

CLASSIFICATION OF ECG ST EVENTS
AS ISCHEMIC OR NON-ISCHEMIC
USING RECONSTRUCTED PHASE SPACES

by

Michael W. Zimmerman, B.S.

A Thesis submitted to the Faculty of the Graduate School,
Marquette University,
in Partial Fulfillment of
the Requirements for
the Degree of
Master of Science

Milwaukee, Wisconsin
May 2004

Acknowledgement

Special thanks go to my advisor, Dr. Richard Povinelli. Without his teachings, comments, and continued support my master's degree work would not have been possible. I have learned so much about conducting research, electrical engineering, and life in general. He has made my time at Marquette as a graduate student both educational and entertaining.

Special thanks also go to my parents for the countless hours they spent reading thesis drafts and giving advice. I would also like to thank the members of Knowledge and Information Discovery Laboratory who read paper drafts, provided assistance, and listened to many presentations.

Finally, I would like to extend a thank-you to the Milwaukee Foundation for their funding of my research through the Frank Rogers Bacon Research Assistantship.

Table of Contents

1	INTRODUCTION	1
1.1	PROBLEM DESCRIPTION	2
1.2	THESIS OUTLINE	2
2	CARDIAC BACKGROUND	4
2.1	CARDIAC FUNCTION	4
2.1.1	<i>Cardiac Mechanical System</i>	<i>4</i>
2.1.2	<i>Cardiac Electrical System</i>	<i>5</i>
2.2	ELECTROCARDIOGRAM	9
2.3	ISCHEMIA	9
2.4	ST EVENTS	11
3	RESEARCH BACKGROUND	14
3.1	CURRENT CLINICAL METHODS	14
3.1.1	<i>Coronary Angiography</i>	<i>14</i>
3.1.2	<i>Echocardiograms</i>	<i>16</i>
3.1.3	<i>Exercise Testing</i>	<i>17</i>
3.1.4	<i>ST Level Alerts</i>	<i>18</i>
3.2	STATE OF ECG RESEARCH IN DETECTING MYOCARDIAL ISCHEMIA	19
3.2.1	<i>Neural Network Classifier</i>	<i>19</i>
3.2.2	<i>Threshold Classifier</i>	<i>21</i>
3.3	CHALLENGE DESCRIPTION	22
4	NONLINEAR SIGNAL PROCESSING THEORY	24

4.1	RECONSTRUCTED PHASE SPACES	24
4.2	GAUSSIAN MIXTURE MODELS	28
4.3	NAÏVE BAYES CLASSIFIER	30
5	DATA AND METHODS	32
5.1	DATA.....	32
5.2	PRE-PROCESSING	34
5.3	CLASSIFICATION ALGORITHM	38
6	EXPERIMENTS	42
6.1	TEN-FOLD CROSS VALIDATION	42
6.2	PARAMETER DETERMINATION	43
6.3	BASELINES.....	47
6.4	NORMALIZING DATA EXPERIMENT	48
6.5	SUB-BAND FILTERING	50
6.6	THREE CLASS EXPERIMENT	54
6.7	MULTIPLE BEAT CLASSIFICATIONS.....	57
6.8	RESULTS COMPARISON	60
7	CONCLUSION	63
7.1	FUTURE WORK	64
	BIBLIOGRAPHY.....	66
	APPENDICES	71

Table of Figures

Figure 2.1 – The mechanical and electrical components of the heart	5
Figure 2.2 – Electrical signal with important points labeled	6
Figure 2.3 – Einthoven’s triangle showing how leads I, II, and III are recorded	9
Figure 2.4 - Example of ST deviation calculation.....	12
Figure 2.5 - Definition of ST Event	13
Figure 3.1 - Figure showing how catheter is inserted for coronary angiography.....	16
Figure 3.2 - Example of echocardiogram with ECG beat along bottom [28].....	17
Figure 3.3 - ECG waveforms that show indication of ischemia during exercise testing [31].....	18
Figure 3.4 - Flowchart of Neural Network Learning.....	20
Figure 3.5 - Flowchart of Neural Network classification.....	20
Figure 3.6 – Langley et al threshold classification algorithm	21
Figure 3.7 – Definition of ST Event	23
Figure 4.1 - An ST segment and T wave with 0 delay, 5 time step delay, and 10 time step delay.....	26
Figure 4.2 - The reconstructed phase space of a set of ST segments	26
Figure 4.3 - The Reconstructed Phase Space of Figure 4.2 with overlay of Gaussian Mixture Models	28
Figure 5.1 - Flowchart of pre-processing stage	34
Figure 5.2 - Description of how ST-T complex (segment in box) is extracted from all signals and several beats surrounding the labeled event	35
Figure 5.3 - Describes how J-point is found in presence of noise on leads one and two ..	36

Figure 5.4 - Description of how signals are filtered	37
Figure 5.5 - Flowchart of algorithm used in training.....	39
Figure 5.6 - Flowchart of algorithm used in classification.....	40
Figure 6.1 - Power spectral density of ECG ST segment in Training Set with frequency boundaries of 8.75Hz and 32.0Hz labeled	52
Figure 6.2 - Power spectral density of ECG ST segment in Training Set with frequency boundaries of 18.75Hz, 18.75Hz, and 32.0Hz labeled	53
Figure 6.3 - Flowchart of the three-class classifier.....	56
Figure 6.4 - Flowchart of the complete algorithm for classification.....	58
Figure 6.5 - Accuracy results for Test Set	59

Table of Tables

Table 2.1 - Stages of cardiac action with corresponding ECG representation.....	7
Table 3.1 - Classification results for Neural Network method applied to the European ST Database.....	21
Table 3.2 - Accuracy results for the Langley et al method.....	22
Table 5.1 - Breakdown of number of events for the Training and Test portions of the database.....	33
Table 6.1 - Accuracy results for combinations of dimension and number of mixtures.....	45
Table 6.2 - Accuracy results for combinations of dimension and number of mixtures.....	45
Table 6.3 - Accuracy results for different lag values.....	46
Table 6.4 - Accuracy results for the Langley et al method.....	47
Table 6.5 - Accuracy results for Neural Network algorithm	48
Table 6.6 – Ten-fold cross validation confusion matrix for normalizing experiment on Training Set	49
Table 6.7 – Ten-fold cross validation accuracy results for normalizing experiment on Training Set	50
Table 6.8 – Ten-fold cross validation confusion matrix for 2 boundary sub-band experiment on Training Set	53
Table 6.9 – Ten-fold cross validation confusion matrix for 3 boundary sub-band experiment on Training Set	54
Table 6.10 – Ten-fold cross validation accuracy results for sub-band experiments on Training Set	54

Table 6.11 – Ten-fold cross validation confusion matrix for 3 class experiment on Training Set	57
Table 6.12 – Ten-fold cross validation accuracy results for 3 class experiment on Training Set	57
Table 6.13 – Ten-fold cross validation confusion matrix for multiple beat experiment on Training Set	59
Table 6.14 – Confusion matrix for 16 beat experiment on Test Set	59
Table 6.15 – Accuracy results for Training Set.....	59
Table 6.16 - Summary of experimental results and baseline accuracies	60

1 Introduction

Heart disease is the leading cause of death in the United States. The most common form of heart disease is myocardial ischemia [1]. This is a condition in which portions of the cardiac tissue are deprived of oxygen. If the deprivation continues for extended periods of time, the cardiac tissue will begin to die. This tissue death is called infarction and is one of several conditions commonly known as a heart attack. Tissue that has died is no longer functional and diminishes the mechanical pumping function of the heart [2, 3].

Early detection of ischemia is crucial because, in most cases, the effects of myocardial ischemia are completely reversible if detected early enough [2, 4, 5]. Currently, many patients are not screened for ischemia due to low test accuracies or invasive and expensive procedures. A more accurate, less expensive, and less invasive screening tool is needed. General screening of patients is vital to preventing myocardial infarction, since ischemia can be present without exhibiting symptoms. Several possible methods have been proposed, which take advantage of easily recorded cardiac electrical signals, but these methods generally have low accuracies and report too many false alarms to be clinically practical.

This thesis presents a novel approach for diagnosing a cardiac patient with myocardial ischemia using recordings of the electrical signals occurring in the heart. This method models the nonlinear components of the cardiac electrical system. Using these models, the algorithm generates a statistical likelihood that is used to determine whether a patient is experiencing ischemia. This approach is beneficial because the required recordings can be obtained without any invasive medical procedures.

1.1 Problem Description

The purpose of this work is to develop an algorithm for the classification of cardiac electrical signals as indicative of ischemia or not. The classifier is supplied with the time in which the electrical signals indicate the possibility of ischemia. The goal of this work is to create an algorithm with high levels of accuracy and reduced levels of false alarms.

The implemented algorithm is completely patient independent, which means that the learning data is independent of the testing data. The data used in this thesis is from the Long Term ECG Database from PhysioNet [6]. This database provides labels for the start time of an ST event, so the developed algorithm need not detect the start of an event automatically. Given an event start time, the algorithm must determine whether the event is caused by ischemia.

1.2 Thesis Outline

This thesis is divided into seven chapters. Chapter 2 provides a background of how the cardiac electrical system works. Chapter 2 also goes into detail about what ischemia is and what effects it can have on the cardiac system. Additionally, the chapter defines ST events as they apply to this body of work.

Chapter 3 gives a background of the current methods used by physicians to diagnose ischemia and the methods that researchers have recently studied. This chapter also provides a description of the Computers in Cardiology Challenge.

Chapter 4 discusses the nonlinear signal processing methods that are used in this thesis to classify ST events. This chapter goes into detail about reconstructed phase spaces (RPS), Gaussian mixture models (GMM), and the Naïve Bayes classifier.

In chapter 5, the Long Term ST Database that is used in this work is discussed. Additionally, the ten-fold cross validation technique, which is used in several of the experiments, is explained. Finally, this chapter provides a discussion of the overall classification framework that is used in the experiments.

Chapter 6 describes the experiments conducted as part of this research. Each experimental section explains how the experiments were conducted, the accuracy results, and the meaning of the results.

Finally, Chapter 7 gives a discussion of conclusions drawn from this work and offers suggestions for future work.

2 Cardiac Background

The cardiac muscle (heart) is the center of the cardiovascular system. This muscle pumps life-sustaining blood to the entire body. The blood supplies oxygen and nutrients to the body's organs so that they can perform their designated functions.

The heart is controlled by a very precise electrical system. This system regulates the mechanical pumping action of the heart so that the entire cardiovascular system can function properly. If a problem occurs in the electrical system of the heart, it can have devastating effects for the entire body.

2.1 Cardiac Function

The function of the cardiovascular system is to supply oxygen to the organs of the body. Blood is the body's medium for transporting oxygen to the organs. The muscular pump, known as the heart, pumps blood throughout the body. A complex set of arteries, vessels and capillaries connect the heart to the entire body.

2.1.1 Cardiac Mechanical System

There are four pumping chambers in the heart: the left and right atria and the left and right ventricles (see Figure 2.1). The purpose of the atria is to receive blood from the body; the right atrium receives oxygen-devoid blood from the body and the left atrium receives oxygen-rich blood from the lungs. The atria are separated from the ventricles by the tricuspid valve on the right side and the mitral valve on the left side. When these valves are opened, the blood from the atria flows into the ventricles. The ventricles are stronger than the atria because they pump the blood throughout the body. The right ventricle pumps the oxygen-devoid blood to the lungs to absorb oxygen and release

carbon dioxide. The left ventricle pumps the oxygen-rich blood to the body's organs. The right ventricle is regulated by the pulmonary valve and the left ventricle is regulated by the aortic valve [2].

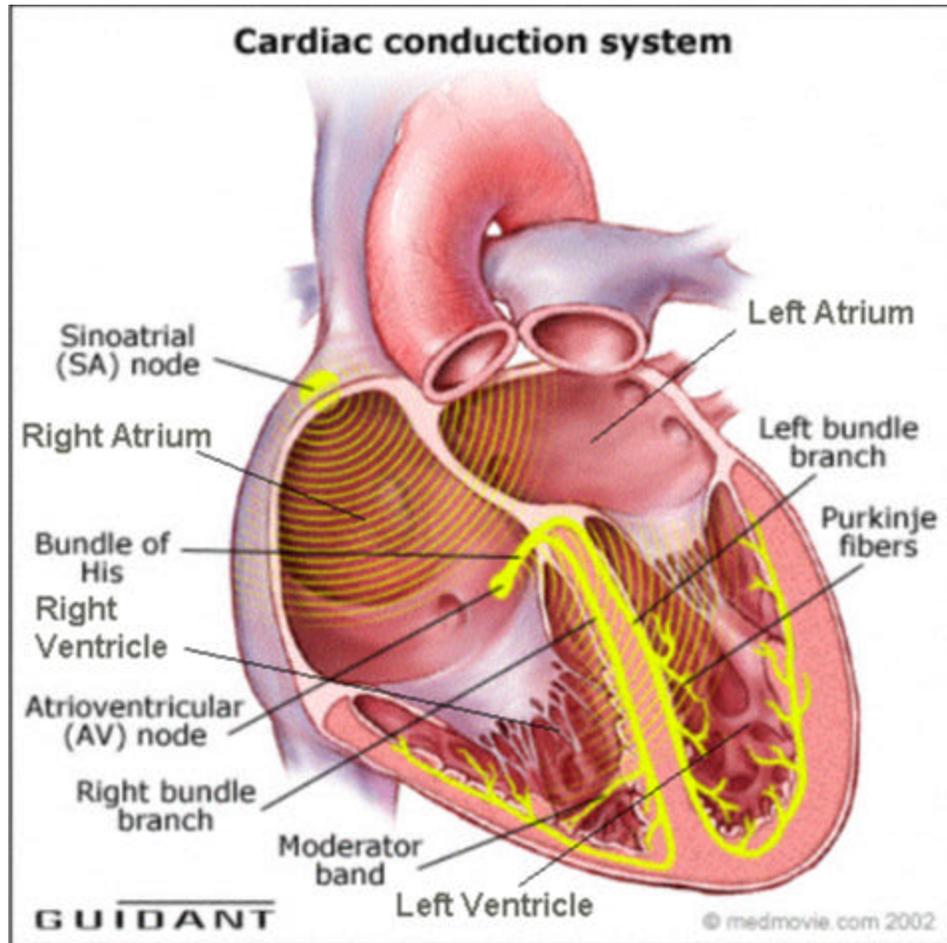


Image used with the permission from Guidant Corporation

Figure 2.1 – The mechanical and electrical components of the heart

2.1.2 Cardiac Electrical System

The mechanical pumping action of the heart results from electrical activation fronts transversing the cardiac tissue. Figure 2.2 shows an example of the electrical signal, also known as the electrocardiogram (ECG), for a single heartbeat. The labels indicate the approximate location of the important waves and components of the ECG

signal. This figure was generated from one of the records contained in the PhysioNet Long-Term ECG Database (see section 5.1 for more information about the database).

