Section Signal Processing Systems
Mekelweg 4,
2628 CD Delft
The Netherlands
https://sps.ewi.tudelft.nl/

SPS-2024-4567862

M.Sc. Thesis

Using our tools backwards, AF detection by confusing time and frequency

Michael Kraaijeveld

Abstract

Atrial Fibrillation or AF is the most common heart rhythm anomaly affecting millions of people. This work explores the possibilities of reinterpreting speech processing techniques for use in atrial fibrillation detection. An existing method of modelling single heartbeat, single lead ECG signals by means of an ARMA model's amplitude response as a time domain signal is implemented. The parameters of the models are then used for AF detection by means of detecting P wave absence. For this detection, the distribution of the P wave associated parameters is compared to a GMM model of normal sinus rhythm beats obtained from a large number of recordings from different sources.

Using our tools backwards, AF detection by confusing time and frequency

Thesis

submitted in partial fulfillment of the
requirements for the degree of

Master of Science

in

Electrical Engineering

by

Michael Kraaijeveld
born in Hardinxveld-Giessendam, The Netherlands

This work was performed in:

Section Signal Processing Systems
Department of Microelectronics
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology

Delft University of Technology

Delft University of Technology
Department of
Microelectronics

The undersigned hereby certify that they have read and recommend to the Faculty of Electrical Engineering, Mathematics and Computer Science for acceptance a thesis entitled **"Using our tools backwards, AF detection by confusing time and frequency"** by **Michael Kraaijeveld** in partial fulfillment of the requirements for the degree of **Master of Science**.

Dated: 9t̂h January 2024

Chairman: _____

dr.ir. R.C. Hendriks

Advisor: _____

dr.ir. J.C. Varon Perez

Committee Members: _____

dr.ir. J.A. Martinez Castaneda

_____

Abstract

Atrial Fibrillation or AF is the most common heart rhythm anomaly affecting millions of people. This work explores the possibilities of reinterpreting speech processing techniques for use in atrial fibrillation detection. An existing method of modelling single heartbeat, single lead ECG signals by means of an ARMA model's amplitude response as a time domain signal is implemented. The parameters of the models are then used for AF detection by means of detecting P wave absence. For this detection, the distribution of the P wave associated parameters is compared to a GMM model of normal sinus rhythm beats obtained from a large number of recordings from different sources.

Acknowledgments

This master thesis required of me to grow as a person, I am very grateful this opportunity was given to me and for the support I have received. From my beloved partner Bart, who put me back on my feet every time I would not. From my parents, knowing me better than I do. From my sisters, by being awesome. From all of my DWH friends, by giving me a sense of belonging and perspective. From my supervisors dr.ir. R.C. Hendriks and dr.ir. J.C. Varon Perez their assistance, patience and enthusiasm during the writing of this thesis was vital.

Michael Kraaijeveld
Delft, The Netherlands
9t̂h January 2024

Contents

List of Figures

List of Tables

1 Introduction

The heart is the single most important muscle found in almost all animals. Its importance has been known about since ancient times as illustrated by how much symbolical meaning it has come to hold in many cultures. One could see this meaning originating in the way we can sense it doing its vital work each and every moment. Even though we have no way of directly controlling its behaviour. Like any living tissue the heart is prone to diseases of various kinds. These are a leading cause of death in developed countries. Some are lethal almost instantly while others may linger for a long time before being noticed. Problems may arise in any part of the organ but of exceptional danger are infarcts and arrhythmia or irregular rhythm. The first is caused by a blocked artery but the arrhythmia are diverse and their causes are not always clear. The most common of these is atrial fibrillation or AF. It can go relatively unnoticed for long periods of time but early diagnosis improves outcomes by a large margin. Luckily it is possible to detect it from outside of the body.

At end of the 19^{th} century it was discovered that the heartbeat could not only be sensed by acoustical means. The movement of the charges inside the cells produced electromagnetic fields that could be measured. By recording these measurements over time by plotting the detected voltage on paper, the regular beating of the heart could be made visible. As this provides a way of closely monitoring the hearts function in a noninvasive way, the electrocardiogram (ECG) was soon recognised as a useful tool for diagnosis.

The usual method of analysis for ECG signals involves a trained human eye looking for features known to be associated with heart disease. A process requiring the time of an expert is always costly, not available everywhere and prone to human error. Automating diagnosis would thus be beneficial especially considering the existence of mobile devices capable of recording ECG signals.

Automated detection requires a model of the ECG signal that can reduce the recorded signal to features usefull for detection of diseases. The finding and extraction of which could also lead to a better understanding of the inner workings of the heart. These features should be informative for both the normal functioning of the heart as well as diseases. The most common abnormality in heart rhythm or arrhythmia is Atrial Fibrillation or AF. The effect of AF on ECG recordings is well described and electrocardiography is the main form of diagnosis for this condition. This diagnosis is mostly done by doctors.

In this work a method is proposed for detecting Atrial Fibrillation from the morphology of the ECG signal using techniques inspired by speech processing. Chapter 2 provides an overview of the physiology of the heart and of typical ECG recordings. Chapter

3 considers methods for modelling ECG signals. In Chapter 4 the chosen model is explored and implemented and a detector based on this model is proposed in Chapter 5. The performance of this detector is assessed in Chapter 6. Discussion and suggested future work is presented in Chapter 7 and conclusions are drawn in Chapter 8.

2 The Heart

The heart is on organ built almost entirely out of muscle. Its purpose is to pump blood through the body. Blood carries nutrients, energy, oxygen and almost everything else the other organs in the body need to function. The pumping action of the heart is achieved by contracting chambers, hollow spaces within the organ surrounded by muscle tissue. There are four chambers in total having one of two functions, both functions are thus represented twice within the organ. The first of these are the atria. Here, blood is collected from the main veins. When a heartbeat is triggered, the atria contract to pump blood into the next chambers, the ventricles. After the atria have provided the blood to be pumped, the ventricles contract to push it into the arteries. The ventricles are surrounded by valves to prevent the blood from flowing back into the atria or out of the arteries. Since the bloodstream is a closed system this last contraction also pushes blood into the atria so the heart is ready for the next cycle. The make up of the organ and this cycle is illustrated in Figure 2.1 This cycle is repeated at least two to four billion times from the formation of the organ until death if all goes well. Assuming average resting heart rate [1] and average global life expectancy in 2019 (73.3 years).

The muscles in the heart are activated by electrical impulses. These impulses largely originate from the nervous system present on the organ itself. The brain does not need

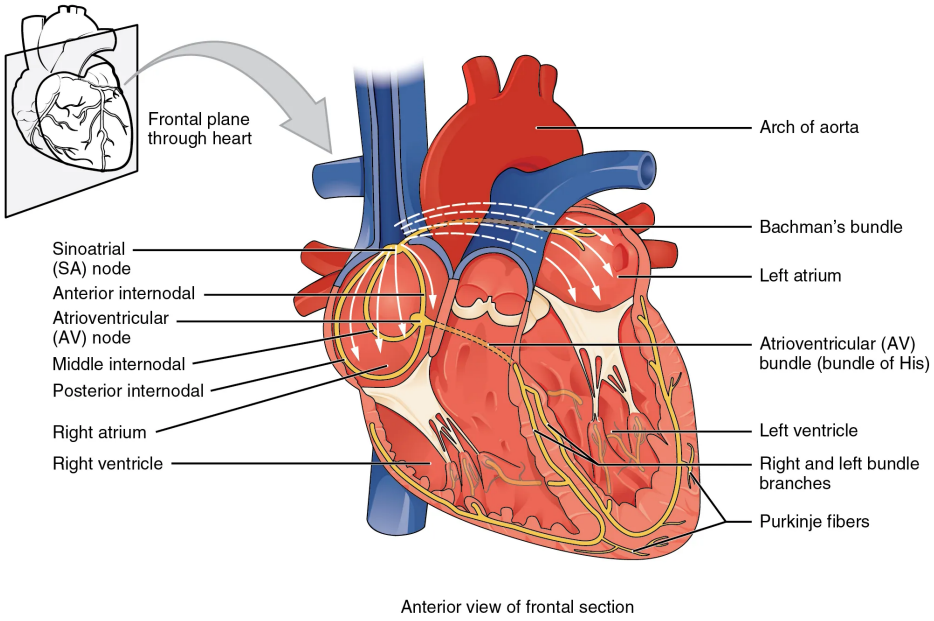

Figure 2.1: Anatomy of the human heart [2]

