



Theses and Dissertations

2019-09-01

The Application of Synthetic Signals for ECG Beat Classification

Elliot Morgan Brown
Brigham Young University

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>

BYU ScholarsArchive Citation

Brown, Elliot Morgan, "The Application of Synthetic Signals for ECG Beat Classification" (2019). *Theses and Dissertations*. 8116.

<https://scholarsarchive.byu.edu/etd/8116>

This Thesis is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact scholarsarchive@byu.edu, ellen_amatangelo@byu.edu.

The Application of Synthetic Signals for ECG Beat Classification

Elliot Morgan Brown

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Master of Science

Jeffrey Humpherys, Chair
Christopher Grant
Tyler Jarvis

Department of Mathematics
Brigham Young University

Copyright © 2019 Elliot Morgan Brown

All Rights Reserved

ABSTRACT

The Application of Synthetic Signals for ECG Beat Classification

Elliot Morgan Brown
Department of Mathematics, BYU
Master of Science

A brief overview of electrocardiogram (ECG) properties and the characteristics of various cardiac conditions is given. Two different models are used to generate synthetic ECG signals. Domain knowledge is used to create synthetic examples of 16 different heart beat types with these models. Other techniques for synthesizing ECG signals are explored. Various machine learning models with different combinations of real and synthetic data are used to classify individual heart beats. The performance of the different methods and models are compared, and synthetic data is shown to be useful in beat classification.

Keywords: ECG, synthetic data, SMOTE, signals, classification, machine learning, neural networks

ACKNOWLEDGMENTS

I'd like to thank all my math teachers from middle and high school for helping me develop a love for math. My professors at BYU, especially those in ACME, for shifting my interest toward the applied side of things. Lonette Stoddard for rapid e-mail communication and solving all of my problems. My advisor, Jeff Humpherys, for providing me with an interesting problem to solve, letting me adapt it, and working with me up until the very last minutes. And Christopher Grant and Tyler Jarvis for taking time to support me by rounding out my committee.

I'd also like to thank my family. My mom for always being there to listen. My dad for hours of math help and passing on his math gene. My brother-in-law for providing medical resources and answering ECG questions. And I'd like to thank my wife, Jamee, for pretending to care about the intricacies of my research and supporting me throughout the whole process. I couldn't have done this without you.

CONTENTS

Contents	iv
List of Tables	vi
List of Figures	vii
1 Physiology and Electrical Activity of the Heart	1
1.1 Basics of the Cardiovascular System	1
1.2 Electrical Activity of a Cardiac Cycle	2
2 Interpreting the Electrocardiogram	4
2.1 Rate	4
2.2 Rhythm: Arrhythmia	5
2.3 Rhythm: Block	9
2.4 Axis	11
2.5 Hypertrophy	12
2.6 Infarction	13
3 Related Work	14
3.1 ECG Classification	15
3.2 ECG Synthesis	15
4 Synthesis Modeling of ECG Signals	16
4.1 Heart Rate Variability	16
4.2 Dynamical Model on a Circle	17
4.3 Cardiac Conduction Model	22
4.4 Parameter Tuning	25
5 Alternative Synthesis Methods	29
6 Classification of ECG Signals	32
6.1 Data	33

6.2	Experimental Design	34
6.3	Classification	35
7	Conclusion	50
	Bibliography	52

LIST OF TABLES

4.1	Parameters used in (4.2) to create Figure 4.1	18
4.2	Parameters used in (4.8) to create Figure 4.3	22
4.3	Parameters used in (4.10) to create Figure 4.6	25
4.4	Parameters used in (4.11) to create Figure 4.6	26
6.1	XGBoost accuracy results	37
6.2	XGBoost F1 score results	37
6.3	DNN accuracy results	43
6.4	DNN F1 score results	44
6.5	CNN accuracy results	46
6.6	CNN F1 score results	47
7.1	Experiment Comparisons	50

LIST OF FIGURES

1.1	Blood flow of the human heart	1
1.2	Basic diagram of the ECG	2
1.3	Conduction system of the human heart	3
2.1	Sinus Rhythms	5
2.2	Wandering Pacemaker	6
2.3	Escape Rhythms	6
2.4	Junctional Escape Beat	7
2.5	Atrial Bigeminy	7
2.6	Premature Beats	8
2.7	Tachyarrhythmias	8
2.8	Fibrillations	9
2.9	Atrial Flutter	9
2.10	1° Block	10
2.11	2° Blocks	10
2.12	3° Block	11
2.13	Bundle Branch Blocks	11
2.14	Axis Deviations	12
2.15	Left Ventricular Hypertrophy	13
2.16	Infarction	14
4.1	A simple synthesized ECG signal	18
4.2	Breakdown of the 5 Gaussians in Figure 4.1	19
4.3	XYZ leads of a more robust synthetic ECG signal	21
4.4	Synthetic VCG from leads in Figure 4.3	22
4.5	Extended Conduction Model Framework	24
4.6	Conduction signal breakdown	27
4.7	Conduction signal summation	28
4.8	Synthetic cardiac conditions created with (4.8)	29

4.9	Bradycardia created with the extended conduction model	30
5.1	Data perturbation example with $L_1 = 0.08$, $L_n = -0.04$, $b = -0.4$	31
5.2	SMOTE example with $k = 3$	32
6.1	XGBoost baseline confusion matrix	36
6.2	XGBoost synthetic/SMOTE confusion matrix	36
6.3	A simple neural network	39
6.4	DNN architecture	42
6.5	DNN baseline confusion matrix	42
6.6	DNN SMOTE confusion matrix	43
6.7	Convolutional neural network architecture	45
6.8	DNN baseline confusion matrix	45
6.9	DNN synthetic/SMOTE confusion matrix	46
6.10	DNN Synthetic model confusion matrix	48
6.11	Parameterized DNN architecture	49
6.12	Parameterized DNN SMOTE confusion matrix	49

CHAPTER 1. PHYSIOLOGY AND ELECTRICAL ACTIVITY OF THE HEART

This chapter will explain some of the medical terminology and processes required to understand the research. It will address the basic mechanics of a beating heart and how the electrocardiogram (ECG or EKG) reflects those beats under normal circumstances. Dr. Dale Dubin's book [1] was used as the main resource for the information in this chapter.

1.1 BASICS OF THE CARDIOVASCULAR SYSTEM

To better understand electrocardiograms, it is helpful to review what the heart is actually doing at a high level in Figure 1.1 below. Under-oxygenated blood comes through superior (1) and inferior vena cava (2) and into the right atrium (3). The atrium contracts and sends the blood into the right ventricle (4). From there, the blood goes through the pulmonary artery (5) to the lungs where it is oxygenated before heading back to the heart. The blood flows through the pulmonary veins (6) and enters the left atrium (7), which contracts, sending the blood down into the left ventricle (8). The blood then flows through the aorta (9), carrying oxygen throughout the rest of the body, before returning to the vena cava.




Figure 1.1: Blood flow of the human heart

1.2 ELECTRICAL ACTIVITY OF A CARDIAC CYCLE

The basic flow of blood throughout the body, as described in the previous section, is powered by the contractions within the heart. Electrical stimulation is what causes the heart muscle (myocardium) to contract. Heart muscle cells are polarized (negatively charged) at rest. When the myocardium is electrically stimulated, the wave of depolarization causes it to contract. This is then followed by a recovery phase known as repolarization. Both depolarization and repolarization are caused by the movement of electrically charged ions.

The ECG records the electrical activity of contraction of the myocardium. Traditionally, twelve leads are placed on the patient's chest and limbs to monitor electrical activity in various regions. These leads, consisting of positive and/or negative electrodes, are placed on a patient's skin to detect the movement of ions. For example, a wave of positive ions moving toward a positive electrode causes an upward deflection on the electrocardiogram. Ions continue to move throughout the depolarization and repolarization. Any significant electrical activity within this process is recorded on the ECG. The process of depolarization and repolarization is known as a cardiac cycle. The various stages of the cardiac cycle can be seen in the ECG. It is critical to understand each part of the cycle and what causes it.