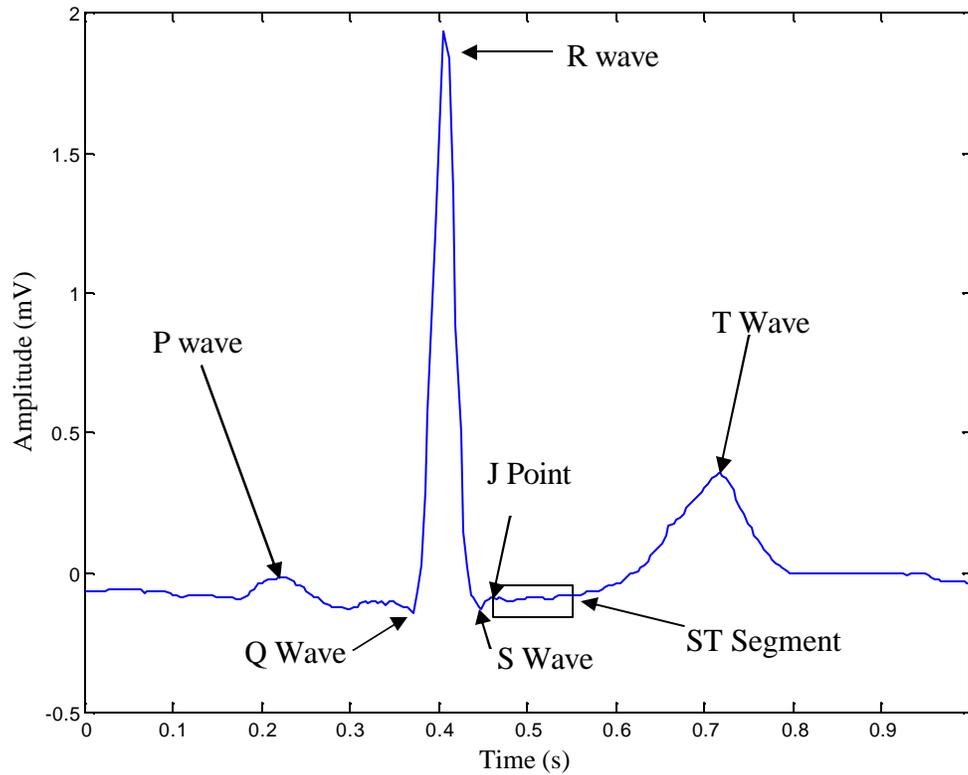


Figure 2.2 – Electrical signal (ECG) with important wave components labeled

The heart tissue experiences a series of stages of electrical depolarization and repolarization that lead to particular muscle contractions. These stages, summarized in Table 2.1, are described in the following paragraphs. The components of the heart discussed throughout this section are labeled in Figure 2.1.

	Electrical Function	Mechanical Function	Electrical Representation
1.	SA Node emits electrical pulse		
2.	Atria depolarize	Atria contract	Start of P Wave
3.	Electrical pulse pauses at AV Node	Blood flows to ventricles	End of P Wave
4.	Pulse travels down His Bundle to Bundle Branches		Q wave
5.	Atria repolarize while ventricles depolarize	Atria relax, Ventricles contract pumping blood to lungs and body	R and S wave
6.	Ventricles repolarize	Ventricles relax	T wave

Table 2.1 - Stages of cardiac excitation with corresponding ECG representation

The first stage of a heartbeat begins when the sinoatrial (SA) node of the heart depolarizes. During this stage, the right atrium is filled with oxygen-devoid blood that has returned from the circulatory system and the left atrium is filled with oxygen-rich blood that has returned from the pulmonary circulation. The SA node, located on the posterior wall of the right atrium, is the pacemaker of the heart, depolarizing at regular time intervals to ensure proper pacing. In a normal heart, the rate at which this node emits pulses is directly correlated to the amount of work that the heart is doing. As the body works harder and requires more oxygen-rich blood, the SA node increases its pace to satisfy the demand [2].

The electrical impulse from the SA node causes the upper portion of the heart, called the left and right atrium, to depolarize. This depolarization causes the atria to contract forcing the blood from these chambers downward into the large lower portion of the heart, called the ventricles. The corresponding component of the electrical signal is the P wave. As soon as the atria have completely contracted, they begin to repolarize in preparation for the next beat. The electrical signature of the repolarization is not

discernable in the electrical signal because it occurs at the same time as the ventricular contraction, which yields the large QRS complex [2].

Following the depolarization of the atria, the depolarizing wavefronts signals converge at the atrioventricular (AV) node. The AV node serves two very important purposes. Its first purpose is to bridge the electrical signal from the atria to the ventricles. The second purpose is to slow the electrical depolarization to allow the blood to completely flow from the atria to the ventricles. The electrical depolarization propagates from the AV node to the His Bundles, which are located at the base of the ventricles. The His Bundles lead to the bundle branches and then into the Purkinje fibers, which rapidly spread the depolarizing wavefront across both ventricles. This electrical signal moves rapidly across the ventricular tissue causing the muscles of the ventricles to pump the blood to the rest of the body. The right ventricle pumps the oxygen-devoid blood to the pulmonary system for oxygenation. The left ventricle pushes the oxygen-rich blood to the circulatory system to bring oxygen to the body. Following the depolarization and contraction, the ventricles begin to repolarize to prepare for the next cycle [2].

As long as the heart is operating properly, the process described above repeats rhythmically with a natural variability. Generally, the resting heartbeat of a healthy person is about 60-80 BPM [7]. If the electrical system of the heart does not properly function, the heart's rhythm can become abnormal. This directly affects the heart's ability to supply blood to the entire body. To monitor for such problems, physicians record and analyze the cardiac electrical signals using the electrocardiogram (ECG).

2.2 *Electrocardiogram*

The electrocardiogram (ECG) is a recording of the electrical signals that control the cardiac function. By analyzing these signals, clinicians can monitor the rhythmic function of the cardiac system.

In order to record an electrocardiogram (ECG) from a patient, a number of electrodes are placed on the patient's chest. The number of electrodes used can vary between two and fourteen depending on what parts of the heart the physician would like to focus the examination. Figure 2.3 describes how the three most common ECG leads (I, II, and III) are recorded [2]. The ECG measures the change in the electrical potential across the electrodes. The recorded potential is converted to a waveform after signal filtering and amplification. In this work, each waveform will be referred to as a lead. The Long Term ST database, which is used in this work, provides either two or three leads for each patient.

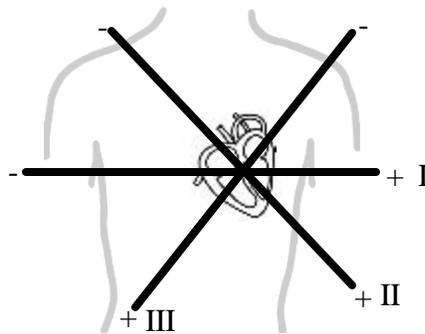


Figure 2.3 – Einthoven's triangle showing how leads I, II, and III are recorded

2.3 *Ischemia*

Ischemia is a condition in which the heart requires more oxygen than the body is able to supply. This can result from an increased demand for oxygen or a decrease in the

supply of oxygen. A decrease in oxygen supply can be caused by an artery blockage, blood clots, and spontaneous artery spasms. Blockages of a coronary artery, which supply oxygen rich blood to the heart, are known as atherosclerosis. This condition is caused by a buildup of fat and cholesterol, known as plaque. In some cases a plaque buildup will cause a crack in the walls of the arteries. At the site of the crack a blood clot may form, further restricting blood flow. In the case of an artery spasm, known as angina, an artery may spontaneously and temporarily contract restricting or stopping blood flow. When any of these three conditions occur, the heart is deprived of the oxygen it needs to function [8].

Quite often, the causes of ischemia may be temporary. As blood flow to the heart is restored, the cardiac oxygen perfusion (absorption) returns to normal and the tissue returns to a healthy state. However, approximately six hours of oxygen deprivation is all that is needed for the tissue to sustain irreversible injury. This injured cardiac muscle tissue becomes scar tissue, which leads to a diminished pumping capability and modified electrical propagation characteristics. If the infarction (tissue death) becomes severe enough, the heart will no longer be able to adequately supply blood to the entire body. The death of heart tissue, commonly known as a heart attack, often leads to strokes and patient death [3, 4].

Myocardial ischemia is commonly seen as a warning sign of cardiac problems. It is beneficial for patients and doctors to be able to recognize the signs and symptoms of ischemia before permanent tissue death has occurred. While ischemia often results in chest pain, there are many cases where the patient experiences no pain. This is known as silent ischemia. Because of silent ischemia, it is important to screen for ischemia even in

the absence of symptoms, especially for patients who are at risk due to family history or health problems. If the warning signs are recognized, the patient can take steps to prevent a heart attack [9]

2.4 ST Events

The literature has established that there is a strong correlation between elevation and depression of the ST portion of the ECG signal and cardiac problems related to ischemia and infarction [10]. In 1920, Pardee first claimed that ST elevation was a sign of ischemic problems [10]. According to Fozzard and Janse, the abnormality is due to the way that ischemic tissue conducts electricity [11, 12]. During the period that the tissue is polarized, current flows from the ischemic tissue towards the normal tissue. When the portion of the heart affected by ischemia begins to depolarize, the current flow switches from the normal tissue to the ischemic tissue. This current flow causes the abnormality of the ST segment during the period in which it is normally a flat waveform [11-14]. It is generally accepted that an ST deviation or elevation greater than 1 millivolt may indicate the presence of ischemia. It is important to note that many other conditions and circumstances can cause this elevation and depression [15].

Based on the research previously described, many systems monitor the elevation or depression of the ST segment. When significant changes of the ST segment are noted, the device records the time of occurrence. Using this technique, this system can provide a simple way to look for ischemic events. However, many false alarms are present because other occurrences such as former heart damage and position changes of the patient can cause similar ST changes. Additionally, it is not possible to attach ECG leads in the same position for every test. This leads to variation in the recorded signals [16].

To compare records from multiple patients, a patient independent function of the ST level is used. The standard value, called the ST deviation, is calculated as the difference between the ST level and an established ST reference function (shown in Figure 2.4). PhysioNet (an organization for the research of physiological signals) has defined the ST reference function. The baseline or starting value of the reference function is chosen as the stable ST level near the beginning of the record. Next, an expert must annotate the record labeling points of ST deviation that are not clinically relevant. For this work, clinically irrelevant points are periods where the ST deviation is not due to ischemia, heart rate-related changes, and noisy ST events. For these clinically irrelevant ST deviations, the ST reference function is designed to compensate for the deviation so that the ST deviation shows no change. The resulting ST deviation can easily be compared with the standard deviation of other records calculated using the same method [17, 18].

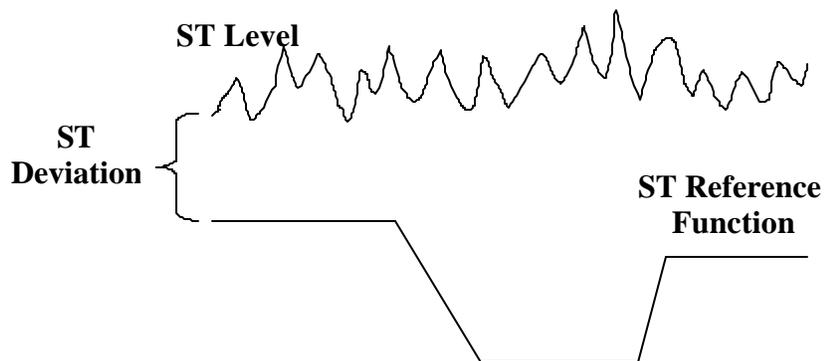


Figure 2.4 - Example of ST deviation calculation

Each ST level monitoring device has specific level criteria that it uses to define an ST event. The creators of the PhysioNet Long-Term ST Database, which is used in this research, defined an ST event using the three criterion shown in Figure 2.5. First, the ST

event begins when the ST deviation exceeds $50\mu\text{V}$. Second, the ST deviation must exceed $100\mu\text{V}$ for at least 30 seconds during the event. Third, the ST event ends when the ST deviation drops below $50\mu\text{V}$ for at least 30 seconds. All three criteria must be met in the described order for an event to be labeled [17, 18].

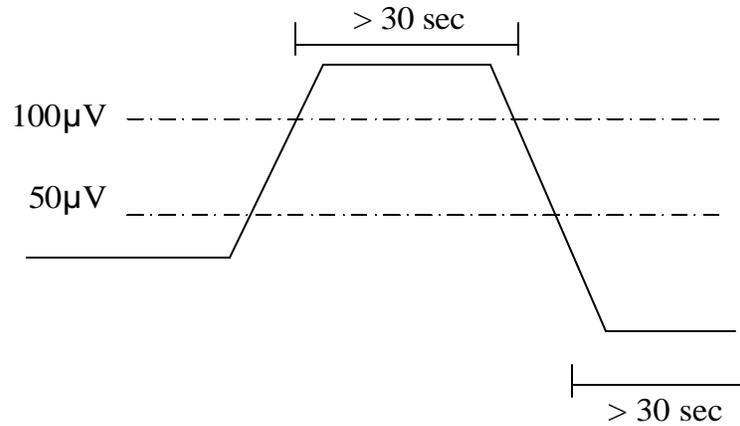


Figure 2.5 - Definition of ST Event

3 Research Background

There are several methods currently used by clinicians to monitor and detect ischemia in cardiac patients. These methods have varying levels of accuracy, cost, and difficulty in administering. Section 3.1 gives an overview and analysis of the methods that are currently being used. Section 3.2 explains some of the methods that are currently being researched to solve the ischemia classification problem. The last section of this chapter provides an overview of the 2003 Computers in Cardiology Challenge, which prompted the research of this thesis.

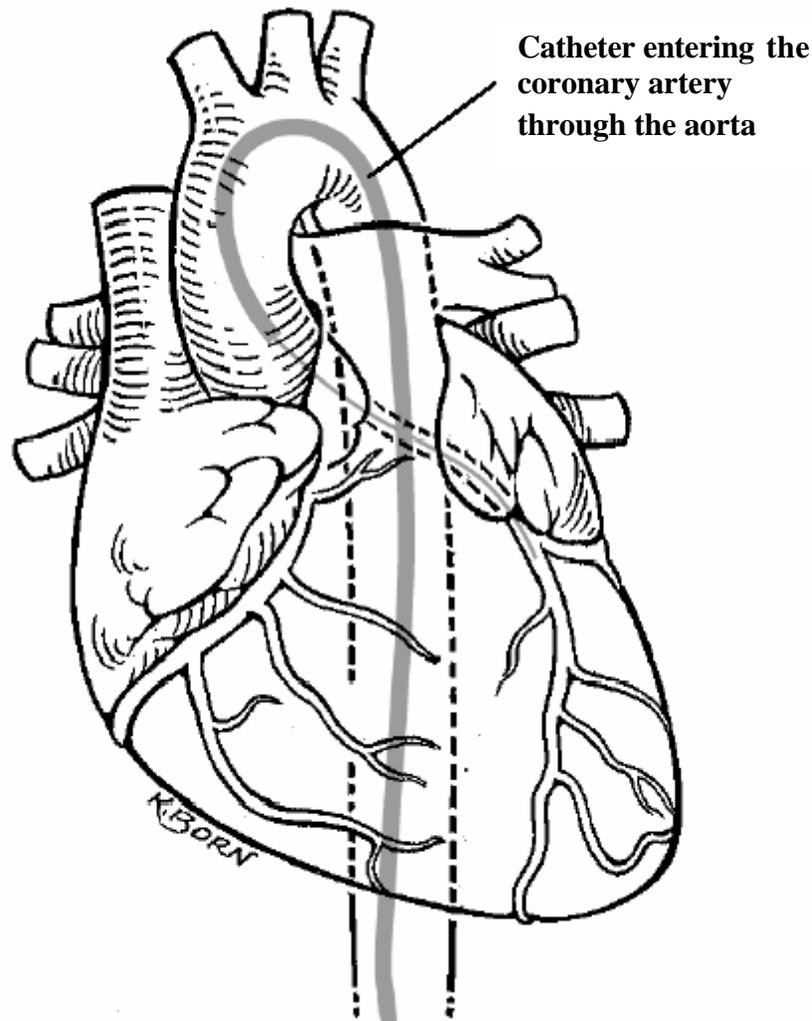
3.1 *Current Clinical Methods*

The procedures available to physicians for diagnosis of myocardial ischemia range from very accurate but expensive and often invasive to relatively inexpensive but indecisive. Each method has its benefits and drawbacks, so physicians must select the appropriate method on a patient-by-patient basis. Frequently used methods include coronary imaging (angiography) [19-22] and imaging using echocardiograms [23], which are both very accurate. The use of ST (the waveform segment connecting the S wave to the T wave in the cardiac electrical signal) event alerts [16] and exercise stress testing [24] are examples of lower cost methods that have lower levels of accuracy.