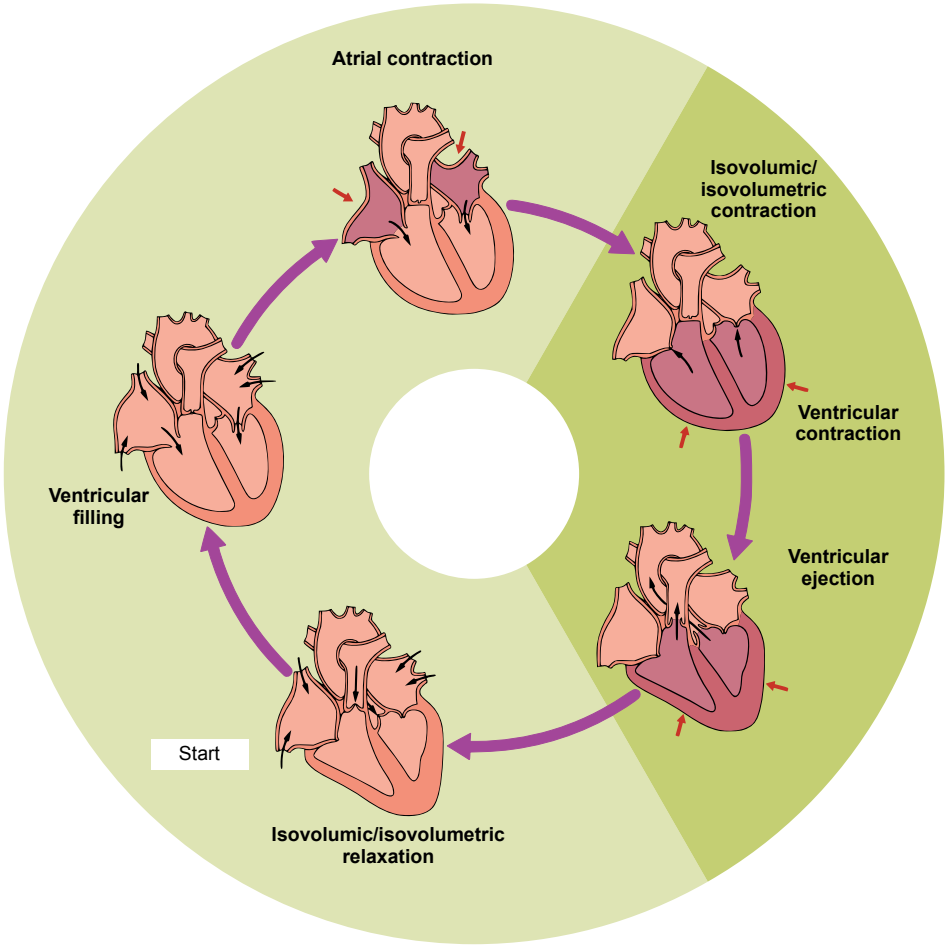

Figure 2.2: Phases of the cardiac cycle, adapted from [2]

to send any information to keep the heart pumping. The origin of these impulses is the so-called SA node, also known as the internal pacemaker. When it fires, signals travel via the atria, which then contract, to the AV node. The AV node acts as a delay for the signals between the atria and the ventricles. After the delay, the AV node fires and the signals are further distributed to the ventricles, causing them to contract. This is the end of the cycle which is restarted with the firing of the SA node, it does this by itself. The heart's nerve system is also illustrated in Figure 2.1.

As it is the main transportation mechanism for energy, the demand for blood is not constant and has to react fast to changes in the use of energy. These changes can be direct and voluntary such as when we start moving, anticipatory such as in an adrenaline rush, or by the internal state of the body. The regulation of heart rate comes from the autonomous nervous system. It influences the rate of firing of the SA node via neurotransmitters

2.1 Arrhythmia

The workings described above are the normal situation when all is working properly. Also known as the Normal Sinus Rhythm or NSR. In case of disease however, these events can be severely disturbed. The most common of these is Atrial Fibrillation (AF) which will also be the main focus of this work.

In Normal Sinus Rhythm, the signals travel over the atria and ventricles as a smooth wavefront. This is not a wave in the electromagnetic sense, the muscle tissue is not a simple conductive surface, it is made up of cells. The cells do also not just conduct the signal sent through them. Instead they respond by sending out a signal themselves. Each cell can thus be seen as a small amplifier reinforcing the signal. If all cells are healthy this constitutes a uniform wavefront resulting in proper contraction and in case of the atria giving a clear signal to the AV node.

If some of the tissue is not healthy however, the wavefront can be disturbed. Under conditions resulting in a chaotic partial contraction or fibrillation. When this happens in the ventricles where the main pumping action occurs the flow of blood is greatly affected and will lead to loss of consciousness and death. When fibrillation occurs in the atria, Atrial Fibrillation or AF occurs. In this case the atria contract chaotically and less blood is pumped to the ventricles. Another effect of chaotic conduction is the lack of a clear signal to the AV node. Instead of a clear front of signals arriving, it now receives a lot of chaotic smaller ones. This causes it to fire into the ventricles at random. Causing the contraction of them to occur normally but at an increased and irregular rate.

2.2 ECG measurements

A nerve cell relaying a signal depolarises, meaning the potential over its cell wall becomes less negative by letting positively charged ions enter. After this has happened, these ions are slowly pumped out again to repolarise the cell. If this happens for many cells at the same time, this can be sensed from outside of the body by measuring potentials. This is true for all muscles, but when staying still the heartbeat is the main signal. The standardised method for monitoring the heart this way is via an electrocardiogram (ECG). As this was invented before computers, the original recording method is on paper being moved under an actuated needle, nowadays this is done digitally.

2.2.1 Typical ECG signals

Electrocardiograms are a useful diagnostic tool as they separately show atrial and ventricular activity. Under healthy conditions they are separated in time and have a distinct morphology as the ventricles are larger and contract more forcefully. The timings and shapes of parts of the ECG can tell a lot about the heart's functioning while the measurement is not invasive.

Under Normal Sinus Rhythm, an ECG signal typically looks like Figure 2.3. We see the different stages in the heart's cycle pass by one by one. The atrium contracting

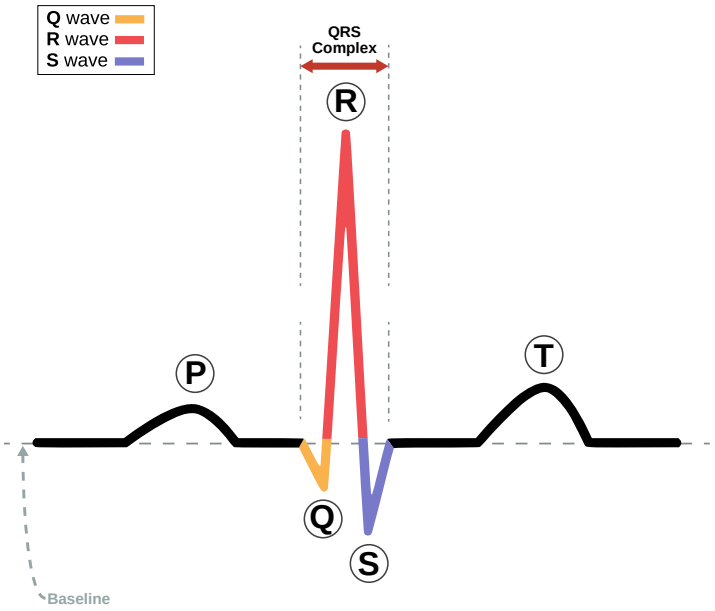

Figure 2.3: A schematic normal heartbeat ECG signal in Lead I

generates the P wave, it is caused by the depolarisation of the nerve cells in the atrium. The repolarisation of these cells is generally invisble as it overlaps with the next stage, the contraction of the ventricles. This main pumping action generates the QRS complex associated with the cells in the ventricles contracting. It changes sign multiple types due to the path the signal takes over the heart. When the contraction is done the ventricles repolarise producing the T wave.

In case of Atrial Fibrillation, the typical ECG changes considerably, as there is no coherent contraction of the atria anymore the P wave disappears. These can be replaced by f(ibrillatory) waves which are very small in amplitude [3]. As the AV is triggered randomly, QRS complexes appear without regular rhythm. This is illustrated in 2.4. If AF is to be detected, both of these changes (absence of P waves and irregular rhythm) can be utilised.

2.2.2 The twelve leads

The typical ECG parts as described above can differ in the height and sign depending on where on the body it is measured. As the potential is measured between different sets of points, the relative direction of the moving charges to those points changes. In standard ECG this is utilised to view the heart from different angles. The standardised connection is depicted in Figure 2.5. Lead I, focused on in this work, can be seen in the top left corner of Figure 2.5a.

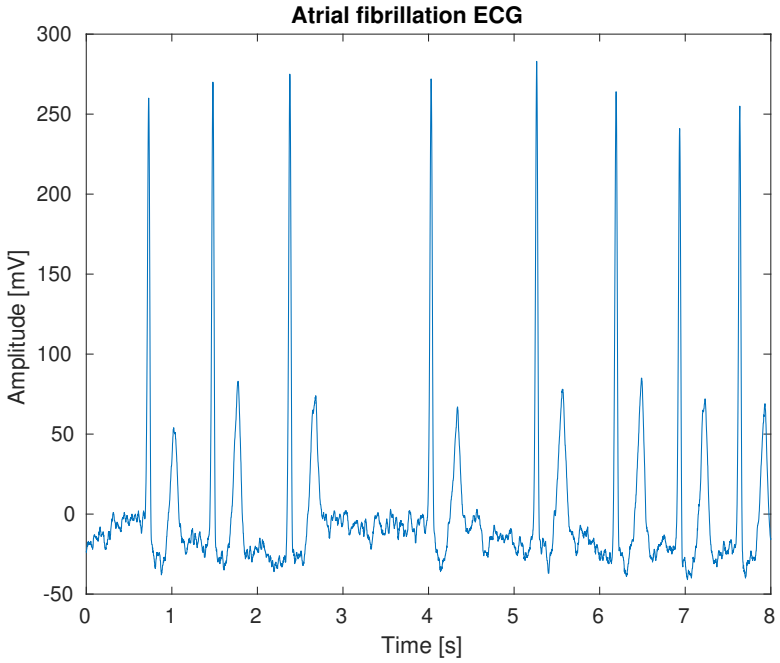

Figure 2.4: Example ECG recording of Atrial Fibrillation [4]

