE:U
LOCATION:        3RD-LEFT-INTERSPACE                    with SIG:3 and FRE:U
RADIATION:       From 3RD-LEFT-INTERSPACE to PRECORDIUM     with SIG:3 and FRE:U
```

- *Example 2*

```
Question? ( H for options, Q to quit)
DESCRIBE PURE-SYSTOLIC

RULE-NUMBER:    E4
RULE-NAME:      PURE-SYSTOLIC
RULE-NATURE:    NON-TERM
RULE-DOCUMENTATION:
"IF input pattern(s) is not empty
   AND does not include EARLY-DIASTOLIC-DECRESCENDO
   AND MID-DIASTOLIC-DECRESCENDO-CRESCENDO
   AND MID-DIASTOLIC-DECRESCENDO-CRESCENDO-WITH-OPENING-SNAP
THEN move AR, PR, MS, MFM, TS and TFM to *dormant-dbase*"
```

A query of the format:
WHY [ABNORMALITY1] OVER [ABNORMALITY2] [LIST-ORDER] is given below.

```
Question? ( H for options, Q to quit)
WHY 1 OVER PR

[ i.e. WHY is AORTIC-REGURGITATION
ranked OVER
PULMONIC-REGURGITATION
in the FIRST list?]

[Interpretation of SIG:     1: Almost pathognomonic
                            2: Important, strongly weighted evidence for the correspoding
                               abnormality
                            3: Negligible, characteristic but not a discriminating finding
 Interpretation of FRE:     A: Always present
                            U: Usually present
                            S: Sometimes present [Rare] ]

The following attribute(s) of AORTIC-REGURGITATION is(are) more consistent with the input
than PULMONIC-REGURGITATION:

HEART-SOUNDS:   A2-DECREASED                      with SIG:2 and FRE:U
                A2-INCREASED                      with SIG:3 and FRE:S
                S3                                with SIG:3 and FRE:S
with SCORE      7.00
vs.
HEART-SOUNDS:   P2-INCREASED                      with SIG:2 and FRE:U
with SCORE      0.00

LOCATION:       AORTIC                            with SIG:2 and FRE:U
                PULMONIC                          with SIG:3 and FRE:S
with SCORE      7.00
vs.
LOCATION:       PULMONIC                          with SIG:2 and FRE:U
                AORTIC                            with SIG:3 and FRE:U
with SCORE      4.90

The following attribute(s) of PULMONIC-REGURGITATION is(are) more consistent with the input
than AORTIC-REGURGITATION:

None.

The total-score for AORTIC-REGURGITATION is 9.10 more than PULMONIC-REGURGITATION.


INPUT
INTENSITY:      Grade 2
QUALITY:        BLOWING
MANEUVER:
PATTERN:        MID-DIASTOLIC-DECRESCENDO-CRESCENDO
HEART-SOUNDS:   A2-DECREASED
                SYSTOLIC-CLICK
                S4
LOCATION:       AORTIC
RADIATION:      From AORTIC to MITRAL
```

And finally, whenever an invalid question is asked, the following complaint will be shown to the user.