3.1.1 Coronary Angiography

The most accurate method to detect myocardial ischemia is coronary angiography [22]. For this procedure, an incision is made near the patient's groin. A catheter is inserted into an artery through this incision. The catheter is then maneuvered through the artery to the location that the physician would like to examine. Typically, the catheter

will be used to examine coronary arteries (arteries that supply blood to the heart muscle), the heart's valves, and the inside of the cardiac chambers. Figure 3.1 shows how a catheter may be maneuvered to examine a coronary artery. Once the catheter is in place, an X-ray sensitive dye is released and X-rays are taken. The X-rays show highlighted images of the area where the dye was released. This allows a cardiologist to look for constriction of arteries or blockages due to fatty buildup known as plaque. These conditions restrict blood flow to the heart and lead to ischemia. Several injections and subsequent X-rays may be required for the physician to complete the examination. This method is utilized because of its high level of accuracy with sensitivity between 90% and 98% and specificity between 95% and 98% [21]. The drawbacks are the level of intrusion into the patient's body, the requirement that the patient abstain from food and drink for 6 to 8 hours, and the average cost of \$17,598 [19-22].



Reproduced with permission

<http://www.americanheart.org/downloadable/heart/1046715993947WhatIsaCoronaryAngiogram.pdf>

© 2004, Copyright American Heart Association

Figure 3.1 - Figure showing how catheter is inserted for coronary angiography

3.1.2 Echocardiograms

A second highly accurate method of myocardial ischemia detection is the use of echocardiograms [23]. An ultrasound transducer emits ultrasound waves that echo off the internal organs of the body. The transducer receives the echoes and the echocardiogram machine converts the received data into 2D and 3D moving images. A physician can then use these images to examine abnormalities in the cardiac tissues. While this system

cannot view the blood flowing through the heart, as the coronary angiography can, abnormalities of the heart's motion can indicate ischemic conditions [25]. The sensitivity for echocardiography for detection of ischemia varies between 80% and 90% and the specificity varies between 75% and 100% depending on the extent of the ischemic changes to the cardiac tissue [26]. This method of detection is relatively expensive at between \$250 and \$1000 and requires a physician to examine the results [27].

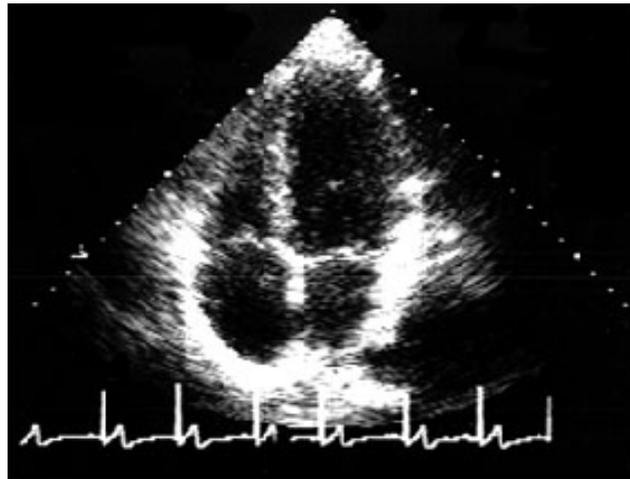


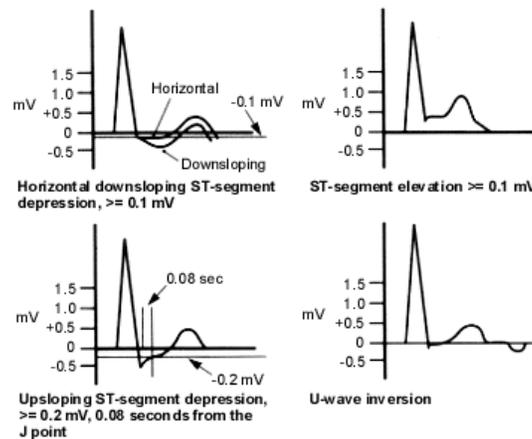
Figure 3.2 - Example of echocardiogram with ECG beat along bottom [28]

3.1.3 Exercise Testing

Exercise testing of a patient is a method of stress testing the functions of the heart [24].

The patient walks on a treadmill or rides a stationary bicycle with an ECG monitor attached to their chest. This is an effective way of forcing the cardiac muscle to pump blood more rapidly thus requiring more oxygen to function. If there are blood flow restrictions to the cardiac muscle, they will become more pronounced as the heart works harder. The recorded ECG will show evidence of these changes as a greater ST deviation than normal [24]. Figure 3.3 shows four ECG waveforms from exercise testing that exhibit signs of ischemia. The advantages related to this method of testing are the ease at

which it is administered and its low cost of only \$200 to \$300. The drawbacks are: it requires a patient to endure the intense physical test that may be difficult depending on other health conditions, it requires a physician to carefully analyze the recorded ECG, and it is much less accurate than other methods with a sensitivity of 68% and a specificity of 77% [29, 30].



© 1999, Copyright American Academy of Family Physicians

Figure 3.3 - ECG waveforms that show indication of ischemia during exercise testing [31]

3.1.4 ST Level Alerts

The final method discussed here is the use of ST level alerts [15]. This is an easy, but imprecise way of detecting the possibility of myocardial ischemia. Because of high sensitivity but low specificity, the results of this diagnostic tool must be verified by another detection method. Due to physiology changes caused by the oxygen deprivation of the cardiac tissue, the ST segment of the ECG signal can become abnormal (see section 2.4). Using long-term (24-48 hour) ECG recording devices, the ECG signals are collected, and the ST events (a period when the ST segment is abnormal according to specific measures, for more information see section 2.4) are noted. A cardiologist later reviews the events and the patient's cardiac history to determine if ischemia is likely. If

the test indicates the presence of ischemia, further tests, such as those discussed above, can be conducted to verify the preliminary diagnosis. While this method is more cost effective with a cost under \$250 and easy to administer, it has lower accuracy with a sensitivity around 46% and a specificity around 91% and positive results can be confused with other causes of ST level changes [27, 32-34]. It is also possible that the patient will experience no ischemic episodes during the time that the ECG is recorded, even though myocardial ischemia exists.

3.2 State of ECG Research in Detecting Myocardial Ischemia

Significant research has been undertaken to develop a more accurate, less invasive, and less expensive method for detecting myocardial ischemia. Much of this research focuses on the use of ECG signals. These methods build models or use thresholds of the ST deviation to determine if a patient's ECG signal might indicate ischemia. The following section provides two examples of these methods.

3.2.1 Neural Network Classifier

Maglaveras et al [35] have investigated a method for ischemia detection that uses supervised neural networks. The first step in this algorithm was to extract the ST segment (section of the ECG waveform between the S wave and the T wave) from the ECG waveform (see Figure 3.4). The developers chose to use 160ms (40 samples at a sampling rate of 250Hz) following the J-point as the length of the ST segment. Using the ST segment and the ST reference function from PhysioNet [17], the ST deviation was calculated (see section 2.4). The 40 ST deviation samples were grouped into 20 consecutive pairs. Each of the two paired values were averaged leaving 20 values. These values were the inputs to a three layer neural network. The neural network consisted of an

input layer with 20 neurons (one for each input), a hidden layer with 10 input neurons (10 was chosen empirically based on best accuracy results and training time), and an output layer with two neurons. The two output neurons each provide a value between zero and one, which is rounded to either zero or one. The four possible outputs are: (1,1) unclassifiable, (1,0) ST depression, (0,1) ST elevation, (0,0) normal beat.

Since one of the developer's goals was to create an algorithm that could be quickly trained, an adaptive back propagation algorithm was used to train the classifier. This method of classification adjusts the neural network weights so as to minimize the mean square error between the output and the target vectors of the training set.

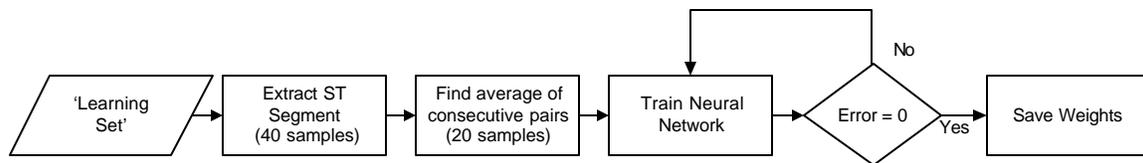


Figure 3.4 - Flowchart of Neural Network Learning

The classification procedure for this algorithm follows the same ST segment extraction and averaging of the training process. Once the beats are extracted, they are input into the neural network. If the output layer produces (1,0) or (0,1) the beat is classified as ischemic, otherwise it is classified as non-ischemic.

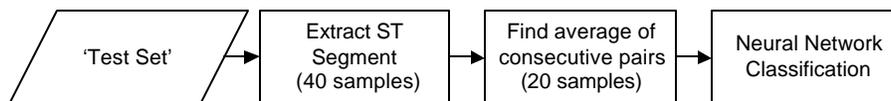


Figure 3.5 - Flowchart of Neural Network classification

The classifier was trained and tested using two exclusive sets of records from the European ST Database. The results are shown in Table 3.1. The paper did not provide

specificity and no confusion matrix was provided so, positive predictive accuracy is offered [35].

Sensitivity	Positive Predictive Accuracy
73.0%	69.5%

Table 3.1 - Classification results for Neural Network method applied to the European ST Database

3.2.2 Threshold Classifier

A second technique for classification developed by Langley et al [36] uses the characteristics of the ST deviation to classify ST events (ST deviation and events are explained in section 2.4). The event classification is determined by several threshold criteria. Figure 3.6 shows the criteria that must be met for an ST event to be classified as ischemia. The event begins when the ST deviation first exceeds $50\mu\text{V}$. The ST deviation must then rise above $100\mu\text{V}$ and stay above that threshold for at least 30 seconds. The event ends when the ST deviation drops below $50\mu\text{V}$ for at least 30 seconds.

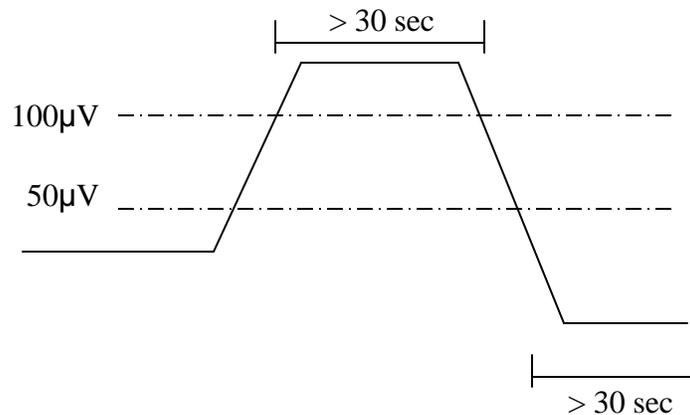


Figure 3.6 – Langley et al threshold classification algorithm

	Sensitivity	Specificity	Accuracy
Training Set	99.0%	88.8%	91.1%
Test Set	99.0%	93.3%	95.6%

Table 3.2 - Accuracy results for the Langley et al method

The data used for this study was the Long Term ST Database. When run on the ‘Training Set,’ the classification algorithm had a nearly perfect sensitivity of 99.0% but the specificity was 88.8%. When the algorithm was applied to the Test Set the sensitivity remained at 99.0% but the specificity increased to 93.3% [36].

3.3 Challenge Description

The Computers in Cardiology Challenge (CINC) is an annual contest in which members of the research community compete to solve a problem which has to date not sufficiently been solved [6]. The 2003 challenge was to design an algorithm which could classify ST events as ischemic or non-ischemic using only the data of a long term two or three lead ECG (the electrocardiogram is a plot of the electrical fields occurring in the cardiac muscle, see section 2.1) signal. The challenge coincided with the release of the Long Term ST database from PhysioNet [18]. This database was created to aid in the design of algorithms that detect ischemia using ECG signals. Forty-three (the Training Set) of the 86 records (entire dataset) contained in the database were released for use in algorithm design. The other 43 records (Test Set) were withheld to be used for scoring the algorithms.

For the CINC challenge, significant ST events are determined by PhysioNet [17] using the criteria shown in Figure 3.7. The onset of an event occurs when the ST deviation exceeds $50\mu\text{V}$. The ST deviation must then exceed $100\mu\text{V}$ for at least 30 seconds. The end of the episode is reached when the ST deviation remains below $50\mu\text{V}$

for at least 30 seconds. After the ST episodes were located in the data using this criteria, independent experts met to reach a consensus on labeling each event as caused by ischemia or one of three types of non-ischemic factors. The three types of non-ischemic classification are axis shift, conduction change, and heart rate-related ST shift.

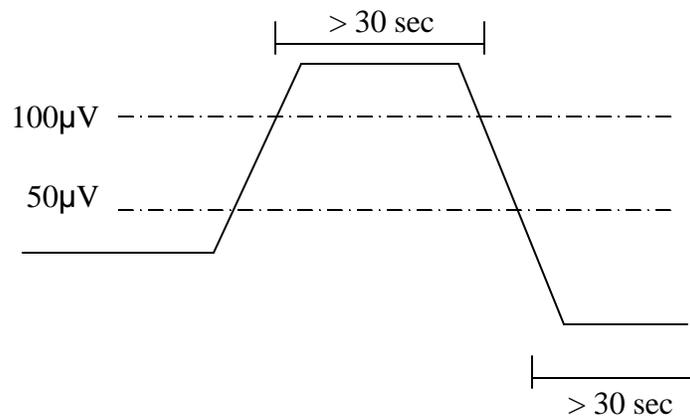


Figure 3.7 – Definition of ST Event

For each record (a record consists of two or three signal long-term ECG recordings for a single patient, see section 5.1) in the training set two ST event label files are provided. The .EPI file contains the event index, the time when the event begins, and the channel that the experts believe will most likely show evidence of ischemia. The event time label indicates the time before the first beat with an ST level. The second file, the .EPR file, contains all of the same information plus a label for the event as ischemic or non-ischemic. The .EPI files are provided for the Test Set, but the .EPR files are not.

4 Nonlinear Signal Processing Theory

The following chapter describes the nonlinear signal processing methods used in this thesis to model and classify the ECG ST events. The first section describes the reconstructed phase space method used to transform the time series data into the time embedded phase dimension. The next section describes the Gaussian mixture model (GMM) method, which is used to model the constructed phase spaces. The final section introduces the Naïve Bayes classifier, which is used to classify signals into classes using the learned GMMs.