Figure 1.2: Basic diagram of the ECG

When a heart is pacing normally, the SA Node is responsible for initiating depolarization. At regular intervals, waves of Na^+ ions leave the SA Node and stimulate both atria, causing them to contract. This spread of positive ions produces an upward P wave on the ECG. As depolarization continues, the wave enters the AV Node, which is the pathway for electrical conduction between the atria and the ventricles. In the AV Node, depolarization slows down and is carried by Ca^{++} ions. This slow-down allows for blood to pass from the atria to the ventricles before depolarization continues. This pause can be seen on the ECG as a flat baseline immediately after the P wave. It is sometimes called the PR segment. The P wave together with the PR Segment make up the PR interval. (In general, a segment is a piece of baseline, whereas an interval contains at least one wave.)

Once depolarization passes through the AV Node, it rapidly makes it way through the His Bundle and then is split into the Left and Right Bundle Branches using fast-moving Na^+ ions. These bundles are comprised of Purkinje fibers. The depolarization of this His-Purkinje (HP) or ventricular conducting system has little direct effect on the ECG, but the resulting depolarization of the ventricular myocardium registers on the electrocardiogram as the QRS complex. The upward wave is the R wave, and the other two waves of the complex are distinguished by their relative position to the R wave. During ventricular depolarization, atrial repolarization (restoration of negative charge) is also occurring. Its effect is sometimes referred to as a Ta wave, but its relatively small electrical impact is masked by the QRS complex and does not appear on the ECG.

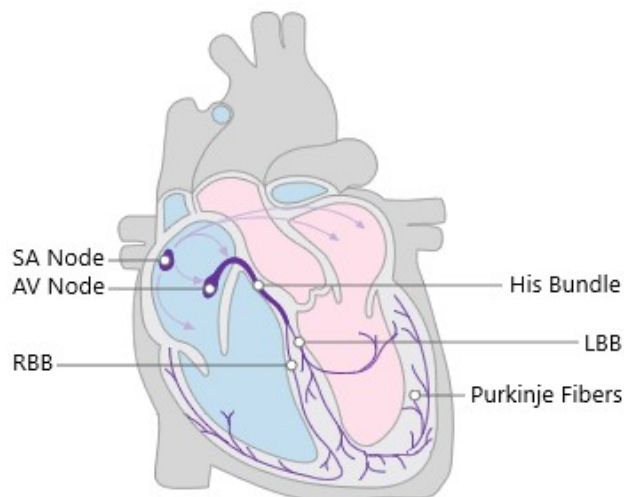


Figure 1.3: Conduction system of the human heart

After the QRS complex, there is a slight baseline plateau. This is known as the ST segment, which represents the beginning of ventricular repolarization. When the speed of repolarization becomes more rapid, the ECG records a broad T wave. K^+ ions are used in the repolarization of both the ST segment and the T wave. Ventricular contraction starts with QRS complex and continues until the end of the T wave. Due to its clinical significance in measuring repolarization, this piece of the cycle is given its own name: the QT interval. After the T wave, there is a brief baseline segment before the next P wave appears to signal the start of another cardiac cycle. This cycle, which typically lasts less than a second, is repeated continuously.

CHAPTER 2. INTERPRETING THE ELECTROCARDIOGRAM

Each part of the ECG signal, including the rate and rhythm, can supply useful information about a patient's overall health. The ECG provides details about the positioning and size of the heart as well as the movement of blood through the coronary artery. In this chapter, we explore how to use the ECG to determine rate, rhythm, axis, hypertrophy, and infarction. Again, [1] is the main source of information, and figures come from [2] with permission.

2.1 RATE

Determining heart rate from the ECG is fairly straightforward as the rate is measured in cycles per minute. On traditional red-line ECG paper, each heavy red line represents 0.2 seconds. Under normal circumstances, the Sinus Rhythm maintains a rate of 60-100 beats per minute (bpm). This is known as Normal Sinus Rhythm (NSR) and is a result of a regular functioning SA Node. At times, the SA Node beats outside of this normal range. When there are less than 60 bpm, this is known as bradycardia. This results from parasympathetic excess and can be found in professional athletes when resting. On the other end of the spectrum, tachycardia is when the heart beats at a rate of more than 100 bpm. Sympathetic stimulation from exercise is what causes this increase in heart rate. In either case, the SA Node is still responsible for each beat.

When the SA Node fails to beat, the heart has multiple backups in place. These fail-safes are known as automaticity foci and are located throughout the heart. Foci found in the atria each pace at their own inherent rate between 60 and 80 bpm. Foci pace at 40-60 in the AV Junction and




Figure 2.1: Sinus Rhythms

20-40 bpm in the ventricles. Faster pace-making activity suppresses slower ones in what is known as overdrive suppression. This creates a hierarchy of pacemakers within the heart which eliminates competition between them. For example, a junctional focus will only assume pace-making activity if the SA Node and all atrial foci fail to keep the heart beating.

2.2 RHYTHM: ARRHYTHMIA

Arrhythmia (or Dysrhythmia) is any rhythm disturbance or variance from NSR. The most common example of this is known as Sinus Arrhythmia. Though completely normal and very minimal, the process of respiration does have an effect on the heart's rhythm. There is a slight increase in heart rate when inhaling (sympathetic stimulation) as well as slight decrease when exhaling (parasympathetic inhibition). Other variations from NSR may be cause for concern and are more prominent on the ECG.

2.2.1 Irregular Rhythm. An irregular rhythm is when there is a lack of constant duration between paced cycles. This can occur when pacemaker activity periodically switches from the SA Node to a nearby atrial focus. This is known as a Wandering Pacemaker when the heart rate is still in the normal range. On the ECG, the P wave is replaced with a different looking P' wave when the SA Node is not pacing. Each focus will produce a distinct P' wave. It creates an irregular ventricular rhythm (observed in the QRS complex) and has cycles of varying length. If the heart rate increases to more than 100 bpm, this is known as Multifocal Atrial Tachycardia (MAT) and is a result of multiple atria foci resisting overdrive suppression to pace together.

Atrial Fibrillation is when multiple atrial foci are firing continually. These foci are completely insensitive to overdrive suppression and create chaotic atrial spikes on the ECG. Since no single impulse can completely depolarize the atria, Atrial Fibrillation also produces an irregular ventricular




Figure 2.2: Wandering Pacemaker

rhythm. This will be covered more in depth later on.

2.2.2 Escape. When there is a pause in SA Node pacing, an automaticity focus can escape overdrive suppression. In the case of Sinus Arrest, the SA Node ceases pacing completely and a backup focus must *escape* and take over pace-making responsibility. This is known as an escape rhythm. When an atrial focus is pacing, the ECG will show a pause followed by a switch from typical P waves to P' waves and a slower overall cycle, since the atrial focus paces at 60-80 bpm. If atrial foci are unable to pace, a junctional focus will take over at 40-60 bpm. A Junctional Escape Rhythm mainly conducts to the ventricles, which creates a series of QRS complexes. If the focus depolarizes the atria from below, this is known as retrograde depolarization, which produces an inverted P' wave on the electrocardiogram before the typical QRS complex. If there is a severe conduction block or complete failure of all automaticity foci above the ventricles, a ventricular focus will begin to pace. This creates very large QRS complexes and a rate of 20-40 bpm, which is unsustainable.



Figure 2.3: Escape Rhythms

Occasionally, an unhealthy SA Node will temporarily fail because of a transient Sinus Block. When the SA Node misses a beat, it creates a pause. If long enough, an automaticity foci will escape to emit an escape beat, after which the SA Node may resume pacing. As with escape rhythm, an escape beat can originate from the atria, AV Junction, or ventricles. In each case the ECG will register the same signal as the escape rhythm but only until the SA Node resumes pacing

(often a single cycle). At that point, the ECG will revert back to a normal rhythm.




Figure 2.4: Junctional Escape Beat

2.2.3 Premature Beat. Whereas escape beats indicate an issue with the SA Node, premature beats are a result of an irritable automaticity focus firing spontaneously. A premature beat comes from one of the three typical regions and often depolarizes the SA Node (or dominant automaticity center). This resets the cycle of the automaticity center so that pacing will resume one cycle after the premature beat. On occasion, a pattern of normal cycles and premature beats can occur. When each normal cycle is followed by a premature beat, this is called Bigeminy. When every two normal cycles are followed by a premature beat, it is known as Trigeminy.