```
Question? ( H for options, Q to quit)
HOW

Oops. I can't answer this.
Please retype your question.
Enter H for a list of the valid question formats.
Enter N for a list of the murmur names and their abbreviations.
```

Typing "Q" to the query-prompt will end the session invoked by the DIAGNOSE operation and return to the top level of the **Murmur Clinic**. The system can then be exited via the EXIT option on the command-menu-pane, if desired.

Chapter 7

Discussion

7.1 An Evaluation of Murmur Clinic

The approach we have adopted in the design and implementation of the Murmur Clinic system has eased the development process. Although not as compact as a conventional rule-based system, the modularity and transparency provided by our design enable straightforward improvement and enhancement of the system. Since the diagnostic procedure is explicitly encoded as discrete conceptual structures, the reasoning process of a physician can be closely simulated, and any deficiency in the formalism can be detected and corrected.

The evidence-records have allowed us to handle complications with multiple murmurs. While the current mechanism is not perfect, we are confident that we have provided the basic means to handle the major difficulties. We do not foresee refinement of the mechanism as a main obstacle in our future work.

During the development process, we have experimented with various applications of Artificial Intelligence techniques, and have learned much about the diagnostic reasoning process of expert physicians. We have built an operational system with a certain amount of expertise in auscultation interpretation. There still exist a number of serious limitations which prevent the system from being sufficiently competent from a professional point of view. We present a few of these problems here.

One of the major deficiencies of our system is the incompleteness of the knowledge base. The encoded causal abnormalities in the database, input manifestation description options, and the collection of judgemental evidence represented in the system are all by no means exhaustive, although the most representative cases are included. Even if we consider the amount of knowledge captured in the system as adequate, we do not have enough skills to distinguish effectively the most important information on the subject from the less important ones.

A second major deficiency of the system lies in the diagnostic algorithm. Although we have captured some of the most important details involved in the reasoning process of a physician, formalization of the algorithm is far from complete. In dealing with classification and confirmation of the abnormalities, for example, the notion of "negative evidence" is captured only in the meaning of absolute inconsistency, i.e. presence of a negative evidence would exclude a particular abnormality from diagnostic consideration, without taking into account how "negative" the evidence is. A formal model which can handle the positive and negative evidences involved in a separate manner consistent to the human cognitive reasoning process would need to be established, or else the algorithm would remain "fuzzy" in nature. Furthermore, at present the system is only capable of performing tentative diagnosis. A real expert, however, performs tentative diagnosis only when a definitive one cannot be made, and knows when he cannot make a diagnosis at all. Our algorithm has not been able to capture these capabilities so far.