4.1 Reconstructed Phase Spaces

In a time series, such as the ECG signal, it is sometimes necessary to search for patterns not only in the time series itself, but also in a higher dimensional transformation of the time series. The reconstructed phase space (RPS) is an example of such a transformation. An RPS is an n dimensional space in which a signal is plotted against time-delayed versions of itself. Each point in the phase space is calculated according to:

$$\mathbf{x}_n = [x_n \quad x_{n-t} \quad \cdots \quad x_{n-(d-1)t}], \quad n = (1 + (d-1)t) \dots N, \quad (4.1)$$

where N is the dimension of the space, t is the time delay, and d is the dimension. The entire phase space is generated by:

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_{1+(d-1)t} \\ \mathbf{x}_{2+(d-1)t} \\ \vdots \\ \mathbf{x}_N \end{bmatrix} = \begin{bmatrix} x_{1+(d-1)t} & \cdots & x_{1+t} & x_1 \\ x_{2+(d-1)t} & \cdots & x_{2+t} & x_2 \\ \vdots & & \ddots & \\ x_N & \cdots & x_{N-(d-2)t} & x_{N-(d-1)t} \end{bmatrix} \quad (4.2)$$

In 1980, Takens proved that a reconstructed phase space is topologically equivalent to the original state space if the embedding dimension is $2d + 1$, where d is the state dimension of the original system [37]. Sauer et al generalized this requirement by showing that if the dimension of the RPS is greater than two times the box counting dimension of the original system, topological equivalence holds [38]. The existence of topological equivalence means that the reconstructed phase space transformation is unique or one-to-one. In classification problems, this theory is important because the reconstructed phase space transformation does not destroy the dynamical information of the system.

Figure 4.1 demonstrates how a signal is delayed by consecutive time lags. For purposes of visualization the signals have been separated along the y-axis. As the figure shows, x_n is not delayed, x_{n-5} is delayed by 5 samples, and x_{n-10} is delayed by 10 samples. The points labeled with boxes are all at the same sample index after the delays. Each time index is transformed into one point in the embedded space; so the three labeled points would create one point in a three dimensional RPS. Figure 4.2 shows an example of a two-dimensional reconstructed phase space. The signal x_n is first delayed by 25 samples. Then the point from each time index of x_n is plotted against the point from the same time index of the delayed signal. For example $x_n(1)$ is plotted against $x_{n-25}(1)$.

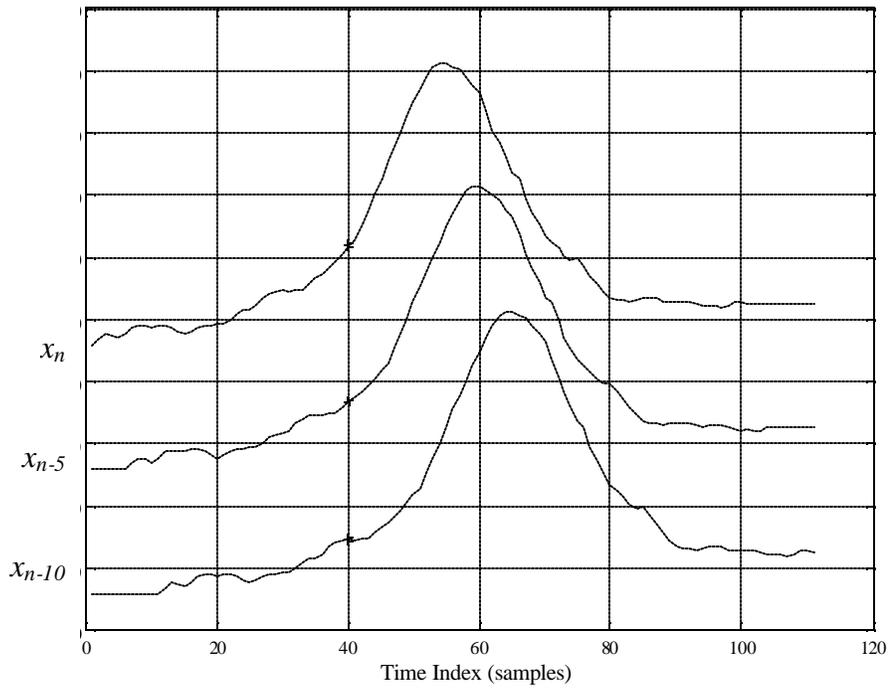


Figure 4.1 - An ST segment and T wave with 0 delay, 5 time step delay, and 10 time step delay (shown on the vertical axis).

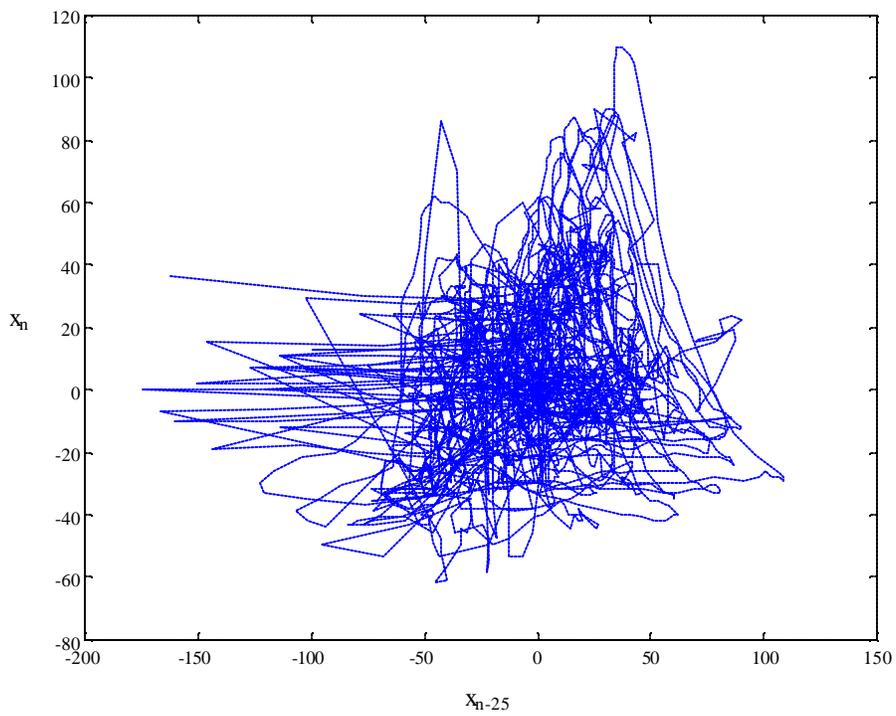


Figure 4.2 - The reconstructed phase space of a set of ST segments

In order to create a reconstructed phase space, the dimension and the lag need to be determined. The proper determination of these parameters can have a significant effect on the classification performance. Selecting values that are too high adds complexity and increases the amount of data needed for an experiment. If the values selected are too low, the complex nature of the problem may not be captured in the phase space.

Two common methods used to estimate the proper lag are: the first minimum of automutual information and an empirical method. The automutual information method determines the information that is shared between consecutive integer lag choices [39]. The first minimum of the mutual information is used as an estimate of the optimal lag. A second method estimates the lag empirically. The lag is selected by examining the results of multiple classification experiments. In each experiment a different lag is studied and the classification accuracy is determined. The lag that provides the highest accuracy is the one that is selected.

The process for choosing the dimension is similar to that of determining the lag. The dimension can be selected using the false nearest neighbors or by an empirical method. The false nearest neighbors are points in a phase space of n dimension which are near each other, but are not near each other in a phase space of dimension $n + 1$. The number of points falsely near each other indicates whether a higher dimension should be used. A threshold on the percentage of false nearest neighbors is used to select the dimension. The empirical method operates the same as the lag determination. Multiple experiments are conducted and the experiment with the highest accuracy indicates the dimension to be chosen.

4.2 Gaussian Mixture Models

A Gaussian mixture model (GMM) is a set of N multidimensional Gaussian distributions. Figure 4.3 gives an example set of eight GMMs on top of the two-dimensional ECG ST segment RPS from Figure 4.2 (the figure has been zoomed to make it easier to see). The principal axes of each ellipse are labeled for each GMM. The set of GMMs approximately models the distribution of the data.

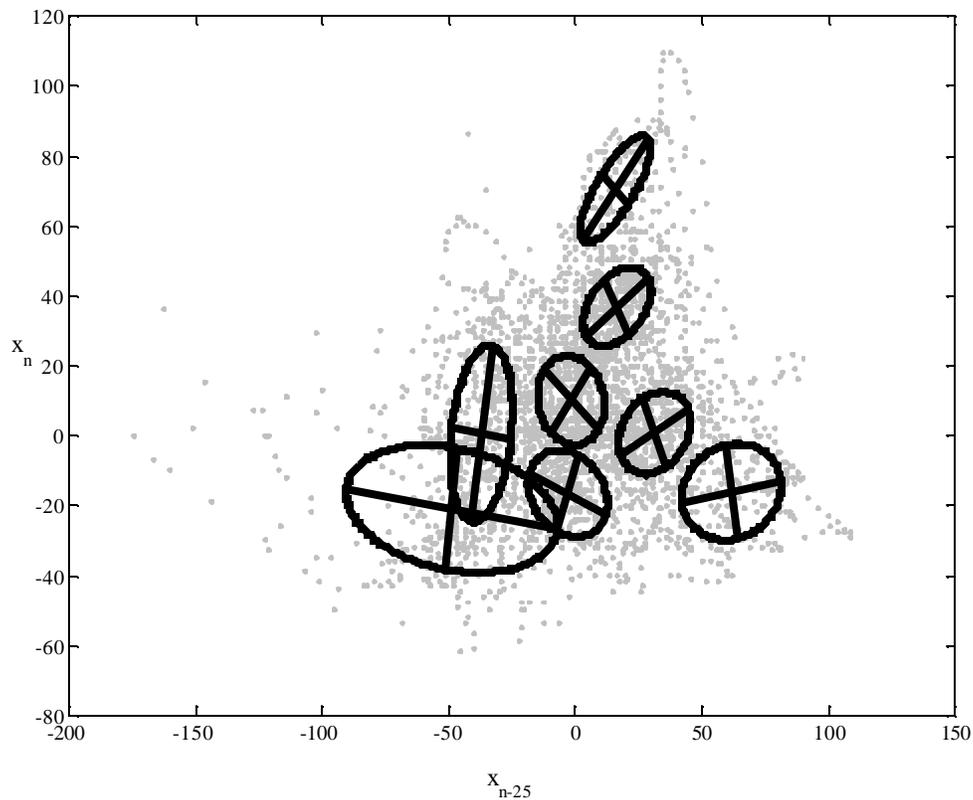


Figure 4.3 - The Reconstructed Phase Space of Figure 4.2 with overlay of Gaussian Mixture Models

The GMM is defined as:

$$p(\mathbf{x}) = \sum_{m=1}^M w_m p_m(\mathbf{x}) = \sum_{m=1}^M w_m N(\mathbf{x} | \boldsymbol{\mu}_m, \mathbf{S}_m), \quad (4.3)$$

where w_m is the mixture weights ($\sum w_m = 1$), M is the number of mixtures, and $N(\mathbf{x}; \boldsymbol{\mu}_m, \mathbf{S}_m)$ is a normal distribution with mean $\boldsymbol{\mu}_m$ and covariance matrix \mathbf{S}_m . Using multiple Gaussian distributions, any distribution of data can be modeled [40]. The number of distributions to use in a GMM model is a parameter dependent on the complexity of the distribution of the data to be modeled. This number can be determined empirically, by running multiple classifications that utilize different numbers of GMMs. The classification with the highest accuracy is the one that is selected.

The weights, means, and covariance matrix of the GMM are estimated using the Expectation Maximization (EM) algorithm [41]. The method begins with initial values for each parameter, then iterates through the available data to find the Maximum Likelihood (ML) estimate. The formulas used for the estimation are:

$$\begin{aligned}
 \mathbf{m}'_m &= \frac{\sum_{n=1+(d-1)t}^N (p_m(x_n) x_n)}{\sum_{n=1+(d-1)t}^N p_m(x_n)}, \\
 \Sigma'_m &= \frac{\sum_{n=1+(d-1)t}^N (p_m(x_n) (x_n - \mathbf{m}'_m)(x_n - \mathbf{m}'_m))}{\sum_{n=1+(d-1)t}^N p_m(x_n)}, \\
 w'_m &= \frac{\sum_{n=1+(d-1)t}^N p_m(x_n)}{\sum_{n=1+(d-1)t}^N \sum_{m=1}^M p_m(x_n)},
 \end{aligned} \tag{4.4}$$

where M is the number of mixtures, p_m is the probability distribution function, and N is the number of points in the signal [41].

4.3 Naïve Bayes Classifier

The Naïve Bayes classifier [42] uses the principles of Bayes' law to compute the class with the highest conditional likelihood. Bayes' law states:

$$p(class | \mathbf{x}_n) = \frac{p(\mathbf{x}_n | class)p(class)}{p(\mathbf{x}_n)}, \quad (4.5)$$

where $p(class | \mathbf{x}_n)$ is the conditional probability of a specific *class* given a point in the RPS. Since the $p(\mathbf{x})$ term is simply a normalizing value, it is left out of the equation when comparing class probabilities. So, the Naïve Bayes classifier becomes:

$$class = \underset{class_i \in AllClasses}{\operatorname{argmax}} p(class_i) p(\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_n | class_i), \quad (4.6)$$

where *AllClasses* is the set of possible classes and $class_i$ is the i^{th} class.

In order to calculate $p(\mathbf{x}_1, \mathbf{x}_2 \dots \mathbf{x}_n | class_i)$ it would be necessary to have a large set of training data for every set of RPS points and class combination. In most cases, the available data will not be sufficient. For this reason, the Naïve Bayes classifier relies on the assumption that the features are conditionally independent with regard to class [43]. This means that the probability that a set of features occurs given a class, is simply the product of the each of the conditional probabilities, $p(\mathbf{x}_n | class_i)$. The Naïve Bayes classifier becomes:

$$class = \underset{class_i \in AllClasses}{\operatorname{argmax}} p(class_i) \prod_{n=1+(d-1)t}^N p(\mathbf{x}_n | class_i), \quad (4.7)$$

where N is the number of points in the RPS and d is the dimension of the RPS. The conditional probabilities ($p(\mathbf{x}_n | class_i)$) can be learned from the training data of the

experiment. In this way, the class of a given RPS is compared with the learned models and then classified.

5 Data and Methods

This chapter describes the dataset and the methods used in the proposed algorithm. The first section describes the Long-Term ST Database, which contains the ECG signals used for training and classification. The second section describes the pre-processing techniques that are necessary to prepare the data for the classifier. The final section outlines the procedure used for training and classification with the proposed algorithm.

5.1 Data

The dataset used in conducting the experiments described in this research is the Long-Term ST Database from PhysioNet. The database contains 86 Holter ECG recordings from 80 independent patients. Holter recordings are ECG recordings that are recorded using portable recording devices, generally taken over a long period. These recordings were selected from the Holter libraries at Beth Israel Deaconess Medical Center in Boston, Physiolab (Laboratory of Biosignal Processing) of the Institute of Clinical Physiology in Pisa, Brigham and Womens Hospital in Boston, and the Zymed company. These records vary in length from 20 to 24 hours. Each record contains either two or three ECG leads. The records are digitized at 250 Hz with 12 bit resolution [17, 18].