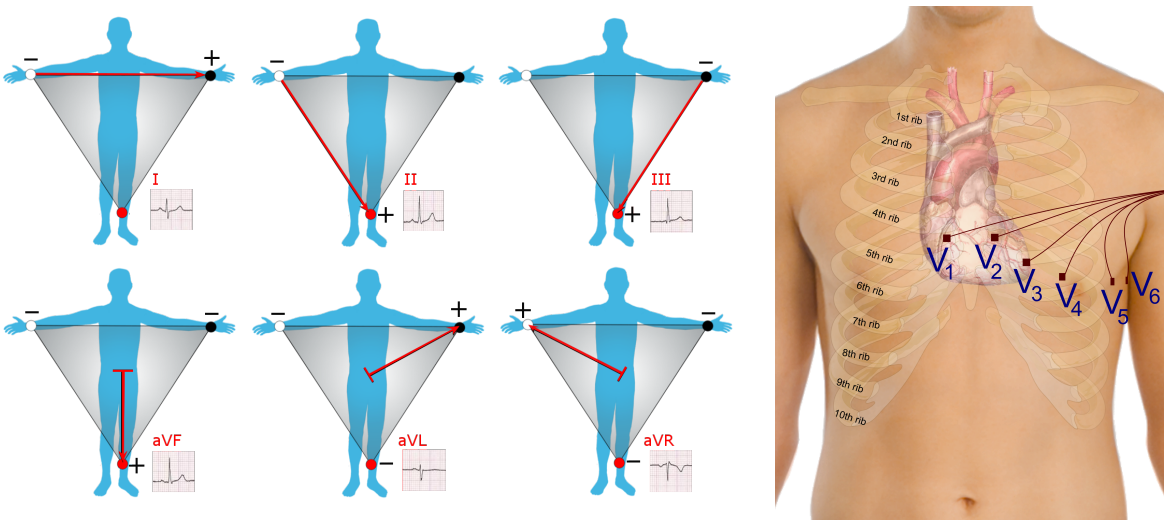

(a) Limb leads [5]

(b) Precordial leads

Figure 2.5: Illustration of the twelve standard ECG lead connections

2.3 Research goal

This work will focus on using an interpretable model for ECG signals based on common techniques as used in speech processing for AF detection. The parameters of the model should have a simple link to the modelled signal or heart physiology. The focus will be on the morphology of the signal as methods based on the rhythmic irregularity

of ventricular beats have already proven performant while detecting the absence of P waves proves more difficult [6]. The intended application is automated detection in a non-clinical environment. Therefore, the work will utilise one defined lead of ECG as full multi-lead recordings are unpractical for this purpose. In the next chapter, possible methods for modelling ECG signals are explored.

3 Modelling of ECG signals

Electrocardiography is an old technique and also a standard procedure executed on many patients. As a result, much work has been done in extracting useful information from it. Classically, this is focused on providing doctors with normal ranges of signals and signs of diseases to look out for [1]. ECG analysis used to be done on paper in the early days but as with all recording methods a lot of work has been done with digital signal processing for a wide variety of goals. Some examples include compression for storage [7], the creation of simulated signals [8], extracting physiological information not directly related to the heart [9], and actual diagnosis algorithms [6]. A lot of these methods rely on modelling the signal in some way. Either by simulating heart activity or only by directly replicating the signal. Some methods use a combination of factors or features to extract information, especially the ones relying on machine learning. These features range from the aforementioned timings and amplitudes also used by doctors, statistical features [10], morphological features [11] and features based on wavelet analysis [12]. The scope of this work will not include neural networks as the stated goal is to have an interpretable model of lower complexity. The features and classification techniques used in machine learning methods could however be used as inspiration. Before the features are extracted some preprocessing is typically done to remove noise or disturbances related to other muscles moving. This varies from using simple filtering to remove global regions of the frequency spectrum associated with unwanted signal to wavelet techniques [13].

As consecutive heartbeats look alike, it would make sense to treat them as a quasi-periodic signal to some degree. We can however not assume full periodicity since the heart rate is variable, moreover this variance exists from the smallest possible timescale of one beat to another as it is modulated by breathing [14]. An ECG model would thus require separating the morphology of the beat from the exact timing. In the field of speech modelling much success has been achieved by AR modelling [15][16] which can do just this by decorrelating the speech signal. What is left is a decomposition of the speech signal into an error signal and a filter modelling the spectral shape of the speech fragment. If a vowel was spoken the error signal resembles a pulse train corresponding to the pitch of the voice. Since the ECG signal would seem to resemble the voiced pattern, it might be worth exploring what the same techniques could achieve. Figure 3.1 illustrates the similarity.

If we feed a short ECG recording into such a model however, it turns out that the parts that make up the ECG signal, the P-, QRS-, T- and/or U waves are not modelled very well by assuming a single initiating event, even for very high model orders. This is illustrated by Figure 3.2. The speech residual shows very clear peaks of which one occurs per period whereas the ECG version shows multiple per heartbeat which seem

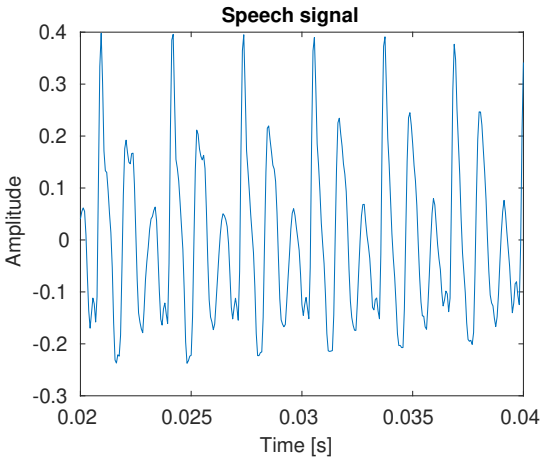

(a) Example voiced speech signal

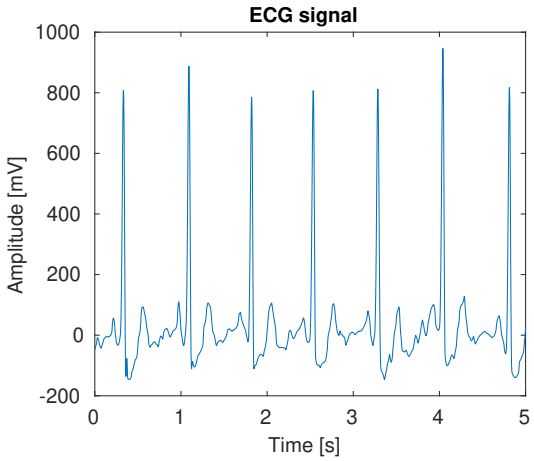

(b) Example ECG signal

Figure 3.1: Comparison of signals

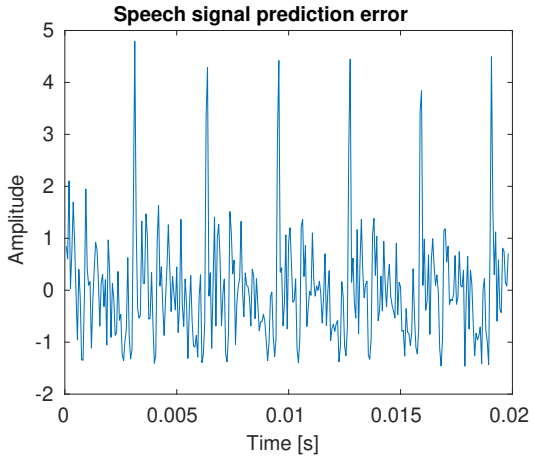

(a) Voiced speech prediction error

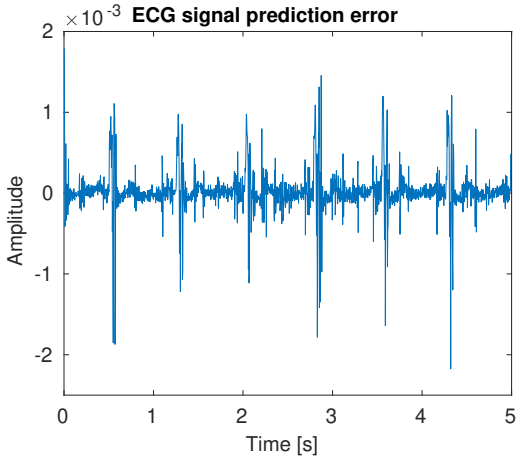

(b) ECG prediction error

Figure 3.2: Comparison of prediction errors for the signals in Figure 3.1

to correspond to the wave components.

We would thus need to model these wave components separately, which would mean they need to be extracted and or subtracted from the original signal. This presents an interesting problem in itself as the signal that is recorded by an ECG system is a superposition of multiple events as explained in the previous chapter. We are seeing moving charges in two muscle groups. Therefor delimiting a wave and setting it to zero or interpolating between its beginning and end would not be a good basis to model reality. A somewhat more plausible extraction might be achieved by using a template morphology for the wave components as we have a concept what these model and what they look like. An interesting solution is found in [17], where the modelling of the partial ECG waves is actually done by a modified form of ARMA modelling. There are more performant modelling methods available in terms of error metrics, for example

[8] which uses a mixture of Gaussian functions. However this requires more complex optimisation methods and relies on already delineated component waves. Furthermore a Gaussian is a strictly positive function while the ECG wave components are not. Since the method discussed in [17] already provides a full pole-zero based model of ECG signals and has not been explored for AF detection this will be done in this work. In the following chapter its workings are explained thoroughly as the method is unusual in its application of ARMA modelling, after this explanation its output parameters will be investigated with respect to their usefullness in AF detection.

