

Automatic detection of Atrial Fibrillation

by

Greg Hayes

A thesis
submitted to the Victoria University of Wellington
in fulfilment of the
requirements for the degree of
Master of Science
in Electronic and Computer Systems.

Victoria University of Wellington
2014

Abstract

Atrial Fibrillation is an abnormal arrhythmia of the heart and is a growing concern in the health sector affecting 1% of the population. The incidence of atrial fibrillation increases with age and has been found to be more detrimental to long term cardiac health than previously thought. Sufferers are five times more likely to experience a stroke than others. Often, atrial fibrillation is asymptomatic and is frequently discovered only when a patient visits a hospital for other reasons. The detection of paroxysmal atrial fibrillation can be difficult. Holter monitors are used to record the ECG over long periods of time, but the resulting recording still needs to be analysed. This can be a time consuming task and one prone to errors. If a miniature, low-power, wearable device could be designed to detect and record when a heart experiences atrial fibrillation, then health professionals would have more timely information to carry out better, more cost effective courses of treatment. This thesis presents progress towards development of such a device.

Atrial fibrillation is characterised by random RR interval, missing P wave and presence of atrial activity. The detection of the P wave and atrial activity can be unreliable due to low signal levels and differences in waveform morphology between subjects. The random RR interval appears to be a more reliable method of detection. By analysing the ECG signal in both the frequency and time domains, feature sets can be extracted for the detection process. In this research, the Discrete Wavelet Transform is used to generate several sub-bands for analysing waveform morphology, and a number of RR interval metrics are created for analysing the rhythm. All features are further processed and presented to a support vector machine

classification stage for the ultimate detection of atrial fibrillation. Forty eight files from the MITDB database of the PhysioNet online ECG repository were downloaded and processed to form separate training and testing data sets. Overall classification accuracy for normal sinus rhythm was 93% sensitivity and 95% specificity, and for atrial fibrillation, 95% sensitivity and 93% specificity.

These results were found to be sensitive to the ECG morphology of the individual subjects. This means that the system either needs to be trained on a greater number of ECG morphologies or perhaps trained on the morphology of the individual under investigation. Putting this issue aside, the research to date shows that it is reasonable to expect a small, low powered, wearable device, to be capable of automatically detecting when a heart experiences atrial fibrillation.

Acknowledgments

I would like to acknowledge the support of my research supervisor Dr Paul D. Teal for his guidance during this research project. Your help was invaluable and very much appreciated. Thank you Paul.

Contents

1	Introduction	1
1.1	The Electrocardiogram	2
1.2	History and development	2
1.3	Current technology	3
1.4	Goals	4
2	Cardiology	5
2.1	Cardiac Physiology	6
2.2	ECG recording	8
2.3	Changes to ECG due to arrhythmia	11
2.4	Changes to ECG due to other causes	13
2.5	Atrial Fibrillation	14
3	Mathematical Background	17
3.1	Mathematical Morphology	17
3.2	Wavelet Transform	19
3.3	Hilbert Transform	27
3.4	Teager-Kaiser Energy Operator	28
3.5	Support Vector Machine	32
4	Research	36
4.1	Arrhythmia	37
4.2	Signal pre-processing	39

<i>CONTENTS</i>	iii
4.3 QRS complex detection	42
4.4 Rhythm analysis	44
4.5 Waveform morphology analysis	49
4.6 Classification and detection of rhythm type	54
5 Results and Discussion	57
5.1 Datasets and testing	57
5.2 Results	58
5.3 Discussion	62
6 Implementation	64
6.1 Hardware	64
6.2 Software	66
7 Conclusions	71

Chapter 1

Introduction

Medicine has advanced considerably over the past century, but surprisingly there is one diagnostic tool that has remained little changed in that time: the electrocardiogram (ECG). The ECG is used to measure the electrical activity of the heart and, along with 100 years of accumulated knowledge, is still an extremely valuable tool for diagnosing heart complaints.

With an ageing population and changing lifestyle, society is now facing increased health issues caused by over-eating, poor diets and little physical activity [7]. These issues can show up in heart related problems such as cardiovascular disease and cardiac arrhythmia [16]. Current practices can pick up cardiac problems when symptoms present, but this can be too late for some subjects that have asymptomatic cardiac problems. Often cardiac issues are discovered as a by-product of a hospital visit [16, 22]. One cardiac arrhythmia problem that is often discovered by accident is Atrial Fibrillation (AF) [16]. The purpose of this research is to design a real time detector of AF that is useful in detecting asymptomatic AF and that can help management of known cases.

1.1 The Electrocardiogram

The electrocardiogram, or electrocardiograph, is the output of an instrument that measures the electrical activity of the heart muscle. Electrodes are placed on predetermined points of the thorax, in such a way as to detect the electrical current generated by the heart muscle as it goes through its pumping cycle [2]. The signals generated are used to create a graph showing a waveform pattern. Knowledge that has been gained over the last century is used to diagnose heart problems by interpreting the morphology of the waveform produced.

1.2 History and development

As long ago as the late 17th century, scientists have been aware of the relationship between electricity and the activation of muscles in the body. First insights into the physical reaction of a muscle to an electrical stimulus were observed in 1668 by Jan Swammerdam [1], although it is unlikely he understood that it was electricity causing the reaction. It wasn't until the late 18th century, when electricity was better understood, that the effects of electricity on muscle activity was more widely known and the ability to measure electrical activity became possible.

In the beginning of the 19th century, new devices for measuring electrical current flow were developed including the galvanometer named after Italian researcher Luigi Galvani [1]. The first ECG type machines using the galvanometer were developed in the early 1800's. However, it wasn't until 1887 that physiologist Augustus Desiré Waller first recorded the electrical activity of the human heart using such a galvanometer [1]. He continued his work and coined the phrase 'electrogram' referring to the action of recording the electrical activity of the heart. This phrase was later modified by Einthoven [1] to 'electrocardiogram'. Einthoven developed the string galvanometer in 1903, an improved form of the device, and started

its manufacture by Edelmann and Sons of Munich, and subsequently, the Cambridge Scientific Instrument Company of London. Surprisingly, the first instance of an ECG being transmitted over telephone wires was accomplished in 1905 by Einthoven [1]. Vacuum tubes were used to amplify the ECG signal in 1928, and 1949 saw the advent of the Holter monitor, a 'portable' backpack version of the ECG monitor developed by Norman J. Holter [1]. Weighing 34 kgs, it was carried around by the patient to record cardiac activity. With the discovery of semiconductor transistors, many new devices have been developed making the recording and analysis of the ECG commonplace in medicine. Modern integrated electronic circuits have enabled the measuring of the ECG to be easy, accurate and affordable. New systems and technologies are now being developed with these modern devices.

1.3 Current technology

Current development in semiconductor electronics has produced at least two examples of integrated circuits (IC) that are designed specifically for ECG recording. These are the Texas Instruments ADS119x/129x [19] and Analogue Devices ADAS1000 family of ECG analogue front end ICs [11]. Both units have the necessary analogue and digital circuitry to detect, amplify and digitise the ECG signal, and offer very small form factor along with low power consumption necessary for miniature portable devices. Coupled with an appropriate microcontroller and associated software, they have the ability to detect, record and interpret the ECG waveform.

There are many ECG monitoring devices currently available, ranging from small portable Holter monitors to very large, sophisticated patient monitoring suites that monitor multiple patients, carrying out real time ECG interpretation. A few examples of Holter monitors include the Philips DigiTrak XT [29], Midmark IQholter [27] and the Schiller Medlog AR12plus

[35]. Examples of large monitoring suites are the Surveyor Central by Mortara using the S12/S19 bedside monitors along with the X12+/T12S wireless transmitters and VERITAS ECG algorithms [28] and the Philips ECG system using the TC70 ECG monitors with wireless communications, combined with the DXL-16 ECG algorithm [30].

Using advances in mathematics and digital signal processing, automatic interpretation of the ECG has become feasible. Philips have developed the computer based 'DX algorithms' [30] to automatically interpret ECG data and similarly Bionet have developed the 'EKG Plus II' algorithm [3].

1.4 Goals

In this thesis, it is shown that the ECG signal can be used as the basis for fully automatic detection of atrial fibrillation. The fundamental signal processing steps to achieve this are presented, including Mathematical Morphology, the Hilbert Transform, the Discrete Wavelet Transform and the Teager-Kaiser Energy Operator. It is shown that the ECG associated with AF has distinctive features that allows AF to be distinguished from normal sinus rhythm and from other cardiac arrhythmia.

The thesis starts with a discussion of Cardiology in Chapter 2, followed by Mathematics in Chapter 3. Proposed new methods for the detection of AF in real time are presented in Chapter 4, and Chapter 5 details the tests involved to validate the proposed new methods and discusses results. Chapter 6 describes the target implementation of these methods and finally, concluding remarks and ideas for future research are provided in Chapter 7.

Chapter 2

Cardiology

Introduction

The heart, in simple terms, is a muscle much like other muscles of the body. However, it has some major physiological differences that allow it to carry out its primary function as a pump. It has a very sophisticated electrical system that controls the activities of the cardiac muscle, synchronising the two top chambers (the atria) with the two bottom chambers (the ventricles) to create an efficient pumping mechanism [2]. The pumping cycle is initiated by an electrical impulse generated within the sino-atrial (SA) node causing the atrial muscles to contract, pumping blood from the atria into the ventricles. This electrical impulse travels across the atria in a wavelike fashion and is detected by the atrio-ventricular (AV) node and channelled down the Bundle of His to the ventricles. A small delay in conduction allows the ventricles to fill with blood before contracting and pumping blood to the lungs and the rest of the body [15].

2.1 Cardiac Physiology

Figure 2.1 shows a cut away view of the heart displaying the two top chambers, the Atria, the two bottom chambers, the ventricles, the SA node, the AV node and the Bundle of His including the Left and Right Bundle Branches.

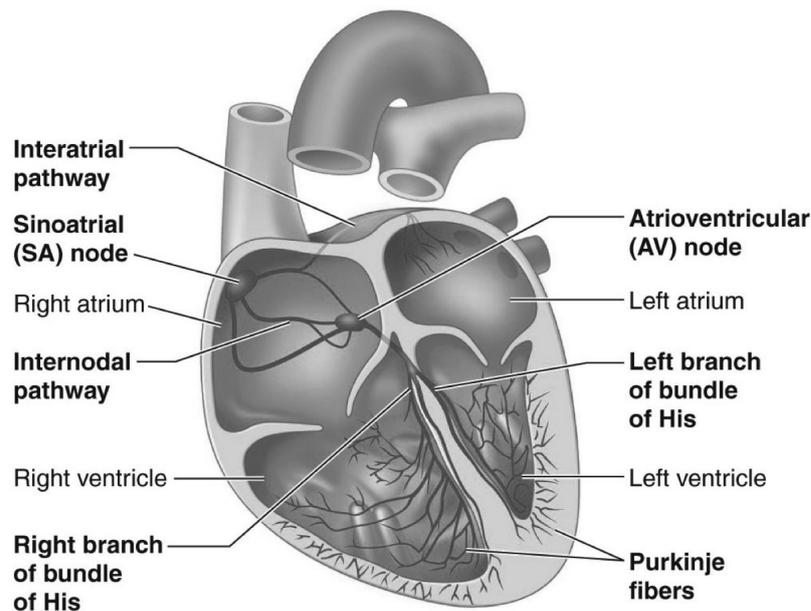


Figure 2.1: Cardiac Physiology. (Courtesy:www.studyblue.com)

Contraction and relaxation of the heart muscle is caused by changes in electric potential between the inside and outside of the muscle cell. This 'action' potential is controlled by ion channels that manipulate the flow of calcium (Ca^{2+}), Sodium (Na^+) and Potassium (K^+) ions into and out of the cell structure [2]. Figure 2.2 (bottom trace, solid line) shows the four phases of action potential as the heart muscle cells contract and relax. In phase 0, rapid de-polarisation occurs where opening of the fast sodium channels causes inward movement of Na^+ ions. Phase 1 sees the sodium channels close. In Phase 2, the long plateau is sustained by inward movement of Ca^{2+} ions through the Calcium channels and outward

movement of K^+ ions through the slow Potassium channels. During Phase 3, re-polarisation occurs by closing the Calcium channels but keeping the Potassium channels open. Phase 4 is the resting phase of the cycle and produces a potential difference of around $-90mV$ [2]. The cycle is triggered by impulses from adjacent cells.

Cells in the SA node and other pacemaker sites operate slightly differently, see the dashed trace in Figure 2.2 . They are able to de-polarise automatically as the action potential in Phase 4 doesn't remain at $-90mV$. Instead the action potential rises slowly and upon reaching a threshold potential of about $-40mV$ triggers the cycle initiating a heartbeat [2].

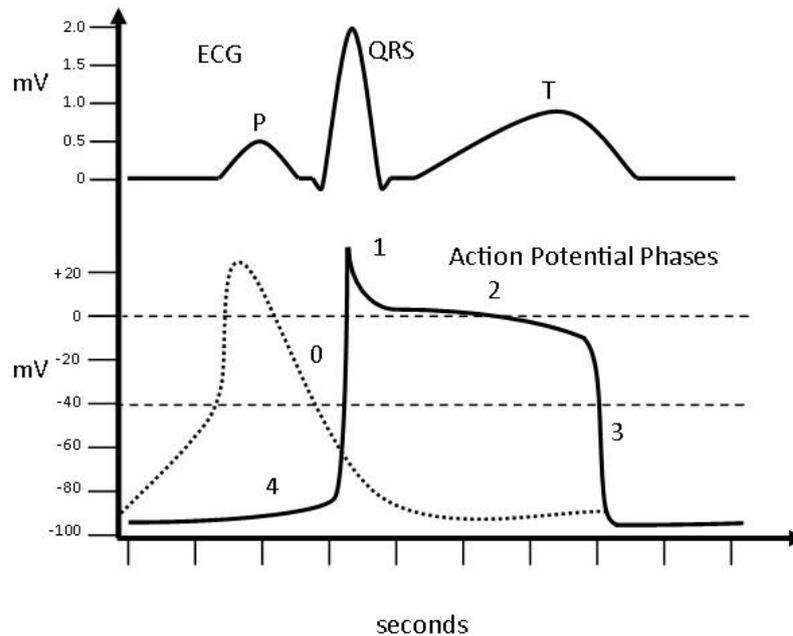


Figure 2.2: Action Potential. The bottom trace shows the action potential of contractile cells (solid line) in the ventricle muscle and pacemaker cells (dashed line) in the atrial muscle, with the top trace showing the resultant ECG.

2.2 ECG recording

The normal ECG waveform is composed of three basic components, the P-wave, the QRS complex and the T-wave (see Figure 2.3). As the electro-chemical processes involved in heart muscle contraction occur, electrical current is passed through body tissue and is detected by ECG electrodes placed on the skin [2].

The P-wave is indicative of the depolarisation of the two top chambers; the Atria. The SA node is responsible for generating the regular electrical impulse that initiates the contraction of the Atria and, consequently, the observable P-wave of the ECG trace. This ‘electrical spark’ from the SA node causes adjacent heart muscle cells to initiate their own depolarisation thus causing an electrical wave to travel throughout the atria triggering more cells to depolarise, manifesting heart muscle contraction. The overall effect is to eject blood from the atria to fill the two bottom chambers, the ventricles.

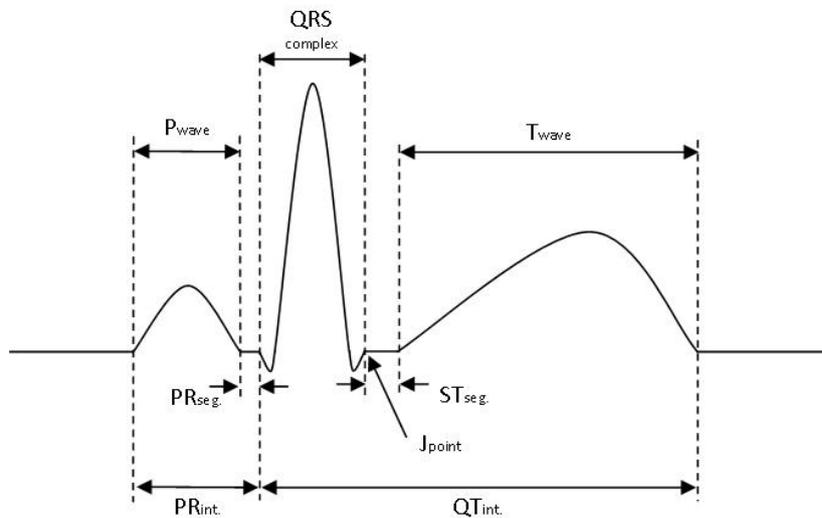


Figure 2.3: ECG waveform showing features used in interpretation.