In order to create a database that accurately models what a classifier might see in a clinical setting, the records chosen for insertion into the database were carefully screened. The records that are included have at least one of the following ST events:

- Ischemic ST episodes
- Non-ischemic ST episodes due to heart rate changes
- Non-ischemic ST shifts due to axis shifts or changes in ventricular conduction

This creates a rich set of clinically relevant scenarios for training and testing [17].

The Long-Term ST Database creators have split the records into two groups. To allow a broad range of researchers to develop algorithms to classify ischemia, one-half of the database is available to the public. This half contains 43 records from 42 of the database's 80 patients. This portion of the database, called the Training Set was to be used to train algorithms for the Computers in Cardiology Challenge described in section 3.3. The second half of the database is available for purchase. This half was purchased for use in this study. The second half, called the 'Test Set,' contains the remaining 43 records from 38 distinct patients. The breakdown of the episodes in the database is shown in Table 5.1. As the table shows, 3266 total episodes were used during the training and testing portions of this research.

	Training Set	'Testing Set'	Total
Ischemic	331	723	1054
Conduction Change	441	454	895
Heart Rate-Related Shift	137	92	229
Axis Shift	383	705	1088
Total	1292	1974	3266

Table 5.1 - Breakdown of number of events for the Training and Test portions of the database

To allow researchers to focus on algorithm development and not data annotation, the database has been filtered and fully annotated. Filtering was conducted to remove the

baseline variability and frequencies over 55Hz (assumed to be noise). Next, the J-point of each beat was determined based on the start of the “most flat” portion of the beat following the R wave. These automatically determined J-Points were edited by experts. Next, a stable ST level near the start of the waveform was used to determine the global ST reference or baseline. Cardiology experts then examined the ST level throughout the signal and marked large ST level changes that are not related to ischemia, heart rate changes or conduction and axis shifts. This allows the ST deviation parameter to reveal important ST level changes. The ST deviation throughout the record is found by subtracting the ST level from the ST reference function. The ST reference function, ST level, and ST deviation are all provided in Long-Term ST database annotation files. The annotation files also provide labels of the R-wave peaks and the ST events (see section 2.4) [17].

5.2 Pre-processing

The Long-Term ST Database, used in this research, contains records that are between 21 and 24 hours long. The algorithm proposed by this thesis uses small segments of the waveform. In order to extract these waveform segments, several steps are needed to process and filter the data. The overall preprocessing stage can be represented as:

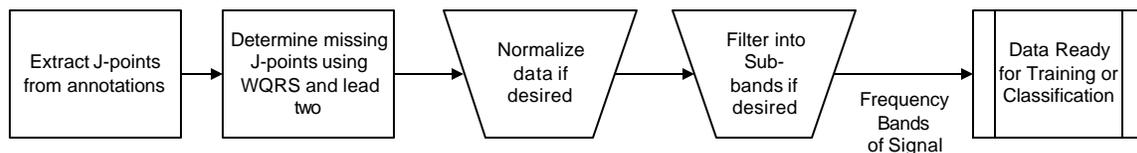


Figure 5.1 - Flowchart of pre-processing stage

The first step in pre-processing is to determine the J-point (point where the ST segments begins – see section 2.1.2) for all ECG beats that are to be studied. For this work, an equal number of beats surrounding the event label are used (N total beats). Figure 5.2 shows how data is extracted from the ECG records. As the figure shows, each record has two to three signals. The algorithm proposed by this thesis extracts the ST segment and T wave from each signal for each of the N beats surrounding the labeled event. The boxes in the figure show the segment of the signal extracted from each beat index. The number of beats used is symmetric on both sides of the event label. The algorithm uses $\frac{1}{2}N$ beats (n_-) before the labeled event time. The ST segment and T-wave are then extracted from the $\frac{1}{2}N$ beats (n_+) after the event time.

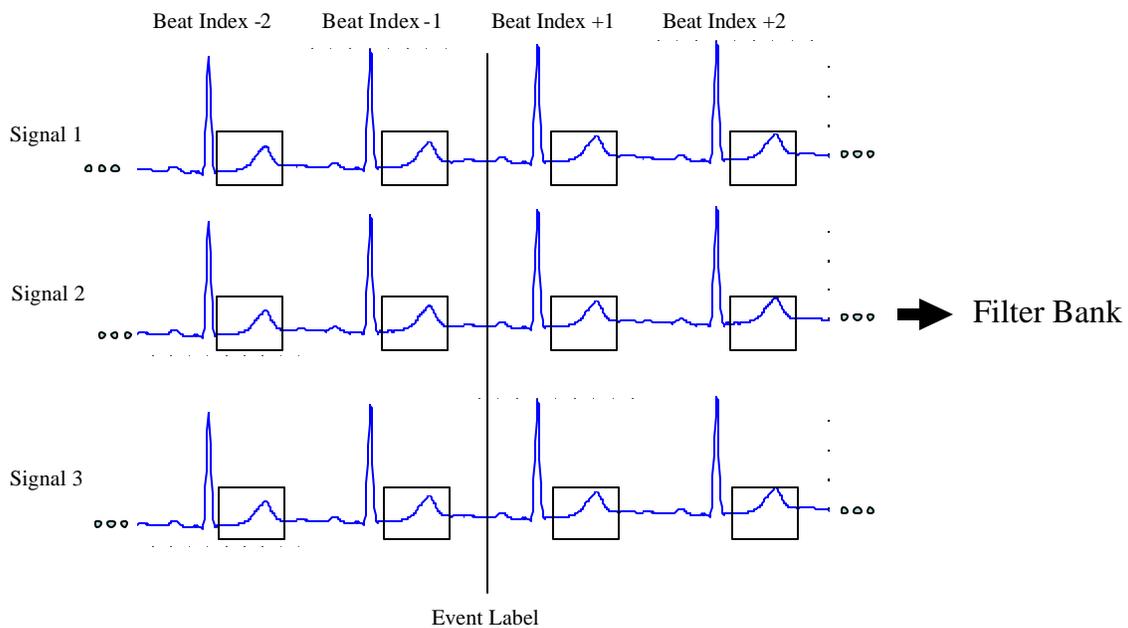


Figure 5.2 - Description of how ST-T complex (segment in box) is extracted from all signals and several beats surrounding the labeled event

The majority of J-points are labeled in the annotation files, but several beats are not labeled due to noise in the ECG signal. Often, localized noise can affect a single lead without affecting the other leads. Because the J-point annotations in the database were computed using only lead one, it is possible to discover more J-points by examining lead two. To attempt to compute the missing J-points, a program called WQRS, provided as a part of the PhysioNet toolkit [44], is used. This program allows the specification of lead two for detection. If using this technique does not provide a J-point for the desired beat, the next J-point forward in time is used, as shown in Figure 5.3.

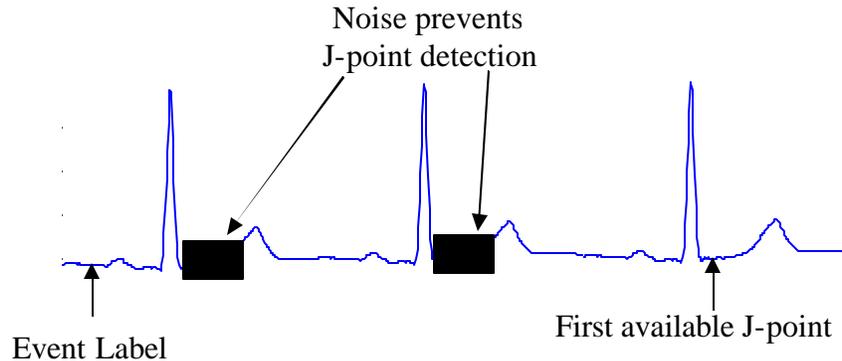


Figure 5.3 - Describes how J-point is found in presence of noise on leads one and two

The detected J-points are used to supplement the annotation files so that the ST segment and T wave can be extracted from all necessary beats.

In order to reduce computational complexity, the end of the T-wave is not calculated for each individual beat. Instead, the 400ms following the J-point is extracted. This time period effectively captures the desired signal portion.

Once the J-points have been calculated and the ST-T segments have been extracted, there are several options available for processing the data. These procedures are not necessary in all cases and vary with each specific experiment. Adjusting the data

to have a zero mean and unit variance in order to normalize the data. To make the data have zero mean, the calculated mean of the data is subtracted from the signal. To make the signal have a unit variance, the signal is divided by the calculated standard deviation of the signal (see section 6.4 for more information on these calculations). Normalizing the signal can help to reduce the effects of baseline and amplitude differences between patients.

A second option for processing the signal is to filter the ST-T complexes into frequency sub-bands. This procedure can help a classifier to focus on important dynamics found in distinct frequency bands of the signal. Figure 5.4 describes the filtering process used. First, the frequency boundaries to be used must be determined. This is done by examining previous research data and a power spectral density plot of the data (see section 6.5 for more details on choosing filter boundaries). The goal is to find frequency bands that will divide the signal into “important” energy bands. Once the filter boundaries have been determined, each of the N beats of the two to three signals for each event is filtered into bands using an infinite impulse response (IIR) filter. This creates $n \times N \times f_n$ discrete filtered signals, where n is the number of signals in the record and f_n is the number of frequency bands used.

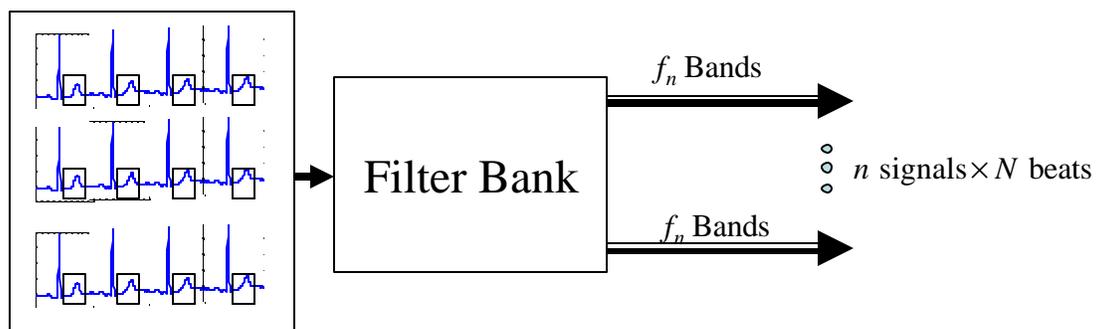


Figure 5.4 - Description of how signals are filtered

5.3 Classification Algorithm

This section describes the basic framework of the classification algorithm used in this research. The research conducted for this thesis focuses on the ST segment and the T wave as indicators of ischemia. Research has shown that deviations in these areas can be indicators of ischemia or, in later stages, infarction [2, 10-12, 14, 15]. Current techniques use these deviations directly to diagnose ischemia, but they provide too many false positives. The proposed method uses the ST segments and T wave in a novel way to improve detection accuracy.

The training algorithm flow chart is shown in Figure 5.5. After the pre-processing described in the previous section, a set of f_n frequency bands are generated for each of ECG lead. This means that there is a total of $f_n \times n \times N$ total signals for each ST event, where n is the number of signals in the record and N is the total number of surrounding beats used. The first step in training the classifier is to embed each one of these signals in a five dimensional reconstructed phase space with a time lag of five (see section 6.2 for information on parameter selection). These RPSs are then grouped by class and each matching RPS is overlaid to create a global RPS for that class. This results in $f_n \times n \times N \times M$ reconstructed phase spaces, where f_n is the number of frequency bands, n is the number of signals in the record, N is the number of beats used, and M is the total number of classes used.

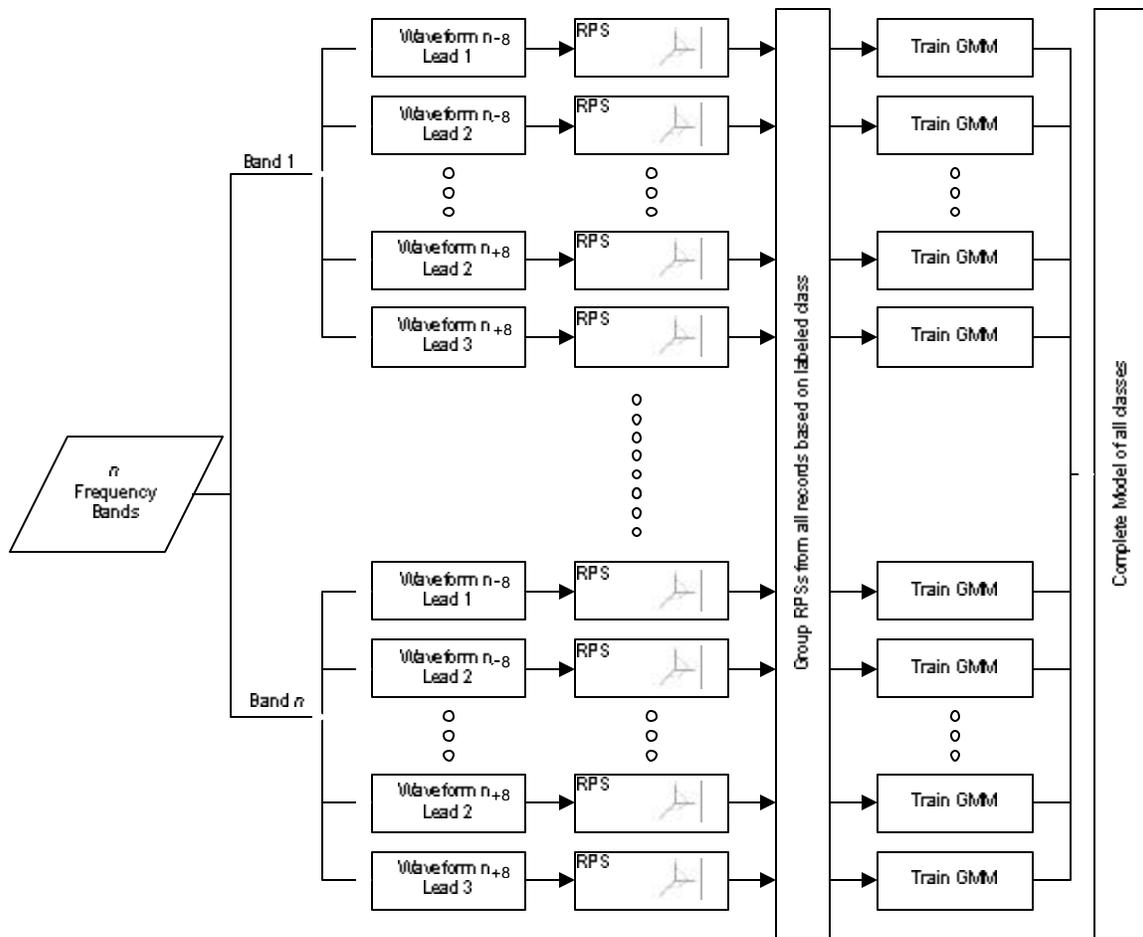
The next step in the training process is to learn the GMM model parameters. One GMM with 25 (see section 6.2 for information on parameter selection) mixtures is trained