4 The Murthy Model

In this chapter an existing modelling method for a single heartbeat, single lead ECG recording with regards to its use as an input for an atrial fibrillation detector is explored. The ECG model found by Murthy [17] proposes to use the impulse response of a second order linear filter as a basis for modelling the ECG component waves (P, QRS, T). The intuition behind this being that the ECG signal morphology resembles the frequency response of such filters. We are thus going to approach the time domain ECG signal as if it was the amplitude frequency response of some linear filter.

The modelling of a single component wave goes as follows. Suppose, for example, we start with a recording of a single P wave denoted by $x(k)$ of length N, we do not view this as a time domain signal but as the Discrete Cosine Transform (DCT) of some impulse response $x(n)$. The DCT is used instead of a Fourier transform as the ECG recording has no imaginary part. The impulse response $x(n)$ is thus given by the inverse DCT in Equation 4.1 with N the length of the recording and δ_{k1} the Kronecker delta.

$$x(n) = \sqrt{\frac{2}{N}} \sum_{n=1}^{N} x(k) \frac{1}{\sqrt{1+\delta_{k1}}} \cos\left(\frac{\pi}{2N}(2k-1)(n-1)\right) \tag{4.1}$$

The P wave is then modelled with a linear second order filter $H(z)$ defined by

$$H(z) = \frac{G(1-b_1 z^{-1})(1-b_2 z^{-1})}{1 - 2r\cos\Theta z^{-1} + r^2 z^{-2}} \tag{4.2}$$

with impulse response $h(n)$ and where we then want to find G,b_1,b_2,Θ and r, such that

$$\sum_{n=0}^{N-1} |x(n) - h(n)|^2$$

is minimised. The DCT of $h(n)$ then replicates the original P wave recording $x(k)$ while being specified by only a few parameters.

The basic building block of the model is thus a second order filter with z-transform (4.2) with Θ and r the angle and magnitude of the conjugate pole pair. Using this building block or fractional component, the common shapes of the ECG wave components can be modelled. The QRS complex generally requires two of these sections as its morphology is more complex. As a full single heartbeat recording consists of multiple component waves (eg. P wave, QRS complex), we would thus need multiple

of these sections added together. These then from a general ARMA(N,N) model as in Equation 4.3. The original authors do not force the model for a full heartbeat to be organised in second order sections, hence, they use.

$$H(z) = \frac{b_0 + b_1 z^{-1} + \cdots + b_N z^{-N}}{a_0 + a_1 z^{-1} + \cdots + a_N z^{-N}} \tag{4.3}$$

To find the correct a and b coefficients, common tools for ARMA modelling can be used. The standard approach when applied to ordinary time domain signals for such a model would be to use a correlation function to optimise a prediction error. In this case however there is no prediction happening. All that is needed is for the model to match the transfer function as closely as possible. The method used by the original authors is the Steiglitz-McBride method [18]. This works by iteratively solving Equation 4.4, where $h(n)$ is the modelled transfer (the impulse response of Equation 4.5) and $x(n)$ the IDCT of the input ECG signal.

$$\min_{a,b} \sum_{i=0}^{N} |x(n) - h(n)|^2 \tag{4.4}$$

This was verified to indeed lead to good results compared to a linear prediction approach. The algorithm is fed the inverse DCT of a single heartbeat ECG signal, a pseudo-impulse response to match with the model. The full heartbeat is thus processed in one step. As the representation also includes a gain term the model is not sensitive to inter-beat changes in amplitude. This is beneficial as signal amplitude varies between and within recordings depending on hardware and patient parameters.

4.1 Implementation

The ECG data used in this work originates from the PhysioNet's [19] 2017 [20] and 2020 [4] Computing in Cardiology Challenges. These data sets combined provide a total of 51629 recordings from different sources. The 2017 challenge consists of single lead (Lead I) recordings and the 2020 part consists of full 12-lead recordings of which Lead I will be examined for consistency. Recording times differ between 6 and 60 seconds. Sampling rates differ from 300 to 1000 Hz. The St. Petersburg [21] recordings are not included as these are 30 minutes long, this would make the definition of AF used inconsistent as this is labelled per recording and thus could contain both normal and arrhythmic beats. This omits 75 recordings. The remaining data set is used in full and thus contains a large amount of patients and multiple recording methods. All recordings are high pass filtered at 0.5 Hz to dampen baseline wander over inter-beat timescales.

As the model is intended for use on single heartbeats the recordings have to be split up into single beat windows. To do this, the QRS complex is chosen as the timing reference as it is the best detectable part in a given heartbeat for both NSR and AF. The QRS

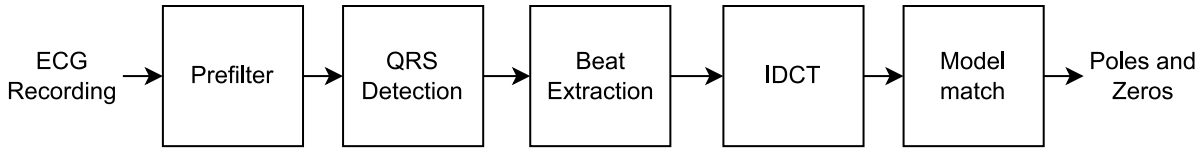

Figure 4.1: Flow diagram of the modelling process

detection is done using a wavelet based approach found in [22] as implemented in ECG-kit [23]. The method works by detecting signatures as produced by a QRS complex over several scales of a wavelet transformation.

When the QRS complex has been located, the window for that heartbeat is determined by fixed time offsets wide enough to encapsulate most normal heartbeats plus some slack, -300 ms and +420 ms relative to the detected QRS peak [1]. No window function is applied as the signal in the pseudo time domain is not of interest. By having all windows correspond to the same amount of time the heartbeats and obtained models can be easily compared later. For high rates this does introduce some overlap. As the window represents a fixed amount of time the input sampling frequency does not affect the placement of the poles. The single beats are then transformed by the IDCT and the model is fit. The main output of interest are the coefficients of the model viewed as poles and zeros. The full process is illustrated in Figure 4.1

As stated in the previous section, the main shapes of the ECG would suggest the usage of 4 second-order filter sections, amounting to 8 poles and zeros needed. In addition, two are supplemented by the original authors to also capture the baseline, which is not generally straight and at zero, variations in the QRS complex, and artefacts. The full model thus becomes of the form in Equation 4.5.

$$H(z) = \frac{b_0 + b_1 z^{-1} + \cdots + b_{10} z^{-10}}{a_0 + a_1 z^{-1} + \cdots + a_{10} z^{-10}} \tag{4.5}$$

The components modelling these artefacts and baseline are clearly distinct from the components modelling the ECG wave components negating the need for complex pre-processing. Figure 4.2 shows an example heartbeat modelled in this way, the locations of the poles and zeros are also displayed. It is clear that the full model indeed converges on multiple second order filters as the poles appear in conjugate pairs. The angles of these poles with the real axis correspond to the time axis of the output. This means they can be matched with their associated ECG wave component. The second order sections can be extracted again from the full model by performing a partial fraction decomposition on the full model and then using only one pole pair to generate an output. Figure 4.3 shows the DCT transforms of only the second order section transfer functions associated to the partial wave components and the difference signal with the original input recording. That is, after the model for the full recording is fit. It is split up into the individual poles and zeros by partial fraction decomposition. The conjugate pole pairs associated with the P, QRS or T component as highlighted in Figure 4.2b are then added together and the DCT of only these poles and zeros is plotted. For reference the output is also subtracted from the original ECG recording to show these filters indeed model the partial waves.

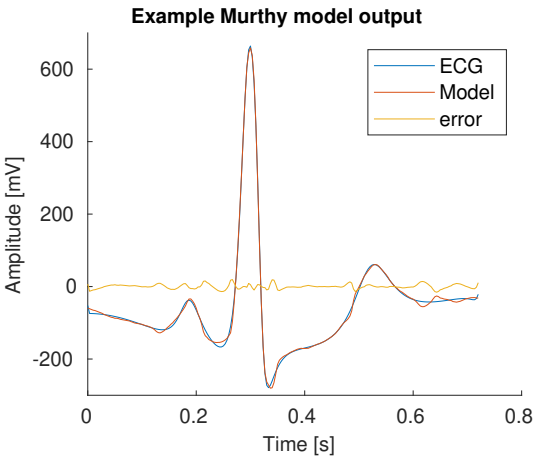

(a) Example full model output

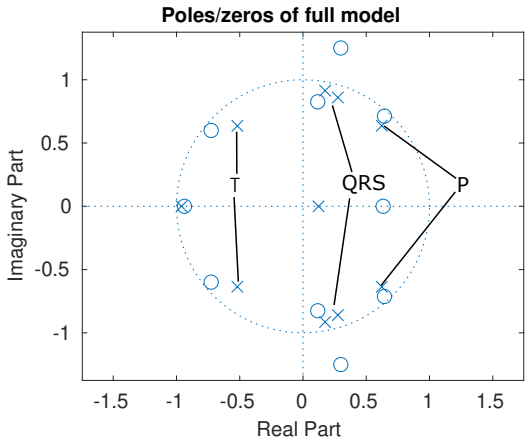

(b) Pole/zero plot, with x's denoting poles and o's denoting zeros

Figure 4.2: An example normal rhythm heartbeat model from a Lead I recording and its associated poles and zeros

As stated, the pole locations are directly interpretable. The pole angle with the real axis corresponds to the location of a peak in the signal and the distance to the unit circle is indicative of the height. The zeros have the same interpretation for their angle and magnitude but for valleys. The pole locations are also constrained to the inside of the unit circle, this eases the building of a detector. Therefore the focus of the next chapters will be to build a detector based on the pole locations.