The performance of the system is further handicapped by the fact that the user may not be able to provide exact information on the findings for an accurate diagnosis because of the limited manifestation description options provided for him.

And last but not the least, there is currently no way, except through the programmer, that the user can change or update the knowledge captured in the system.

7.2 Future Directions

It is obvious that a considerable amount of work still has to be done on the system to achieve the level of expertise we desire. Some of the major issues we envision requiring substantial improvements are discussed below.

The first principle area that future work on the system should concentrate on is further formalization of the reasoning algorithm, especially in the parts dealing with interpretation of the evidences, and uncertainty management in general. In particular, the quantitative model we have adopted should be rigorously tested for its feasibility, and appropriate changes have to be introduced if necessary. The current diagnostic procedure should also be upgraded to perform definitive diagnosis after some extensive testing of the algorithm. Tentative diagnosis should be made only when a definitive one cannot be made, and the system should be aware when it is unable to produce a diagnosis.

A second area is the refinement and expansion of the knowledge base. This would require extensive work done on interrogating human experts, which in turn involves a lot of knowledge engineering techniques.

Another area that can be improved is the accuracy of the input information. Problems in this area arise because, as mentioned in the previous subsection, only a limited number of options are provided for the user to describe the auscultatory findings. This is particularly true for the murmur *pattern* attribute. As the display area on the screen is limited, a feasible alternative to trying to display all the options at the same time is to provide the options in different categories sequentially. As for the *pattern* attribute, more accurate input can be obtained by allowing the user to tailor his input, for example, he can choose to place the opening snap or the peak of a murmur at a particular position. In this case, however, we would have to compromise the accuracy of the input information with a much extended knowledge base to keep track of minute details. Moreover, a feature that would allow the user to modify or update the knowledge captured in the system should be incorporated. This feature, however, should be carefully implemented so as not to cause any inconsistency that the system maintainer is unaware of.