Figure 2.5: Atrial Bigeminy

Both atrial and Junctional foci are commonly irritated by adrenaline, an increase in sympathetic stimulation, the presence of stimulants such as caffeine and cocaine, or an excess of digitalis or toxins. On the ECG, a Premature Atrial Beat (PAB) produces a P' wave earlier in the cycle. Sometimes, it is early enough that it is on the peak of the previous T wave, creating an unusually tall T wave. If the AV Node is not depolarized by the PAB, there may be no QRS complex or T wave. Premature Junctional Beats (PJB) produce a premature QRS complex and may record an inverted P' wave just prior if there is retrograde depolarization of the atria. If ventricular depolarization is not complete due to PAB or PJB, aberrant ventricular conduction produces a widened QRS complex.

Irritable ventricular foci are mainly caused by hypoxia (low O_2) but also by hypokalemia (low K^+) and various pathologies and medications. When a ventricular focus beats spontaneously, the ECG shows a very large QRS complex of opposite polarity, which is often referred to as a Premature Ventricular Complex (PVC). If the PVC is delayed enough, it can fall on a T wave and lead to dangerous arrhythmias. This irregularity is due to the non-simultaneous depolarization of the ventricles. Unlike with PAB and PJB, the SA Node is not depolarized, which means the cycle is not reset. There is a brief pause after the PVC, but this is due to arrival of a P wave before the ventricles have repolarized. It is also possible for a ventricular focus to suffer from entrance block and not be irritable. The focus ignores overdrive suppression and beats at its own rate while the SA Node continues to pace, producing a dual rhythm with PVCs.




Figure 2.6: Premature Beats

2.2.4 Tachyarrhythmia. When an irritable focus paces very rapidly, this creates a tachyarrhythmia. If the heart rate is between 150-250 bpm, this is known as Paroxysmal Tachyarrhythmia. The ECG can be used to determine the origin and probable cause. Paroxysmal Atrial Tachyarrhythmia (PAT) creates P' waves and is usually sign of digitalis excess or toxicity. PAT with AV Block has rapid, spiked P' waves with 2:1 ratio to QRS complexes. Paroxysmal Junctional Tachyarrhythmia (PJT) can record an inverted P' before, during, or after the QRS complex. It is caused by stimulants or premature beats. Paroxysmal Ventricular Tachyarrhythmia (PVT or simply VT) is a sign of hypoxia. It creates a rapid series of PVCs on the ECG.



Figure 2.7: Tachyarrhythmias

If an irritable focus paces at 250-350 bpm, it is called a flutter. An atrial flutter produces a signal with a saw-tooth baseline on the ECG. A ventricular flutter records smooth sine waves and often leads to fibrillation. Fibrillation is when the focus paces at 350-450 bpm. In reality, multiple foci are pacing at the same time and so the actual heart rate is not a true rate. Atrial fibrillation has a very wavy baseline without P waves, as well as an irregular QRS complex. Ventricular fibrillation has no identifiable waves or patterns.



Figure 2.8: Fibrillations



Figure 2.9: Atrial Flutter

2.3 RHYTHM: BLOCK

Blocks slow down or completely prevent the conduction of depolarization. A block can develop in the SA Node, AV Node, or Bundle Branches.

2.3.1 Sinus Block. This is when an unhealthy SA Node momentarily fails to pace for at least one cycle. This creates a pause on the ECG. As discussed previously, this pause could lead to an escape beat. Either way, after the Sinus block, normal pacing resumes with a P wave identical to the one before the block.

2.3.2 AV Block. A block in conduction from the atria to the ventricles is known as AV block. AV block prolongs the delay between atrial and ventricular depolarization. This, in turn, increases the duration of the PR interval to more than 0.2 seconds. There are three different degrees of AV

blocks. When the length of the prolonged PR interval is consistent and the P-QRS-T sequence is unaltered, first degree AV block is occurring.



Figure 2.10: 1° Block

Second degree AV block allows atrial to ventricle depolarization to conduct sometimes but not always. When the depolarizing P wave is blocked, there is no associated QRS complex. Wenckebach blocks are second degree blocks of the AV Node. The AV node is progressively blocked in a series of cycles until the final P wave is totally blocked. This creates a consistent P:QRS ratio of $n : n - 1$ on the ECG. Second degree blocks of His Bundle or Bundle Branches are called Mobitz. They produce a P:QRS ratio of $n : 1$, where a normal P-QRS-T cycle is preceded by cycles with lone P waves.



(a) Wenckebach, 3:2



(b) Mobitz, 2:1

Figure 2.11: 2° Blocks

A third degree AV block is a complete block of conduction from the atria to the ventricles. The P wave has no QRS response, and the PR interval is totally variable. The SA Node continues to pace the atria, but the ventricles are independently paced by either a junctional or ventricular focus, depending on the location of the block. If a junctional focus is pacing, a normal QRS complex will be present. When a ventricular focus paces, the ECG will reflect wide, PVC-like complexes.

2.3.3 Bundle Branch Block. BBB is caused when there is a block in the Right or Left Bundle Branch. It retards depolarization to its respective ventricle. Since the depolarization of the right and left branches typically occur simultaneously, this delay produces a wide QRS complex with two peaks, one for each ventricle. The first one to depolarize creates the R wave; the second creates an




Figure 2.12: 3° Block

R' wave. When the Left Bundle Branch is blocked, there is a small valley between the R and R' waves. When the right is blocked, the valley dips well below baseline.




Figure 2.13: Bundle Branch Blocks

2.4 AXIS

Depolarization of the ventricles begins in the endocardium (innermost layer of tissue lining the chamber) and proceeds outward through the ventricular wall. Each vector of depolarization is normal to the endocardium and has magnitude relative to the thickness of the wall through which it travels. The Mean QRS Vector is the sum of the smaller vectors with origin at the AV Node. Since the left ventricle is thicker, the Mean QRS vector should point down and to the patient's left.

If the Mean QRS Vector does not point in this direction it can indicate various issues. If there is a rotation or deviation in the heart, the Vector will indicate the severity and direction. If one of the ventricles is enlarged (hypertrophy), the Vector will point more toward the hypertrophied side. If there is a necrotic area of the heart, it will not depolarize, so the Mean QRS Vector will point away from that area.

As previously stated, there are 12 different leads used in a standard ECG. The lead on the left arm (lead I) is positive and horizontal. The QRS complex as seen in lead I should be positive. This

indicates that the the Mean QRS Vector is pointing somewhat to the left. If the QRS complex is negative, this indicates that there is a Right Axis Deviation (RAD) Similarly, the QRS complex on the positive AVF lead in the left foot should also be positive. This means the Vector is pointing downward. If the QRS complex is negative, there is Left Axis Deviation (LAD). By checking both these leads, one can determine whether or not the Vector is in the correct quadrant. Since the Mean QRS Vector is actually three-dimensional, chest leads can also be used to ensure the heart is properly rotated, but the heart is not able to rotate significantly in the horizontal plane.




Figure 2.14: Axis Deviations

2.5 HYPERTROPHY

Hypertrophy is an increase in the muscle mass of the heart. It can increase both the thickness of the wall and the volume of the chamber. Both anomalies can be diagnosed using the ECG. Because the lead V_1 is closest to the atria, it can provide the best information about atrial enlargement, where both the wall and chamber are enlarged. A diphasic (positive and negative) P wave is the best indicator of atrial enlargement. If the initial phase of the wave is larger than the other, the enlargement is on the right side. If the second phase dominates, it is on the left.

Signs of ventricular hypertrophy are found in the QRS complex. If the R wave is large in V_1 , it is on the right. If the S wave is large in V_1 and the R wave is large in V_5 , then it is on the left. An inverted T wave may also be present. Often, ventricular hypertrophy is caused by a ventricle that is straining due to some added resistance. Ventricular strain produces a depressed, convex ST segment.




Figure 2.15: Left Ventricular Hypertrophy

2.6 INFARCTION

Infarction occurs when a coronary artery is completely occluded. The infarcted area of the myocardium does not receive any blood; it suffers necrosis and cannot depolarize. The ECG can show which artery is occluded, reveal blocks caused by infarction, and show if a coronary vessel is narrowed. Indicators of ischemia, injury, and necrosis can help identify infarction.

2.6.1 Ischemia. Reduced blood flow is known as ischemia. In the heart, this typically leads to hypoxia since the heart is not getting the oxygen it needs. On the ECG, ischemia produces an inverted T wave, which is typically symmetrical. Ischemia may lead to infarction, but it can also cause chest pain.