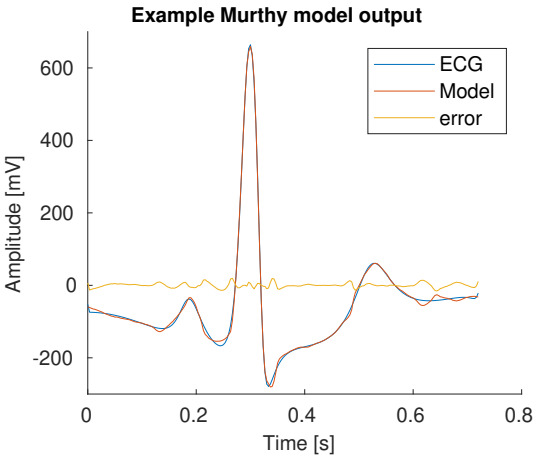

(a) Example full model output

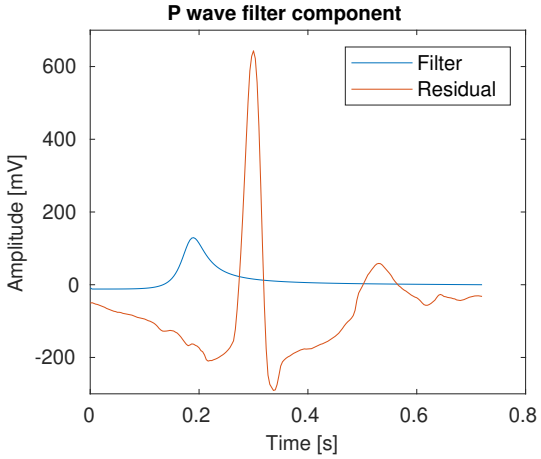

(b) P wave component

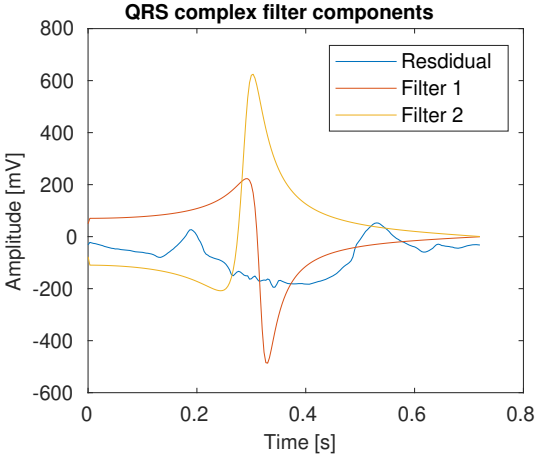

(c) QRS wave components

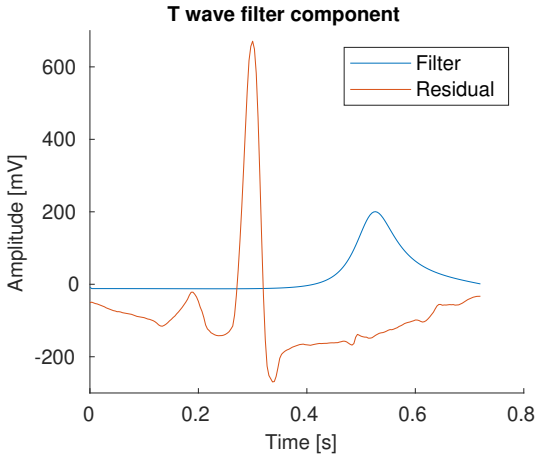

(d) T wave component

Figure 4.3: Example Murthy model output split into single second order filters

5 AF detection

In this chapter the use of the Murthy model for detecting Atrial Fibrillation is explored. First, the location of the output poles is compared over a large dataset. Then, a method is proposed to model the distribution of these poles after which a classifier is devised to perform atrial fibrillation detection.

5.1 Murthy model output for a large dataset

As we now have a model that can describe a single heartbeat with a few parameters that are closely linked to the component waves. It would make sense to start investigating how this model responds on average over multiple heartbeats, patients and datasets, and with the occurrence of arrhythmia. Since the model is set up to centre around the QRS complex and is agnostic to amplitude and sampling frequency. It can compare inter patient data from multiple sources.

Figures 5.1a and 5.1b show the differences in pole distributions between a normal sinus rhythm and atrial fibrillation recordings. In the NSR plot we can clearly see the P waves, QRS complex and T waves represented as three distinct clusters close to the unit circle. We are only viewing the positive imaginary axis as the poles of interest form conjugate pairs and are thus mirrored in the real axis. The QRS complex is the most concentrated as this is the feature the frames are centred on. Because of the variance in timing of the other wave components these pole distributions are more spread out but they are still visible. Not all poles are within the indicated areas as the number of poles estimated is higher than necessary for modelling only the wave components, they are mainly there to ensure good model fit in the presence of artefacts and because the QRS complex sometimes is better modelled with an extra set of poles, indicated by the small secondary cluster towards the origin from the main QRS pole location. The excess poles also cluster near ± 1 as the frame does not necessarily begins with a flat line centred at zero or the frames may overlap.

Looking at the pole distribution for the atrial fibrillation case we see a clear difference. As expected the region where the P wave cluster normally resides is now much less densely populated. This also means more poles are 'freed' as there is no P wave to model. The QRS complex associated double poles move closer together (this is more clearly visible in the histogram view of Figure 5.2a). The T wave region is more spread out as the variation in timing increases. Because of the fixed frame lengths, some frames might contain more than one of each wave component as AF leads to periods of rapid ventricular contraction.

Since the absence of P waves is the most apparent indicator for AF and the chaotic

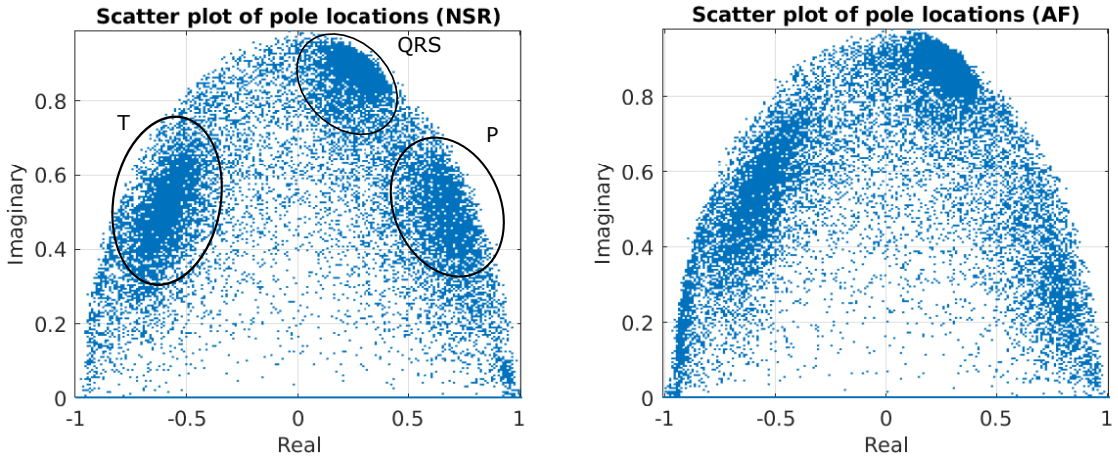

(a) Scatter plot of pole locations for normal heartbeats with component wave regions indicated

(b) Scatter plot of pole locations for Atrial Fibrillation heartbeats

Figure 5.1: Pole locations for the two different heart rhythms

rate can already be well detected by monitoring RR intervals, the vanishing P wave poles are a good candidate for building an AF detector based on the model described. In the following sections one such a detector is explored.

5.2 Pole distribution based AF detection

To utilise the information captured by the poles, a classifier has to be devised to decide whether a recording falls into the NSR or AF category. As the pole locations are described by complex numbers in which the magnitude represents peak size and the angle timing, these variances are correlated in Cartesian coordinates apart from possible correlations between size and timing themselves. Conceptually we would expect to be able to describe the pole distribution by a few locations corresponding to the mean component pole location, in addition to there being a variance around these points caused by noise, timing variations and variability in the wave components itself. If a distribution can be modelled on the data, a threshold could be set above which a pole is deemed to be in the P wave region. To get an idea about the kind of distribution that could be used here the pole locations displayed in Figure 5.1 are given again in Figure 5.2 as histograms.