We mentioned earlier in this paper that the system does not require any pathophysiologic knowledge underlying the manifestations. This is actually not exactly true. We, as human beings, always make better judgement when we have more information at our disposal. This is also true for the system. Pathophysiologic knowledge is needed sometimes for making accurate differentiation of similar manifestations ascribed to different abnormalities. This kind of knowledge has to be encoded explicitly as evidence rules for the system to recognize. Although as much judgemental knowledge as possible can be encoded into the current system, it would be much more efficient to incorporate a pathophysiologic model, to which the diagnostic procedure can turn to when the encoded judgemental knowledge is insufficient for making a decision. By arranging the knowledge base into a hierarchy, the system is shielded from dealing with detailed pathophysiologic knowledge unnecessarily, but is assured of its availability.

And finally, an interface is needed for coupling the **Murmur Clinic** system to the rest of the Heart Failure Program. At the moment, however, we are yet to come up with an approach in which the coupling can be done.

Chapter 8

Conclusions

In this paper we have discussed the design and implementation of the Murmur Clinic, a cardiac auscultation expert system. A description of the problem domain and the motivations behind our approach is also presented. At present we have an operational system which shows a promising competence in its performance. We have identified the major areas in which improvements are required, and the refinement process is actually underway at the completion of this paper.

The **Murmur Clinic** system is currently on-line for demonstration and testing. The system can be invoked on any Symbolics 3600 Lisp Machines by the following procedure:

1. Load the file "Z:>LEONG>MURMUR>MAKE-MURMUR-CLINIC" [1].
2. Type the expression (make-murmur-clinic) to the *Lisp Listener*, and wait for the source files to be loaded.
3. Type the expression (visit-murmur-clinic) to the *Lisp Listener* to invoke the program.
4. Enjoy your visit!

[1] This file is mounted on server MIT-ZERMATT, net-address: ZERMATT.LCS.MIT.EDU.

Appendix A

THE ALGORITHM

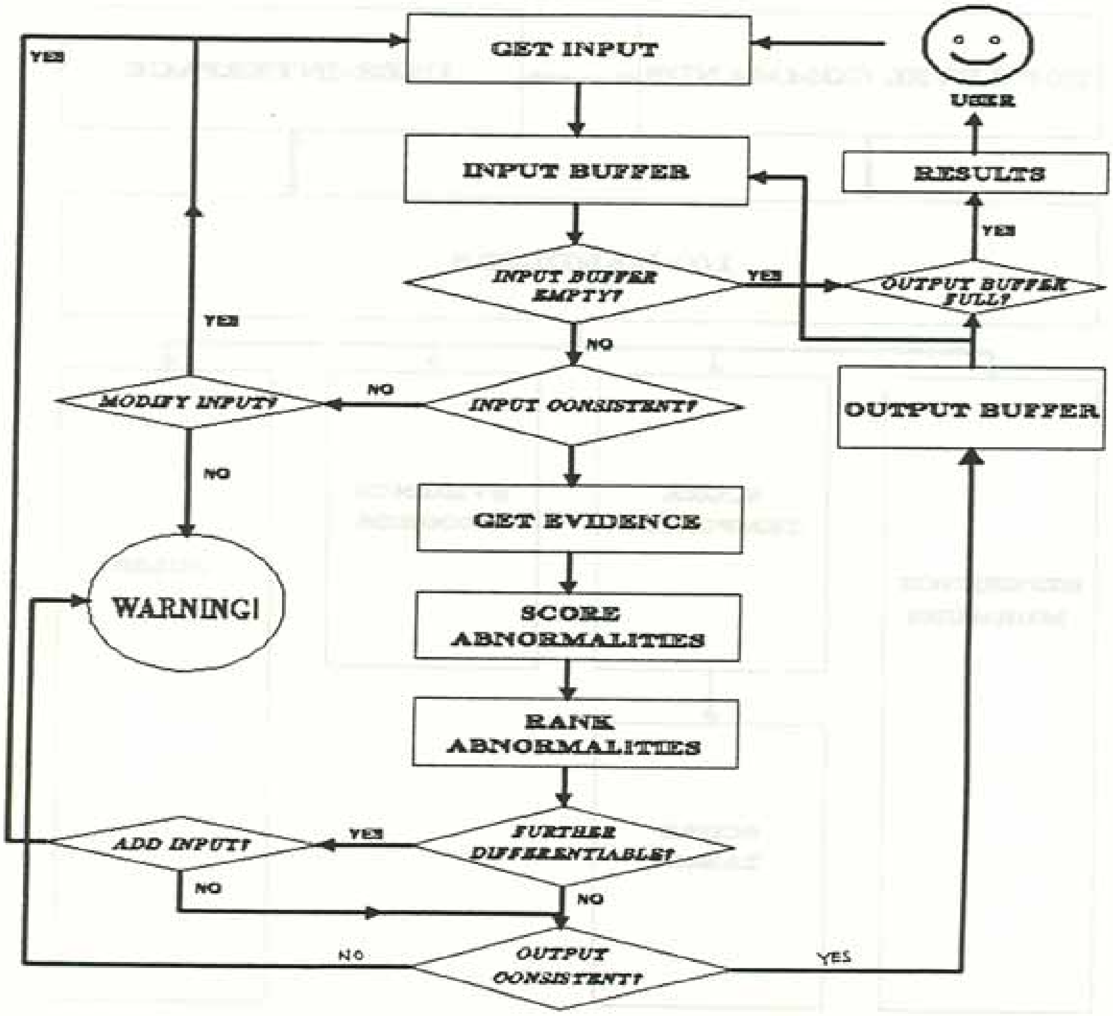

Appendix B

THE SYSTEM LAYOUT

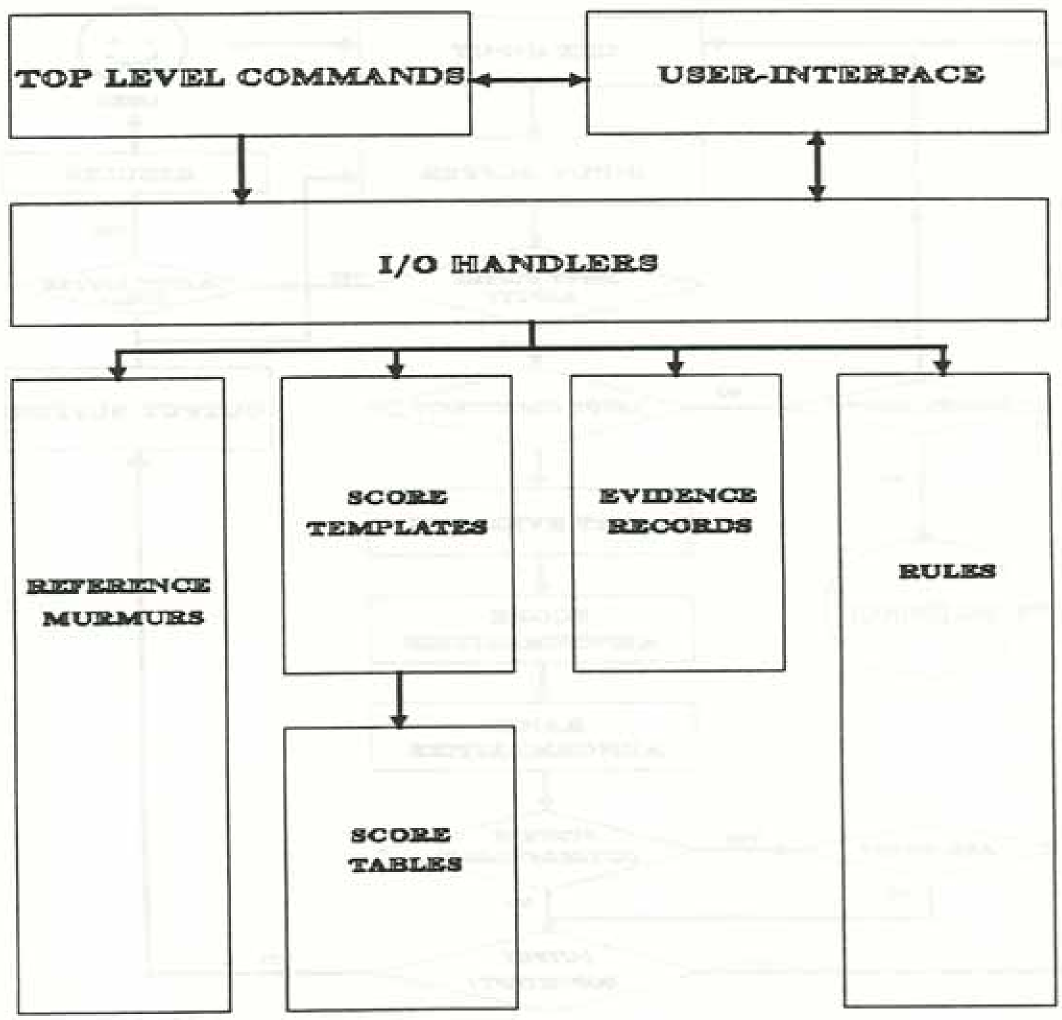

Appendix C

THE REFERENCE MURMURS

PULMONIC-REGURGITATION		
INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
QUALITY:	BLOWING	with SIG:2 and FRE:U
	HIGH	with SIG:2 and FRE:U
	RUMBLE	with SIG:3 and FRE:S
MANEUVER:	INCREASED with INSPIRATION	with SIG:2 and FRE:U
	UNCHANGED with INSPIRATION	with SIG:3 and FRE:S
PATTERN:	EARLY-DIASTOLIC-DECRESCENDO	with SIG:2 and FRE:A
HEART-SOUNDS:	P2-INCREASED	with SIG:2 and FRE:U
LOCATION:	PULMONIC	with SIG:2 and FRE:U
	AORTIC	with SIG:3 and FRE:U
RADIATION:	From PULMONIC to LEFT-STERNAL-BORDER	with SIG:2 and FRE:U
	From PULMONIC to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to LEFT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to MITRAL	with SIG:3 and FRE:U