The QRS complex is made up of three distinct waveforms, the Q, R and S waves, but not all are necessarily present on the ECG recording. The

QRS complex, in whatever form, is the observable indication of the depolarisation of the ventricles. The electrical stimulus for this action is derived from the impulse generated by the SA node being detected by the AV node and transmitted via the the Bundle of His for distribution throughout the ventricles. There is an inherent delay in the conduction time of the impulse through the Bundle of His that allows the ventricles to fill before contracting and pumping blood to the rest of the body. This conduction delay can be observed in the ECG waveform as the time between the P-wave and the onset of the QRS complex, called the PR segment (PR_{seg} in Figure 2.3). The T-wave is indicative of the re-polarisation of the ventricles. The ECG waveform shown in Figure 2.3 is commonly referred to as Normal Sinus Rhythm (NSR), i.e., a normal rhythm initiated by the SA node. The re-polarisation of the atria occurs concurrently with ventricular de-polarisation and is therefore not visible on the ECG.

To measure the electrical activity of the heart, it is standard practice to use the 12 Lead ECG and place the electrodes in the positions depicted in Figure 2.4. Nine electrodes are used to generate the 12 Lead ECG. Three electrodes, Right Arm (RA), Left Arm (LA) and Left Leg (LL) are used to generate Leads I, II, III and augmented Leads aVL, aVF, aVR using the following formulae.

$$V_I = V_{LA} - V_{RA}$$

$$V_{II} = V_{LL} - V_{RA}$$

$$V_{III} = V_{LL} - V_{LA}$$

$$V_{aVR} = V_{RA} - V_{\frac{1}{2}(LA+LL)}$$

$$V_{aVL} = V_{LA} - V_{\frac{1}{2}(RA+LL)}$$

$$V_{aVF} = V_{LL} - V_{\frac{1}{2}(RA+LA)}$$

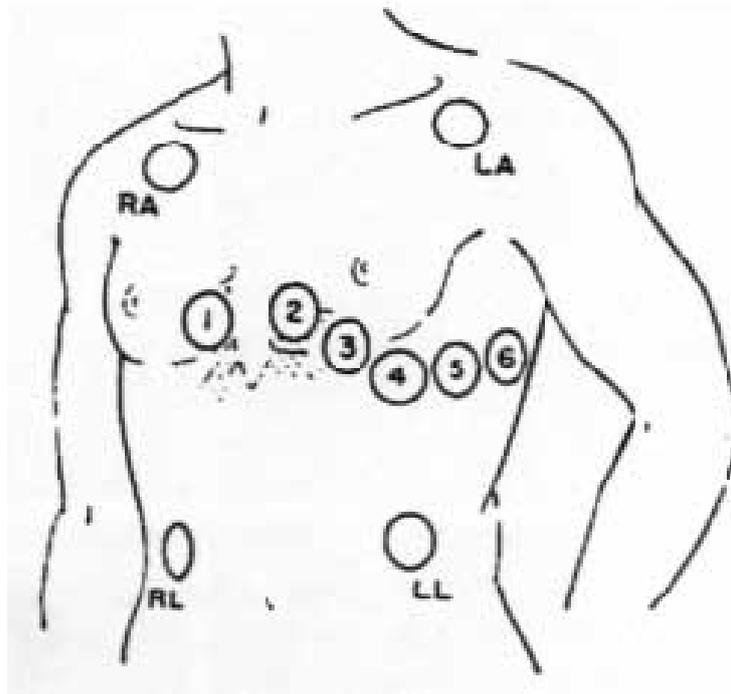


Figure 2.4: Placement of electrodes for 12 Lead ECG.

(Courtesy <http://nuclearcardiologyseminars.com>)

The electrical centre of the heart is generated by averaging the signal of all three electrodes and is called the Wilson Central Terminal (WCT).

$$V_{WCT} = \frac{1}{3}(V_{RA} + V_{LA} + V_{LL})$$

Effectively, the three electrodes placed on the skin give a total of six views of the hearts electrical activity called Leads. Because of the placement of the three electrodes, a two dimensional view of the electrical activity of the frontal or anterior plane is created [2]. A highly skilled cardiac specialist can interpret the angles of propagation of the various ECG morphologies and determine normal or abnormal electrical activity. The same is true for the six Pre-cordial Leads, $V_1 \dots V_6$ in the horizontal, or transverse plane measured with respect to the WCT.

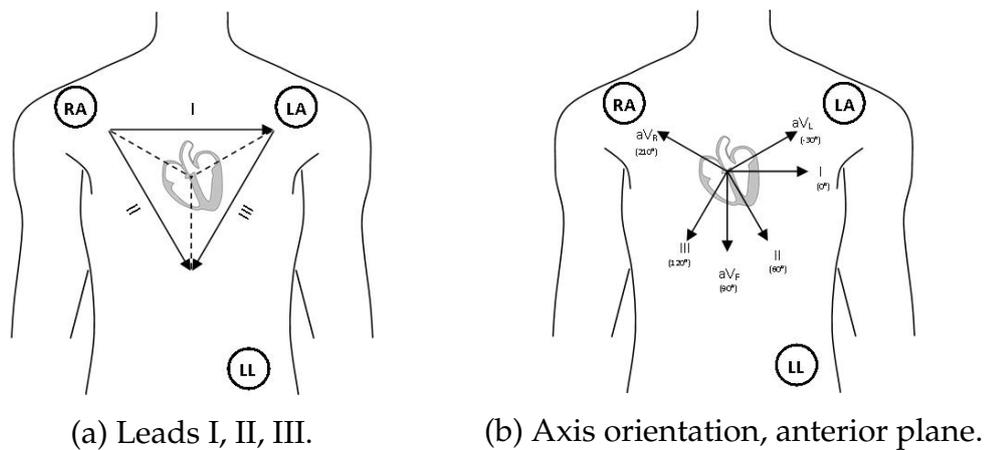


Figure 2.5: Six Lead ECG showing electrode connection and resultant six Lead axis orientation.

2.3 Changes to ECG due to arrhythmia

Irregular heart beat, known as arrhythmia, can be caused by a number of conditions. Following is a brief description of various arrhythmias grouped by their origin.

- Sinus: The sino-atrial (SA) node, is a group of specialised cells that reside in the right atrium. Their biological make up causes them to automatically depolarise and re-polarise at about 80 times per minute and it is this action that starts the electrical sequence of events producing the pumping action of the heart [18]. The timing of the electrical pulses created by the SA node can be altered by the sympathetic and parasympathetic nervous systems allowing external control of the heart rate. When functioning correctly, the heart rate caused by the SA node is called 'normal sinus rhythm' [2].

Sick sinus syndrome occurs when the SA node fires at a much slower rate or fails to fire at all [15]. This results in bradycardia and sinus pauses.

- **Atrial:** Imperfections in the atria can cause a number of arrhythmia problems such as atrial fibrillation, atrial flutter and Premature Atrial Contractions (PAC) [2]. In AF, the SA node loses focus and the regular contractions of the atria are reduced to random activity causing irregular heart beat. With atrial flutter, the SA node again loses focus, but the atria beat at a regular, much faster rate than SA node generated impulses. Premature atrial contractions occur when the atria contract prematurely, unsynchronised with the SA node impulses. These may be lone occurrences or if more than two consecutive PACs occur then this is called atrial tachycardia.
- **Junctional:** The combination of the AV node and the Bundle of His form what is referred to as the junction between atria and ventricles and any heartbeat initiated by electrical activity in this region is called Junctional arrhythmia. The QRS complex looks normal, but the P-wave is abnormal and is inverted due to depolarisation of the atria being propagated in the opposite direction [18].
- **Ventricular:** Ventricular arrhythmia is caused by impulses created from within the ventricles and is characterised by a wide, bizarre shaped QRS complex due to the irregular speed and direction of propagation of depolarisation [15].
- **Atrio-ventricular blocks:** AV blocks are characterised by either prolonged conduction through the AV node or partial or full conduction block [18]. First degree AV block occurs when there is a prolonged PR interval but otherwise normal beat sequence. This arrhythmia is often asymptomatic. Second degree AV block is an intermittent block of conduction via the AV node and third degree AV block is full, permanent block of conduction with no electrical impulses from the atria being passed to the ventricles. In this case the atria and ventricles beat independently and at their own respective rates.

- Bundle Branch blocks: Bundle Branch block is described as a conduction blockage of the Bundle of His either just below the AV node or within either the left or right bundle [2]. Because of the block the depolarisation of the ventricles takes longer and so the QRS complex becomes wider and abnormal in morphology.

2.4 Changes to ECG due to other causes

Other changes to the ECG morphology can be caused by various conditions of the heart. A description of some of these follows, but it is outside the scope of this research to study them in depth.

- Pre-excitation: On rare occasions secondary conduction pathways exist between the atria and ventricles, short circuiting the AV node/ bundle conduction pathway [2]. This shows up on an ECG as a widened QRS complex having a distinctive 'delta' wave on the leading edge of the complex.

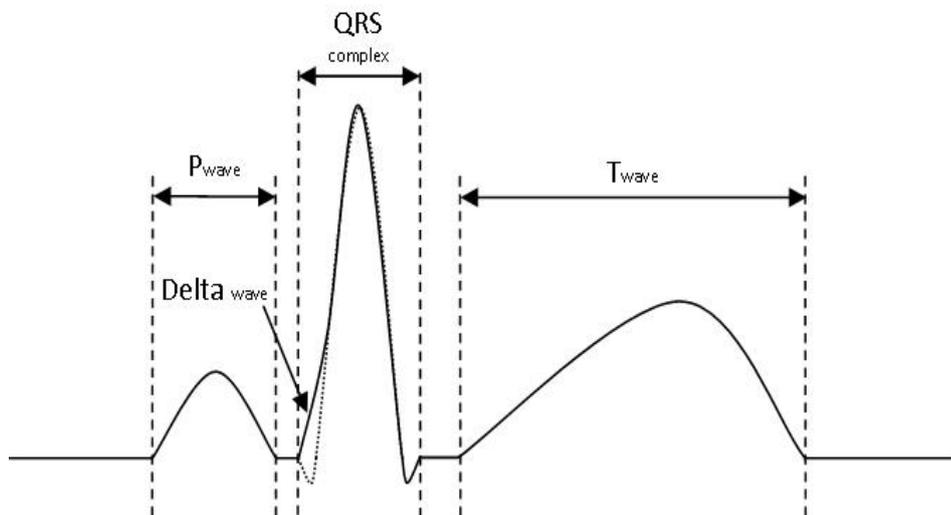


Figure 2.6: An example of the Delta wave caused by pre-excitation of the ventricles.

- **Chemical imbalances:** As previously discussed, the heart muscle undergoes contraction and relaxation due to the flow of chemical ions into and out of the heart muscle cells. These are ions of calcium, sodium and potassium, and any deficiency or excess of these electrolytes can cause changes to the cardiac action and thus the ECG morphology [18].
- **Pulmonary disease:** Disease of the lungs causes strain on the heart that in time effects the ability of the heart to function correctly [15].
- **Myocardial disease:** Often with heart diseases, blood pressure increases causing the enlargement (increased volume) and hypertrophy (increased wall thickness) of the atria and ventricles. This results in an associated change to the ECG morphology [18].
- **Myocardial damage:** Lack of blood supply to the heart muscle (the myocardium) causes damage that results in changes to the ECG morphology [18]. An experienced cardiac specialist is able to determine how extensive the damage is and sometimes the location of the damage. The damage to the heart muscle caused by low blood perfusion can be temporary (myocardial ischemia), or permanent (myocardial infarction) [2].

2.5 Atrial Fibrillation

The focus of this research is Atrial Fibrillation (AF). AF is seen in 1.0% of the total population and accounts for 11% of all arrhythmias seen in patients presented to intensive care units [23]. Its prevalence increases with age and is higher in the western world compared with that in Asia. AF is a condition of the heart whereby the two top chambers of the heart, the atria, are not beating in unison with the two bottom chambers of the heart, the ventricles. More specifically, the Sino-atrial node (the hearts

natural pacemaker) loses focus and the regular contractions of the atria are reduced to random activity causing irregular heart beat. The resulting changes in fluid dynamics can create blood clots and studies show that patients with AF are five times more likely to suffer from strokes [46] than have without. In a study by Katz and Pick referred to by [23], 50,000 consecutive patients presented to hospital over a 25 year period showed that the prevalent cardiac arrhythmias were Sinus tachycardia (18.7%), Sinus arrhythmia (16.6%), Premature systoles (14.3%) and Atrial fibrillation (11.7%). In another study [33], 133 consecutive patients with arrhythmias were monitored. Of those the most frequent arrhythmias found were Ventricular tachycardia (48.6%) and Atrial fibrillation (29.8%)

The detection of AF is no trivial task. Electrocardiograph devices print out a real time waveform for diagnostic purposes and these are interpreted by highly trained cardiac specialists. However AF can often be paroxysmal and not always picked up when an ECG is recorded [25]. Devices are available, such as the Holter monitor, that record the ECG over long periods, but the data collected still has to be analysed by a cardiac expert. This is a lengthy process and one prone to errors. Systems are now available that automatically interpret ECG tracings saving time, but these are expensive desktop units or computer based software.

Diagnosis of Atrial Fibrillation from ECG relies on detection of the following three indicators [8]

1. Missing P-wave.
2. Presence of abnormal atrial activity.
3. Highly variable RR interval.

Contractile cells (those that contract on de-polarisation) can change and become inappropriate pacemaker cells, causing erratic contractions of the atria. This causes the P-wave to disappear and abnormal atrial activity appears on the ECG. Instead of one pacemaker, (i.e., the SA node generating regular impulses) a number of sites cause multiple wave fronts that travel throughout the atrial mass. These haphazard waves occur too

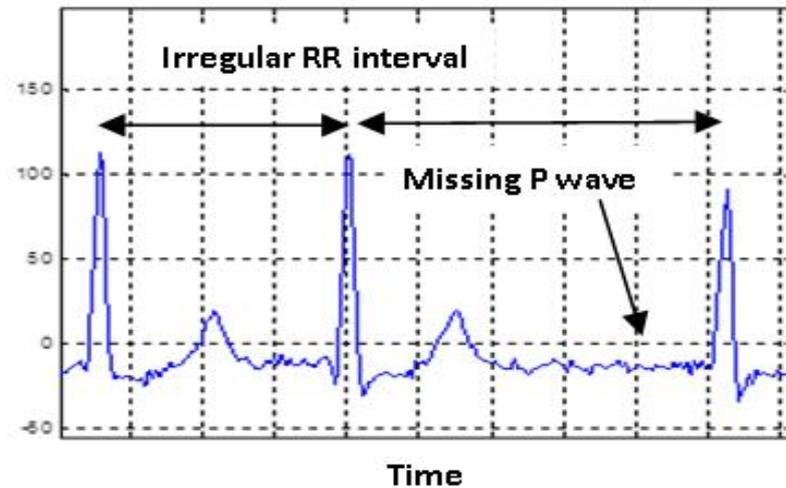


Figure 2.7: ECG morphology: Atrial fibrillation showing irregular RR interval and P-wave replaced by random atrial activity.

quickly for synchronised atrial contractions and thus produce the highly variable RR intervals observed during AF (see Figure 2.7).

In this chapter the physiology of the heart has been examined and AF described. In the following chapter a number of mathematical processes that are useful in the analysis the of ECG signal are presented.

Chapter 3

Mathematical Background

In the previous chapter on cardiology, the ECG waveform was discussed and arrhythmia, in particular atrial fibrillation, was described. Characteristics of Normal Sinus Rhythm and Atrial Fibrillation were also summarised. What was not mentioned was that various ECG waveforms have certain characteristics both in the time and frequency domains. It is these characteristics that are to be extracted and analysed using digital signal processing techniques, allowing various arrhythmias to be detected and distinguished. This is described in more detail in Chapter 4. In this chapter, various mathematical processes are introduced that are suitable for the construction of reliable, efficient algorithms for AF detection.

3.1 Mathematical Morphology

Mathematical Morphology (MM) is the discipline of analysing geometrical shapes using techniques from a number of mathematical fields including Set theory, Lattice theory, Topology and Discrete Mathematics. It originated from work done at the École des Mines de Paris by Jean Serra and Georges Matheron [13] during the mid 1960's. More often associated with image processing, MM is becoming more common in one dimensional signal processing especially where the morphology of the presented signal is

being investigated. In this section we limit our discussion to one dimensional signals only.

Mathematical morphology has been used to remove baseline wander in the ECG signal [34], and as a QRS complex detector [47]. Also, ECG feature extraction and heartbeat classification has been implemented using mathematical morphology [39]. Two fundamental MM operators are:

Erosion

$$(f \ominus B)_{(n)} = \min_{m=0, \dots, M-1} \left\{ f\left(n - \frac{M-1}{2} + m\right) - B_{(m)} \right\}$$

Dilation

$$(f \oplus B)_{(n)} = \max_{m=0, \dots, M-1} \left\{ f\left(n - \frac{M-1}{2} + m\right) + B_{(m)} \right\}$$

where

$$n = \left\{ \frac{M-1}{2}, \dots, N - \frac{M+1}{2} \right\}$$

and B is a symmetric structuring element of M points having value zero.