2.6.2 Injury. Injury shows the acuteness of an infarct, meaning how recently it occurred. It is characterized by an elevated ST segment. Over time, the elevation decreases, and the ST segment returns to normal. If there is no associated Q wave, this can be a sign of an impending infarct. Pericarditis can also elevate the ST segment, but in that case, it is flat or concave. Digitalis, stress tests, and subendocardial infarction can depress the ST segment. ST depression along with large R waves can be a sign of a posterior infarction.

2.6.3 Necrosis. Significant Q waves on the ECG are produced by an area of dead tissue. They are the main tool for diagnosing myocardial infarction. A Q wave is considered significant when it

lasts more than 0.04 seconds or is at least one third of the overall amplitude of the QRS complex. Examining various leads can provide information about which portion of the ventricle has necrosis.




Figure 2.16: Infarction

2.6.4 Hemiblock. Hemiblocks are blocks of the Left Bundle Branch (LBB) that occur frequently with infarction. Anterior hemiblocks occur in the anterior division of the LBB. They are characterized by LAD, a slightly widened QRS complex, the presence of a Q wave in lead I and a large S wave in lead III. Posterior hemiblocks are rare due to the nature of the posterior division of the LBB. They produce RAD, a slightly widened QRS complex, a large S in lead I and the presence of a Q in III. These blocks can be intermittent or permanent.

Being familiar with normal ECG signals, as well as how they change in the event of cardiac conditions, is important both for improving synthesis and classification. While the ideas behind the synthesis and classification of ECG signals are mathematical, domain knowledge is critical to put those mathematical ideas into context and to obtain the best results.

CHAPTER 3. RELATED WORK

The electrocardiogram has become a critical diagnostic tool since its invention over a century ago. As the previous chapter demonstrates, many cardiac conditions are reflected in the ECG. Trained medical professionals routinely use the ECG to aid them in their care of patients. Many researchers have used machine learning techniques in an effort to assist physicians in their diagnoses. A lack of biomedical ECG recordings have in turn led to the generation of synthetic electrocardiogram signals, often to aid in filtering the raw ECG signal.

3.1 ECG CLASSIFICATION

Many researchers have worked on the classification of heart beats and rhythms. Traditionally, statistical methods such as random forests and support vector machines were used [3]. In order to make such methods successful, various types of preprocessing must be performed. Linear transformations such as principal component analysis or linear discriminant analysis are often used to transform the data. Other techniques include using the DWT (Discrete Wavelet Transform) or Fourier transforms, as well as various filtering methods [4].

Lately, however, neural networks have emerged as the preferred method for classification [5]. Some methods use 1D signals while others use 2D images [6]. While various preprocessing steps may still be taken, neural network models are less dependent on them. (More detail about how neural networks work will be given in Chapter 5.)

A recent Stanford study with access to a novel data set achieved a higher accuracy than that of a cardiologist in rhythm detection with little to no signal preprocessing [7]. While their approach to the model architecture drew inspiration from the latest neural network techniques, a big reason for their success was their data.

Machine learning techniques in general, but especially neural networks rely on large amounts of data. One difficulty with ECG classification is the availability of labelled data. Hospitals all over the world collect ECG data on a daily basis, but it is not available to the general public. And, even if it were, it wouldn't have each peak or rhythm nicely labelled. As more data sets become available to researchers, results will only improve.

Researchers hope that by correctly identifying cardiac conditions, they cannot only aid doctors and nurses, but also provide medical assistance in areas where medical professionals are not readily available. While these algorithms only identify conditions rather than treating them, they could help prioritize needs and provide insight into future treatments.

3.2 ECG SYNTHESIS

Due to the usefulness and scarcity of ECG recordings, researchers have explored synthesizing ECG signals. Many of these are focused on their uses for signal processing techniques with aims to improve feature extraction [8] [9], discriminant analysis [10], and setting benchmarks[11]. Many

researchers use clean synthetic signals before adding noise or using real ECG signals in order to test their new techniques. Other researchers hope to improve medical education [12] by accurately modeling the heart.

Methods vary in their approaches to recreating a realistic signal. A summary of various models may be found in [13], where models are divided into three different types. Some methods focus on creating a close approximation [8] with a system of ODEs on a circle. The other two focus on modeling the actual physiological processes occurring within the heart [12]. Within these, some model the natural pacemakers of the heart, whereas the other models both these pacemakers and the muscles involved. In both the physiological and non-physiological case, varying levels of detail are employed, depending on the final goal of the synthetic signal. One of each of these approaches will be explored in the following chapter.

In addition to methods developed especially for ECG signals, there are many methods designed for general augmentation of data. There are techniques that use a single sample and others that use multiple samples in order to generate synthetic data. A handful of these techniques will also be explored later on in Chapter 5.

CHAPTER 4. SYNTHESIS MODELING OF ECG SIGNALS

As mentioned in the previous section, there are many methods for synthesizing ECG signals. Here, we consider two different models. The first model is chosen for its simplicity and many years of improvement and fine-tuning. The second model is newer and more complex but better approximates the associated physiology. We also explore adding natural variations to the ECG signals with the aim to make them more realistic.

4.1 HEART RATE VARIABILITY

When synthesizing a realistic ECG signal, it is important to consider the natural fluctuation in heart rate due to respiration and other factors. Heart rate can be measured by measuring the RR intervals (the distance between the peaks of two consecutive R waves), and then finding the average. The series of beat-to-beat RR intervals is known as the RR tachogram. Heart rate varies as the balance between sympathetic and parasympathetic shifts.

Spectral analysis of the RR tachogram can be used to estimate the effect of these changes. There are two main frequency bands present in the ECG. The higher frequency band (0.15-0.4 Hz) corresponds to the respiratory cycle [8]. This is the previously mentioned respiratory sinus arrhythmia (RSA). If a patient takes 15 breaths per minute, there will be a 4 second oscillation in the ECG signal, with a corresponding 0.25 Hz peak in the power spectrum. The lower frequency band (0.04-0.15 Hz) similarly generates an oscillation and is most likely influenced by the response of the baroreflex to changes in blood pressure [8].

To replicate these two phenomena, the following bi-modal spectrum may be used [14]:

$$S(f) = \sum_{i=1}^2 \frac{\sigma_i^2}{\sqrt{2\pi}\gamma_i} \exp\left(-\frac{(f-f_i)^2}{2\gamma_i^2}\right), \quad (4.1)$$

where $i = 1$ corresponds to the LF band and $i = 2$ is for the HF. Note that this is the sum of two Gaussians with means f_i and standard deviations γ_i . The power of each band is given by σ_i^2 and the variance is their sum.

To generate a time series $T(t)$ with power spectrum $S(f)$, we take the inverse Fourier transform of a series of complex numbers with amplitudes $\sqrt{S(f)}$ and phases randomly distributed between 0 and 2π [14]. This time series may then be scaled and offset to have a desired mean and standard deviation. Thus, a natural heart rate variability may be incorporated into the model.

4.2 DYNAMICAL MODEL ON A CIRCLE

The model presented in this section will build off of the dynamical model first presented by McSharry et al. in [14]. In this three-dimensional model, let $W = [P, Q, R, S, T]$ the set of named waves on the ECG. Then

$$\begin{aligned} \dot{x} &= \alpha x - \omega y \\ \dot{y} &= \alpha y + \omega x \\ \dot{z} &= - \sum_{i \in W} a_i \Delta\theta_i \exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right) - (z - z_0), \end{aligned} \quad (4.2)$$

where $\alpha = 1 - \sqrt{x^2 + y^2}$, $\Delta\theta_i = (\theta - \theta_i) \bmod 2\pi$, $\theta = \arctan2(y, x)$, and ω is the angular velocity. The locations of five peaks are given by their corresponding angle θ_i . The coefficients a_i and

b_i determine the magnitude and width of each peak respectively. Setting $z_0(t) = A \sin(2\pi f_2 t)$ introduces baseline wander due to respiration. In this model, the output signal $s(t)$ is simply $z(t)$, which is found by integrating (4.2) with a numerical solver. To add in heart rate variability, let $\omega(t) = 2\pi/T(t)$ with $T(t)$ generated by the process described in Section 4.1.