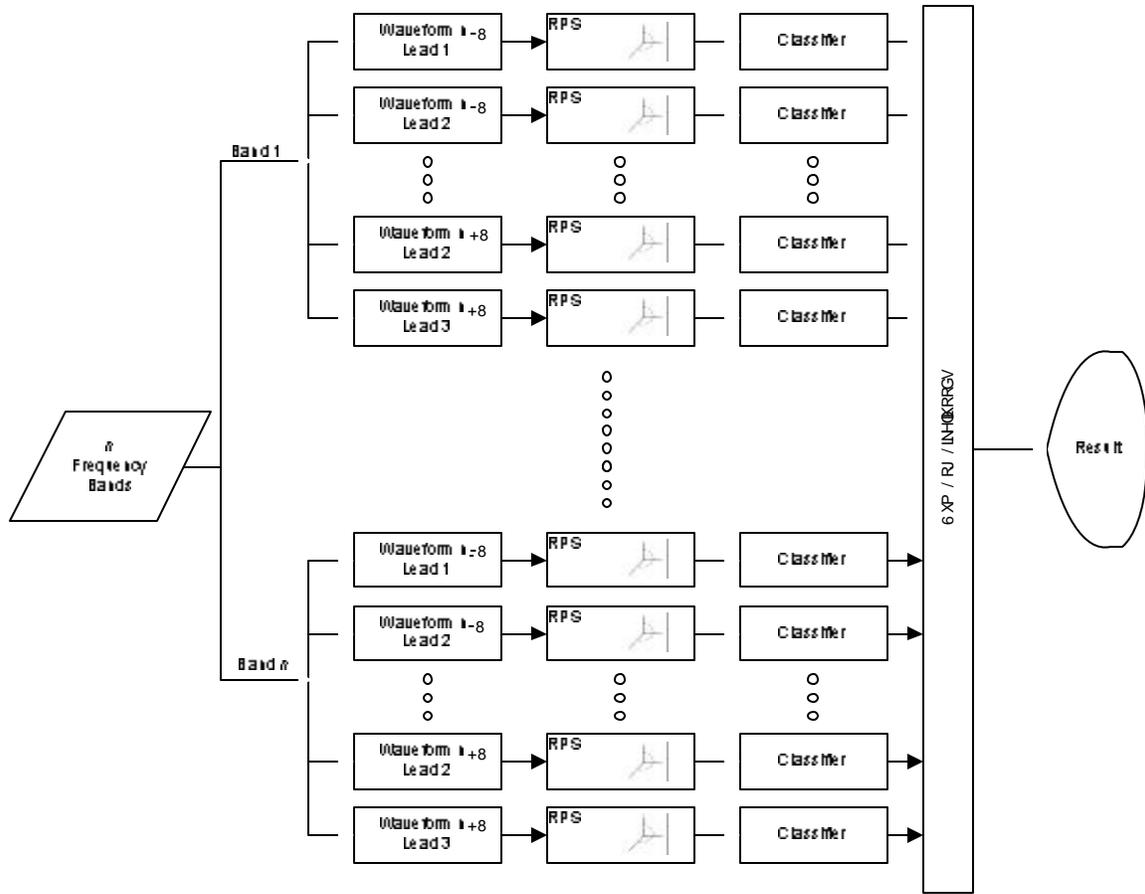
for each of the global RPSs. This creates $f_n \times n \times N \times M$ GMM classifiers or

$f_n \times n \times N \times M \times 25$ total models, where f_n is the number of frequency bands, n is the

number of signals in the record, N is the number of beats used, and M is the total number

of classes used.





$$\sum_{i=1..M} \log_{10} \hat{l}_i \quad (4.8)$$

This generates an overall log-likelihood result for each of the M classes. The class with the highest summed log-likelihood is selected as the estimated class according to:

$$\hat{\mathbf{w}} = \max_{i=1..C} \{\hat{l}_i(\mathbf{x})\}, \quad (4.9)$$

where \mathbf{x} is the test data vector and M is the number of classes [42].

6 Experiments

This chapter describes the experiments used to validate the proposed method for ischemia classification discussed in chapter 5. Each section contains an explanation of the experimental procedure, a table of results, and a discussion of how the results contribute to the overall study.

The first section of this chapter explains a technique used to validate the proposed algorithm. Section 6.2 explains how several important parameters are determined. Section 6.3 defines the set of baselines used for comparison to the results of this study. Section 6.4 evaluates the normalization of the signals input to the classifier. Section 6.5 evaluates the filtering of the signal into frequency sub-bands. Section 6.6 describes a validation of the hypothesis that using three classes (one ischemic and two non-ischemic) will improve the classification accuracy over classifications with two classes. Section 6.7 describes the experiment to validate the use of multiple beats surrounding the labeled ST event. The final section of this chapter provides a summary and discussion of the experimental results.

6.1 *Ten-Fold Cross Validation*

In experiments where the Test Set is not available, it is necessary to use a method of algorithm validation that is data/patient independent, but does not need a separate set of data for testing. The method used in this research is called ten-fold cross validation. The first step in this method is to separate the Training Set into ten groups (folds). The folds are divided to match, as closely as possible, the distribution of events and episodes of the entire ‘Training Set.’ The classifier is trained on nine of the ten folds. Then the

learned models are used to classify the tenth fold. This is repeated 9 more times, each time leaving out a different fold. In the end, 10 patient independent classifications have been conducted. The classification results are combined to calculate the overall accuracies.

Using ten-fold cross validation maintains patient independent classification without the availability of the ‘Test Set.’ This allows for an effective validation of possible algorithms while not analyzing the ‘Test Set.’ Keeping the Test Set aside until the final algorithm validation is important in making sure that the algorithm is not tailored to the data of the ‘Test Set.’ Thus, ten-fold cross validation is an accurate way of simulating ‘Test Data’ results when a testing dataset is not available [42, 45].

Implementing this procedure increases the likelihood that if the developed algorithm has high classification accuracy on the ‘Test Set,’ it will be able generalize to other data sets. It also allows results from classifications run on the Training Set to generally be compared to the Test Set results.

6.2 *Parameter Determination*

In order to create an accurate classifier of the ST event data, the parameters of the reconstructed phase space (RPS) and Gaussian mixture model (GMM) are adjusted for the data that is to be classified. The three parameters determined are the lag of the reconstructed phase space, dimension of the phase space, and the number of GMMs used for each class. There are several methods for estimating these parameters, refer to section 4.2 for more details.

In this work, the values for the model parameters are determined empirically. A simple version of the classification algorithm described in chapter 5 is run for different

combinations of parameters to determine which one provide the best overall accuracy. No normalization or sub-band filtering is done. Ten-fold cross validation is used to take full advantage of the training data.

Table 1 shows the results of a classification series over different numbers of mixtures. In this experiment, a lag of three and a dimension of three are artificially chosen. The results in the table below show that both 5 and 25 mixtures provide high levels of accuracy. Because the accuracies are only 0.7% different, a future experiment will use both values to determine the overall best.

Number of Mixtures	Dimension	Lag	Sensitivity	Specificity
2	3	3	67.2%	66.3%
3	3	3	67.4%	70.1%
4	3	3	68.9%	70.0%
5	3	3	70.0%	72.8%
6	3	3	68.9%	70.6%
7	3	3	59.9%	69.3%
8	3	3	65.9%	70.0%
9	3	3	68.5%	71.1%
10	3	3	68.6%	70.4%
11	3	3	68.5%	68.8%
12	3	3	65.4%	70.9%
13	3	3	64.1%	70.1%
14	3	3	65.0%	70.6%
15	3	3	67.1%	70.2%
16	3	3	62.3%	71.1%
17	3	3	66.1%	69.6%
18	3	3	66.0%	69.8%
19	3	3	66.1%	71.2%
20	3	3	66.6%	68.2%
21	3	3	65.3%	71.4%
22	3	3	66.0%	70.1%
23	3	3	67.1%	69.6%
24	3	3	68.4%	70.9%
25	3	3	69.9%	72.1%
26	3	3	68.9%	71.3%
27	3	3	68.1%	70.3%
28	3	3	66.6%	70.7%
29	3	3	66.2%	66.3%

30	3	3	68.1%	70.1%
----	---	---	-------	-------

Table 6.1 - Accuracy results for combinations of dimension and number of mixtures

Table 2 shows the results of a classification run with multiple values for the dimension and the two different values for the number of mixtures. This classification will be able to show which value for the mixtures combined with a value for the dimension provides the best results.

Dimension	Number of Mixtures	Lag	Sensitivity	Specificity
1	5	3	60.5%	71.8%
2	5	3	66.1%	71.9%
3	5	3	65.9%	71.4%
4	5	3	66.4%	71.8%
5	5	3	53.9%	69.9%
1	25	3	61.2%	71.1%
2	25	3	66.6%	71.5%
3	25	3	68.4%	71.9%
4	25	3	69.9%	71.8%
5	25	3	70.1%	72.2%

Table 6.2 - Accuracy results for combinations of dimension and number of mixtures

Table 3 shows the results of an experiment run with multiple values for the lag. Again, the experiment uses the structure discussed in chapter 5. 25 mixtures and a dimension of five is used based on the results of the previous two series of classifications. The results show that a lag of five provides the best results.

Dimension	Number of Mixtures	Lag	Sensitivity	Specificity
1	25	5	65.7%	71.6%
2	25	5	65.0%	68.6%
3	25	5	65.6%	73.3%
4	25	5	67.9%	70.8%
5	25	5	70.8%	74.7%
6	25	5	69.6%	72.5%
7	25	5	68.3%	72.3%
8	25	5	66.1%	72.0%
9	25	5	62.8%	73.9%
10	25	5	56.7%	72.3%
11	25	5	59.4%	72.1%
12	25	5	53.4%	73.3%
13	25	5	53.0%	71.0%
14	25	5	51.9%	70.3%
15	25	5	57.0%	72.1%
16	25	5	55.5%	72.0%
17	25	5	52.3%	71.7%
18	25	5	56.8%	70.9%
19	25	5	53.7%	70.9%
20	25	5	54.5%	72.4%

Table 6.3 - Accuracy results for different lag values

Experimentation has shown that the values that provide the best overall accuracy are a lag of 5, dimension of 5, and 25 mixtures in the GMM. The experiments in the following sections use these values. It should be noted, however, that the variation between consecutive runs of the classification with like parameters is similar to the variation among runs with different parameters. This is indication that the results may merely indicate the presence of noise and not conclusive evidence of the optimal parameter selections. It is likely that any selection of parameters from the above tables will provide similar accuracy results.

6.3 *Baselines*

Two baselines will be used to examine the experimental results. The first baseline is the Langley et al method [36] described in section 3.2. This algorithm uses threshold levels of the ST deviation to classify ECG signals. The results when the algorithm is run on the Training Set are 99.0% sensitivity and 88.8% specificity. When the algorithm is run on the ‘Test Set,’ the sensitivity stays at 99.0% and the specificity actually increases to 93.3%.

	Sensitivity	Specificity	Accuracy
Training Set	99.0%	88.8%	91.1%
Test Set	99.0%	93.3%	95.6%

Table 6.4 - Accuracy results for the Langley et al method

The second baseline experiment is a neural network method. This classification method follows the approach of [35], but applies the algorithm to the data of the Long-Term ST database. The algorithm was originally implemented using the European ST database. The classifier had a sensitivity of 73.0% and specificity of 69.45%.

The first step in this neural network algorithm is to extract the needed ST segments from the data. The data used is the 40 samples (160ms) of ST deviation following the J-point for the beat labeled as the event start. The number of samples is reduced to 20 by averaging every two samples to reduce the effects of noise. These deviation values are input into a three layer backpropagation neural network with 20 input neurons. This neural network has one hidden layer with 10 neurons and an output layer with two neurons. The output of the neural network gives a value between 0 and 1. When the values are rounded (.5 rounds up to 1) the neural network becomes a binary

classifier where 1 represents ischemia and 0 represents non-ischemic causes of the ST event.

	Sensitivity	Specificity	Accuracy
Test Set	71.3%	68.7%	69.2%

Table 6.5 - Accuracy results for Neural Network algorithm

The neural network classifier is trained on the Training Set and tested on the Test Set of the Long-Term ST database. The resulting sensitivity was 71.3% and the specificity was 68.7%. These accuracy results are similar to those that were achieved on the European ST database.

6.4 Data Normalization Experiment

The unique nature of an individual's cardiovascular system leads to ECG recordings with varied amplitudes and wave characteristics. Due to this, the data varies greatly from patient-to-patient and hour-to-hour for a specific patient. This variation is especially prevalent when the ECG waveform is affected by muscle artifact (distortion that appears in the ECG signal). In order correct to for these effects, the data must be adjusted to have standard statistical parameters.

This experiment uses the 400 ms following the ST event (which captures the ST segment and T wave) as labeled by the annotation file. The data is extracted from the lead labeled by the annotations as indicating the ST event. The mean of the signal is calculated and subtracted so that the signal is zero meaned. Next, the standard deviation of the signal is found using:

$$\mathbf{s}_x = \left(\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2 \right)^{\frac{1}{2}}, \quad (5.1)$$

where

$$\bar{x} = \frac{1}{n} \sum_{i=1}^n x_i \quad [1]. \quad (5.2)$$

The signal is then divided by the standard deviation to give a unit variance. The complete normalization is characterized by:

$$x' = \frac{(x - \mathbf{m}_x)}{\mathbf{s}_x}. \quad (5.3)$$

This normalization provides a set of signals with the same first and second order statistics.

Next, the training data is transformed into a reconstructed phase space (see section 6.2 for information on the parameter values used). The RPSs are combined into classes based on the label assigned to them by the database annotations. All of the RPSs labeled as ischemic are concatenated to form one global phase space. The same procedure is applied for the non-ischemic phase spaces. Once the class sets are created, they are modeled with Gaussian mixture models. This creates two trained statistical classification models.

Signal classification is done by comparing the learned models with the test data to be classified. In order to utilize all available data, ten fold cross validation is implemented (see section 6.1 for more information).

	Classified As	
	Ischemic	Non-Ischemic
Ischemic	222	109
Non-Ischemic	460	501

Table 6.6 – Ten-fold cross validation confusion matrix for normalizing experiment on Training Set

	Sensitivity	Specificity	Accuracy
'Learning Set' – Ten Fold Cross Validation	67.1%	47.9%	56.0%
Non-normalized results	70.8%	74.7%	73.7%

Table 6.7 – Ten-fold cross validation accuracy results for normalizing experiment on Training Set

The above results show that, for this experiment, normalizing the data actually causes a decrease in sensitivity from 70.8% to 67.1% and a decrease in specificity from 74.7% to 47.9% when compared with the same experiment done without normalizing the data (see section 6.2). This result is not entirely unexpected as it is thought that there may be dynamics of the signal that are lost during normalization. The experiment confirms that there are features of the non-normalized signal that have indications of ischemia. When the data is normalized, these features are lost and the classifier is unable to distinguish between classes. It is clear that normalization is not a good approach for the proposed classification method.

6.5 Sub-Band Filtering

This section describes an experimental approach known as sub-band filtering. In this method, the signal is filtered into discrete frequency bands using a set of IIR filters. Each band is treated as a separate signal and the class likelihoods are calculated. The likelihoods derived from the GMMs are joined together using a summation method called fusion. This classification method is currently being used with high levels of success to classify other cardiac malfunctions in ECG signals [46]. Two versions of implementing this method are described in this section.

Selection of filter boundaries is important for this approach to yield high accuracies during classification. The goal is to select boundaries that separate the signal into sub-bands containing band-independent information. Analyzing the power spectral density (PSD) of the ST segments, like the one shown in Figure 6.1, and taking into account the results from [46], two sets of boundaries are determined.

The first set creates three sub-bands with two boundaries at 8.75Hz and 32.0Hz. Figure 6.1 shows that the first frequency band captures the large energy of the signals low frequency components. The value of 32.0Hz because it has been shown that frequency bands above this frequency are due to noise and not the ECG signal. Consequently, the band above 32 Hz is dropped, leaving a total of two frequency bands.