More complex operators can be constructed from erosion and dilation. Following are two very common operators:

Opening

$$f \circ B = f \oplus B \ominus B$$

Closing

$$f \bullet B = f \ominus B \oplus B$$

In simple terms, opening can be understood as sliding the structuring element beneath the signal producing the highest points reached by any part of the structuring element, effectively suppressing peaks. With a longer structuring element, more of the peak is removed; with a shorter structuring element, less of the peak is removed. Closing can be understood as sliding the structuring element above the signal producing the lowest points reached by any part of the structuring element, effectively suppressing valleys. In the same way as described above for the opening operator, the closing operator removes valleys to a greater or lesser extent, see Figure 3.1

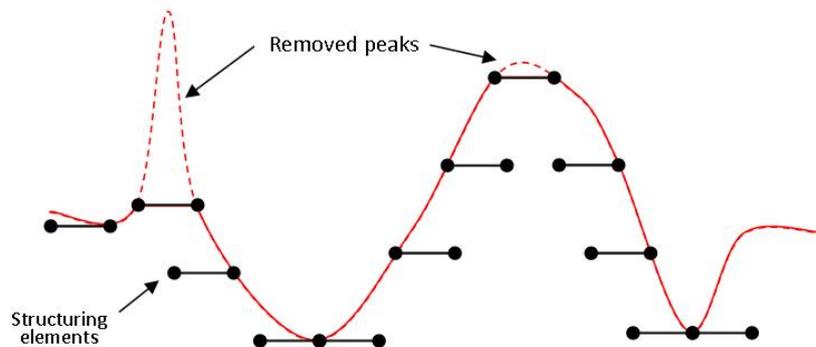


Figure 3.1: Opening operation showing elements and removed peaks.

By combining operators, more sophisticated signal manipulation can be accomplished. The use of Mathematical Morphology for baseline removal is explained in more detail in section 4.2.

3.2 Wavelet Transform

The Fourier Transform is ubiquitous in modern signal processing applications. Originally devised by Joseph Fourier in 1822 for the analysis of heat flow, it is now used extensively in signal processing to resolve the frequency composition of electronic signals. Its ability to transform a sig-

nal represented in the time domain to the same signal represented in the frequency domain has given scientists and engineers a valuable tool for signal and systems analyses. However, its main drawback has been its inability to separate signals in the frequency domain, with respect to time. This is due to the infinite support of the basis functions, sine and cosine. By using the Short Time Fourier Transform this limitation has been partially overcome but unfortunately it has fixed frequency and time resolution [6]. A more natural process would provide good time resolution for high frequencies and good frequency resolution for low frequencies. The Wavelet Transform overcomes the deficiencies of the Fourier Transform in this respect by separating out the frequency spectrum of a signal with respect to time, and providing appropriate scaling in time and frequency domains.

3.2.1 Brief history

Work done by Alfred Haar (orthonormal basis functions, 1909) and Alberto Caldéron (harmonic analysis, 1964) paved the way for further research into wavelets. In the late 1970's Jean Morlet, a French geophysicist, proposed a method of analysing signals that had short duration high frequency components, and long duration of low frequencies components, by using dilation and compression of a prototype function to create a set of windowing functions, in this case a Gaussian window [32]. Working with Alex Grossmann, a theoretical physicist, they formalised the transform and inverse transform in 1980. It was Yves Meyer that noticed the similarity of work done by Morlet and Grossmann to that of Caldéron and consequently constructed new wavelets with better localisation properties. More research by Ingrid Daubechies and Stéphane Mallat lead to the development of the discretised version of the continuous wavelet transform. It was latter works of Daubechies that laid the foundations of modern wavelet theory and created a set of useful signal processing tools to be used by scientists, engineers and economists. [10, 32]

3.2.2 Continuous Wavelet Transform

Essentially the prototype wavelet function is a filter like any other, however it is used as a template to create a family of filters for resolving different frequency components within a signal.

The Continuous Wavelet Transform (CWT) takes the form

$$\Upsilon(s, t) = \int f(t) \Psi_{s,t}^*(t) dt,$$

when the wavelet function Ψ is

$$\Psi_{s,\tau}^*(t) = \frac{1}{\sqrt{s}} \psi\left(\frac{t - \tau}{s}\right)$$

where signal $f(t)$ is transformed into a two dimensional function Υ , with τ and s being the translation and scaling factors respectively.

Using dilation and compression of the prototype (the mother wavelet), a series of wavelets is created, each inheriting the fundamental properties of the mother wavelet, but having differing frequency and time domain characteristics. For these relationships to exist the mother wavelet must have the following properties.

Admissibility:

A function that is square integrable and satisfies the following admissibility condition can be used to decompose and reconstruct a signal without losing any information

$$\int_{-\infty}^{\infty} \frac{|\Phi(\omega)|^2}{|\omega|} d\omega < +\infty$$

where $\Phi(\omega)$ is the Fourier Transform of $\psi(t)$.

This also means that the Fourier Transform of $\psi(t)$ is zero, i.e. vanishes, at zero frequency.

$$|\Phi(\omega)|^2 |_{\omega=0} = 0$$

Wavelets have a band-pass like spectrum. A zero at zero frequency, means that the average value of the wavelet is zero, i.e.,

$$\int_{-\infty}^{\infty} \Psi(t) dt = 0$$

and shows that the wavelet is oscillatory in nature.

Regularity:

It is an advantage for the wavelet to be local in the time and frequency domains. Although this is not necessary, it is an important property. The time-bandwidth product of the wavelet transform is the square of the signal (viz. a one-dimensional signal transforms into two dimensions and a two dimensional signal transforms into four dimensions), which is undesirable. To overcome this adverse effect, other conditions are imposed on the wavelet functions so they decrease rapidly with decreasing scaling factor s . Vanishing moments are a measure of quality of the wavelet function and provides reduced computation time and allows detection of higher order derivatives. The moments of a wavelet M_p are defined as

$$\int_{-\infty}^{\infty} t^k \psi(t) dt = 0, \quad 0 \leq k < P$$

$$M_p = \int_{-\infty}^{\infty} t^p \psi(t) dt$$

For a wavelet having P vanishing moments, the wavelet function Ψ , is orthogonal to polynomials up to degree $P - 1$, and the detail signal produces no energy for polynomials up to degree $P - 1$.

In summary, admissibility gives a wave like function and regularity gives a smooth function with fast decay, i.e., local in both the time and frequency domains. However, the Continuous Wavelet Transform has three

major disadvantages for use in real world applications. Firstly, having the wavelet function continuously shifted over the whole signal results in much redundant information. Secondly, by having a continuously variable wavelet function, an infinite number of wavelets is created. Thirdly, most of the functions of the wavelet transform can only be calculated numerically and have no analytical solution. To overcome these issues the Discrete Wavelet Transform was devised. [9, 31, 37, 38, 41, 42, 43]

3.2.3 Discrete Wavelet Transform

The Discrete Wavelet Transform has wavelets that are scaled and translated in discrete steps, effectively making it a piecewise continuous function. The discrete wavelet function Ψ takes the form

$$\Psi_{s,\tau}^*(t) = \frac{1}{\sqrt{s^j}} \psi\left(\frac{t - k\tau s^j}{s^j}\right)$$

where $\tau = 1$, $j = 0..n$ and $s = 2$, giving the *dyadic* implementation of the transform, i.e. dyadic sampling of the frequency and time axis (see Figure 3.2).

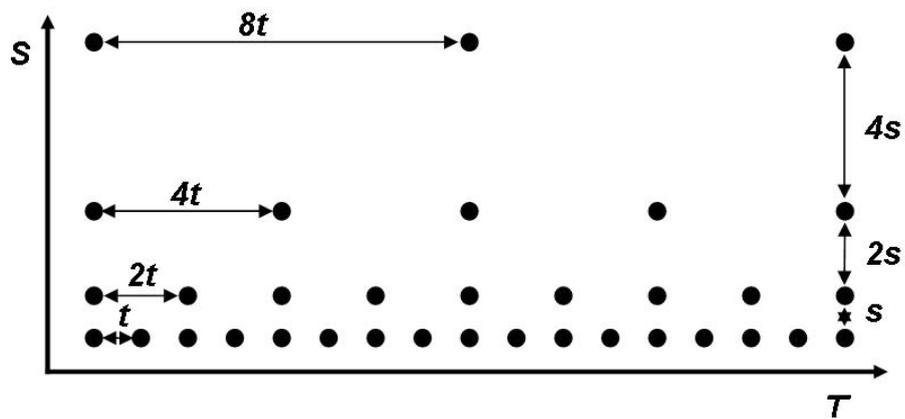


Figure 3.2: The dyadic grid, time-scale representation.

For perfect reconstruction, Ingrid Daubechies [9] showed that it is a

necessary and sufficient condition that the energy of the wavelet lie between two positive bounds

$$A\|f\|^2 \leq \sum_{j,k} |\langle f, \Psi_{j,k} \rangle|^2 \leq B\|f\|^2$$

where $\|f\|^2$ is the energy of $f(t)$, $0 < A < B < \infty$ and A, B are independent of $f(t)$. The set $\Psi_{j,k}$ is then said to be a frame. If $A = B$, then the frame is said to be tight and the discrete wavelets have similar properties to an orthonormal basis. When $A \neq B$, perfect reconstruction can still be achieved, but the decomposition and reconstruction wavelets are dissimilar.

The wavelet effectively acts like a band-pass filter and subsequent dilations produces a family of band-pass filters, each filter having half the bandwidth of the previous. Because the ratio of bandwidth to centre frequency remains constant, the resulting family of filters is called a constant Q filter bank.

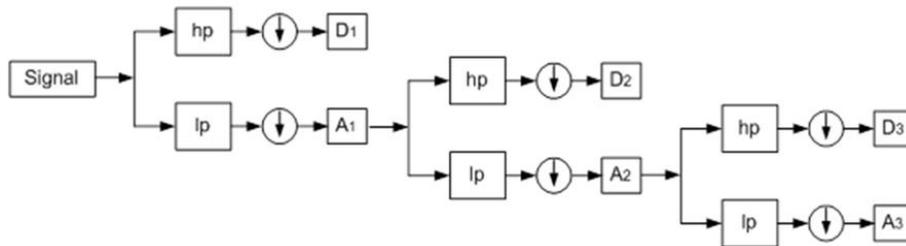


Figure 3.3: The dyadic implementation of the DWT filter bank showing high pass (wavelet) and low pass (scaling) filters with following decimators. $D_1..D_n$ are the detail outputs which are effectively the outputs of a series of band-pass filters, each having half the bandwidth and half the number of coefficients as the previous.

3.2.4 Filter design

Once the scaling function (ϕ , the decomposition low pass filter) has been designed, it is relatively easy to create the wavelet function (ψ , the decomposition high pass filter) by creating a filter that is the Quadrature Mirror Filter (QMF) of the scaling function. Figures 3.4–3.5 show the resulting filters for decomposition and reconstruction of a signal in the time domain and frequency domain, respectively. The wavelet function (decomposition high pass filter) is defined as

$$g_1[n] = (-1)^n g_0[N - 1 - n]$$

being the QMF of g_0 (the decomposition low pass filter). The reconstruction low pass filter is defined as

$$h_0[n] = g_0[N - 1 - n]$$

being the mirror of g_0 , and the reconstruction high pass filter

$$g_1[n] = (-1)^n g_0$$

is the quadrature of g_0 .

In the Z domain the respective filters become

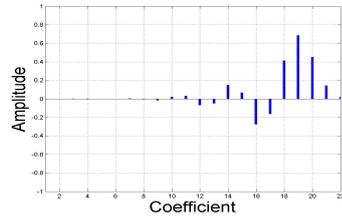
$$G_1[n] = Z^{-N} G_0(-Z^{-1})$$

$$H_0[n] = Z^{-N} G_0(Z^{-1})$$

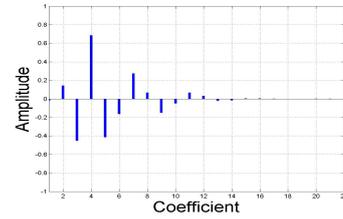
$$H_1[n] = G_0(-Z^1)$$

see Figure 3.5.

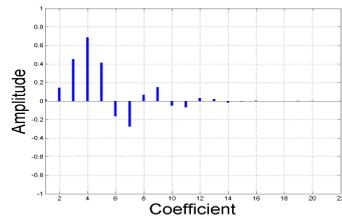
The Discrete Wavelet Transform is used in this research for filtering the ECG signal into a series of frequency sub-bands for further analysis. See sections 4.4 and 4.5



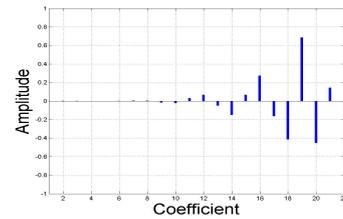
(a) Decomposition low pass filter



(b) Decomposition high pass filter

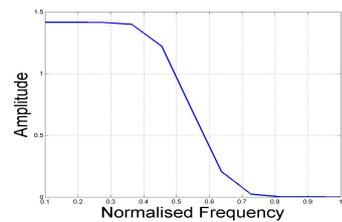


(c) Reconstruction low pass filter

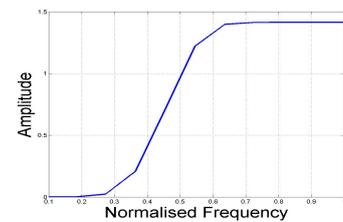


(d) Reconstruction high pass filter

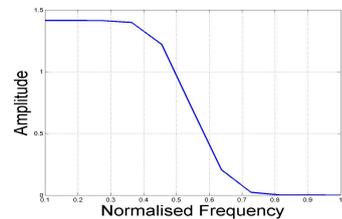
Figure 3.4: Daubechies db11 wavelet filter construction (Time domain). Quadrature mirror left to right, mirror (time reversal) top to bottom.



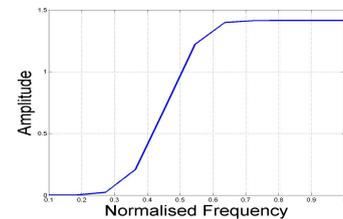
(a) Decomposition low pass filter



(b) Decomposition high pass filter



(c) Reconstruction low pass filter



(d) Reconstruction high pass filter

Figure 3.5: Daubechies db11 wavelet filter construction (Frequency domain). Quadrature mirror left to right, mirror (time reversal) top to bottom.

3.3 Hilbert Transform

The Hilbert Transform is used in many signal processing applications including envelope detection, magnitude estimation, decay time estimation and propagation delay measurements, and also finance and economic analysis. It has also been used as an edge detector in 2 dimensional images [26]. The transform defines $h(t)$, an analytic signal, i.e., a signal having no "negative" frequencies [20], from a real valued time signal deriving $g(t)$ so that $h(t) = f(t) + jg(t)$. The transform takes the form

$$g(t) = \frac{1}{\pi} \int_{-\infty}^{\infty} \frac{f(\tau)}{t - \tau} d\tau$$

However, as can be seen the integral is improper with a singularity at $t = \tau$. It is more properly defined as

$$g(t) = \frac{1}{\pi} \text{P.V.} \int_{-\infty}^{\infty} \frac{f(\tau)}{t - \tau} d\tau$$

where P.V. is the Cauchy's Principle Value i.e., the transform can be written as

$$g(t) = \lim_{\epsilon \rightarrow 0} \left[\int_{-\infty}^{t-\epsilon} \frac{f(\tau)}{t - \tau} d\tau + \int_{t+\epsilon}^{\infty} \frac{f(\tau)}{t - \tau} d\tau \right]$$

The Hilbert Transform can be considered as the convolution of the real valued signal with $\frac{1}{\pi t}$, namely $g(t) = f(t) * \frac{1}{\pi t}$ which has a singularity at $t = 0$.

By obtaining the Hilbert Transform of an input signal, the amplitude of a signal $f(t)$ can be found using

$$A(t) = \sqrt{f(t)^2 + g(t)^2}$$

where $g(t)$ is the Hilbert Transform of $f(t)$, and the corresponding phase can be found using

$$\theta = \tan^{-1} \left(\frac{g(t)}{f(t)} \right)$$

A number of methods can be used to obtain the Hilbert Transform of a signal. One method is taking the Fourier Transform of signal $f(t)$, setting all negative frequencies in the frequency domain to zero, then taking the inverse Fourier Transform. However a simpler approximation can be achieved by creating an All-pass filter with a -90° phase shift. The Hilbert Transform must exhibit three important characteristics [24], namely

- Continuous and Differentiable
- Linear, i.e., multiplying by a constant retains the phase and frequency.
- Harmonic correspondence. For a simple sinusoid, the amplitude and frequency should be retained, i.e., the Hilbert Transform of $a \cos(\omega t + \theta)$ is $a \sin(\omega t + \theta)$ for any a, ω, θ .