The type of distribution depicted and described can be modelled using a bivariate Gaussian Mixture Model with N components all having a mixing proportion m_i, mean μ_i and covariance matrix Σ_i as defined in Equation 5.1 where $\mathbf{x}$ is the two-dimensional vector of the real and imaginary component of a pole location.

$$f(\mathbf{x}) = \sum_{i=1}^{N} m_i \frac{1}{2\pi\sqrt{\det(\Sigma_i)}} e^{\frac{-(\mathbf{x}-\mu_i)^T \Sigma_i^{-1} (\mathbf{x}-\mu_i)}{2}} \tag{5.1}$$

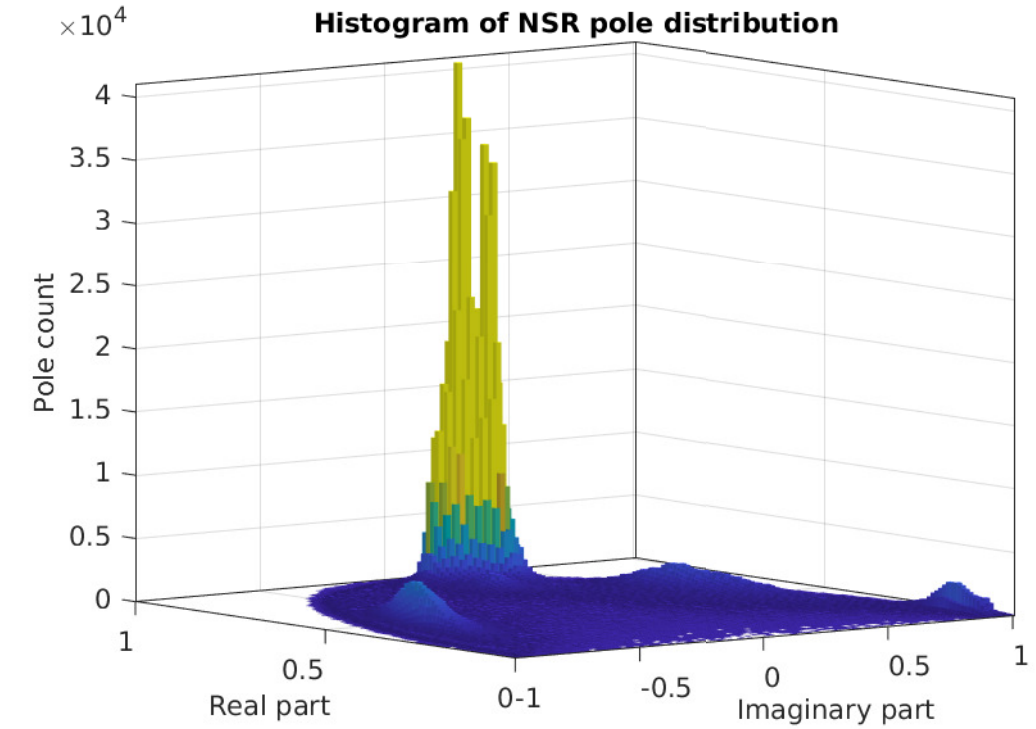

(a) Histogram of pole locations for Normal Sinus Rhythm heartbeats

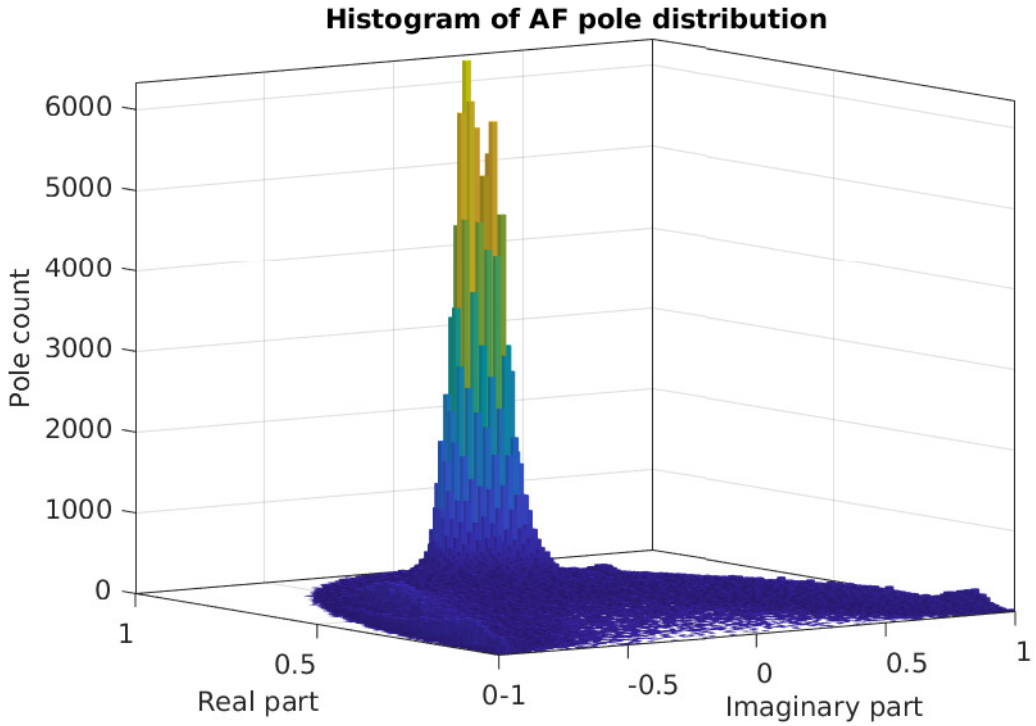

(b) Histogram of pole locations for Atrial Fibrillation heartbeats

Figure 5.2: Histogram view of pole distributions

The means, variances and mixing proportions are found using the Expectation-Maximisation algorithm with starting points determined by the kmeans++ algorithm in MATLAB. For the NSR case around 20 components are needed to capture the full complexity of the distribution. This is higher than one would expect if only modelling the component wave cluster. However the presence of the extra poles and the closely spaced double cluster for the QRS complex necessitates these extra components. This number may be reduced by not considering poles we do not find to be informative such as those closer to the origin or the real axis, the number of components however is not a problem for the intended use case as the distribution only has to be fit once. The GMM fit on the pole locations for a training set of NSR recordings can be seen in Figure 5.3.

The main difference between a normal heart rhythm and AF, when looking at pole locations, is the absence of P waves. Corresponding to an absence of poles in that region of the unit circle. Figure 5.2b largely confirms this. For a successive number of AF heartbeats it is unlikely that a majority of them will have poles in this region. The detector will thus be based on a model of the normal sinus rhythm pole distribution and will classify a recording as AF if there are not enough poles in the P wave region. The size of the region is determined by a threshold on the estimated Gaussian Mixture Model. A single pole is therefore classed as being in the P wave region if

$$f(\mathbf{x}) > t$$

The optimal value for the threshold t will be experimentally determined in the next chapter. The P wave region is constrained by a box around the lowest contour for which the P wave and QRS complex regions are disjoint, any other poles are not used, see Figure 5.4 for an illustration.