The second frequency boundary set uses three boundaries to create three bands after the highest frequency band is dropped. Figure 6.2 shows the boundaries that are chosen on the power spectral density plot. The three boundaries are 8.75Hz, 18.75Hz, and 32.0Hz. The 8.75Hz boundary is again chosen to capture the high energy in a low frequency band. The 18.75Hz boundary separates the mid-frequency band into two approximately equal bands. This should allow a classifier to better model this portion of the signal. The final boundary, at 32.0Hz, again is used to separate the drop the high frequency noise.

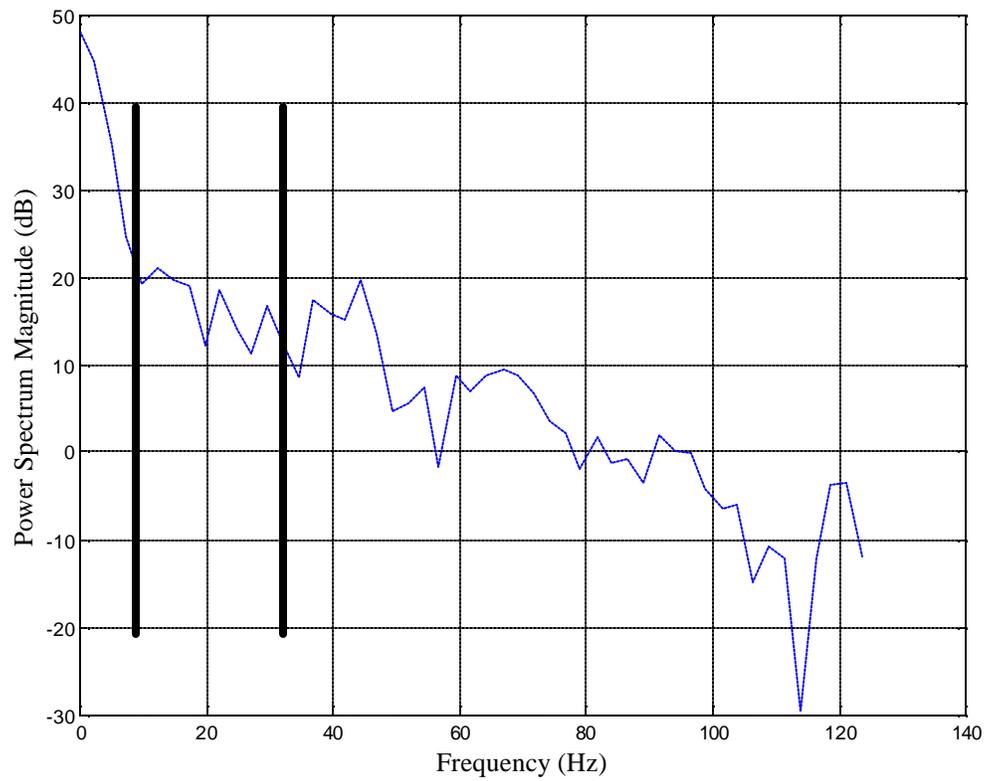


Figure 6.1 - Power spectral density of ECG ST segment in Training Set with frequency boundaries of 8.75Hz and 32.0Hz labeled

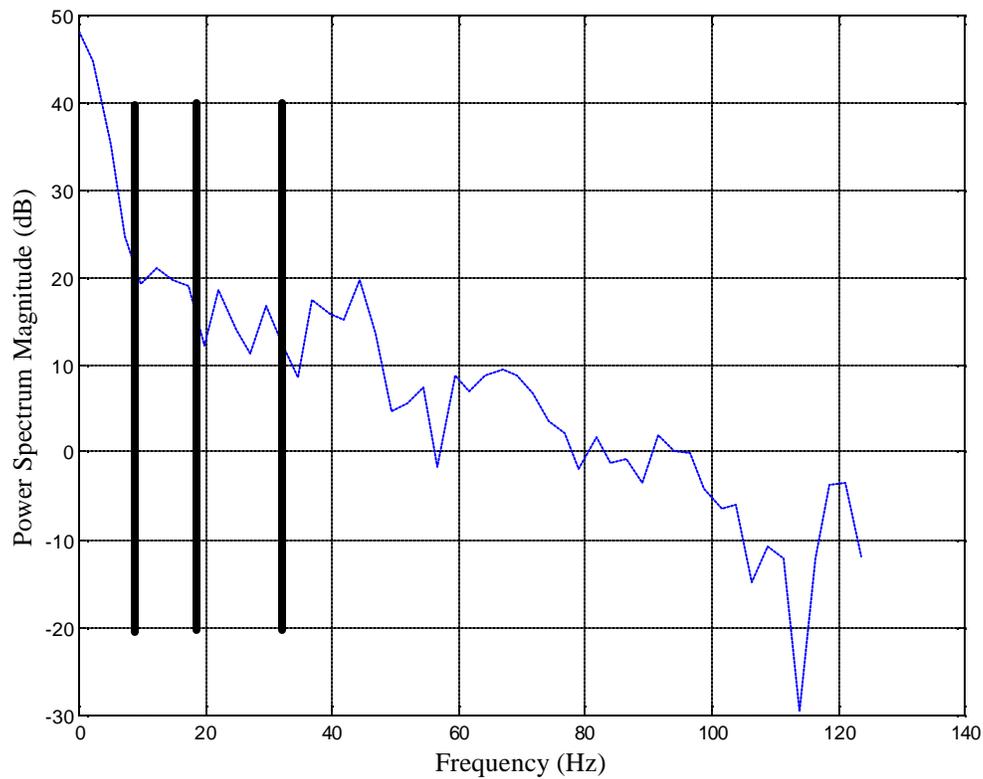


Figure 6.2 - Power spectral density of ECG ST segment in Training Set with frequency boundaries of 18.75Hz, 18.75Hz, and 32.0Hz labeled

Once the signals are separated into frequency bands, GMM models are trained for each combination of class and frequency band. The ST events to be classified are similarly filtered into frequency bands and then compared with the learned models. The Naïve Bayes classifier method is used to classify the events into classes. The results of a ten-fold cross validation implementation of this method are shown below.

	Classified As	
	Ischemic	Non-Ischemic
Ischemic	151	180
Non-Ischemic	570	391

Table 6.8 – Ten-fold cross validation confusion matrix for 2 boundary sub-band experiment on Training Set

	Classified As	
	Ischemic	Non-Ischemic
Ischemic	175	156
Non-Ischemic	600	361

Table 6.9 – Ten-fold cross validation confusion matrix for 3 boundary sub-band experiment on Training Set

	Sensitivity	Specificity	Accuracy
2 Boundaries	45.6%	40.7%	42.0%
3 Boundaries	52.9%	37.6%	41.5%
Non-banded	70.8%	74.7%	73.7%

Table 6.10 – Ten-fold cross validation accuracy results for sub-band experiments on Training Set

The results show that both the sensitivity and specificity are lower than the baseline for both filter-boundary choices. The first set of filter bands has a slightly higher specificity than the second, while the second set of filter bands has a slightly higher sensitivity than the first. These results are somewhat unexpected, as this method has had much better accuracy when detecting other arrhythmia events [46]. One hypothesis is that ST-T complex is too short to split into effective frequency bands. Other experiments using sub-band filtering have used the entire beat, which has several frequency components that correspond to different stages of cardiac function.

6.6 Three Class Experiment

This experiment confirms the hypothesis that breaking the non-ischemic class into two classes will improve classification accuracy. By splitting the classes into more precise models, the classifier can better learn the distribution of each specific class. After the classification is completed, all of the non-ischemic classes are folded together to get the overall results. This means that all events that are classified as any type of non-

ischemic cause are considered to be correct, even if they are not classified as the correct non-ischemic cause.

The ST events in the Long Term ST database have been classified into four classes: ischemic, axis shift, conduction change, and heart rate-related shifts. The ischemic class is all events that have been determined to be caused by ischemia. The axis shift class represents events that correspond to a non-ischemic ST level shift caused by a change of the physical path of conduction with respect to the leads placed on the patient's chest. The conduction change class corresponds to a non-ischemic ST shift caused by a change in the cardiac conduction pattern. The non-ischemic heart rate-related class corresponds to a change in the ST level due to a change in heart rate. The conduction change data is distributed between two records in the 'Training Set.' One record contains 429 conduction change events and a second record contains 12 events. This distribution demands a combination of two classes to allow for ten-fold cross validation. Because the axis shift and conduction change classes both represent ST shifts as opposed to the ST episodes of the heart rate-related class, they are treated as one class.

In order to use all available data, in this experiment, all available leads (signals) are used (in previous experiments only the lead identified by the annotations was used). If the ST event is visible on the other leads of the data, they will be helpful in making the classification. To better model each specific lead, separate models of the three classes are created for each lead. This increases the total number of classes from three to nine, since there are up to three leads per record.

To train the classifier, each signal of the training set is embedded into a reconstructed phase space. This means that the ST segment for every lead at the labeled

event beat is embedded. Then the RPSs for each lead and class combination are overlaid to make one global RPS. For example, all segments from lead 1 and the ischemic class are combined. Next, one GMM classifier is trained for each global RPS. This creates a total of $N \times L$ classifiers, where N is the number of classes and L is the maximum number of leads for any segment.

To validate the classifier, ten-fold cross validation is used. All of the leads of the event are converted into their own RPS. The RPS is then compared with the matching lead of each class. This returns a log-likelihood that the lead belongs to that class. The log-likelihoods are then summed across all of the leads for each class to get an overall likelihood that an event belongs to a certain class. The class with the highest likelihood is selected as the estimated class. Figure 6.3 gives an overview of this classifier. Table 6.11 and Table 6.12 provide the results of ten-fold cross validation using this classifier.

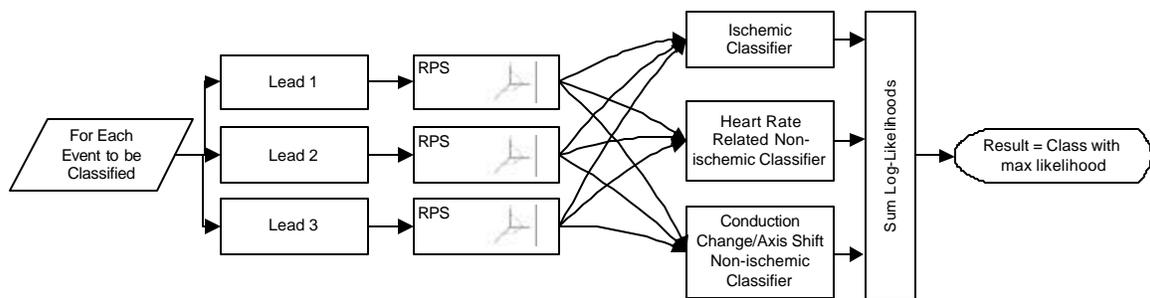


Figure 6.3 - Flowchart of the three-class classifier

	Classified As		
	Ischemic	Conduction Change/Axis Shift	Heart Rate-Related Shift
Ischemic	261	47	23
Conduction Change/Axis Shift	120	662	42
Heart Rate-Related Shift	6	15	116

Table 6.11 – Ten-fold cross validation confusion matrix for 3 class experiment on Training Set

In order to compare these results to other experiments, all of the non-ischemic results (conduction change/axis shift and heart rate-related shift) are folded into one non-ischemic class for the accuracy calculations.

	Sensitivity	Specificity	Accuracy
Training Set – Ten Fold Cross Validation	78.9%	86.9%	84.8%
Single Class Experiment	70.8%	74.7%	73.7%

Table 6.12 – Ten-fold cross validation accuracy results for 3 class experiment on Training Set

These results show that converting to three classes and using all available leads gave a very significant improvement in results over previous experiments. These are the results that were expected since more data was being used in modeling and classification and the data was grouped more logically. Clearly, the classifier is much better at distinguishing the non-ischemic data when it is broken down into two available classes.

6.7 Multiple Beat Classifications

In order to take advantage of the ischemic condition's progression over time, the proposed classifier uses the 16 beats surrounding the labeled event. In the Long Term ST database, the ST level must drop below $50\mu\text{V}$ to be labeled as an ST event. By examining the eight beats before the event occurs, it may be possible to find indications of ischemia

or other changes when the ST level is beginning to drop. Additionally, to be labeled as an event, the ST changes must occur for at least 30 seconds [2]. This means that ST changes are still apparent during this time. By modeling and classifying using the eight beats following the event, it may be possible to develop a more accurate classification method.

During the training phase, the ST segment and T Wave (400 ms following labeled J-Point) are extracted from the beats before and after the labeled ST event of the training set. Each signal in the set is transformed into a reconstructed phase space (every lead for every beat) with dimension five and lag five. Each of the RPSs with a matching class, lead, and waveform index are overlaid to create global phase spaces. A GMM with 25 mixtures is then learned for each of the global phase spaces. This creates $N \times L \times M$ classifiers, where N is the number of classes, L is the maximum number of leads in any record, and M is the number of beats used. In this study $N=3$, $L=3$, $M=16$, so there are 144 total classifiers trained.

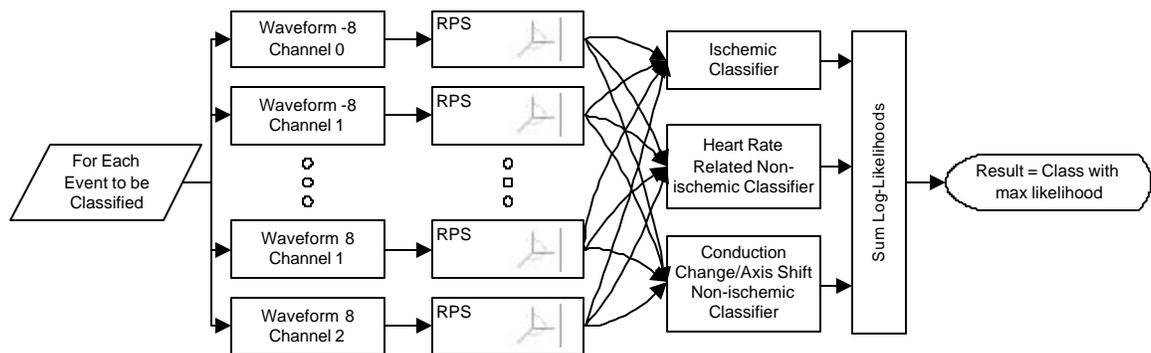


Figure 6.4 - Flowchart of the complete algorithm for classification

Figure 6.4 shows a flowchart of the process used to classify an ST event using the proposed algorithm. The ST segment for each lead of each of the eight beats before the event and the eight beats after the event must be extracted. Each combination of lead and

waveform index is then compared with the same lead-index combination for the three possible classes. The log-likelihoods for each class are then summed across the lead-index combinations to get an overall class likelihood. The class with the highest likelihood is selected.