This original model is fairly intuitive. The x and y values trace the circle as z leaves the baseline at each peak. The time it takes to trace that circle depends on the angular velocity, which is a function of the heart rate. Note that z is simply a sum of five Gaussian functions with amplitude, mean, and standard deviation (a_i, θ_i, b_i) as input. By examining an actual ECG, appropriate values for a_i, θ_i, b_i , and ω may be found to generate a synthetic ECG signal. The result of this model can be seen in 4.1 with parameters given in Table 4.1.

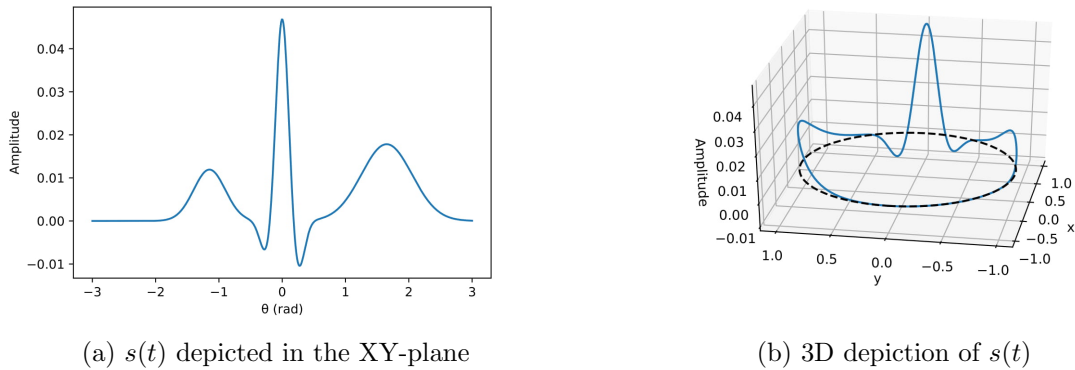


Figure 4.1: A simple synthesized ECG signal

Index	P	Q	R	S	T
θ_i	-1.15	-0.25	0.00	0.25	1.65
a_i	1.20	-5.00	30.0	-7.50	0.70
b_i	0.25	0.10	0.10	0.10	0.40

Table 4.1: Parameters used in (4.2) to create Figure 4.1

A few simple modifications proposed in following years can be utilized to improve this model. If we wish for z to be measured in millivolts (mV), and we assume that both the θ_i and b_i are in radians, with time measured in seconds, then simple dimensional analysis of the equation for \dot{z} in (4.2) without baseline wander shows that




Figure 4.2: Breakdown of the 5 Gaussians in Figure 4.1

$$\begin{aligned}
\frac{mV}{s} = [\dot{z}] &= - \sum_{i \in W} [a_i] [\Delta\theta_i] \left[\exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right) \right] \\
&= [a_i] [\Delta\theta_i] \\
&= [a_i] \cdot \text{rad}.
\end{aligned} \tag{4.3}$$

This implies that $[a_i] = mV/(\text{rad} \cdot s)$. In order to make the units more interpretable and to simplify the dimensions of a_i , we add a ω/b_i^2 term as a factor in the summation, which has units $1/(\text{rad} \cdot s)$. With this modification, the a_i terms are now measured in millivolts. This modification was first introduced in [15] with alternate reasoning.

The model can further be simplified by a change of variables from Cartesian coordinates to polar [16]. Differentiating the equation $r^2 = x^2 + y^2$ and making appropriate substitutions for x , y , and their derivatives gives an equation for \dot{r} .

$$\begin{aligned}
\dot{r} &= \frac{x\dot{x} + y\dot{y}}{r} \\
&= \frac{r \cos \theta (\alpha r \cos \theta - \omega r \sin \theta) + r \sin \theta (\alpha r \sin \theta + \omega r \cos \theta)}{r} \\
&= \alpha r \cos^2 \theta - \omega r \sin \theta \cos \theta + \alpha r \sin^2 \theta + \omega r \sin \theta \cos \theta \\
&= \alpha r \\
&= (1 - r)r.
\end{aligned} \tag{4.4}$$

Similarly, the derivative of $\tan \theta = \frac{y}{x}$ with respect to t shows that

$$\begin{aligned}
\dot{\theta} &= \frac{x\dot{y} - y\dot{x}}{r^2} \\
&= \frac{r \cos \theta (\alpha r \sin \theta + \omega r \cos \theta) - r \sin \theta (\alpha r \cos \theta - \omega r \sin \theta)}{r^2} \\
&= \alpha \sin \theta \cos \theta + \omega \cos^2 \theta - \alpha \sin \theta \cos \theta + \omega \sin^2 \theta \\
&= \omega.
\end{aligned} \tag{4.5}$$

Note that neither this equation for $\dot{\theta}$ nor the equation for \dot{z} in (4.2) depend on r . Furthermore, (4.4) indicates that \dot{r} has no dependence on other variables. As such, it provides no necessary information. Thus, the dynamical system simply becomes:

$$\begin{aligned}
\dot{\theta} &= \omega \\
\dot{z} &= - \sum_{i \in W} \frac{a_i \omega}{b_i^2} \Delta \theta_i \exp \left(-\frac{\Delta \theta_i^2}{2b_i^2} \right) - (z - z_0),
\end{aligned} \tag{4.6}$$

where $-\pi \leq \theta < \pi$, $\Delta \theta_i = (\theta - \theta_i) \bmod 2\pi$, and ω is the angular velocity. This system has the benefit of fewer equations, with the equation for θ now representing the angular position of each peak.

To improve the model further, note that \dot{z} in (4.2) is the sum of five Gaussian functions, one for each peak. This number is intuitive, but in order to synthesize the asymmetrical T wave an additional Gaussian must be used [17]. To better replicate the P wave, the uncommon U wave, or other abnormalities, even more Gaussians are required [11]. In fact, an arbitrarily close approximation can be achieved by using a sufficient number of Gaussian functions [18]. To implement this change, we instead take \dot{z} to be the sum of n Gaussians, with $n \geq 5$.

Finally, the system can be expanded to generate multi-lead signal by letting z be a d -dimensional vector. The parameters θ , a , and b are likewise expanded to be $d \times n$ matrices Θ , A , and B . If $d = 12$, a full 12-lead ECG will be generated. An alternate approach to generating 12 leads, is also explored in [11]. They propose setting $d = 3$ in order to generate the 3 XYZ Frank leads. The 12 leads are then found by the following equation:

$$ECG(t) = H \cdot R \cdot \Lambda \cdot s(t) + W(t), \tag{4.7}$$

where $\text{ECG}(t)$ is the 12×1 final vector, $s(t)$ is the 3×1 vector result of the system of ODEs, H is the 12×3 Dower transformation matrix (which accounts for the change in body mass), Λ is a 3×3 diagonal matrix used to scale each of the directions, R is the 3×3 rotation matrix, and $W(t)$ adds noise. Due to variations in body mass and electrode placement, however, there is no universal H , that can be reliably used in this way. It may be found using a least squares approach when both the 12 traditional leads as well as the three VCG leads are known, but that solution varies from person to person.

Nevertheless, allowing for more Gaussians and for z to be a vector of \mathbf{z} of desired length, yields the final version of the dynamical model presented in this paper:

$$\begin{aligned} \dot{\theta} &= \omega \\ \dot{\mathbf{z}} &= - \sum_{i=1}^n \frac{A_i \omega}{B_i^2} \Delta \Theta_i \exp\left(-\frac{\Delta \Theta_i^2}{2B_i^2}\right) - (\mathbf{z} - \mathbf{z}_0), \end{aligned} \quad (4.8)$$

which is robust enough to model many cardiac conditions with multiple leads. Figure 4.3 shows the result of using (4.8) with the parameters in Table 4.2.