The Hilbert transform is used in this research as a signal energy detector and its use is described in more detail in section 4.3.

3.4 Teager-Kaiser Energy Operator

The energy content of a signal is an important measure for signal processing purposes. Signal energy is the measure of energy of a signal that does not extend to $\pm\infty$, whereas the power of a signal is the energy per unit time of a signal that does extend to $\pm\infty$. Signal energy is defined as

$$E_s = \int_{-\infty}^{\infty} |x(t)|^2 dt$$

whereas signal power is

$$P_s = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{-\frac{T}{2}}^{+\frac{T}{2}} |x(t)|^2 dt$$

Signal energy described above is not strictly the same as the measurement of energy in the physics sense. For simple harmonic motion the kinetic energy for a signal $A \cos(\omega n + \theta)$ is proportional to the velocity squared of the object i.e.,

$$K(t) = \frac{1}{2}mA^2\omega^2 \sin^2(\omega t - \phi)$$

and hence the average kinetic energy is proportional to the square of both the frequency and the amplitude. i.e.,

$$K(t) \propto A^2\omega^2 \quad (3.1)$$

It is this 'physical' paradigm that has given rise to the Teager-Kaiser Energy Operator (TKEO) for signal energy detection.

3.4.1 The continuous operator

The Teager-Kaiser energy operator [12, 24, 40] was derived from work by Teager and later Kaiser [21]. In the continuous domain the operator is defined as

$$\Psi\{x(t)\} = \dot{x}(t)^2 - x(t)\ddot{x}(t)$$

The effectiveness of this operation can be demonstrated on an example signal $x(t) = A \cos(\omega t)$

The TKEO is then

$$\begin{aligned} \Psi\{x(t)\} &= (-A\omega \sin(\omega t))^2 - (A \cos(\omega t))(-A\omega^2 \cos(\omega t)) \\ &= A^2\omega^2 \sin^2(\omega t) + (A^2\omega^2 \cos^2(\omega t)) \\ &= A^2\omega^2(\sin^2(\omega t) + \cos^2(\omega t)) \\ &= A^2\omega^2 \end{aligned}$$

3.4.2 The discrete operator

The discrete operator takes the form of

$$\Psi\{x[n]\} = x[n]^2 - x[n-1]x[n+1]$$

For example, if

$$x[n] = A \cos(\Omega n + \theta), \text{ and } \Omega = \frac{2\pi f}{F_s}$$

and f is the analogue frequency, F_s the sampling frequency, and θ is the phase, then

$$\begin{aligned} x[n] &= A \cos(\Omega(n) + \theta) \\ x[n-1] &= A \cos(\Omega(n-1) + \theta) \\ x[n+1] &= A \cos(\Omega(n+1) + \theta) \end{aligned}$$

Using identities

$$\begin{aligned} \cos(\alpha + \beta) \cos(\alpha - \beta) &= \frac{1}{2} [\cos(2\alpha) + \cos(2\beta)] \\ \cos(2\alpha) &= 2 \cos^2(\alpha) - 1 = 1 - 2 \sin^2(\alpha) \end{aligned}$$

then it follows that

$$\begin{aligned} x[n-1]x[n+1] &= (A \cos(\Omega(n-1) + \theta))(A \cos(\Omega(n+1) + \theta)) \\ &= A^2 \cos^2(\Omega n + \theta) - A^2 \sin^2(\Omega) \end{aligned}$$

It can be seen that the first term is in fact $x[n]^2$, and so

$$\begin{aligned} x[n]^2 - x[n-1]x[n+1] &= A^2 \cos^2(\Omega n + \theta) - (A^2 \cos^2(\Omega n + \theta) - A^2 \sin^2(\Omega)) \\ &= A^2 \sin^2 \Omega \end{aligned}$$

For small, $\Omega \sin \Omega \approx \Omega$ and hence $A^2 \sin^2 \Omega = A^2 \Omega^2$

So

$$\Psi\{x[n]\} = x[n]^2 - x[n-1]x[n+1]$$

can be used as the discrete TKEO

3.4.3 Cross Teager-Kaiser Energy Operator

The Cross Teager-Kaiser Energy Operator (CTKEO) is a non-linear measure of the interaction between two (real) signals and can be used as a measure of how much of one signal is present in the other. [4, 5]. The CTKEO takes the form

$$\Psi\{x(t), y(t)\} = \dot{x}(t)\dot{y}(t) - x(t)\ddot{y}(t)$$

or

$$\Psi\{y(t), x(t)\} = \dot{y}(t)\dot{x}(t) - y(t)\ddot{x}(t)$$

Note that

$$\dot{x}(t)\dot{y}(t) - x(t)\ddot{y}(t) \neq \dot{y}(t)\dot{x}(t) - y(t)\ddot{x}(t)$$

and hence the operator is not commutative. The operator may be viewed as a non-linear version of cross correlation and used in the same context.

The Teager Energy Operator is used in this research as a simple signal energy detector and its use is described in more detail in section 4.3.

3.5 Support Vector Machine

The Support Vector Machine is a learning machine that can be used as a binary classifier to determine whether a set of input data belongs to known class or not. Originally conceived by Vapnik and Cortes [36], the SVM is a classifier that possesses a high level of generalisation, accuracy comparable to other classifiers, and ease of implementation as an online classifier in the target device. As a linear classifier, the SVM calculates the weighted sum of a function relating the input data to support vectors to determine whether the input data is of one class or another. The SVM takes the form

$$\hat{y} = \text{sign}\left(\sum_{i \in s} \alpha_i y_i \text{K}(s_i, x) + b\right), \quad (3.2)$$

where \hat{y} is the output class, y are the input classes, α the weight of the support vector, and b the bias. The kernel function K is a measure of similarity between the support vectors, s_i and the input data, x . For a linear classifier this is the inner product of the support vector, s_i and the input data, x .

$$\text{K}(s_i, x) = \langle s_i, x \rangle$$

Figure 3.6 shows the linear separation of two data sets, class 1 and class 2. The samples of data lying on the dashed lines are the 'support vectors' and the central dividing line (a hyperplane in higher dimensional examples), partitions the two classes. The distance between the two dashed lines (the margin) is the maximum distance between the two classes. During training, the classifier maximises this distance by minimising the following

$$\tau_{(w)} = \frac{1}{2} \|w\|^2$$

This creates a 'hard margin' whereby no points are allowed to occur inside, or beyond, the margin. w is the distance between the hyperplane and the sample point. However, classifying with such tight constraints can be difficult. For classifiers that are not so rigid, a 'soft margin' can be instigated

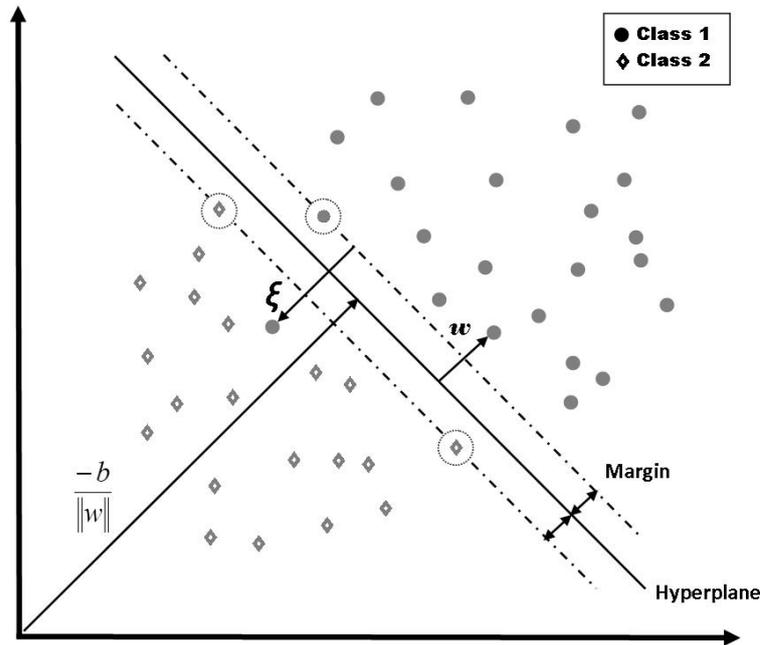


Figure 3.6: SVM showing linear separation of two classes with points on dashed lines being support vectors.

whereby points are allowed inside, or beyond, the margin and therefore open to misclassification. During training, the classifier minimises the following

$$\tau_{(w,\epsilon)} = \frac{1}{2} \|w\|^2 + \frac{C}{M} \sum_{i=1}^M \xi_i,$$

subject to

$$y_i = (w^T x_i + b) \geq 1 - \xi_i, \quad \xi_i \geq 0.$$

where, w is the distance between the hyperplane and the sample point, M is the number of data points, ξ is the loss function (see Figure 3.6) and C regulates the "capacity" of the learning i.e., the trade off between the margin and the loss. By altering C the loss function will have a greater or lesser effect on overall training. If C is large, the system over-trains and will not generalise well to unseen data.

The inner product is a suitable kernel for linearly separable data, however for high dimensional data, other kernels can be utilised that take advantage of what is called the ‘kernel trick’. This allows the use of a kernel function that carries out the inner product of the two data sets in a higher dimension, without having to map the input data to that higher dimension. An example of this trick is the Radial Basis Function

$$\mathbf{K}(x_i, x) = e^{-\left(\frac{\|s_i - x\|^2}{2\sigma_0^2}\right)}.$$

To explain this concept, consider the Taylor series expansion of the exponential function. This expansion has an infinite number of coefficients, and the series is effectively the inner product of the two data sets. Thus, the result is the equivalent of calculating the inner product with infinite dimensions weighted in a particular way. This kernel maps the Euclidean distance of zero (when perfectly matched) to one, and a non-zero Euclidean distance (when badly matched) to a number less than 1. $[0 \dots \infty] \mapsto [1 \dots 0]$.

Other suitable kernels include the Polynomial kernel:

$$\mathbf{K}(s_i, x) = (s_i^T x + b)^d$$

the Laplacian kernel

$$\mathbf{K}(s_i, x) = e^{-\left(\frac{\|s_i - x\|}{2\sigma_0^2}\right)}.$$

and the Cosine kernel

$$\mathbf{K}(s_i, x) = \frac{a_i}{(\|s_i - x\| + a_i^2)^{\frac{1}{2}}}$$

For this research, the Radial Basis Function has been used for classification. The chapter on implementation describes how the kernel function is implemented in the target device. Training and testing of the classifiers has been carried out using this kernel.

In this chapter, a number of mathematical processes have been discussed that may be of use fulfilling the goal of developing algorithms to detect atrial fibrillation. The next chapter discusses methods of carrying out the detection process by using the mathematics discussed above.

Chapter 4

Research

Introduction

This chapter describes how, with knowledge gained from chapters 2 and 3, the ECG signal can be analysed to automatically detect AF. The premise is that AF can be accurately and reliably detected by extracting feature sets using digital signal processing techniques, and classifying using a Support Vector Machine, all on a low power, miniature ECG monitoring device.

A number of avenues of research were followed, some that showed promise while others were discarded as being too difficult to implement on the target device. During this phase of research, it was always remembered that any methods of feature extraction or classification chosen, needed to be implemented on a resource constrained microcontroller. By applying the previously discussed mathematical processes, it was hoped that new methods of signal feature identification could be found. Firstly, timing of the RR interval and, secondly, the frequency spectrum of the ECG signal, were analysed by using appropriate mathematical methods previously described. The waveform and rhythm feature sets created were then brought together in the last stage of research, classification, to see if detection of AF was realisable. Finally, implementing the desired detection scheme in the target device's native assembly language was carried

out.

4.1 Arrhythmia

To effectively distinguish between arrhythmia types an understanding of how each type of arrhythmia originates is necessary. AF is a chaotic firing of a number of abnormal pacemakers in the atrial mass. This causes random firing of the AV node and consequently random activation of the ventricles. Therefore, the RR interval is random. Contrast this with other types of arrhythmias such as bigeminy or trigeminy where irregular RR intervals are generated by Premature Atrial Contractions (PAC) or Premature Ventricular Contractions (PVC) that have more predictable timing.

Bigeminy is defined as a PAC or PVC every second beat while trigeminy is a PAC or PVC every third beat. See Figure 4.1 for a visual aid to the three types of arrhythmias, along with normal sinus rhythm. As can be seen, the RR intervals of AF are random, but the waveform type is considered normal, i.e., the QRS complex and T wave are of NSR morphology, even though the P wave is missing. Compare this with ventricular bigeminy or trigeminy. In these instances only the 'N' type beats have normal NSR morphology, but the premature beats have QRS complexes and T waves that differ considerably from NSR morphology. Also, the timing of ventricular bigeminy or trigeminy is far more consistent. The n-n, n-v and v-n intervals are consistent and curiously the time period n-v-n is approximately equal to $2n$. This is due to the ventricles needing time to re-polarise between beats and their inability to be reactivated by the first normal beat immediately following the premature beat. In the case of atrial bigeminy or trigeminy, the morphology of the normal and premature beat are the same, i.e., normal. This creates a problem in determining the exact type of rhythm.

As discussed in Section 2.5 on cardiology, key indicators for AF are:

1. Missing P-wave.

2. Presence of abnormal atrial activity.
3. Highly variable RR interval.

Detection of the presence, or absence, of the P wave is problematic. There can be large morphological differences between subjects ECGs, and the amplitude and shape can vary widely depending on lead placement and electrode impedance. The same applies to the detection of abnormal atrial activity. The key indicator to AF then, is the highly variable RR interval. As discussed above, there are a number of other arrhythmia that can cause variable RR intervals so methods to distinguish between them need to be devised. This thesis looks at using only the RR interval variations for detecting AF. Future work will include P wave and atrial activity detectors to determine the effect they have on overall accuracy and reliability of AF detection.

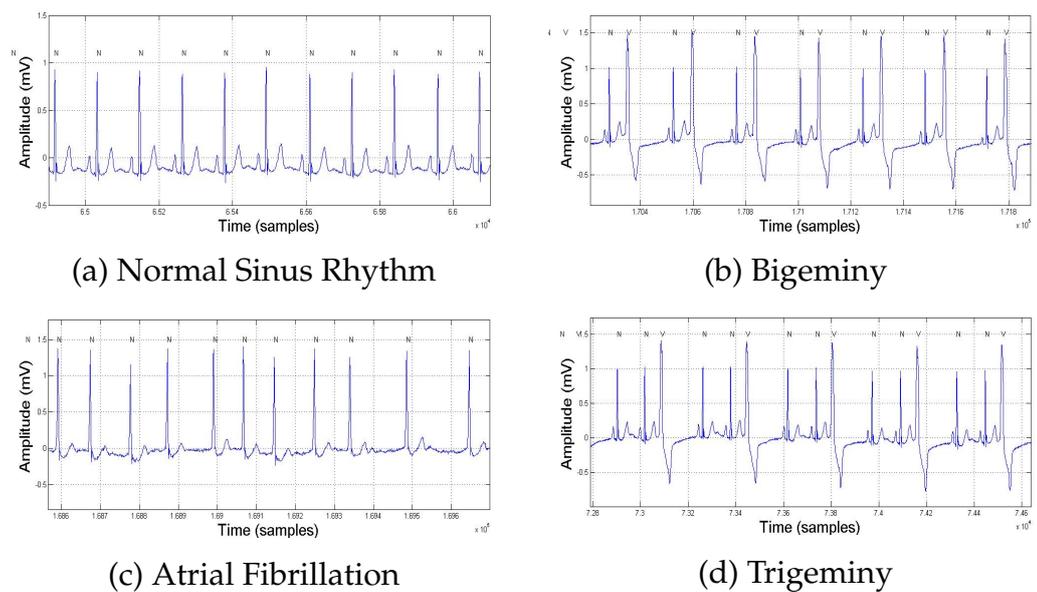


Figure 4.1: Rhythm types showing (a) NSR with P-wave present, (b) Bigeminy with a premature beat for every normal beat, (c) AF with random RR intervals and absent P wave, and (d) Trigeminy with premature beat for every two normal beats.

With the above knowledge, it is proposed that the AF detection scheme include two important feature sets. Firstly, the timing of the RR intervals will be analysed to determine their randomness, and secondly, the type of QRS complex will be found in order to more accurately decide on the type of arrhythmia being presented. A classification system is proposed to decide whether the heart is in NSR or experiencing some form of arrhythmia, such as AF.

Five areas to be investigated are:

- Signal pre-processing
- QRS complex detection
- Rhythm analysis
- Waveform morphology analysis
- Classification and detection of rhythm type

The following sections describe each of the areas of research.

4.2 Signal pre-processing

The accurate, reliable detection of the QRS complex, subsequently referred to as R peak detection, is paramount for accurate measurement of the RR interval. All statistics concerning RR interval and arrhythmia classification will be carried out on information derived from R peak detection. Consequently, a reliable, accurate R peak detector must be utilised. This exposes a number of problems that need to be resolved before the RR interval can be calculated. Firstly, short bursts of noise can emulate the high frequency content of the QRS complex and trigger false detections, and secondly, baseline wander, or movements in the underlying DC component of the ECG signal, can create signal threshold problems. Preconditioning of the ECG signal to remove low frequency baseline wander and

higher frequency noise needs to be carried out. Two methods of baseline wander removal are presented here.