	Classified As		
	Ischemic	Conduction Change/Axis Shift	Heart Rate-Related Shift
Ischemic	266	48	17
Conduction Change/Axis Shift	112	693	19
Heart Rate-Related Shift	3	15	119

Table 6.13 – Ten-fold cross validation confusion matrix for multiple beat experiment on Training Set

	Classified As		
	Ischemic	Conduction Change/Axis Shift	Heart Rate-Related Shift
Ischemic	604	70	49
Conduction Change/Axis Shift	244	813	102
Heart Rate-Related Shift	6	6	80

Table 6.14 – Confusion matrix for 16 beat experiment on Test Set

Beats Before/Beats After	Sensitivity	Specificity	Accuracy
2	77.9%	79.3%	78.9%
4	79.2%	81.2%	80.7%
6	79.7%	85.4%	83.9%
8	80.4%	88.0%	86.1%

Table 6.15 – Accuracy results for Training Set

Beats Before/Beats After	Sensitivity	Specificity	Accuracy
2	77.0%	78.8%	78.2%
4	80.0%	80.0%	80.0%
6	80.9%	80.1%	80.4%
8	84.6%	85.5%	85.3%

Figure 6.5 - Accuracy results for Test Set

These results show significant improvement from the previous methods. The training data overall accuracy improved by 1.3% over the three class experiment. The method's accuracy was 12.1% better than the neural network baseline method, but was still 14.3% lower than the Langley et al method. The ability of the method to generalize to unseen data is evident from the high accuracy level in both ten-fold cross validation and classification on a never-before-seen test data set.

6.8 Results Comparison

Table 6.16 shows a summary of all of the experimental results for this thesis. It also shows the two baseline classification accuracies for comparison.

Experiment	Sensitivity	Specificity	Accuracy
Baseline 1: Langley Classifier ¹	99.0%	88.8%	91.1%
Baseline 1: Langley Classifier ²	99.0%	93.3%	95.6%
Baseline 2: Neural Network ¹	71.3%	68.7%	69.2%
Simple Experiment from Parameter Determination ¹	70.8%	74.7%	73.7%
Normalizing Data 2 Bands ¹	67.1%	47.9%	56.0%
Filter Banks 2 Bands ¹	45.6%	40.7%	42.0%
Filter Banks 3 Bands ¹	52.9%	37.6%	41.5%
Three Classes ¹	78.9%	86.9%	84.8%
Multiple Beats ¹	80.4%	88.0%	86.1%
Multiple Beats ²	84.6%	85.5%	85.3%

Table 6.16 - Summary of experimental results and baseline accuracies

The best classification results come from the final experiment, which uses three classes and eight beats on either side of the labeled event. This experiment has similar accuracies for both the 'Learning Set' ten-fold cross validation and the 'Test Set.' It is expected that these accuracies would be the same because both sets are classified using

¹ Ten-fold cross validation on 'Learning Set'

² Test Set

an independently trained classifier. While the experiment was able to classify more accurately than the Neural Network classifier, it was not able to exceed the overall accuracy of the Langley classifier. The proposed method was able to match the specificity of the Langley method when operating on the 'Training Set.' The sensitivity, however, is considerably lower for both the Training Set and the 'Test Set.'

There are several possibilities for the cause of the difference in classification accuracy from the Langley method. First, it is possible that the durations of ST deviation levels are more important than the shape and trajectory of the ST segment itself. This would explain why a method based on thresholds of ST deviation level would have a much higher accuracy. This is the most damaging hypothesis to the proposed algorithm, since it does not rely on time durations at all. A second hypothesis is that different RPS parameters and numbers of surrounding beats should be used. The parameters should be empirically found using the final classification method developed. It may be possible to find a combination which yields better classification accuracies.

The greatest contribution to classification accuracy comes from splitting the non-ischemic class into two sub-classes. Splitting the classes improves overall accuracy by 11.1%. Even if the proposed algorithm is not used in the future, the use of sub-classes could be applied to other classification approaches. There appears to be distinct differences between classes that a generalized classifier is not able to recognize.

It is interesting to note that the multiple beat classifier has a higher specificity than sensitivity. This is contrary to the operation of most other classification methods. This means that the algorithm would have more missed alarms than false alarms. On its own, this would not be an appropriate method for clinical use, since patients would not be

notified of important health conditions. It may be possible, however, to combine this method with a high sensitivity, low specificity method to achieve a balanced overall accuracy.

7 Conclusion

A novel method for the classification of myocardial ischemia has been presented in this thesis. Several experiments, which analyze the effectiveness of the proposed method, have been presented and discussed. The method presented captures dynamical data present in the ST segment of the ECG waveform by embedding the signal in a reconstructed phase space. This approach is different from other approaches, which have been presented in the past. Because the approach is unique, it lends itself well to the possibility of combination with other methods. It may be possible to increase classification accuracy by finding a classifier that has a mutually exclusive set of errors from the proposed algorithm.

A second strategic part of this algorithm is how it takes advantage of an ST event's progression over time by using the ST segments surrounding the event start time. Since ischemia is a condition that generally develops over time, it may be possible to recognize the signs of a developing cardiac condition. This is especially helpful in distinguishing ischemia from sudden axis shifts and conduction changes. This could likely be a contributing factor as to why the specificity of the algorithm is higher than the sensitivity.

Finally, the algorithm also increases accuracy by breaking the non-ischemic class into two sub-classes. This procedure introduced an 11.1% increase in accuracy over the simple experiment run during parameter determination. It is a very significant increase for this type of experiment. Clearly, the classifier is more much more accurate when it is able to model the heart rate-related events separately from the axis shifts and conduction changes.

While the algorithm was not able to get a higher accuracy than the best baseline method (Langley et al method), it was able to achieve significantly higher accuracy than the Neural Network. The accuracy results of this classification approach are promising, but this method is not yet ready for clinical use. In order to be accepted for use, more research must be done to improve sensitivity and specificity.

7.1 Future Work

While this classification algorithm does not have accuracies that exceed current methods, it is important to note that it has been under development for a relatively short time. There are many possibilities for future research that may be able to generate an increase in classification accuracy. The future areas generally fall into two categories: development of new feature sets and combination with other classification methods.

The classification algorithm described herein used the ST segment based on its popularity and the accepted idea of injury current. It may be possible, however, to find other features that may help to improve the classification accuracy. For example, it may be possible to use other portions of the ECG waveform besides only the ST segment and T wave. On a related note, it may be helpful to generate separate models for both the ST segment and T wave so that a classifier can search for details in how each wave is affected by ischemia. In addition, since it is known that one of the non-ischemic classes corresponds to a change in the heart rate, it is likely that incorporation of the patient's heart rate changes would be a helpful feature.

Another possible classification improvement might come from examining different combinations of the beats surrounding the event label. Currently, a symmetric number of beats (same number before and after the event) are used. Based on research

regarding the length of ischemia development and the average duration, it might be valuable to adjust the beat distribution. It may also be helpful to examine beats further into the future, since ischemia labeled in the Long-Term ST database must be at least 30 seconds (approximately 40 beats) in duration (based on a heart rate of 80 BPM).

The second category for future research work is in combining this classification method with other methods to increase classification accuracies. An especially promising algorithm for combination is the Langley method described in this thesis. This method uses characteristics of the ST deviation to classify. Because the algorithm proposed in this thesis uses the dynamical information of the ST segment and T wave, the errors from both methods may not be the same. Since the Langley classifier has such a high level of sensitivity and a lower specificity, it erroneously classifies non-ischemic beats as ischemia. The most probable way to implement this combination is to reclassify the events that the Langley classifier calls ischemic. Additionally, since the Langley classifier labels all heart rate-related shifts as ischemia, use of the heart rate's deviation may be beneficial.

Bibliography

- [1] "Preventing Heart Disease and Stroke," Centers for Disease Control and Prevention, 2004.
- [2] G. S. Wagner, *Marriott's Practical Electrocardiography*, 9th ed. Baltimore, MD: Williams & Wilkins, 1994.
- [3] *Prevention and rehabilitation in ischemic heart disease*. Baltimore: Williams & Wilkins, 1980.
- [4] *Principles of critical care*. New York: McGraw-Hill, Health Professions Division, 1998.
- [5] "Pathophysiology and therapeutics of myocardial ischemia," presented at A.N. Richards Symposium, Philadelphia, 1976.
- [6] "Computers in Cardiology Challenges," vol. 2003. Cambridge, MA, 2003.
- [7] P. D. Richard E. Klabunde, "Cardiovascular Physiology Concepts," 2002.
- [8] S. Garas, "Coronary artery spasm," vol. 2004: US National Library of Medicine, 2002.
- [9] T. H. Institute, "Silent Ischemia," vol. 2004: Texas Heart Institute, 2003.
- [10] H. Pardee, "An electrocardiographic sign of coronary artery obstruction," *Arch Int Med*, vol. 26, pp. 244-257, 1920.
- [11] H. Fozzard and D. DasGupta, "ST-segment potentials and mapping. Theory and experiments," *Circulation*, vol. 54, pp. 533-537, 1976.

- [12] M. Janse and A. Kleber, "Electrophysiological changes and ventricular arrhythmias in the early phase of regional myocardial ischemia," *Circulation Research*, vol. 49, pp. 1069-1081, 1981.
- [13] T. Fujimoto, T. Peter, H. Hamamoto, and W. J. Mandel, "Electrophysiological Observations During the Spontaneous Initiation of Ischemia-Induced Ventricular Fibrillation," *American Heart Journal*, vol. 105, pp. 189-200, 1983.
- [14] E. Braunwald and P. Maroko, "ST-segment mapping. Realistic and unrealistic expectations," *Circulation*, vol. 54, pp. 529-532, 1976.
- [15] J. M. Foody, *Preventive cardiology : strategies for the prevention and treatment of coronary artery disease*. Totowa, N.J.: Humana Press, 2001.
- [16] B. Drew and M. Krucoff, "Multilead ST segment monitoring in patients with acute coronary syndromes: a consensus statement for healthcare professionals," *Am J Crit Care*, vol. 8, pp. 372-388, 1999.
- [17] F. Jager, A. Taddei, G. Moody, M. Emdin, G. Antolic, R. Dorn, A. Smrdel, C. Marchesi, and R. Mark, "Long-term ST database: a reference for the development and evaluation of automated ischaemia detectors and for the study of the dynamics of myocardial ischaemia," *Medical and Biological Engineering and Computing*, vol. 41, pp. 172-182, 2003.
- [18] F. Jager, A. Taddei, M. Emdin, G. Antolic, R. Dorn, G. Moody, B. Glavic, A. Smrdel, M. Varanini, M. Zabukovec, S. Bordigiago, C. Marchesi, and R. Mark, "The Long-Term ST Database: A Research Resource for Algorithm Development and Physiologic Studies of Transient Myocardial Ischemia," *Computers in Cardiology 2000*, vol. 27, 2000.

- [19] *Coronary Heart Disease (Arteriosclerotic Coronary Artery Disease; Ischemic Heart Disease)*, 37 ed: Stamford: Appleton & Lange, 1997.
- [20] HCUPnet, "Healthcare Cost and Utilization Project." Rockville, MD: Agency for Healthcare Research and Quality, 2003.
- [21] S. M. M.D., "Aortic Angiography," in *Family Practice Notebook: Family Practice Notebook*, LLC, 2000, pp. 4416.
- [22] D. A. Kohli, "CT Coronary Angiograms," *Und J Raiol Imag*, vol. 13, pp. 11, 2003.
- [23] V. Fuster, *Hurst's The heart*, 10th ed. New York: McGraw-Hill Health Professions Division, 2001.
- [24] *Cardiac Nursing*, 4th ed. Philadelphia: Lippincott Williams & Wilkins, 2000.
- [25] M. D. Thippeswamy H. Murthy, "Echocardiogram," in *Medical Encyclopedia*. Washington D.C.: National Institutes of Health, 2002.
- [26] A. Støylen and K. Bjørnstad, "Stress echocardiography."
- [27] A. S. Bensky, W. Covitz, and R. H. DuRant, "Primary care physicians' use of screening echocardiography," *Pediatrics*, vol. 103, pp. e40, 1999.
- [28] heartfailure.org, "Echocardiogram," heartfailure.org, 2004.
- [29] R. Gianrossi, R. Detrano, D. Mulvihill, K. Lehmann, P. Dubach, A. Colombo, D. McArthur, and V. Froelicher, "Exercise-induced ST depression in the diagnosis of coronary artery disease. A meta-analysis," *Circulation*, vol. 80, pp. 87-98, 1989.
- [30] S. V. Williams, S. D. Fihn, and R. J. Gibbons, "Guidelines for the management of patients with chronic stable angina: diagnosis and risk stratification," *Ann Intern Med*, vol. 135, pp. 530-47, 2001.

- [31] M. D. D. M.D., "Ordering and Understanding the Exercise Stress Test."
Greenville, NC: Academy of Family Physicians, 1999.
- [32] K. Channer and F. Morris, "ABC of clinical electrocardiography: Myocardial ischaemia," *Bmj*, vol. 324, pp. 1023-6, 2002.
- [33] J. Edhouse, W. J. Brady, and F. Morris, "ABC of clinical electrocardiography: Acute myocardial infarction-Part II," *Bmj*, vol. 324, pp. 963-6, 2002.
- [34] F. Morris and W. J. Brady, "ABC of clinical electrocardiography: Acute myocardial infarction-Part I," *Bmj*, vol. 324, pp. 831-4, 2002.
- [35] N. Maglaveras, T. Stamkopoulos, C. Pappas, and M. Strintzis, "Use of neural networks in detection of ischemic episodes from ECG leads," *Neural Networks for Signal Processing*, pp. 518-524, 1994.
- [36] P. Langley, E. Bowers, M. D. J Wild, J. Allen, A. Sims, N. Brown, and A. Murray, "An Algorithm to Distinguish Ischaemic and Non-ischaemic ST Changes in the Holter ECG," presented at Computers in Cardiology, Thessaloniki, Greece, 2003.
- [37] F. Takens, "Detecting strange attractors in turbulence," presented at Dynamical Systems and Turbulence, Warwick, 1980.
- [38] T. Sauer, J. A. Yorke, and M. Casdagli, "Embedology," *Journal of Statistical Physics*, vol. 65, pp. 579-616, 1991.
- [39] H. Kantz and T. Schreiber, *Nonlinear time series analysis*. Cambridge: Cambridge University Press, 1997.
- [40] L. Liu and J. He, "On the use of orthogonal GMM in speaker recognition," presented at ICASSP, Tempe, AZ, 1999.

-
- [41] T. K. Moon, "The Expectation-Maximization Algorithm," in *IEEE Signal Processing Magazine*, 1996, pp. 47-59.
- [42] T. M. Mitchell, *Machine Learning*. New York: McGraw-Hill, 1997.
- [43] P. Domingos and M. Pazzani, "On the Optimality of the Simple Bayesian Classifier under Zero-One Loss," *Machine Learning*, vol. 29, pp. 103-130, 1997.
- [44] G. B. Moody, "WFDB Applications Guide," vol. 2003, 10th ed. Cambridge, MA, 2003.
- [45] S. S. Haykin, *Neural networks : a comprehensive foundation*, 2nd ed. Upper Saddle River, N.J.: Prentice Hall, 1999.
- [46] F. M. Roberts, R. J. Povinelli, and K. M. Ropella, "Rhythm Classification Using Reconstructed Phase Space of Signal Frequency Subbands," presented at Computers in Cardiology, Thessaloniki, Greece, 2003.

Appendices

Marquette University

This is to certify that we have examined
this copy of the
master's thesis by

Michael W. Zimmerman, B.S.

and have found that it is complete
and satisfactory in all respects.

The thesis has been approved by:

Richard J. Povinelli Ph.D. - Department of Electrical Engineering, Thesis Director

Michael T. Johnson Ph.D. - Department of Electrical Engineering

Kristina M. Ropella Ph.D. - Department of Biomedical Engineering

Approved on