Figure 4.3: XYZ leads of a more robust synthetic ECG signal




Figure 4.4: Synthetic VCG from leads in Figure 4.3

Index	1	2	3	4	5	6	7	8	9
θ_i^x	-1.09	-0.83	-0.19	-0.07	0.00	0.06	0.22	1.20	1.42
a_i^x	0.03	0.08	-0.13	0.85	1.11	0.75	0.06	0.10	0.17
b_i^x	0.09	0.11	0.05	0.04	0.03	0.03	0.04	0.60	0.30
θ_i^y	-1.10	-0.90	-0.76	-0.11	-0.01	0.07	0.80	1.58	2.90
a_i^y	0.04	0.02	-0.02	0.32	0.51	-0.32	0.04	0.08	0.01
b_i^y	0.07	0.07	0.04	0.06	0.04	0.06	0.45	0.30	0.50
θ_i^z	-1.10	-0.93	-0.70	-0.40	-0.15	0.10	1.05	1.25	1.55
a_i^z	-0.03	-0.14	-0.04	0.05	-0.40	0.46	-0.12	-0.20	-0.35
b_i^z	0.03	0.12	0.04	0.40	0.05	0.05	0.80	0.40	0.20

Table 4.2: Parameters used in (4.8) to create Figure 4.3

4.3 CARDIAC CONDUCTION MODEL

One drawback to the model presented in 4.2 is that the model has no direct physical correlation. The model is simply a sum of Gaussians that does a nice job of fitting the signal, especially as the number of Gaussians increases. The parameters of each Gaussian are unrelated to what is actually happening in the cardiovascular system. Since the ECG is already an abstraction of the actual conduction occurring within the heart, (4.8) essentially becomes a representation of a representation.

In later papers, models of natural pacemakers are presented [19]. This idea is built upon by accounting for signals from cardiac muscles as well [13]. These models attempt to replicate the building blocks of the ECG signal, which actually removes a layer of abstraction. In [12], the

authors build upon their model presented in [13] to create a better ECG signal. Finally, in [20] that model is combined with the RR tachogram used in [8] to add heart rate variability.

The natural pacemakers considered are the SA Node (SA), the AV Node (AV), and the His-Purkinje system (HP). The equations for these natural pacemakers with $i = 1, 2, 3$ corresponding to SA, AV, and HP respectively, are given by

$$\ddot{x}_i = -F(x_i)y_i - G(x_i) + K_i (y_{i-1}^{\tau_i} - y_i), \quad (4.9)$$

where $F(x_i) = a_i (x_i^2 - u)$ correspond to the damping terms and $G(x_i) = f_i x_i (x_i + d)(x_i + e_i)$ are the harmonic forces. The value K_i is the corresponding coupling constant between the i th node and the previous one, so that $K_2 = K_{SA:AV}$ and $K_3 = K_{AV:HP}$, while K_1 is set to zero. The notation $y_i^{\tau_i} = y_i(t - \tau_i)$ represents the function y_i evaluated at time $t - \tau_i$, where τ_i follows the same convention as the K_i as time delayed coupling. (Note: all other instances of y_i and x_i are evaluated at time t .) Finally, the values $a_i, f_i, u, d,$ and e_i are parameters of each oscillator.

Writing (4.9) in two-dimensional form and making substitutions for F and G yields:

$$\begin{aligned} \dot{x}_i &= y_i \\ \dot{y}_i &= -a_i (x_i^2 - u) y_i - f_i x_i (x_i + d_i)(x_i + e_i) + K_i (y_{i-1}^{\tau_i} - y_i). \end{aligned} \quad (4.10)$$

As mentioned in [12], the natural pacemakers have a much smaller effect on the ECG signal than the actual cardiac muscles do. They propose the use of modified FitzHugh-Nagumo oscillators [21] to model depolarization and repolarization of both the atrial (AT) and ventricular (VN) muscles:

$$\begin{aligned} \dot{z}_j &= k_j (-c_j z_j (z_j - w_{j1})(z_j - w_{j2}) - b_j v_j - g_j v_j z_j + I_j) \\ \dot{v}_j &= k_j h_j (z_j - v_j). \end{aligned} \quad (4.11)$$

Here, $j = 1, \dots, 4$ correspond to AT depolarization (P wave), AT repolarization (Ta wave), VN depolarization (QRS complex), and VN repolarization (T wave) respectively. The term $I_j = C_j \tilde{y}_j H(\tilde{y}_j)$ indicates the magnitude of the stimulation that couples SA to AT and HP to VN. This is accomplished by letting C_j be the coupling coefficients and letting $H(\tilde{y})$ be the Heaviside step function with $\tilde{y}_1 = y_1, \tilde{y}_2 = -y_1, \tilde{y}_3 = y_3,$ and $\tilde{y}_4 = -y_3$. The remaining parameters $k_j, c_j, w_{j1}, w_{j2}, b_j, h_j,$ and g_j affect the rest state, the excitability, the duration of the action potential, the

excitation threshold and excited state of each oscillator.

Notice that the equation for \dot{z}_j found in (4.11) depends on the values of y_i found in (4.10). The solutions z_j are then combined to generate the final ECG waveform:

$$ECG = z_0 + z_1 - z_2 + z_3 + z_4, \quad (4.12)$$

where z_0 shifts the baseline value of the signal. Here, the value z_2 is subtracted because its polarity is known to be opposite of z_1 [20].

In order to extend this idea, the RR tachogram (Section 4.1), baseline wander, and additive noise may be added into the model [20]. To incorporate the RR tachogram, the value of f_1 can be made to vary with the time series $T(t)$. Baseline wander can be added as before by allowing z_0 to vary periodically as in (4.2), and Gaussian noise is easily added. Weighting coefficients β_1, \dots, β_4 are also included to vary the contribution of each component. By using different values of β_j , one can generate a 12-lead ECG. The ECG signal then becomes:

$$ECG = z_0 + \beta_1 z_1 - \beta_2 z_2 + \beta_3 z_3 + \beta_4 z_4, \quad (4.13)$$

with the overall process for generating this signal shown in Figure 4.5.

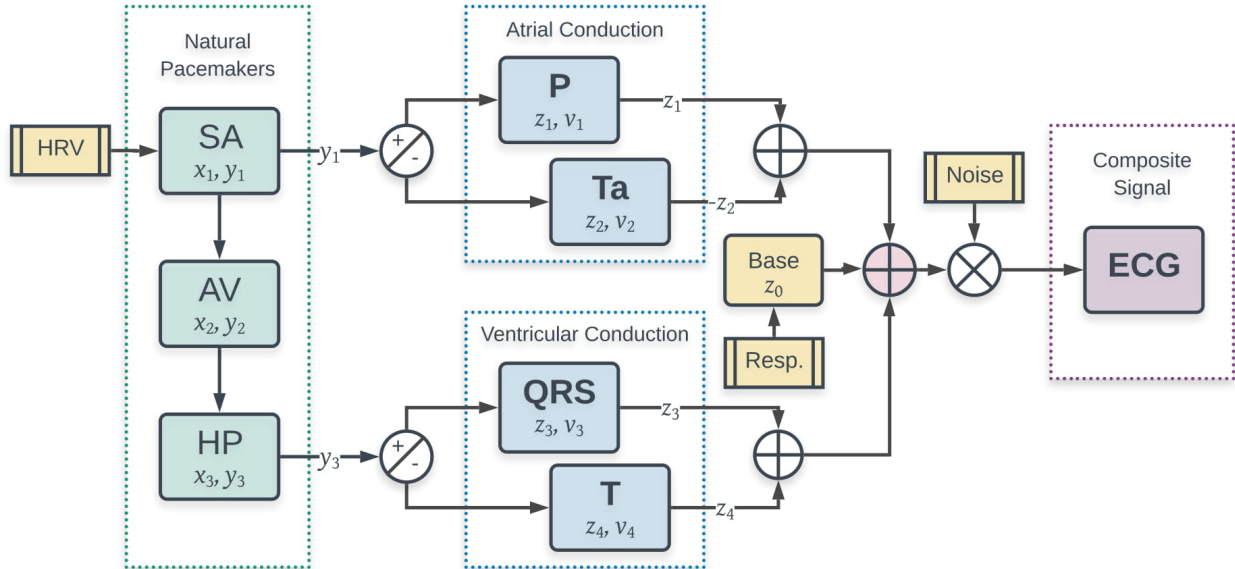


Figure 4.5: Extended Conduction Model Framework

This extended conduction model includes the variability originally introduced in the circular

Parameter	Description	Value
a_1	Nonlinear damping coefficients	40
a_2		40
a_3		50
f_1	Parameters defining the amplitude of a pulse	22.0
f_2		8.4
f_3		1.5
u	Damping coefficients controlling excitation threshold and excited state	0.69
d	Position control of rest state/dynamics (hyperbolic saddle)	3.0
e_1	Position control of hyperpolarization of excitation (stable node)	22.0
e_2		8.4
e_3		1.5
K_1	Coupling coefficients	0
K_2		22.0
K_3		22.0
τ_1	Time delay	0
τ_2		0.092
τ_3		0.092

Table 4.3: Parameters used in (4.10) to create Figure 4.6

dynamical model [14] to a model with a more physiological backing. It is more likely to gain insights from the parameters themselves if they have a physiological meaning. Furthermore, parameters could be tuned according to changes in the actual sub-processes that lead to the ECG, rather than modifying the reflection of those changes directly in the ECG itself. An example of an ECG of a normal patient is shown in Figures 4.6 and 4.7.