4.2.1 Baseline wander removal: Mathematical morphology.

As briefly discussed in section 3.1, mathematical morphology can be used to remove baseline wander. The algorithms are easy to implement and are effective in removing baseline wander and sudden changes in signal bias. Baseline wander removal, shown in Figure 4.2, is accomplished by applying the following operations

$$V_b = V_s - \frac{((V_s \circ B_1 \bullet B_3) + (V_s \bullet B_2 \circ B_3))}{2}$$

where V_b is the signal with baseline removed, V_s is the original signal and B_1, B_2, B_3 are appropriate structuring elements, in this case 31, 17 and 24 sample periods in length, respectively. The size of the structuring elements B effects the degree of signal peak and valley removal thus affecting the value of the final signal peak and valley amplitudes. See section 3.1 for a more detailed description of how the structuring element length affects peak and valley removal.

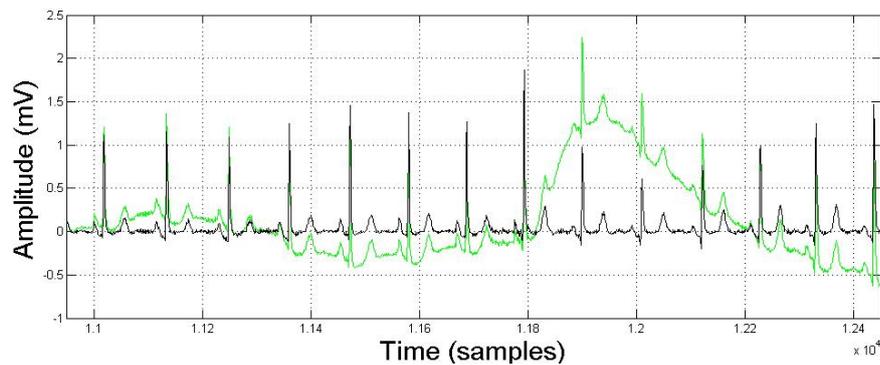


Figure 4.2: Baseline wander removal using morphological mathematics.

4.2.2 Baseline wander removal: Discrete Wavelet Transform.

Another method of baseline removal is to remove the low frequency components causing the variations in the baseline. The Discrete Wavelet Transform as discussed in section 3.2 is used for this purpose. At a sampling frequency of 128s/s, the first six sub-band outputs of the DWT are details D1 [32–64 Hz], D2 [16–32 Hz], D3 [8–16 Hz], D4 [4–8 Hz], D5 [2–4 Hz] and D6 [1–2 Hz]. By eliminating frequencies below 1Hz, significant portions of baseline wander are removed. It is also convenient at this time to suppress noise in the signal. By excluding DWT details D1, much noise and also 50/60 Hz hum is removed from the R peak detection process, increasing reliability and accuracy. By using only details D2–D5 of the DWT, baseline wander and high frequency noise is removed from the ECG signal making it more suitable for the next stage of processing i.e., QRS detection.

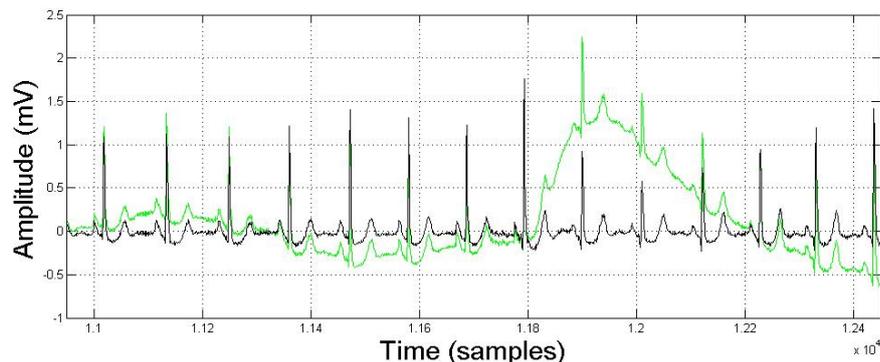


Figure 4.3: Baseline wander removal using discrete wavelet transform.

It is proposed here to use the DWT method of baseline wander removal, as it has a number of benefits. First, not only is the baseline wander removed, so is higher frequency noise and 50/ 60 Hz hum. Second, all DWT details D1–D6 produced, are available for further processing. Also, being a linear

transform, the ECG signal can be easily reconstructed if required.

4.3 QRS complex detection

As mentioned in the previous section on signal pre-processing, the accurate, reliable detection of the QRS complex (referred to as R peak detection), is critical for all subsequent rhythm calculations. After the removal of unwanted signal anomalies such as baseline wander and noise, the RR interval can be evaluated with far more confidence. Two methods of R peak detection are discussed, both using mathematical processes to find the amplitude peak of the QRS complex. They are the Hilbert Transform (section 3.3) and the Teager-Kaiser energy operator (section 3.4) .

4.3.1 R peak detection: Hilbert Transform.

The Hilbert transform is applied to DWT details D2–D5 to obtain a set of imaginary signals. The real and imaginary signals are then used to find the respective signals energy using $A(t)^2 = V_R(t)^2 + V_I(t)^2$. All four signal energy levels are then added, with appropriate scaling, to form a signal (Rpeak), for locating the QRS complex.

4.3.2 R peak detection: Teager-Kaiser Energy Operator.

A second method of R peak detection is to apply the TKEO to DWT details D2–D5, providing an estimate of each signals energy. This is a relatively efficient process compared to applying the Hilbert Transform. Again, all four signal energy levels are added, with appropriate scaling, to form the Rpeak signal to locate the QRS complex.

It is proposed here to use the Hilbert Transform method of R peak detection. Although this method is more complex compared to using the TKEO, the imaginary signal produced will be useful for waveform feature creation as discussed in section 4.5.

4.3.3 R Peak location.

The location of the QRS complex is now found by a simple peak detector algorithm applied to the Rpeak signal. By finding the 'difference of the sign of the difference' of the Rpeak signal, a simple, effective R peak detector is created with positive one indicating positive peaks, negative one indicating negative peaks and zero indicating no peak detected. From this, positive peak positions can be determined, see Figure 4.4.

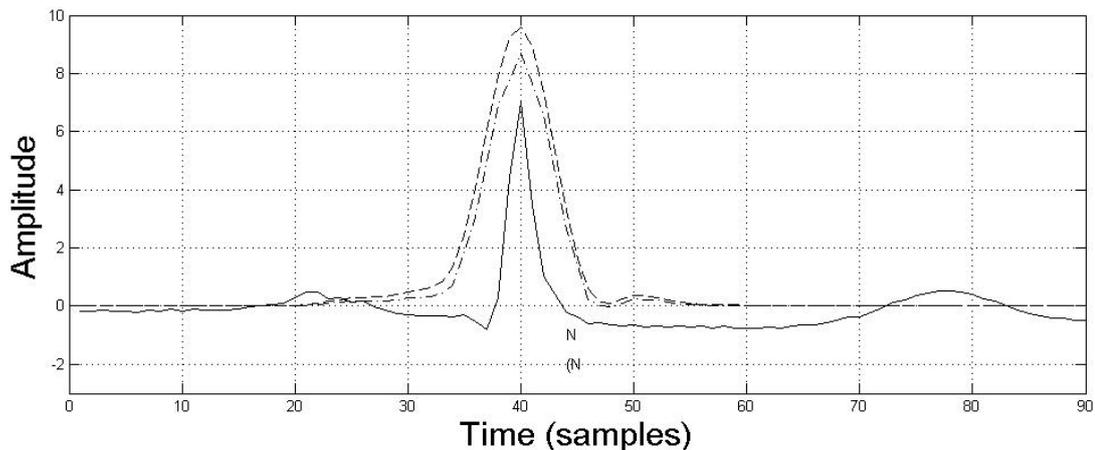


Figure 4.4: R peak detection showing ECG signal (—), D2 peak (--) using Hilbert Transform, and D2 peak using TKEO (-·-).

4.3.4 R Peak validation.

For all positive peak locations, the amplitude of the R peak is taken and if it exceeds a predetermined threshold, the peak is considered to be a valid R peak for timing purposes. Having a constant threshold value for R peak detection is a simple approach. A more sophisticated approach is to use adaptive thresholding. In this way, timing of subsequent heart beats can be determined statistically. For instance, a premature beat can occur no less than about 200ms after a previous beat, so it is unlikely a beat will occur within that time frame because the ventricles do not have time to re-polarise; the threshold can be raised during that period. Also, if a certain rhythm is in progress, then it is more likely that that rhythm will be sustained, so any beats detected outside the current rhythm pattern are more likely to be invalid R peaks. The R peak threshold can be adapted, as the rhythm changes over time, to more accurately detect R peaks and reject noise and motion artefact, see Figure 4.5. By combining the inverse of the ever changing rhythm PDF, Figure 4.7, with the PDF mask, Figure 4.6, an adaptive threshold is created

Up to this point, this chapter has concentrated on the preconditioning of the ECG signal ready for rhythm detection. Baseline wander removal and noise reduction has been discussed, and methods of R peak detection have been proposed. Next, methods of arrhythmia detection are examined and final arrhythmia classification is discussed.

4.4 Rhythm analysis

After preprocessing and QRS complex detection, the RR interval metrics can be evaluated. Figure 4.7 shows the Probability Density Functions

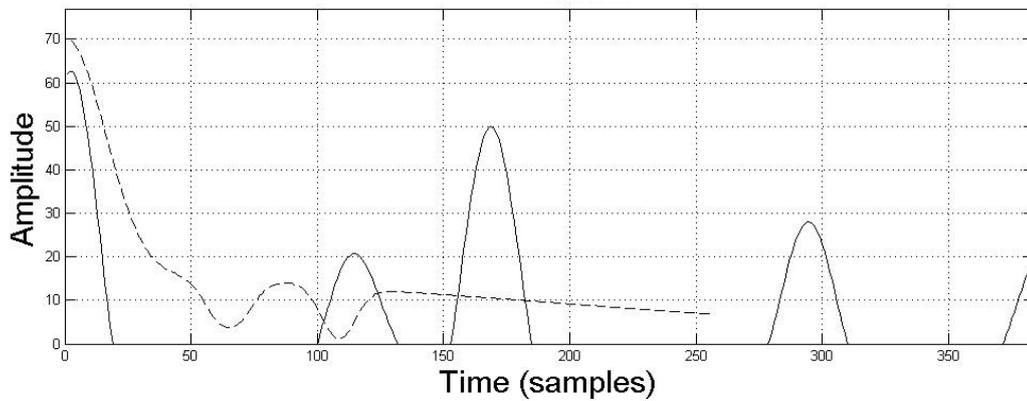


Figure 4.5: R peak adaptive threshold showing R peaks (—) and threshold (---). R peaks greater than threshold will be detected as valid R peaks, those below are ignored. This example shows two dips in the threshold due to the presence of Bigeminy.

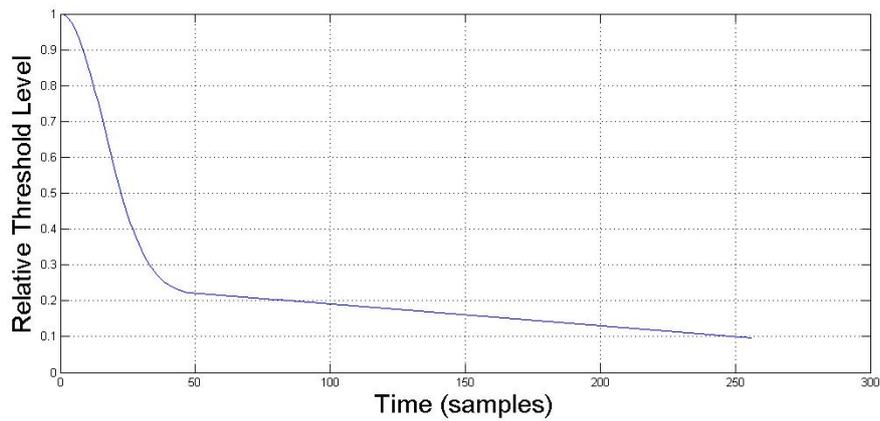


Figure 4.6: R peak adaptive threshold mask. The closer to the previous R peak (to the left in this diagram), the greater the threshold and the less likely a beat will be detected.

(PDFs) of four rhythms. In these examples, six consecutive RRint samples are taken and smoothed with an eight point Gaussian window. NSR is

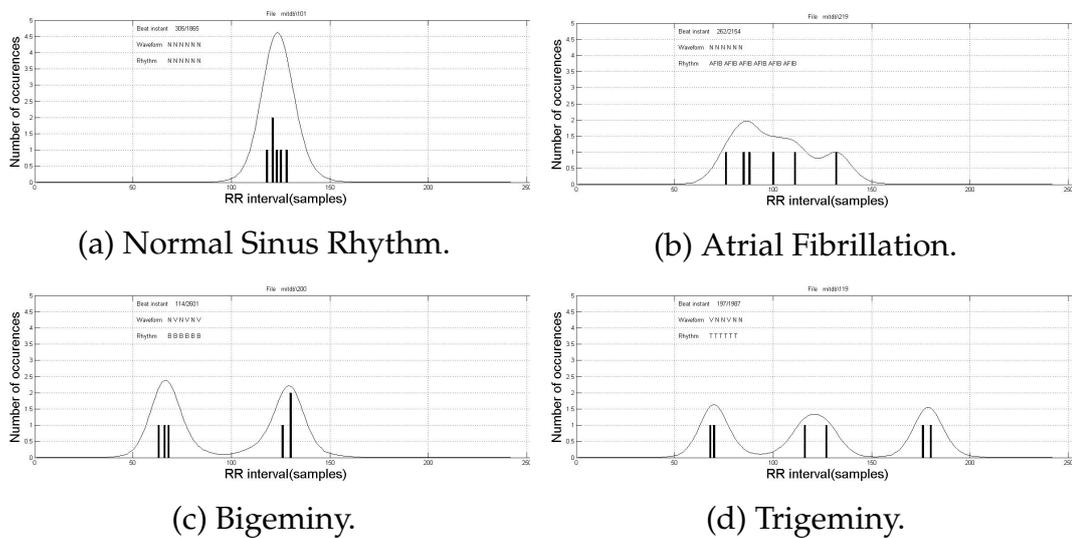


Figure 4.7: Various Rhythm PDF's using a 6 sample frame (bars) and smoothed with 8 point Gaussian window (—) .

shown as a narrow, high peak distribution similar to a Gaussian curve. AF is seen as a wider, more flat distribution, whereas bigeminy and trigeminy have a two and three bump distributions, respectively. Although these plots are useful for visualising the different arrhythmia, simpler metrics need to be obtained for implementation. Heart rate variability is calculated by simply taking the difference (dRR_{int}) of consecutive RR intervals. This signal shows interesting characteristics. Normal sinus rhythm should have a small amount of variability indicating the constant shift in control of the heart rate by the opposing sympathetic and parasympathetic nervous systems [2]. No heart rate variability signifies problems with these autonomous nervous systems [2]. If the variability is high, then arrhythmia is present. By creating a number of RR interval metrics, the type of arrhythmia becomes detectable.

Plotting the dRR_{int} signal, Figure 4.8, shows interesting relationships between consecutive RR intervals for different types of rhythm. NSR for instance, has low variability whereas bigeminy shows large, symmetrical

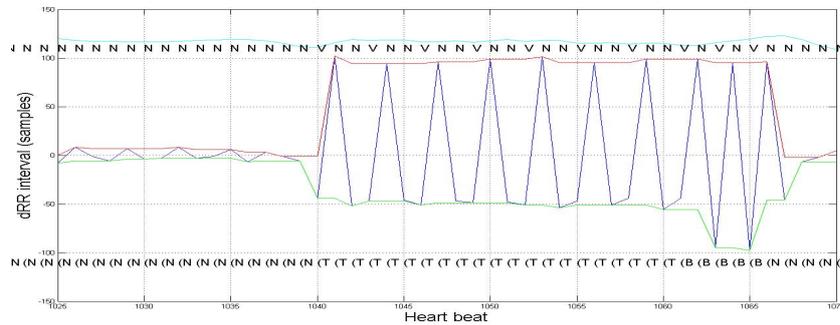


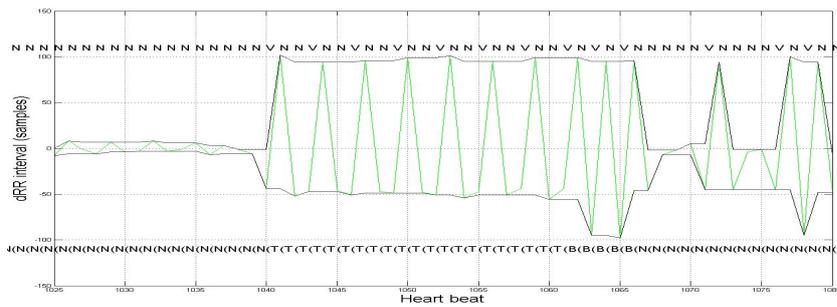
Figure 4.8: RR interval difference showing episodes of NSR, bigeminy and trigeminy. The positive and negative magnitudes are tracked using a three sample MM operator.

positive and negative magnitudes. Contrast that with trigeminy having asymmetrical positive and negative magnitudes, usually in a ratio of 2:1. AF, on the other hand, has larger variability in positive and negative magnitudes. These differences may be useful in determining the type of arrhythmia present.