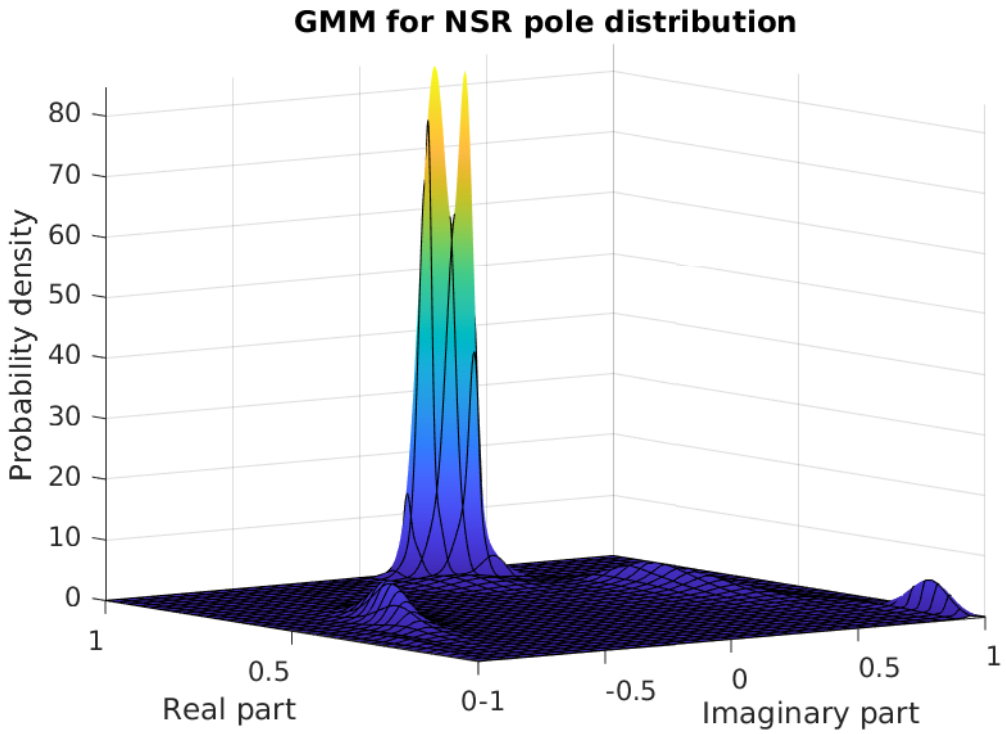

Figure 5.3: Gaussian Mixture Model of Normal Sinus Rhythm pole distribution

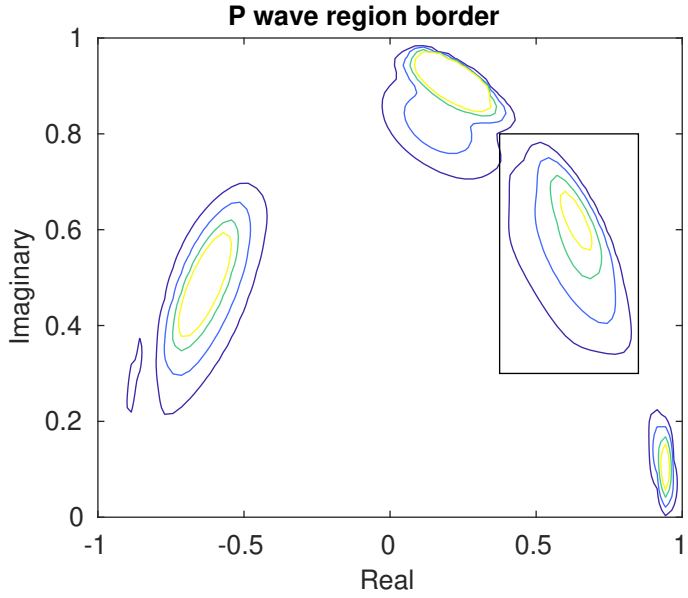

Figure 5.4: Contour plot of Gaussian Mixture Model over the training set of NSR beats

6 Results

For evaluating the performance of the detector constructed in the previous chapter, accuracy is measured in terms of the sensitivity and false alarm rate of the classifier such that both error types are made explicit. The influence of the classifier parameters is made visible by plotting both performance indicators for each combination. This gives a multidimensional receiver operating characteristic (ROC) showing the sensitivity of the classifier to these parameters, the performance and the optimal operating point in one figure. To make sure the results are valid for unseen data the datasets are split into a test and a training set containing all available AF records and an equal amount of unseen NSR recordings. The test set contains the full Georgia set of [4] not used in fitting the GMM.

When the detector is run over the test set. The ROCs in Figure 6.1a can be obtained by varying the proportion of beats with detected absent P waves and the threshold for the GMM model.

The curve for the GMM threshold where the ROC gets closest to the point of perfect performance (0 flase alarm rate, 100% sensitivity) is plotted again in Figure 6.1b to indicate an optimal curve for this detector.

Other methods utilising only the morphology of the signal are rare but do exist in [24] and [25]. In [24], the recordings used are not publicly available. In [25], the MIT-BIH

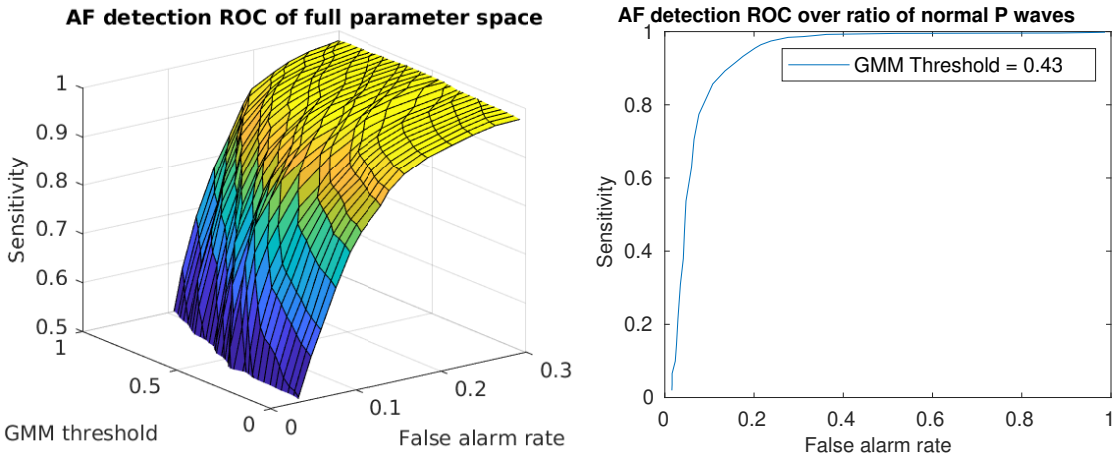

(a) ROC for the full parameter space

(b) Optimal ROC with respect to GMM threshold

Figure 6.1: Detector performance given as ROC curves over the GMM threshold and fraction of non AF P waves detected.

	Sensitivity	Specificity
Slocum et al. [24]	68.3%	87.8%
Ladavich et al [25]	88.9%	90.7%
Proposed	89.3%	86.5 %

Table 6.1: Comparison of P wave absence detection methods

database is used which contains long recordings of few patients, instead of the short recordings of many patients used in this work. As shown in Table 6.1, the performance of the proposed method is better than the former on sensitivity and similar to the latter. Worthy of note is that [25] employs a very similar method of detection by training a GMM on features extracted from the P wave. The difference is in the features which are based on average values of small time windows of the input signal. This approach is more limited than the Murthy model as it captures a downsampled version of the signal and may not map as well to other sections of the ECG with higher frequency content such as the QRS complex. The performance still being similar may be explained by the different means of classification. This performance is such that using only this detector no reliable classification can be performed for any real world application. It could however be used in conjunction with other methods such as RR interval variability detection to add certainty about the cause of that variability. Be it AF or a different arrhythmia that typically show irregular ventricular beats. The method is also light on computing requirements as the modelling relies only on a single DCT transform and a few rounds of Steiglitz-McBride. The QRS detection used is more intense as it is based on a wavelet transform but lighter methods exist.

7 Discussion

This thesis presents a methodology for detecting AF using a model of which the parameters are interpretable in terms of the ECG signal itself. It thus represents the signal as features comprehensible by both computers and humans. It is, however, not a model of the heart itself, the parameters extracted are not related to physical quantities. The definition of AF is not consistent because of the data set used, the presence of AF is labelled per recording but the recordings differ in length. The method is fairly robust as it can deal with different sampling frequencies and amplitudes. However, because of the nature of the ARMA model and the fitting algorithm used, the poles and zeros are not guaranteed to be globally optimal. It could be argued it would be better to only use an AR model as only poles are used in the classification. Long term changes in the signal baseline should also not cause problems due to the amplitude invariance. The achieved detection performance on AF is such that it would be advised to at least combine it with a method based on the variability of the heart rate if it were to be used in a medical device. It is not tested with respect to different arrhythmia and any arrythmia characterised by weaker P waves might also be classified as AF. Although the used features are more flexible than those employed in a similar work [25]. This does not immediately lead to a better performance in AF detection. This may also be explained by differences in the classification algorithm.

7.1 Future work

The presented method shows comparable performance to earlier attempts of detecting P wave absence. The modelling and parameter distribution extraction used in this work can be improved upon and expanded in a number of areas. The most obvious would be to apply it to different kinds of arrhythmia using not only the P wave associated poles but the entire distribution. Special care would have to be taken for arrhythmia where to QRS complex is severely affected as this is the timing reference. Including more pole regions may even help AF detection itself as the QRS complex associated poles move closer together and the T wave region smears out. Another suggestion would be to apply the method to multiple leads for robustness and performance which may also lead to new observations. This would then be helped by incorporating not only the poles but also the zeros as polarity of wave components may change over different leads. They are not constrained to the unit disk but the distribution of phase angles does change between NSR and AF as visible in Figure 7.1. One could also vary the number of model parameters used per heartbeat to reduce the excess amount for some situations. Finally it may be worthwhile to look into applying a similar detector as in [25] to find if the Murthy model indeed brings clear advantages. The GMM may also

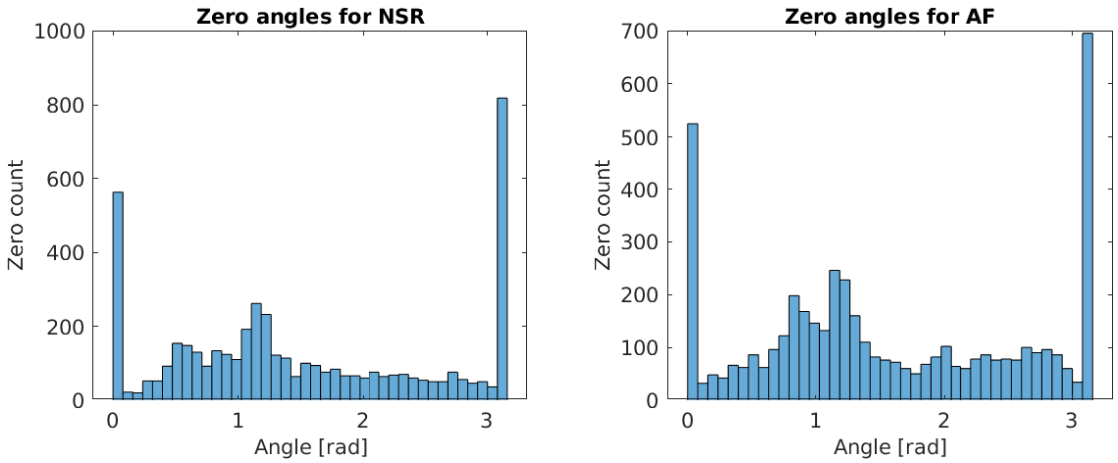

(a) Histogram of model zero angles in normal sinus rhythm recordings

(b) Histogram of model zero angles in atrial fibrillation recordings

Figure 7.1: Zero angles comparison of 400 heartbeats per category

be specified in polar coordinates as this would decouple timing and shape of a pole's influence, the remaining correlation between the two may also give usefull information.