4.4 PARAMETER TUNING

In order to generate realistic signals for patients with various cardiac conditions, the parameters of the models must be tuned appropriately.

4.4.1 Circle Model. The MatLab code presented with [8] on PhysioNet was converted to Python as a base implementation. It was used to generate Figure 4.1. The code was then adapted to handle additional parameters and use the proposed model in (4.8).

The first step in signal synthesis was to synthesize a single RR interval. When focusing on a single beat, it is more practical to integrate the $\dot{\mathbf{z}}$ term in (4.8) with respect to t to get an equation for $\mathbf{z}(t)$:

Parameter	Description	Value
k_1	Scaling coefficients	2e3
k_2		1e2
k_3		1e4
k_4		2e3
c_1	Parameters defining the amplitude of a pulse	0.26
c_2		0.12
c_3		0.12
c_4		0.10
w_{11}	Parameters controlling the excitation threshold	0.13
w_{12}		0.12
w_{13}		0.12
w_{14}		0.22
w_{21}	Parameters controlling the excited state	1.0
w_{22}		1.1
w_{23}		1.1
w_{24}		0.8
b_1	Parameters changing rest state and dynamics	0.0
b_2		0.0
b_3		0.015
b_4		0.0
h_1	Parameters representing excitability and controlling abruptness of activation and duration of action potential	0.004
h_2		0.008
h_3		0.008
h_4		0.008
g_1	Parameters changing rest state and dynamics	0.40
g_2		0.09
g_3		0.09
g_4		0.10
C_1	Coupling coefficients	4e-5
C_2		4e-5
C_3		9e-5
C_4		6e-5

Table 4.4: Parameters used in (4.11) to create Figure 4.6

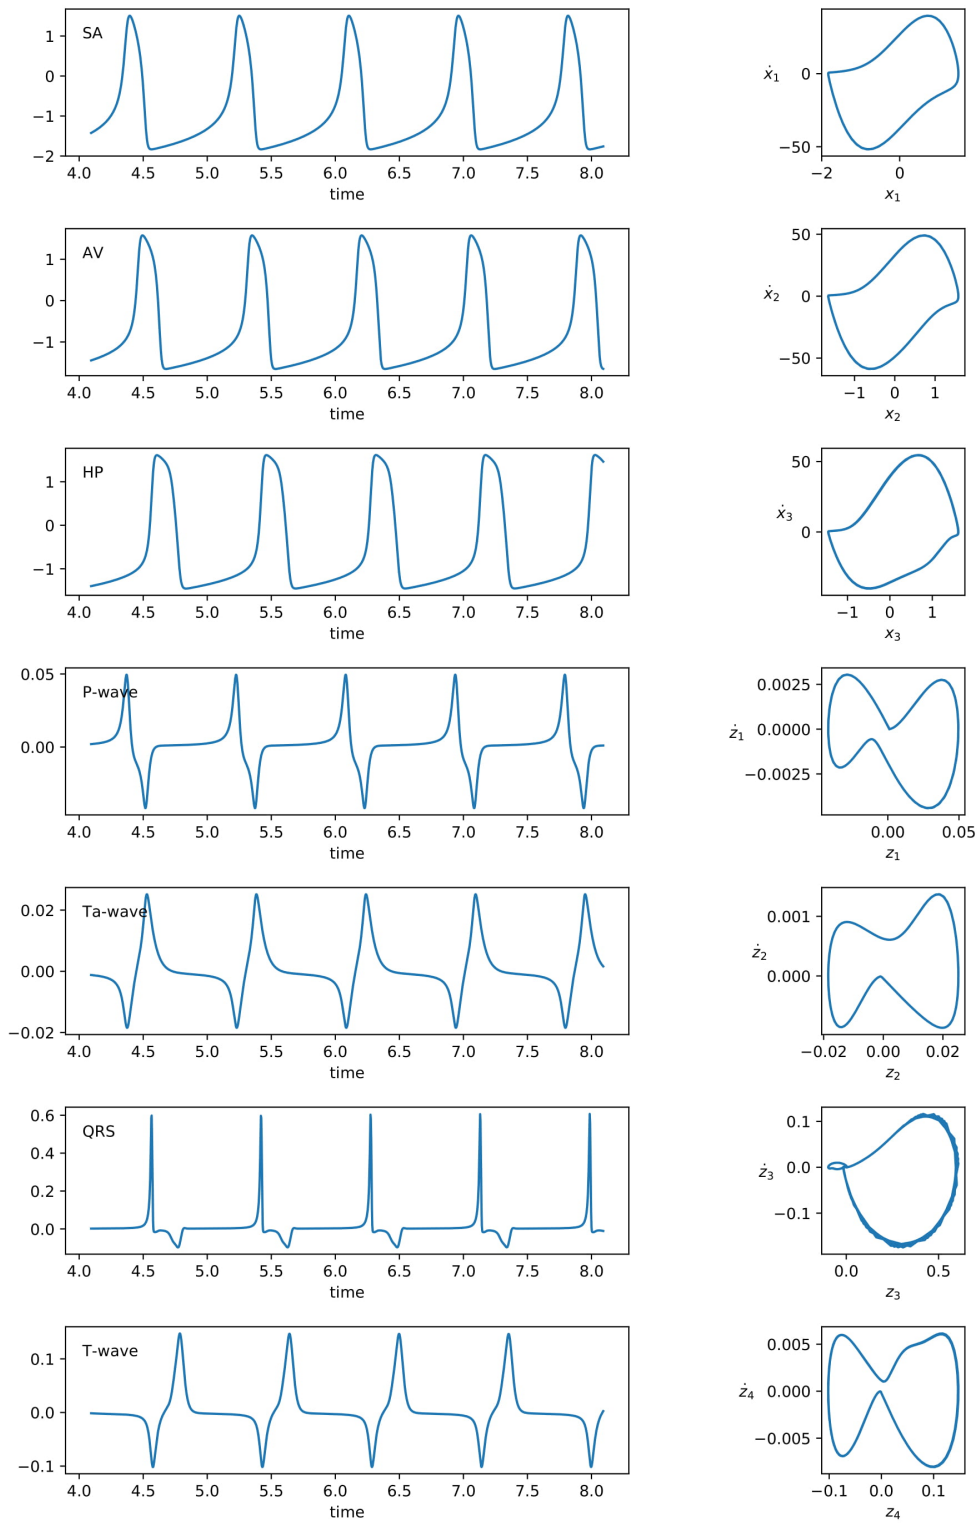


Figure 4.6: Conduction signal breakdown




Figure 4.7: Conduction signal summation

$$z(t) = \sum_{i=1}^n a_i \omega \exp\left(-\frac{\Delta\theta_i^2}{2b_i^2}\right). \quad (4.14)$$

In order to fit a synthetic $z(t)$ to an actual beat of signal $s(t)$, we used (4.14). Once the ideal parameters for a , b , and θ were found, they could then be used in the system of ODEs to generate a synthetic signal of that type of beat. An actual signal may be approximated by minimizing the least squares error of the synthetic and actual signals:

$$\min_{a_i, b_i, \theta_i} \|s(t) - z(t)\|^2. \quad (4.15)$$

The number of Gaussians is not a learnable parameter in this scenario, so this had to be determined by experimentation. For most signals, seven Gaussians seemed to be sufficient to give good enough results without adding an unnecessary amount of parameters to the model.

In addition to fitting actual signals, domain knowledge of various cardiac conditions was used to determine a family of parameters that describe a given cardiac condition, allowing for the synthesis of realistic signals. Examples are shown in Figure 4.8.