PULMONIC-STENOSIS		
INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
QUALITY:	HARSH	with SIG:3 and FRE:U
	MEDIUM	with SIG:3 and FRE:U
MANEUVER:	DECREASED with VALSALVA	with SIG:2 and FRE:U
	INCREASED with INSPIRATION	with SIG:2 and FRE:U
	UNCHANGED with INSPIRATION	with SIG:3 and FRE:S
PATTERN:	MID-PEAK-SYSTOLIC	with SIG:3 and FRE:U
	LATE-PEAK-SYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:	P2-DECREASED	with SIG:2 and FRE:U
	SYSTOLIC-CLICK	with SIG:3 and FRE:U
LOCATION:	PULMONIC	with SIG:2 and FRE:U
	AORTIC	with SIG:3 and FRE:U
	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From PULMONIC to LEFT-STERNAL-BORDER	with SIG:2 and FRE:U
	From PULMONIC to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to LEFT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:U
	From AORTIC to MITRAL	with SIG:3 and FRE:U
PULMONIC-FLOW-MURMUR		
INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
QUALITY:	HARSH	with SIG:3 and FRE:U
	MEDIUM	with SIG:3 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	INCREASED with INSPIRATION	with SIG:2 and FRE:U
PATTERN:	EARLY-PEAK-SYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:		
LOCATION:	PULMONIC	with SIG:3 and FRE:U
	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From 3RD-LEFT-INTERSPACE to LEFT-STERNAL-BORDER	with SIG:3 and FRE:U

AORTIC-REGURGITATION

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
QUALITY:	BLOWING	with SIG:2 and FRE:U
	HIGH	with SIG:3 and FRE:U
MANEUVER:		
PATTERN:	EARLY-DIASTOLIC-DECRESCENDO	with SIG:2 and FRE:A
	EARLY-PEAK-SYSTOLIC	with SIG:3 and FRE:S
HEART-SOUNDS:	A2-DECREASED	with SIG:2 and FRE:U
	A2-INCREASED	with SIG:3 and FRE:S
	S3	with SIG:3 and FRE:S
LOCATION:	AORTIC	with SIG:2 and FRE:U
	PULMONIC	with SIG:3 and FRE:S
RADIATION:	From AORTIC to MITRAL	with SIG:3 and FRE:U
	From AORTIC to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:S
	From PULMONIC to MITRAL	with SIG:3 and FRE:S

AORTIC-STENOSIS

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:2 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	DECREASED with VALSALVA	with SIG:2 and FRE:U
	UNCHANGED with INSPIRATION	with SIG:3 and FRE:U
	DECREASED with INSPIRATION	with SIG:3 and FRE:U
PATTERN:	EARLY-PEAK-SYSTOLIC	with SIG:3 and FRE:U
	MID-PEAK-SYSTOLIC	with SIG:3 and FRE:U
	LATE-PEAK-SYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:	A2-DECREASED	with SIG:2 and FRE:U
	SYSTOLIC-CLICK	with SIG:3 and FRE:U
LOCATION:	AORTIC	with SIG:3 and FRE:A
	MITRAL	with SIG:3 and FRE:U
RADIATION:	From AORTIC to LEFT-NECK	with SIG:2 and FRE:U
	From AORTIC to RIGHT-NECK	with SIG:2 and FRE:U
	From MITRAL to LEFT-NECK	with SIG:2 and FRE:U
	From MITRAL to RIGHT-NECK	with SIG:2 and FRE:U

AORTIC-FLOW-MURMUR

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:2 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:		
PATTERN:	EARLY-PEAK-SYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:		
LOCATION:	AORTIC	with SIG:2 and FRE:A
RADIATION:	From AORTIC to RIGHT-NECK	with SIG:2 and FRE:U

IDIOPATHIC-HYPERTROPHIC-SUBAORTIC-STENOSIS