8 Conclusion

The research goal of this thesis was to use an interpretable model of ECG signals based on common speech processing techniques for AF detection. The parameters of the model were specified to have a direct simple link to the signal or the hearts physiology. AF detection was specified to focus on the morphology of the signal. The method was to be used on single-lead recordings.

The model selected for the ECG was found in literature [17]. It was selected based on its use of ARMA models which is commonly used in speech processing. The use of the ARMA model is different from conventional linear prediction type systems as it models the ECG time signal as if it were a frequency response. This way all parts of the ECG signal can be efficiently modelled without the need for partial wave delineation. It models single heartbeats independently.

The AF detector was then built by observing the differences in pole distributions between normal and AF heartbeats. In these distributions the absence of the P wave in AF can be observed. A GMM model of the normal rhythm pole distribution was fit. If a recording does not have a sufficient number of poles in the P wave region specified by a threshold on the GMM model a recording is classified as AF. This provides a similar performance as earlier methods. Room for improvement is discussed in the previous chapter.

In conclusion it is shown that an AF detector can be built using these techniques and that the model found by Murthy is worth exploring further for different arrhythmia types. ECG signals however are of higher complexity than speech as the partial waves originate from different parts of the heart that are linked but do generate very distinct signals.

Bibliography

[1] J. W. Mason, D. J. Ramseth, D. O. Chanter, T. E. Moon, D. B. Goodman, and B. Mendzelevski, "Electrocardiographic reference ranges derived from 79,743 ambulatory subjects," *Journal of Electrocardiology*, vol. 40, no. 3, pp. 228–234, May 2007, ISSN: 0022-0736. DOI: 10.1016/J.JELECTROCARD.2006.09.003.

[2] J. G. Betts, Young. Kelly A., J. A. Wise, *et al.*, "The Cardiovascular System: The Heart," in *Anatomy and Physiology*, OpenStax, Apr. 2013, ch. 19. [Online]. Available: https://openstax.org/books/anatomy-and-physiology/pages/19-introduction.

[3] L. Sörnmo Editor, *Atrial Fibrillation from an Engineering Perspective* (Series in BioEngineering), L. Sörnmo, Ed. Cham: Springer International Publishing, 2018, ISBN: 978-3-319-68513-7. DOI: 10.1007/978-3-319-68515-1. [Online]. Available: http://link.springer.com/10.1007/978-3-319-68515-1.

[4] E. A. Perez Alday, A. Gu, A. J. Shah, *et al.*, "Classification of 12-lead ECGs: The PhysioNet/Computing in Cardiology Challenge 2020," *Physiological Measurement*, vol. 41, no. 12, Dec. 2020, ISSN: 13616579. DOI: 10.1088/1361-6579/ABC960. [Online]. Available: https://physionet.org/content/challenge-2020/1.0.2/.

[5] N. Patchett, *Limb leads of EKG*, Mar. 2015. [Online]. Available: https://commons.wikimedia.org/wiki/File:Limb_leads_of_EKG.png.

[6] A. Rizwan, A. Zoha, I. B. Mabrouk, *et al.*, "A Review on the State of the Art in Atrial Fibrillation Detection Enabled by Machine Learning," *IEEE Reviews in Biomedical Engineering*, vol. 14, pp. 219–239, 2021, ISSN: 1937-3333. DOI: 10.1109/RBME.2020.2976507. [Online]. Available: https://ieeexplore.ieee.org/document/9016113/.

[7] M. Sabarimalai Sur and S. Dandapat, "Wavelet-based electrocardiogram signal compression methods and their performances: A prospective review," *Biomedical Signal Processing and Control*, vol. 14, no. 1, pp. 73–107, Nov. 2014, ISSN: 1746-8094. DOI: 10.1016/J.BSPC.2014.07.002.

[8] M. A. Awal, S. S. Mostafa, M. Ahmad, *et al.*, "Design and Optimization of ECG Modeling for Generating Different Cardiac Dysrhythmias," *Sensors (Basel, Switzerland)*, vol. 21, no. 5, pp. 1–27, Mar. 2021, ISSN: 14248220. DOI: 10.3390/S21051638. [Online]. Available: /pmc/articles/PMC7956726/%20/pmc/articles/PMC7956726/?report=abstract%20https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7956726/.

[9] Z. Long, G. Liu, and X. Dai, "Extracting emotional features from ECG by using wavelet transform," *2010 International Conference on Biomedical Engineering and Computer Science, ICBECS 2010*, 2010. DOI: 10.1109/ICBECS.2010.5462441.

[10] S. Dilmac and M. Korurek, "ECG heart beat classification method based on modified ABC algorithm," *Applied Soft Computing*, vol. 36, pp. 641–655, Nov. 2015, ISSN: 1568-4946. DOI: 10.1016/J.ASOC.2015.07.010.

[11] T. Ince, S. Kiranyaz, and M. Gabbouj, "A generic and robust system for automated patient-specific classification of ECG signals," *IEEE Transactions on Biomedical Engineering*, vol. 56, no. 5, pp. 1415–1426, May 2009, ISSN: 00189294. DOI: 10.1109/TBME.2009.2013934.

[12] E. Castillo, D. P. Morales, G. Botella, A. García, L. Parrilla, and A. J. Palma, "Efficient wavelet-based ECG processing for single-lead FHR extraction," *Digital Signal Processing*, vol. 23, no. 6, pp. 1897–1909, Dec. 2013, ISSN: 1051-2004. DOI: 10.1016/J.DSP.2013.07.010.

[13] S. Kaplan Berkaya, A. K. Uysal, E. Sora Gunal, S. Ergin, S. Gunal, and M. B. Gulmezoglu, "A survey on ECG analysis," *Biomedical Signal Processing and Control*, vol. 43, pp. 216–235, May 2018, ISSN: 17468108. DOI: 10.1016/J.BSPC.2018.03.003.

[14] A. Kordes, "An Expanded IPFM Model for Heart Rhythm Analysis Detecting Atrial Fibrillation Using a Physiological Model," Ph.D. dissertation, 2021. [Online]. Available: http://repository.tudelft.nl/..

[15] M. R. Schroeder and B. S. Atal, "CODE-EXCITED LINEAR PREDICTION (CELP): HIGH-QUALITY SPEECH AT VERY LOW BIT RATES.," *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pp. 937–940, 1985, ISSN: 07367791. DOI: 10.1109/ICASSP.1985.1168147.

[16] A. S. Spanias, "Speech Coding: A Tutorial Review," *Proceedings of the IEEE*, vol. 82, no. 10, pp. 1541–1582, 1994, ISSN: 15582256. DOI: 10.1109/5.326413.

[17] I. S. Murthy and G. S. Prasad, "Analysis of ECG from Pole-Zero Models," *IEEE Transactions on Biomedical Engineering*, vol. 39, no. 7, pp. 741–751, 1992, ISSN: 15582531. DOI: 10.1109/10.142649.

[18] P. Stoica and T. Soderstrom, "The Steiglitz-McBride Identification Algorithm Revisited-Convergence Analysis and Accuracy Aspects," *IEEE Transactions on Automatic Control*, vol. 26, no. 3, pp. 712–717, 1981, ISSN: 15582523. DOI: 10.1109/TAC.1981.1102679.

[19] A. L. Goldberger, L. A. N. Amaral, L. Glass, *et al.*, "PhysioBank, PhysioToolkit, and PhysioNet," *Circulation*, vol. 101, no. 23, 2000, ISSN: 0009-7322. DOI: 10.1161/01.cir.101.23.e215.

[20] G. D. Clifford, C. Liu, B. Moody, *et al.*, "AF classification from a short single lead ECG recording: The PhysioNet/computing in cardiology challenge 2017," *Computing in Cardiology*, vol. 44, pp. 1–4, 2017, ISSN: 2325887X. DOI: 10.22489/CINC.2017.065-469.

[21] *St Petersburg INCART 12-lead Arrhythmia Database v1.0.0*. [Online]. Available: https://physionet.org/content/incartdb/1.0.0/.

[22] J. P. Martínez, R. Almeida, S. Olmos, A. P. Rocha, and P. Laguna, "A Wavelet-Based ECG Delineator Evaluation on Standard Databases," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 4, pp. 570–581, Apr. 2004, ISSN: 00189294. DOI: 10.1109/TBME.2003.821031.

[23] *ECG-Kit v1.0*. [Online]. Available: https://physionet.org/content/ecgkit/1.0/.

[24] J. Slocum, A. Sahakian, and S. Swiryn, "Diagnosis of atrial fibrillation from surface electrocardiograms based on computer-detected atrial activity," *Journal of Electrocardiology*, vol. 25, no. 1, pp. 1–8, Jan. 1992, ISSN: 0022-0736. DOI: 10.1016/0022-0736(92)90123-H.

[25] S. Ladavich and B. Ghoraani, "Rate-independent detection of atrial fibrillation by statistical modeling of atrial activity," *Biomedical Signal Processing and Control*, vol. 18, pp. 274–281, Apr. 2015, ISSN: 1746-8094. DOI: 10.1016/J.BSPC.2015.01.007.