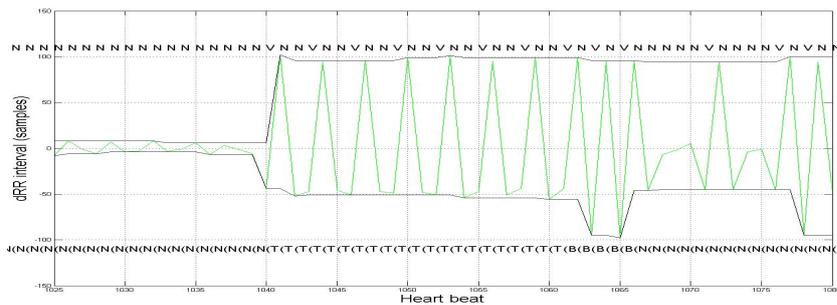
By using a mathematical morphology operator to track the positive and negative amplitudes of dRRint signal, useful indicators of rhythm are created. Bigeminy requires at least a two point structuring element to track the positive and negative dRRint values as it is characterised by two values that alternate between positive and negative. Trigeminy requires at least a three point structuring element as it is characterised by one positive dRRint value followed by two negative values. The dRRint peak tracking for bigeminy is not affected by using a three point structuring element instead of two, consequently a three point structuring element can be used to track both bigeminy and trigeminy, creating a signal $dRRint(3)$. The opening operator is used to track the positive peaks and the closing operator is used to track the negative peaks.

To detect arrhythmia a number of rhythm metrics are calculated. First, a moving average (RRmean) of the RRint is obtained by using a 16 point

Gaussian window applied to the RRint signal. Second, the difference in RR interval (dRRint) is taken and from that an asymmetrical envelope tracker is used to create a signal dRRint(3) as described in the previous paragraph. When interpreting the ECG signal, it is customary not to recognise a change in rhythm until two or more consecutive sequences of the new rhythm are detected. To take this into account, a further simple metric is created using MM operators with a six sample structuring element. This creates a signal that tracks dRRint when more than two sequences of bigeminy or trigeminy occur, and creates a signal dRRint(6) in a manner similar to dRRint(3). In total, seven rhythm features are extracted including RRint, RRmean, dRRint, \pm dRRint(3), \pm dRRint(6).



(a)



(b)

Figure 4.9: RR interval difference showing (a) RR interval difference \pm dRRint(3). (b) RR interval difference \pm dRRint(6).

However, it is not always that clear cut. Bigeminy and trigeminy can be caused by either PAC's or PVC's. It may not be necessary to distinguish between the two for the purposes of detecting AF, but it would be advantageous. PAC's, being premature atrial contractions, show normal sinus morphologies, whereas PVC's, being premature ventricular contractions, show QRS complex and T wave morphologies that are very different from that of sinus derived waveforms. To distinguish the difference between the two, the waveform morphology must be determined. This opens up the possibility of distinguishing between other, more rare types of arrhythmia and making AF determination far more reliable and accurate.

4.5 Waveform morphology analysis

In the previous section the RR interval was used to determine the rhythm of the heart beat and thus distinguish between arrhythmias. However it was noted that the ability to distinguish between different heart beat morphologies would be of great advantage, increasing the accuracy and reliability of AF detection. Following is a discussion on methods for determining the morphology of the heart beat waveform. Two approaches are considered:

- The amplitudes and positions of each of the DWT detail peaks with respect to the R peak location.
- The amplitudes of the Real and Imaginary outputs of the Hilbert Transform for each of the DWT details at specific sample periods from the R peak location.

4.5.1 DWT detail peak amplitude and position

The ECG signal was first decomposed into sub-bands using the discrete wavelet transform. The first six sub-bands were considered adequate as lower frequency bands included unacceptable levels of baseline wander. The peak amplitudes of the six frequency sub-bands were plotted (left column of Figure 4.10), along with the relative position of the peak about the centreline of the detected QRS complex (centre column of Figure 4.10). For interest, a two dimensional image depicting the amplitude vs. position of each peak was created (right column of Figure 4.10)

The left column depicts amplitudes of each type of waveform for DWT details in order of lowest frequency (D6) to highest frequency (D1). The plots demonstrate that differing waveform morphologies have differing frequency spectrum. The position of these peaks with respect to the located R peak is shown in the centre plots. A cursory look at the plots suggests that these may be useful features for an appropriate artificial intelligence classifier to determine waveform morphology.

4.5.2 DWT detail Real and Imaginary amplitudes

By applying the Hilbert Transform to the ECG signal and then using the resulting real and imaginary signals, features for waveform classification can be developed. To gauge the validity of this approach, a number of polar plots of the real vs. imaginary ECG signals were created and shown in Figures 4.12. They show the magnitude vs. phase for each ECG waveform types of Normal, PAC, Left Bundle Branch Block (LBBB), Right Bundle Branch Block (RBBB), and PVC. Differences can be detected in the plots for the five waveform types shown, however, it wasn't considered necessary to convert to polar form for feature generation. The same information is present in the original real and imaginary signals.

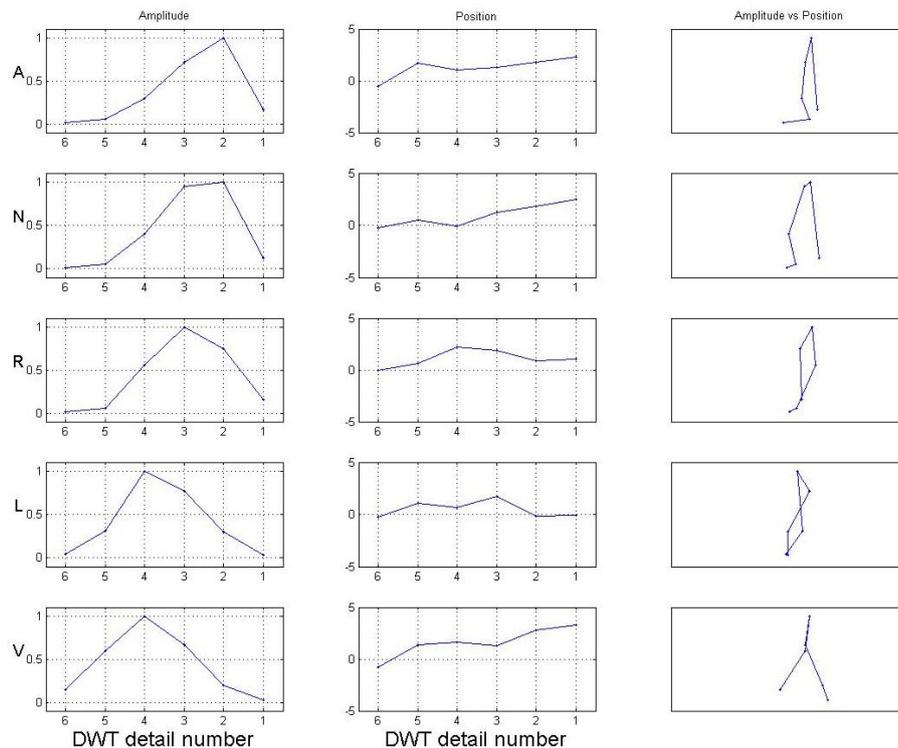


Figure 4.10: Frequency spectrum of 5 waveform morphologies. The left column shows relative peak amplitude and the centre column the relative position of the peak with respect to the R peak centre. The right column is a two dimensional representation of peak amplitude vs. peak position. Morphologies shown are: A (PAC), N (NSR), R (Right bundle branch block), L (Left bundle branch block), V (PVC)

4.5.3 Waveform features extraction.

It was decided to pursue the second option of feature generation, the Hilbert transform. Although more complex than determining the DWT detail peak and location, it was considered more mathematically robust. Values of each of the real and imaginary signals at the R peak location were used for features giving a total of 12 features for classification pur-

poses. However, as will be discussed in the next section, it was found that better classification was accomplished when two values, four sample periods each side of the R peak location were used. This gave a total of 24 features for waveform morphology classification, see Figure 4.11.

Now that the rhythm and waveform features have been extracted from the ECG signal, they are ready to be presented to the classifier section for ultimate rhythm determination.

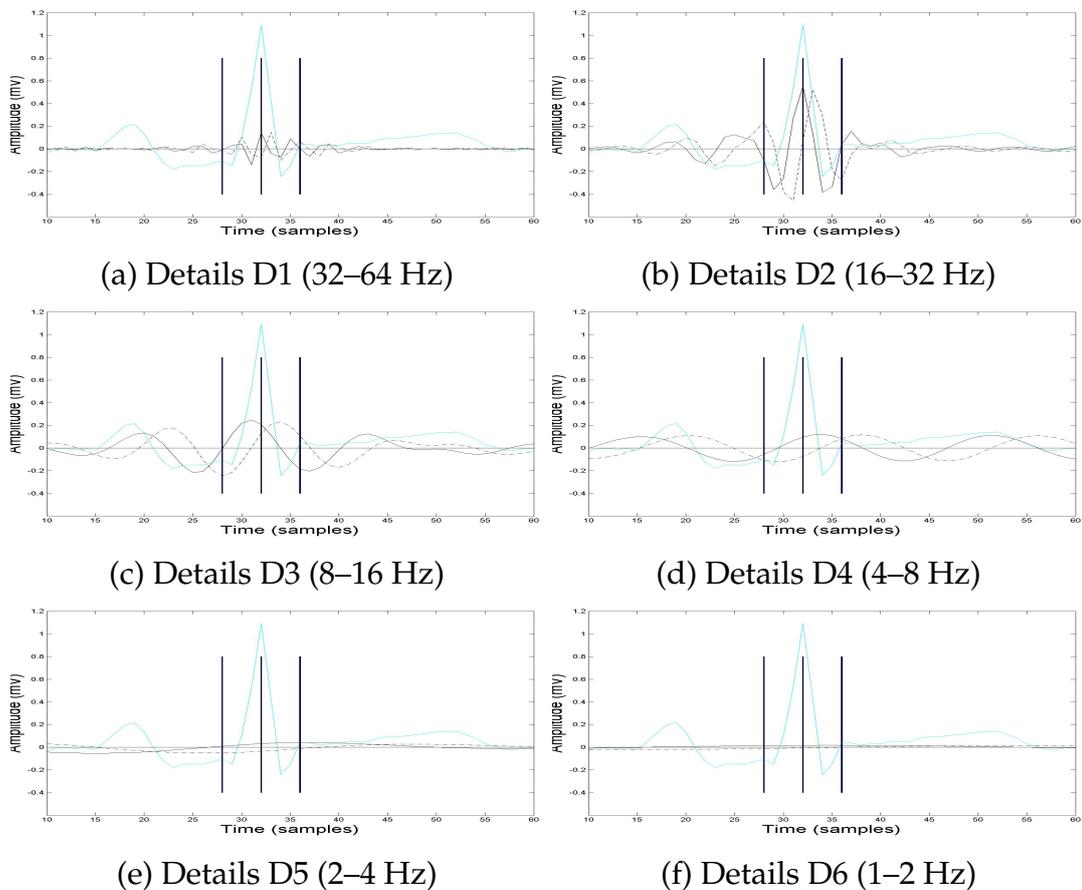
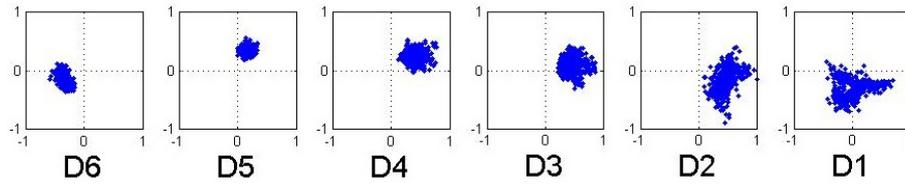
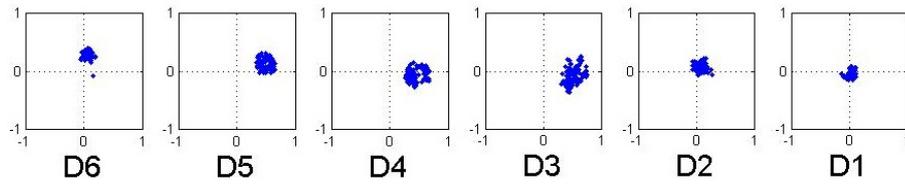


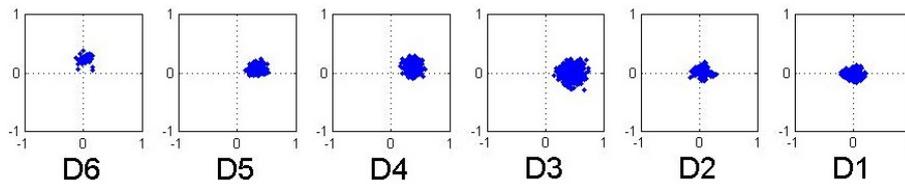
Figure 4.11: ECG waveform features showing ECG signal (light solid), six real (dark solid) and six imaginary details (dark dashed). The black bars show the positions of the feature points four samples either side of the R peak centre.



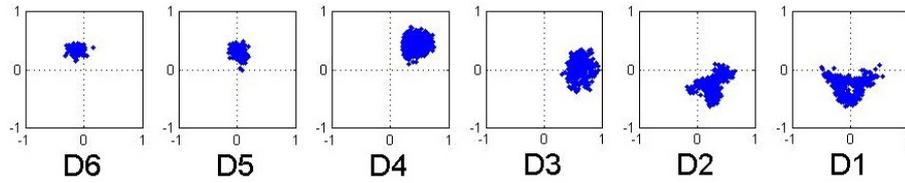
(a) Premature atrial contraction.



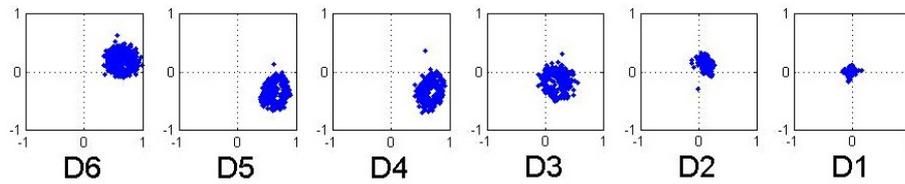
(b) Normal sinus.



(c) Left bundle branch block.



(d) Right bundle branch block.



(e) Premature ventricular contraction.

Figure 4.12: Polar plots derived from the Hilbert transform of the ECG signal, showing relative magnitude and phase with respect to the R peak location for each DWT detail for five ECG morphologies.

4.6 Classification and detection of rhythm type

In the previous two sections, ECG rhythm and waveform feature extraction was described. Here a classification method is proposed to determine the rhythm of the signal.

4.6.1 ECG Arrhythmia Classification

In section 3.5, the support vector machine was described in detail and is used here to determine the rhythm of an ECG signal by passing the ECG feature set onto the classifier stage shown in Figure 4.13.

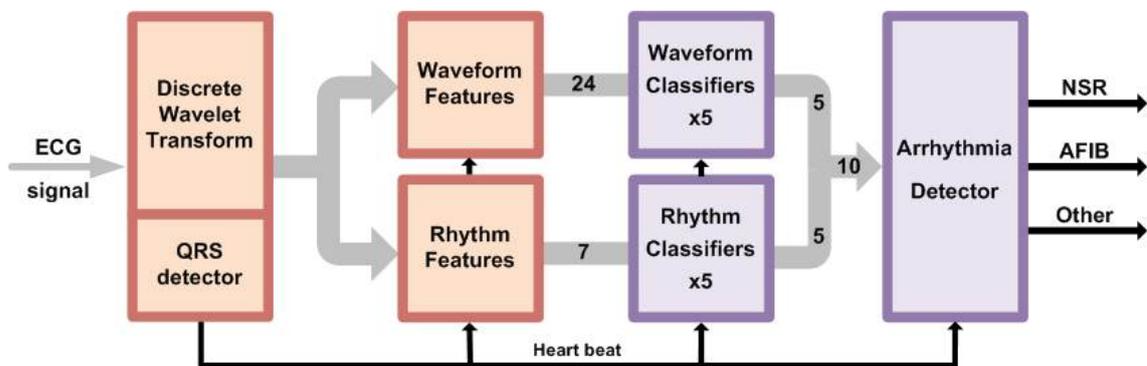


Figure 4.13: Classifier stage showing five Waveform classifiers, five rhythm classifiers and the final arrhythmia classifier.