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:3 and FRE:U
MANEUVER:	DECREASED with SQUATTING	with SIG:2 and FRE:U
	INCREASED with STANDING	with SIG:2 and FRE:U
	INCREASED with VALSALVA	with SIG:2 and FRE:U
PATTERN:	LATE-PEAK-SYSTOLIC	with SIG:3 and FRE:U
	HOLOSYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:	S4	with SIG:3 and FRE:U
LOCATION:	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From 3RD-LEFT-INTERSPACE to PRECORDIUM	with SIG:3 and FRE:U

MITRAL-REGURGITATION

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:3 and FRE:U
	HIGH	with SIG:3 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	INCREASED with SQUATTING	with SIG:2 and FRE:U
	UNCHANGED with VALSALVA	with SIG:2 and FRE:U
	DECREASED with INSPIRATION	with SIG:2 and FRE:U
	UNCHANGED with INSPIRATION	with SIG:3 and FRE:U
PATTERN:	LATE-SYSTOLIC-CRESCENDO	with SIG:2 and FRE:U
	HOLOSYSTOLIC	with SIG:2 and FRE:U
HEART-SOUNDS:	S3	with SIG:3 and FRE:S
LOCATION:	MITRAL	with SIG:2 and FRE:U
	TRICUSPID	with SIG:3 and FRE:U
	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From MITRAL to LEFT-AXILLA	with SIG:2 and FRE:U
	From MITRAL to BACK	with SIG:2 and FRE:U
	From TRICUSPID to LEFT-MIDCLAVICULAR-LINE	with SIG:3 and FRE:U
	From 3RD-LEFT-INTERSPACE to PRECORDIUM	with SIG:3 and FRE:U

MITRAL-REGURGITATION-WITH-PROLAPSED-MITRAL-VALVE

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:3 and FRE:U
	HIGH	with SIG:3 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	INCREASED with SQUATTING	with SIG:2 and FRE:U
	UNCHANGED with VALSALVA	with SIG:2 and FRE:U
	DECREASED with INSPIRATION	with SIG:2 and FRE:U
	UNCHANGED with INSPIRATION	with SIG:3 and FRE:U
PATTERN:	LATE-SYSTOLIC-CRESCENDO-WITH-SYSTOLIC-CLICK	with SIG:1 and FRE:U
	LATE-SYSTOLIC-CRESCENDO	with SIG:2 and FRE:U
	HOLOSYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:	SYSTOLIC-CLICK	with SIG:2 and FRE:U
LOCATION:	MITRAL	with SIG:2 and FRE:U
RADIATION:	From MITRAL to LEFT-AXILLA	with SIG:2 and FRE:U

MITRAL-STENOSIS

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
QUALITY:	RUMBLE	with SIG:2 and FRE:A
MANEUVER:	DECREASED with INSPIRATION	with SIG:2 and FRE:U
PATTERN:	MID-DIASTOLIC-DECRESCENDO-CRESCENDO-WITH-OPENING-SNAP	with SIG:1 and FRE:U
	MID-DIASTOLIC-DECRESCENDO-CRESCENDO	with SIG:1 and FRE:U
	EARLY-DIASTOLIC-DECRESCENDO	with SIG:1 and FRE:S
HEART-SOUNDS:	OPENING-SNAP	with SIG:1 and FRE:U
	S1-INCREASED	with SIG:2 and FRE:U
	P2-INCREASED	with SIG:3 and FRE:U
LOCATION:	MITRAL	with SIG:3 and FRE:U
	TRICUSPID	with SIG:3 and FRE:U
RADIATION:		

MITRAL-FLOW-MURMUR

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
QUALITY:	RUMBLE	with SIG:2 and FRE:U
MANEUVER:		
PATTERN:	MID-DIASTOLIC-DECRESCENDO-CRESCENDO	with SIG:1 and FRE:U
HEART-SOUNDS:	S1-DECREASED	with SIG:3 and FRE:U
LOCATION:	MITRAL	with SIG:3 and FRE:U
RADIATION:		

TRICUSPID-REGURGITATION

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:3 and FRE:U
	SOFT	with SIG:3 and FRE:U
	HIGH	with SIG:3 and FRE:S
MANEUVER:	INCREASED with INSPIRATION	with SIG:2 and FRE:U
PATTERN:	HOLOSYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:		