The classification stage is made up of a number of individual classifiers. For rhythm detection, five classifiers are used to detect the presence of the following rhythms; Normal Sinus, Atrial Fibrillation, Atrial Flutter, Bigeminy and Trigeminy. For waveform detection, another five classifiers are used to detect the presence of the following morphologies; Normal Sinus, Premature Atrial Contraction, Premature Ventricular Contraction, Left Bundle Branch Block and Right Bundle Branch Block. The regression values of these 10 classifiers are then passed as inputs to the final classification stage for ultimate determination of rhythm. The final stage comprises

of two classifiers, one for NSR and the other for AF, with the binary output indication shown in Table 4.1.

NSR	AF	Class
0	0	Other rhythm
1	0	Normal sinus rhythm
0	1	Atrial Fibrillation
1	1	Classification error

Table 4.1: Final Classification.

The idea behind the classification stage design is to provide feature sets related to both the rhythm and the waveform morphology to improve overall classification performance. The design is similar to that of radial basis function neural network [17], the difference being the final layer here is another set of non-linear classifiers instead of the usual linear summation layer. By passing the regression values from the first set of classifiers to the second, the second stage has a set of pertinent features enabling better rhythm determination.

4.6.2 Training and testing procedures

For training and testing of the system classifiers, files from the PhysioNet [14] online repository are used. A number of databases are available for use including the following:

- MIT-BIH Normal Sinus Rhythm Database (NSRDB).
- MIT-BIH Atrial Fibrillation Database (AFDB).
- Long Term Atrial Fibrillation Database (LTAfDB).
- MIT-BIH Arrhythmia Database (MITDB).

All databases have annotations provided to describe each beat (i.e., ECG waveform type) with the MITDB also having rhythm annotations. Because of this, the MITDB is used for all training and testing of the system classifiers. This database includes 48 files of 30 minutes each with a variety of waveforms and rhythms including normal sinus, PAC and PVC waveform morphologies and NSR, AF rhythms. Initial training and testing was carried out at annotated R peak locations rather than at the detected R peak locations. This was done to set a standard for comparison in testing subsequent R peak detectors.

This chapter has described the various methods employed to extract the necessary features required for classifier training and testing. The following chapter describes in more detail how the features were extracted and the testing regimes used. Results of the testing are presented.

Chapter 5

Results and Discussion

In previous chapters, mathematics and methods that form components of a system to detect AF were presented. Here, testing of the proposed system is carried out and the results discussed.

5.1 Datasets and testing

To test the waveform and rhythm classifiers, files from the PhysioNet MIT-BIH arrhythmia database were used. In total the database consisted of 48 files from 48 individual patients, each 30 minutes in duration and having a total of 109,000 beats (approx. 2200 per file). Each file was processed to extract 31 features made up of 24 waveform features and 7 rhythm features,

The waveform features were made up of values extracted from the DWT of the signal. Six frequency bands of the DWT (D1–D6) were used and further processed to produce the 24 features. This processing consisted of generating the Hilbert Transform of each of the DWT details and obtaining the amplitude of the resulting real and imaginary signals at a number of locations either side of the detected R peak. Initial tests showed that only two points at ± 4 sample points from the designated R peak location were required to produce reasonable results. More samples points only marginally increased performance. With two points at ± 4 for each of

the six DWT details for each of the Real and Imaginary signals, giving a total of 24 waveform features.

The rhythm features were generated from the interval between successive R peaks. RR intervals (RRint) of each ECG file were calculated and a 16 point moving average (RRmean) produced. From the first difference (dRRint) of the RR interval, four more features were created by applying an asymmetrical envelope detector using mathematical morphology operators with structuring elements three and six heart beats in length, see section 4.4 for a more detailed explanation. In total, seven rhythm related features were generated including the RRint, RRmean, dRRint, \pm dRRint(3), \pm dRRint(6). A training set of 10,000 randomly chosen beats (waveform and rhythm features sets) were extracted from the 48 individual files and used for training of the classifiers. For testing, the remainder of the 109,000 beats were used.

5.2 Results

Waveform and Rhythm classifiers

Table 5.2 shows the classification ability of the waveform classifiers. The number of support vectors retained ranged from 8% to 23% of those presented for training. Premature Atrial Contractions (the A type waveform) was the most difficult to classify and this is not surprising as the A type waveform morphology is similar to the N type. What was surprising was the relative ease that the other classifiers identified their respective waveform types considering they were set up for one-against-the-rest classification.

Table 5.1 shows the classification ability of the rhythm classifiers. The number of support vectors retained ranged from 15% to 30% of those presented for training. Most show sensitivity and specificity in the high 80% range with NSR being the worst case with sensitivity of 79%.

Rhythm Type		Sensitivity	Specificity	C	σ
Atrial Fibrillation	(AFIB)	0.90	0.90	1	0.5
Atrial Flutter	(AFL)	0.88	0.88	1	1
Bigeminy	(B)	0.88	0.94	1	1
Normal sinus rhythm	(NSR)	0.79	0.89	3	0.5
Trigeminy	(T)	0.89	0.88	1	2

Table 5.1: Rhythm Classification, σ is radial basis function parameter and C is the capacity control parameter, with $0 \leq \alpha \leq C$.

Waveform Type		Sensitivity	Specificity	C	σ
Premature atrial contract.	A	0.74	0.99	6	4
Left bundle branch block	L	0.97	0.99	6	4
Normal sinus beat	N	0.98	0.93	6	4
Right bundle branch block	R	0.98	0.99	6	4
Premature ventricular contract.	V	0.93	0.99	6	4

Table 5.2: Waveform Classification.

The final two classifiers provided the results shown in tables 5.3 and 5.4. Receiver operating characteristic (ROC) curves for the two final classifiers are shown in Figure 5.1. These curves were created by adjusting the bias (b in equation 3.2) during the testing phase. Both exhibit good classification ability. Table 5.3 shows the accuracy of the final classification stage when only the classes of the previous stage are passed as features to the final stage, whereas table 5.4 shows the accuracy when the regression values of the previous stage are passed as features to the final stage. This shows an improvement in overall classification when using the regression values as opposed to just the classes.

Final Rhythm Class		Sensitivity	Specificity	C	σ
Atrial Fibrillation	(AFIB)	0.91	0.91	1	3
Normal sinus rhythm	(NSR)	0.91	0.92	1	3

Table 5.3: Overall Rhythm Classification using classes from previous classification stages.

Final Rhythm Class		Sensitivity	Specificity	C	σ
Atrial Fibrillation	(AFIB)	0.94	0.93	1	3
Normal sinus rhythm	(NSR)	0.93	0.95	1	3

Table 5.4: Overall Rhythm Classification using regression values from previous classification stages.

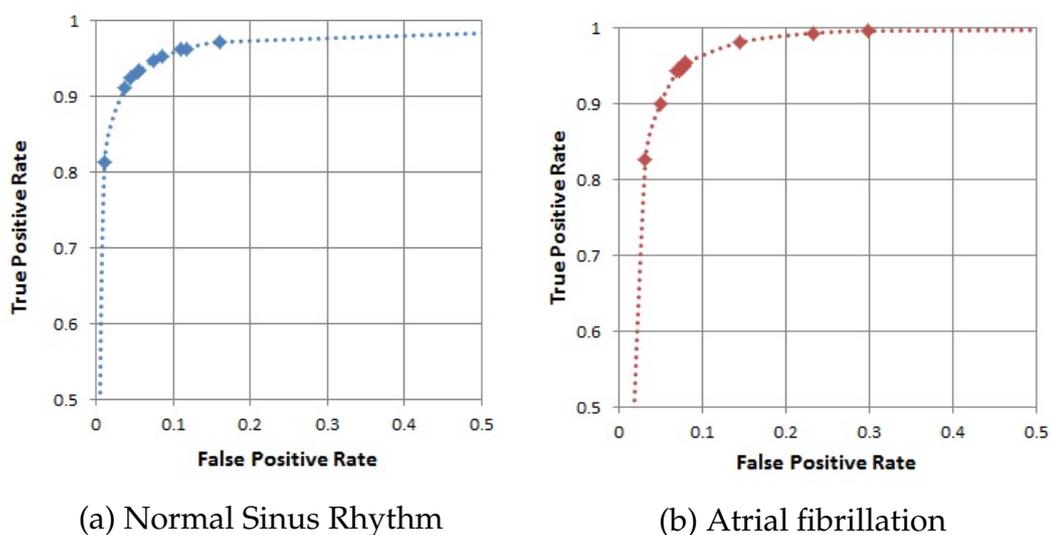


Figure 5.1: Final Classifier ROC curves. The bias (b equation 3.2) was adjusted to generate different points of the ROC curve.

To determine the necessity of the waveform classifiers, tests were carried out without the waveform feature set, using only the seven rhythm features. Final classification accuracy decreased showing that the waveform features do in fact help classification accuracy and reliability, compare tables 5.5 and 5.4

Final Rhythm Class		Sensitivity	Specificity	C	σ
Atrial Fibrillation	(AFIB)	0.92	0.88	1	3
Normal sinus rhythm	(NSR)	0.82	0.88	1	3

Table 5.5: Overall Rhythm Classification using only rhythm features and regression values from previous classification stages.

All the tests above were carried out by randomly choosing features from all 48 patient files for the training and testing of the classifiers. To determine classifier generalisation, further tests were carried out whereby classifier training used random features from 47 of the 48 files, and classifier testing used random features from the remaining unused file, in a leave-one-out training scheme. Tables 5.7 and 5.6 show the classification accuracy of the five rhythm and the five waveform classifiers.

Rhythm Type		Sensitivity	Specificity	C	σ
Atrial Fibrillation	(AFIB)	0.75	0.87	1	0.5
Atrial Flutter	(AFL)	0.73	0.88	1	1
Bigeminy	(B)	0.71	0.90	1	1
Normal sinus rhythm	(NSR)	0.73	0.89	3	0.5
Trigeminy	(T)	0.73	0.89	1	2

Table 5.6: Rhythm Classification using 47 of 48 files for training

Waveform Type		Sensitivity	Specificity	C	σ
Premature atrial contract.	A	0.2	0.96	6	4
Left bundle branch block	L	0.04	0.98	6	4
Normal sinus beat	N	0.81	0.92	6	4
Right bundle branch block	R	0.76	0.95	6	4
Premature ventricular contract.	V	0.76	0.93	6	4

Table 5.7: Waveform Classification using 47 of 48 files for training.

The classifiers show some degradation in accuracy which is expected, but the waveform classifiers now have difficulty in distinguishing premature atrial contractions and left bundle branch block waveform morphologies.

5.3 Discussion

The initial tests showed good classification ability of the AF detection scheme. In all the above tests, the R peak centre used is the annotated R peak associated with each occurrence of the QRS complex. The annotations were provided by PhysioNet and positioned by hand by cardiac experts. The next stage of research will automatically detect the R peak as described in section 4.3, generate the feature sets and carry out the same series of tests for comparison. By doing so, the effectiveness of the overall AF detection scheme will be measured.

The waveform classifiers showed reduced generalisation ability when using 47 of 48 files for training in a leave-one-out training scheme. This could be due to the fact that all QRS complexes of one type (e.g., N), for one patient, are almost exactly the same morphology. In the future, to overcome this issue the waveform classifiers will be trained by using any number of PhysioNet files that have annotated QRS complexes. This will

train the classifier with a more diverse set of waveform morphologies and improve generalisation. Another view of looking at this problem would be to carry out online training of the classifiers to tune them to a specific patients morphologies, but this would require more advanced algorithms and consequently more advanced digital signal processors. The rhythm classifiers showed good classification ability even when trained with 47 of 48 files and tested with the remaining file.

The final classifier stage was tested with two sets of feature sets, one including waveform and rhythm features from the previous classification stage, and the second, using only the rhythm features from the previous stage. The second test showed a small decline in classification accuracy over the first. This begs the question on whether the waveform features are really necessary for AF detection. Further research will likely answer this question, however as discussed in section 4.1, various arrhythmia are associated with particular waveform types.

This chapter discussed the tests carried out to determine the ability of the proposed AF detection scheme to detect AF. The following chapter describes the implementation of this detection scheme on the target device.

Chapter 6

Implementation

This Chapter describes implementation of the proposed AF detection process on the target device. Throughout the research and design of the monitor, attention has always been paid to the end goal of implementing the theory using integer and fixed-point mathematics. Without the accuracy of floating point arithmetic, many of the processes are approximated, thus creating errors that must be managed so that the end result of reliable AF detection is achieved.

6.1 Hardware

Following is a description of the design and implementation of the hardware used in the target device. Key design criteria for hardware implementation include small form factor and low power consumption. However, this is in contrast to having enough processing power to carry out the intensive digital signal processing required for the task at hand. Usually DSPs require high clock speeds with corresponding increased power consumption. DSPs also come with high pin count for use in applications involving many digital/analogue input/outputs. As a compromise solution, a device with a hardware multiplier and low power design with minimal pin count would suffice. For detecting the ECG signal, an appli-

cation specific Analogue Front End (AFE) would be necessary to reduce component count and power consumption. These devices would also add enhanced resources desirable for the application.

6.1.1 ECG AFE

There are two low power, small form factor ECG front end ICs currently available: the Texas Instruments ADS119x/129x and the Analogue Devices ADAS1000 family of ECG ICs. The Texas Instruments ADS1298 was chosen for this research. Briefly, this AFE has all the necessary analogue amplifiers, analogue-to-digital converters and digital circuitry necessary to detect, amplify and digitise the ECG signal. It has eight ECG channels, more than adequate for this research. Digital output is via a three wire SPI port that easily integrates with many micro-controllers. The ability to program the numerous on-board features is very desirable and allows the trade off between circuit complexity and power consumption.

6.1.2 Micro-Controller

There are many micro-controllers on the market that have the abilities required for this research. The TI MSP4305528 microcontroller was chosen for its low power consumption, modest clock speed (25MHz maximum) and internal hardware multiplier capable of carrying out a 32 x 32 bit fixed point multiply and accumulate instruction, which is preferred for digital signal processing. A modest pin count of 64 gives sufficient number of input and outputs for the target device design, and it has features such as two SPI ports, a USB port and a JTAG port for system programming and debugging. The memory mapped hardware multiplier also has the ability to be controlled by a Direct Memory Access controller that allows faster filter convolution processing.

6.1.3 PCB design and layout

Other support electronics include keypad, LCD display, non-volatile flash memory, four digital to analogue converters for debugging, and power supply regulators along with a Li-Pol battery charger. The PCB was designed in house and all manufacturing was carried out by specialised contract manufacturers. During the PCB design phase, particular attention was paid to reducing interference by separating the low level analogue signals from the high noise digital circuitry.

6.2 Software

Following is a description of the implementation of the various software modules in the target device, each performing a specialised mathematical function.

6.2.1 Pre-processing (Discrete Wavelet Transform)

From discussions above on the DWT, see section 3.2, it can be seen that ultimately, the DWT is simply a series of bandpass filters. Implementation is therefore straight forward. Six Finite Impulse Response (FIR) filters are used, each having half the bandwidth of the previous. Filter coefficients were created in Matlab and rounded to provide 16 bit values. Code for the convolution algorithm is relatively simple with a loop that is executed once for every filter coefficient, see Listing 6.1

```

convolve      clr      &MPYS           ;clear h/w multiplier
              clr      &OP2
next_tap     mov      @R8+,&MACS       ;multiply and accumulate
              mov      @R6+,&OP2
              cmp      R6,R7           ;finished?
              jne      next_tap        ;no, get next tap
              mov      &RES1,R10      ;yes, save result
              ret

```

Listing 6.1: Assembly code for FIR filter convolution algorithm.

6.2.2 R peak detector

Each of the FIR filter outputs D2–D5, are passed to the TKEO algorithm thus producing four energy level estimations. These are added, with appropriate scaling, to produce the R peak signal used for determining the location of the QRS complexes. See Listings 6.2 and 6.3

```

TKEO         mov      R5,&MPYS
              mov      R5,&OP2         ;square s(n)
              negate   R6             ;negate s(n-k)
              mov      R6,&MACS       ;multiply & accumulate s(n+k)
              mov      R4,&OP2       ;s(n)^2 - s(n-k)*s(n+k)
              ret

```

Listing 6.2: Assembly code for TKEO.

```

Rpeak      sub    @R9+, Raddr(R9)      ; difference
           and    #8000h, Raddr(R9)  ; sign
           sub    Raddr(R9), Raddr-2(R9) ; difference
           cmp    #08000h, chan1_RPK+8 ; not-peak if 0
           jne    not_peak
           cmp    #1000h, chan1_rpk+8 ; compare to threshold
           jl     not_peak

```

Listing 6.3: Assembly code for R peak detector.

6.2.3 Support Vector Machine

One of the more difficult mathematical processes to implement in assembly is the Gaussian Kernel used by the SVM. There are three common methods of approximation of the exponential function, which is required for the RBF kernel.

- Taylor series expansion
- CORDIC algorithms
- Table lookup

Each has advantages and disadvantages. Taylor series expansions can provide accuracy but at the cost of computation time. CORDIC arithmetic is accurate and efficient. The algorithms were developed by Jack E. Volder [44] and later expanded by John Stephen Walther [45] of Hewlett Packard for the calculation of many common functions such as sine, cosine, and exponential on small, resource constrained computers. These algorithms were designed to use only add/ subtract and left-shift/ right-shift, as hardware multipliers were not available. For the SVM, high accuracy is not required. In simple terms, the kernel maps the squared Euclidean distance $[0, \infty]$ to $[1, 0]$ before further processing. A table lookup suffices as

long as the results are repeatable. Off line training of the SVM was carried out using a lookup table, instead of the built-in exponential function, to emulate the target device.

Another issue that arises using the SVM is that of feature scaling. During off-line training, the training algorithm uses the common practise of normalising all inputs by subtracting the mean and dividing by the standard deviation. On the target device, normalising can be problematic. Firstly, division is computationally expensive and secondly the resulting values still need to be scaled to $[-32768, +32767]$. This can create a loss of resolution. It would be better to not scale the input features, but this also creates its own problems, especially when differing feature sets have different ranges. Consider the waveform features (discrete wavelet transform outputs) having range $[-32768, +32767]$ and the rhythm features (RR interval timing) having range $[0, 256]$. By using different classifiers for different feature sets, each classifier is presented with features that have the same range. The outputs of all the first stage classifiers are in the range of $[-32768, +32767]$, effectively normalising the inputs for the final classifier stage.

The exponential algorithm uses a simple but efficient lookup and interpolate mechanism. The input value, in this case the squared Euclidean distance, is split into two values. The most significant 4 bits (ms4) are used to index into two tables, one providing the slope information (m) and the other the intercept point (b), whereas the least significant 12 bits (ls12) of the input value (x) are used to interpolate between the values,

$$y = mx + b$$

For this application, only values $[0 - -11.09]$ were needed for approximation. The exponential then maps $[0 - -11.09]$ to $[1, 0.5, 0.25 - 0.00001526]$. By scaling up the input values by 5909.28 the input range becomes $[0 - 65536]$ in 16 steps of 4096, and by scaling the output range up by 65535 then the table look up algorithm maps $[0 - 65536]$ to $[65535 - 0]$, all values now within the 16 bit (2^{16}) limitation of the target device, see Figure 6.1.

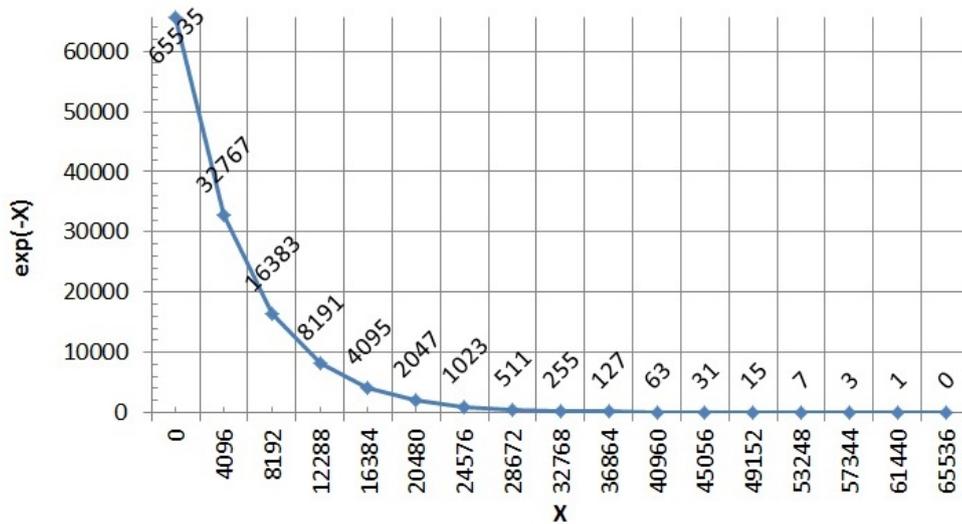


Figure 6.1: Exponential lookup table mapping $[0-65536] \mapsto [65535-0]$.

The exponential look up algorithm simply becomes,

```

expLookUp    mov     R10,R11
              and     #0F000h,R10      ;extract ms4
              swpb   R10
              rrrm   #4,R10
              and     #00FFFh,R11     ;extract ls12 (x)
              mov    #1,&MPYS         ;load 1
              mov    b(R10),&OP2      ;load b
              mov    R11,&MACS         ;load x
              mov    m(R10),&OP2      ;load m
              mov    &RES1,R11        ;result = mx + b
              ret

```

Listing 6.4: Assembly code for exponential lookup.

Chapter 7

Conclusions

A number of mathematical processes were investigated in this research project including mathematical morphology, the discrete wavelet transform, the Teager-Kaiser energy operator and the support vector machine classifier. These are basic digital processing techniques employed in many signal processing applications and have shown to be helpful in this investigation. Could more sophisticated processes such as the Kalman Filter for adaptive filtering of the ECG signal, genetic programming for finding better feature extraction algorithms, and more sophisticated classifiers for accurate determination of rhythms be employed, and can they be implemented in small form factor devices? The use of a low power device such as the MSP430 micro-controller may have been too constraining, and perhaps moving up to a true digital signal processor may produce superior results. This would enable the use of floating point arithmetic in the sensitive area of classifier implementation, and increase throughput with the higher clock speeds. Not only that, by using a higher level language such as C would allow faster and more accurate implementation of the system. However, this would be at the expense of greater power consumption and shorter operating times.

Cardiac Interpretation

Some important problems have been highlighted with regard to automatic detection of AF, and more broadly with regard to automatic interpretation of the ECG signal. A number of ECG files supplied by PhysioNet seem to have inconsistent rhythm annotations. This is in no way a reflection on PhysioNet, but on the way cardiac experts classify the types of beat morphology and rhythm instances. It appears that some arrhythmia can be very difficult to interpret with underlying issues that are hard to distinguish. This became obvious when a section of MITDB record 219 (14:50-15:00) was analysed, see Figure 7.1. To the untrained eye, this section of ECG tracing appears to be bigeminy in nature, but was annotated as normal. Subsequent communications with PhysioNet showed that the interpretation was carried out by two cardiac specialists and that the rhythm annotation was 'agreed' upon due to other complex underlining indicators. However, PhysioNet kindly located a similar example where a number of experts could not agree upon the interpretation. The example provided (<http://ecgguru.com/ecg/bigeminal-rhythm>), see Figure 7.2, emphasised the difficulty in interpreting ECG tracings. Here, four experts couldn't agree upon the cardiac problems presented. In fact, one expert contradicted that of another. If the experts can't agree on the interpretation, it makes it extremely unlikely that an automated system can be designed to do so. How can a classification system be designed, trained and tested if the classifications used for training and testing are inconsistent?

Another issue highlighted in the research is the standard practice of not indicating a change in rhythm unless more than two consecutive beats of the new rhythm have occurred. Waveform and rhythm detectors tuned for waveform morphology and rhythm detection would have to be more complex in structure or a more abstract view of rhythm detection would have to be employed. Perhaps other methods of categorising heart beat waveforms and rhythms need to be devised so hardware and software

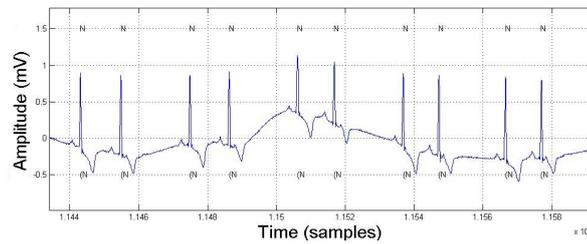


Figure 7.1: MITDB record 219 showing bigeminal rhythm classified as normal sinus rhythm.



Figure 7.2: Example of bigeminal rhythm where experts cannot agree on diagnoses.

algorithms can be more easily defined.

That being said, these two issues also highlight how incredible the human mind is at carrying out interpretation of the ECG tracing. If the human mind can analyse and detect tiny nuances in ECG signals and come up with plausible diagnoses, is it reasonable to expect an artificial intelligence classifier to do the same, especially when the ECG signal from different subjects can look so different? Perhaps a better view of miniature, automatic interpretation devices would be for them to provide pertinent indicators to enhance further decision making by experts i.e., complement the cardiologist, rather than replace. For a small, miniature device, such as the one investigated in this thesis, to ably carry out rhythm detection, perhaps it is enough to detect basic rhythm changes and report these anomalies, thereby notifying cardiac specialists that something is amiss needing further investigation. In that case, the classifier would need to reliably and accurately detect normal sinus rhythm. Anything else would be an anomaly needing attention.

Implementation

One difficulty uncovered by this research is the implementation of the SVM classifier in the target device. By choosing fixed point arithmetic, the realisation of the SVM classifier proved challenging. The issue involved the normalisation of the feature sets before classification. Research into other classifiers and, in particular, distance measuring functions, could resolve this problem. Here the Radial Basis Function was used as the kernel. By using a table lookup scheme, the exponential was approximated with reasonable accuracy. However, is it possible to conceive simpler classifiers that could be more readily implemented in small, resource constrained devices? Or, perhaps, better online normalisation of data could be achieved. Division in small micro-controllers is computationally expensive and would have to be sparingly used. Are there other methods of normalisation that could be employed in such applications? Or could normalisation be eliminated by using new types of classifiers or even the transformation of feature sets into more abstract forms? The structure presented here of using multiple classifiers helps, and has the advantage of increasing classification ability, but at a cost of running 12 classifiers. Is that necessary?

Summary

Even though this research has uncovered some difficulties, it has still shown that complex systems of this nature can be created in small form factor, low power devices. Creating algorithms that approximate complex mathematical processes in an efficient manner is the key. The table lookup scheme depicted in Listing 6.4 is an example. With only ten lines of assembly code and no conditionals, the exponential function is estimated. Also, the FIR filter convolution algorithm, shown in Listing 6.1, with eight lines of code and one conditional, and the Teager-Energy Operator with only six lines

of code, all being compact in size and fast in execution making it plausible to carry out sophisticated mathematical functions in small, low power devices. Although the detection of AF from analysing the ECG signal on such a device is challenging, this thesis shows that it is achievable. However, time is money and is it realistic to do so?

Bibliography

- [1] <http://www.ecglibrary.com/ecghist.html>.
- [2] ACHARYA, SURI, S. K. *Advances in Cardiac Signal Processing*. Springer, 2007.
- [3] BIONET. Bionet ekg plus ii software. <http://www.bionetus.com/product/ekg-plus-ii/>. Accessed 10 January, 2014.
- [4] BOUDRAA, A., BENRAMDANE, S., CEXUS, J., AND CHONAVEL, T. Some useful properties of cross Ψ_B energy operator.
- [5] BOUDRAA, A., CEXUS, J., AND ABED-MERIAM, K. Cross Ψ_B energy operator based signal detection.
- [6] CHUN-LIN, L. A tutorial of the wavelet transform.
- [7] CONEN, D., OSSWALD, S., AND ALBERT, C. Epidemiology of atrial fibrillation. *Swiss Med. Weekly*. 139 (2009), 346–352.
- [8] COUCEIRO, R., CARVALHO, P., HENRIQUES, J., ANTUNES, M., HARRIS, M., AND HABETHA, J. Detection of atrial fibrillation using model-based ECG analysis. *Annual ICPR Int. Conf.* (2008).
- [9] DAUBECHIES, I. *Ten Lectures on Wavelets*. Philadelphia:SIAM, 1992.
- [10] DEBNATH, L. Brief historical introduction to wavelet transforms. *Int, J, Math. Educ. Sci. Technol.* (1998).

- [11] DEVICES, A. Adas1000 low power, five electrode ecg analog front end.
- [12] ELDERY, M., AND SALAMA, E. E.-S. M. A novel energy operator algorithm for voltage envelope tracking. *IEEE, Transaction on Power Systems* (2005).
- [13] GLASBEY, C., AND HORGAN, G. *Image Analysis for the Biological Sciences*.
- [14] GOLDBERGER, A., AMARAL, L., GLASS, L., HAUSDORFF, J., IVANOV, P., MARK, R., MIETUS, J., MOODY, G. B., PENG, C.-K., AND STANLEY, H. PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals. *Circulation*, 23 (2000 (June 13)), e215–e220. Circulation Electronic Pages: <http://circ.ahajournals.org/cgi/content/full/101/23/e215> PMID:1085218; doi: 10.1161/01.CIR.101.23.e215”.
- [15] GOLSCHLAGER, N., AND GOLMAN, M. *Principles of Electrocardiography*.
- [16] G.Y.H. LIP, P. K., AND WATSON, T. Atrial fibrillation, the growing epidemic. *Heart v.93(5) PMC1955544* (2007).
- [17] HALICI, U. *Artificial Neural Networks*.
- [18] HUSZAR, R. *Basic dysrhythmias, interpretation and management*.
- [19] INSTRUMENTS, T. Ads1198 low-power, 8-channel, 16-bit analog front-end datasheet.
- [20] JOHANSSON, M. The hilbert transform. *Master Thesis*.
- [21] KAISER, J. Some useful properties of teager’s energy operators. *Proc.IEEE ICASSP* (1993).
- [22] KANNEL, W., AND BENJAMIN, E. Final draft status of the epidemiology of atrial fibrillation. *Med. Clin. North Am.* (2008).

- [23] KODANI, E., AND ATARASHI, H. Prevalence of atrial fibrillation in asia and the world. *Journal of Arrhythmia* (2012).
- [24] KVEDALEN. Signal processing using the teager energy operator and other nonlinear operators. *Thesis* (2003).
- [25] LIAO, J., KHALID, Z., SCALLAN, C., MORILLO, C., AND O'DONNELL, M. Noninvasive cardiac monitoring for detecting paroxysmal atrial fibrillation or flutter. *Stroke, American Heart Foundation*. (2007).
- [26] LORENZO-GINORI, J. An approach to the 2d hilbert transform for image processing applications.
- [27] MIDMARK. Midmark iqholter monitor. <http://www.midmark.com>. Accessed 13 January, 2014.
- [28] MORTARA. Motara surveyor central. <http://www.mortara.com/products/healthcare/patient-monitoring/surveyor-telemetry/>. Accessed 10 February, 2014.
- [29] PHILIPS. Philips digitrak xt holter monitor. <http://www.healthcare.philips.com>. Accessed 13 January, 2014.
- [30] PHILIPS. Philips dxl 16-lead ecg algorithm. <http://www.healthcare.philips.com>. Accessed 10 January, 2014.
- [31] POLIKAR. The wavelet tutorial.
- [32] POLIKAR, R. The story of wavelets.
- [33] REINELT, P., KARTH, G., GEPPERT, A., AND HEINZ, G. Incidence and type of cardiac arrhythmias in critically ill patients: a single center experience in a medical-cardiological icu. *Intensive Care Med.* (2001) 27:1466-1473 (2001).

- [34] ROMERO, V., ALONSO, D., AND KHALED, N. Ecg baseline wander removal and noise suppression analysis in an embedded platform.
- [35] SCHILLER. Schiller ar12plus. <http://www.schiller-usa.com>. Accessed 13 January, 2014.
- [36] SCHOLKOPF, B., AND SMOLA, A. *Learning with Kernels*. MIT Press, 2002.
- [37] SELESNICK. Wavelet transforms a quick study. *Physics Today* (2007).
- [38] SHUKLA, K., AND TIWARI, A. Efficient algorithms for discrete wavelet transform. *SpringerBriefs in Computer Science* (2003).
- [39] TADEJKO, P., AND RAKOWSKI, W. Mathematical morphology based ecg feature extraction for the purpose of heartbeat classification.
- [40] TOMAR, V., AND PATIL, H. On the development of variable length teager energy operator (vteo). *Interspeech* (2008).
- [41] VALENS. A really friendly guide to wavelets. Tech. rep., 1999.
- [42] VETTERLI. Wavelets and filters banks: Theory and design.
- [43] VIDAKOVIC, M. Wavelets for kids. Tech. rep., Duke University, 1991.
- [44] VOLDER, J. E. The CORDIC trigonometric computing technique. *IRE Trans. Electron. Comput* (1959).
- [45] WIKIPEDIA. Cordic. <http://en.wikipedia.org/wiki/CORDIC>.
- [46] WOLF, P., ABBOTT, R., AND KANNEL, W. Atrial fibrillation as an independent risk factor of stroke: the Framingham study. *Stroke* (1991), 983–988.
- [47] ZHANG, F., AND LIAN, Y. Qrs detection based on multiscale mathematical morphology for wearable ecg devices in body area networks. *IEEE Transactions on Biomedical Circuits and Systems* (2009), 220–228.