LOCATION:	TRICUSPID	with SIG:2 and FRE:U
	MITRAL	with SIG:3 and FRE:U
	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:S
RADIATION:	From TRICUSPID to LEFT-MIDCLAVICULAR-LINE	with SIG:2 and FRE:U
	From TRICUSPID to RIGHT-STERNAL-BORDER	with SIG:3 and FRE:U
	From MITRAL to LEFT-AXILLA	with SIG:3 and FRE:U
	From 3RD-LEFT-INTERSPACE to PRECORDIUM	with SIG:3 and FRE:U

TRICUSPID-STENOSIS

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:S
QUALITY:	RUMBLE	with SIG:1 and FRE:U
MANEUVER:	INCREASED with INSPIRATION	with SIG:2 and FRE:U
PATTERN:	MID-DIASTOLIC-DECRESCENDO-CRESCENDO-WITH-OPENING-SNAP	with SIG:1 and FRE:U
	MID-DIASTOLIC-DECRESCENDO-CRESCENDO	with SIG:1 and FRE:U
HEART-SOUNDS:	OPENING-SNAP	with SIG:1 and FRE:U
LOCATION:	TRICUSPID	with SIG:2 and FRE:U
	MITRAL	with SIG:3 and FRE:U
RADIATION:		

TRICUSPID-FLOW-MURMUR

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:S
QUALITY:	RUMBLE	with SIG:1 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	INCREASED with INSPIRATION	with SIG:2 and FRE:U
PATTERN:	MID-DIASTOLIC-DECRESCENDO-CRESCENDO	with SIG:1 and FRE:U
HEART-SOUNDS:		
LOCATION:	TRICUSPID	with SIG:2 and FRE:U
RADIATION:		

ATRIAL-SEPTAL-DEFECT

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:3 and FRE:U
	SOFT	with SIG:3 and FRE:U
MANEUVER:	DECREASED with VALSALVA	with SIG:2 and FRE:U
	INCREASED with INSPIRATION	with SIG:3 and FRE:U
	INCREASED with SQUATTING	with SIG:3 and FRE:U
PATTERN:	EARLY-PEAK-SYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:	SYSTOLIC-CLICK	with SIG:3 and FRE:U
	S4	with SIG:3 and FRE:U
LOCATION:	PULMONIC	with SIG:2 and FRE:U
	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From PULMONIC to LEFT-NECK	with SIG:3 and FRE:U

VENTRICULAR-SEPTAL-DEFECT

INTENSITY:	Grade 1	with SIG:3 and FRE:U
	Grade 2	with SIG:3 and FRE:U
	Grade 3	with SIG:3 and FRE:U
	Grade 4	with SIG:3 and FRE:U
	Grade 5	with SIG:3 and FRE:S
	Grade 6	with SIG:3 and FRE:S
QUALITY:	HARSH	with SIG:2 and FRE:U
MANEUVER:		
PATTERN:	EARLY-SYSTOLIC	with SIG:2 and FRE:U
	HOLOSYSTOLIC	with SIG:3 and FRE:U
HEART-SOUNDS:		
LOCATION:	3RD-LEFT-INTERSPACE	with SIG:3 and FRE:U
RADIATION:	From 3RD-LEFT-INTERSPACE to RIGHT-STERNAL-BORDER	with SIG:2 and FRE:U
	From 3RD-LEFT-INTERSPACE to PRECORDIUM	with SIG:3 and FRE:U
	From 3RD-LEFT-INTERSPACE to AORTIC	with SIG:3 and FRE:U

Bibliography

[1] Szolovitz, P. "Artificial Intelligence and Medicine", in Szolovitz, P. (Ed), *Artificial Intelligence in Medicine*, pp 1-19, Westview, 1982.

[2] Come, P. C. (MD) "Auscultation of the Heart", in *The Cardiac Physical Examination*, pp 32-52.

[3] Bruns, D. L. "A General Theory of the Causes of Murmur in the Cardiovascular System", Am. J. Med. 27:360, 1959.

[4] Leatham, A. *Systolic Murmurs*, Circulation 17:601, 1958.

[5] Leonard, J. J. et al *Examination of the Heart, Part IV: Auscultation*, pp 38, American Heart Association, 1974.

[6] Ledley, R. S. Lusted, L. B. "Reasoning Foundations of Medical Diagnosis", in *Science*, volume 130, July 3, 1959, pp 9-21, American Association for the Advancement of Science, 1959.