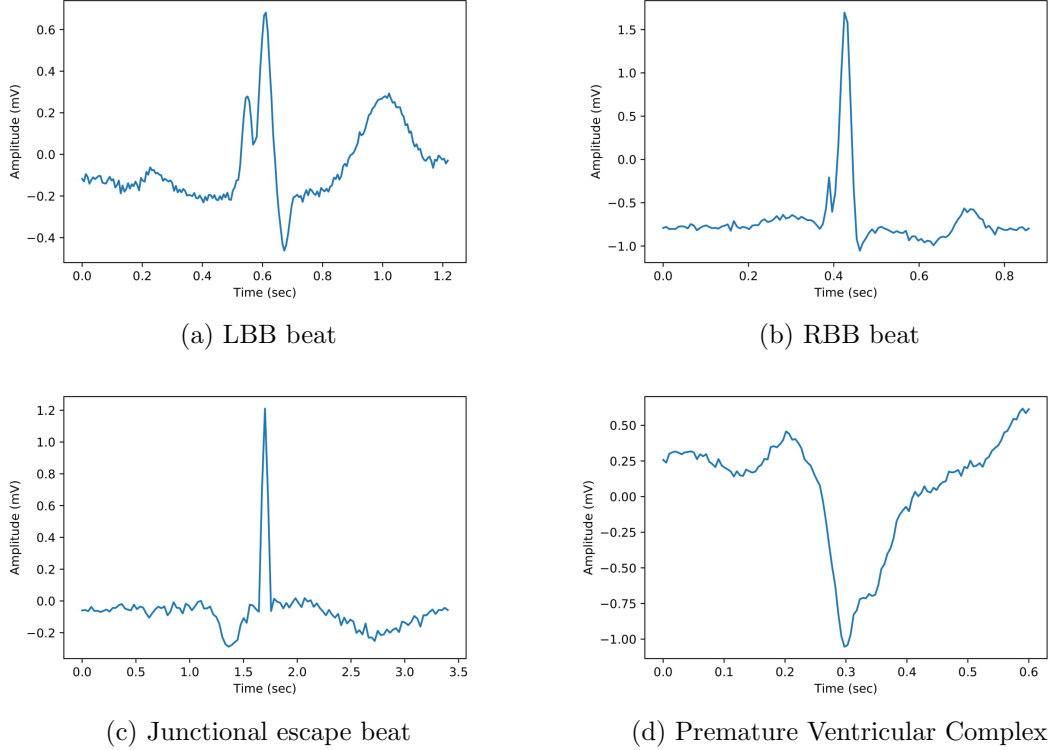


Figure 4.8: Synthetic cardiac conditions created with (4.8)

4.4.2 Conduction Model. The frequencies of the pacemakers are controlled by f_1 . Thus, in order to synchronize the oscillators as the rate of the SA Node changes, the coupling coefficients should be proportional to f_1 . In order to determine the value of f_1 for a given a heart rate, one must find its nonlinear dependence on the frequency f (bps). This can be fit with a binomial function:

$$f_1 = 9.098f^2 + 8.798f - 0.732, \quad (4.16)$$

as found in [20].

Unfortunately, this model was much more difficult to fit, due to the delay and the stiffness and number of the parameters. Besides creating normal signals and altering the heart rate as shown above, other cardiac conditions beyond bradycardia and tachycardia (Figure 4.9) could not be reproduced. Since our goal is to synthesize as many conditions as possible, the dynamical model on a circle was used for all synthesis throughout the remainder of the paper.




Figure 4.9: Bradycardia created with the extended conduction model

CHAPTER 5. ALTERNATIVE SYNTHESIS METHODS

When classification problems have an imbalance among the various classes, machine learning models are prone to classify samples as the most frequent classes. (In the case of heart beat classification, some arrhythmias are more common than others, and all arrhythmias are less common than normal sinus rhythm.) Since class imbalance is a common occurrence in machine learning problems, various methods have been devised to help, and a few include the generation of synthetic signals.

5.0.1 Data Perturbation. For the purposes of this paper, we will define perturbation to be the transformation of the data (signal) without additional data points. This type of augmentation is often used when classifying images. Images are rotated, cropped, dilated, compressed, shifted, and/or blurred to create “new” images. When performing these various transformations, it is important to make sure that the resulting image still has the same class as the original image. This is not generally a concern since no reasonable use of these simple operations is going to make the image of a dog look like a car.

However, when augmenting other types of data, such as ECG signals, one must be careful to preserve the characteristics of the class. By stretching a signal, a normal beat now looks more like

a sign of an irritable focus firing prematurely. Inverting a normal beat creates a signal that looks more like PVC, a rapid series of which would indicate ventricular tachycardia. As a result, only certain minor changes can be made to reliably augment signals.

In this paper, we employ the following perturbation process for a signal S of length n . Two draws L_1 and L_n from a normal distribution with mean 0 and standard deviation 0.1 are used to create endpoints for a sloped line L of length n . A number b is then selected uniformly from the interval $[-0.2, 0.2]$. The new signal S' is created by adding S and L element-wise and then adding b to each element. This results in a signal S' that has a different slope than the original as well as a vertical shift, as seen in Figure 5.1. Neither of these operations should affect the class of a given signal beat, since baseline wander is common in ECGs.




Figure 5.1: Data perturbation example with $L_1 = 0.08$, $L_n = -0.04$, $b = -0.4$

5.0.2 Synthetic Minority Over-sampling. One way to handle class imbalance is to simply under-sample the majority classes. This however, eliminates a lot of useful data, especially when some classes have very few samples. We would rather over-sample the minority classes. The naive approach is to simply copy examples of under-represented samples to improve class balance. This helps the model to not pick favorites but does not add novel data. As a result, it often has only a minor effect on model performance.

The Synthetic Minority Over-Sampling Technique (SMOTE) is a method that uses minority samples to generate similar synthetic samples. Rather than perturbing a given sample, SMOTE uses

similar samples to generate a new one. These synthetic examples are made by creating synthetic examples along the line segments between the original sample and some fraction of its k nearest neighbors. The fraction of segments chosen depends on the amount of over-sampling required, and the ones selected are chosen at random. Synthetic samples are then generated by taking the difference between a given sample and its nearest neighbor. This difference is then multiplied by a random number between 0 and 1 and then added to the original sample. This creates a new sample from a point along the line segment between two specific samples [22]. Figure 5.2 shows an example of the SMOTE process.

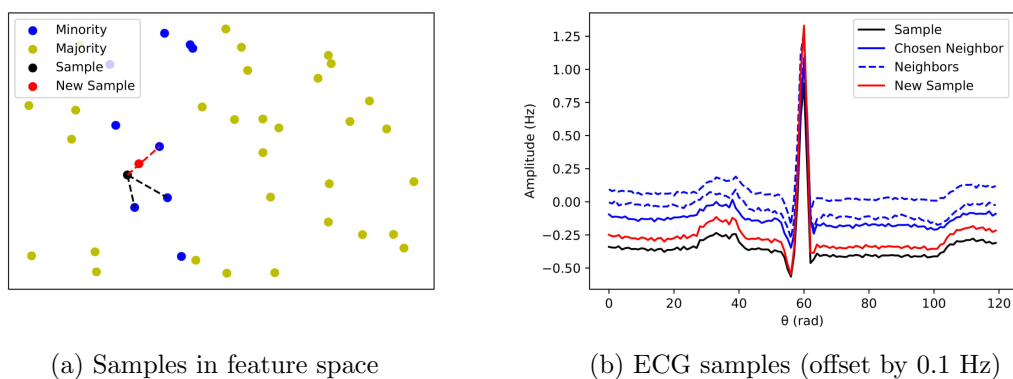


Figure 5.2: SMOTE example with $k = 3$

The synthetic samples generated through SMOTE are hopefully representative of other samples in the class since they are essentially a combination of two real, similar samples. With ECG signals, as long as they have similar neighbors, this technique should preserve the characteristics of the class.

Unlike synthetic models, these methods are not as well suited for generating long ECG signals. Single beats are easily perturbed, but the minor perturbations suggested above would be mostly unnoticed in longer signals. Additional baseline wander and noise could reasonably be added, and horizontal shifts would become valid, but the effects would be minimal. SMOTE poses more problems as neighbors become harder to define. A particular difficulty is the combining of two signals with differing heart rates. These methods would then be less effective for rhythm classification but should be suitable for beat classification. To help compare these various methods, perturbation, naive over-sampling, and SMOTE were all implemented.

CHAPTER 6. CLASSIFICATION OF ECG SIGNALS

Various machine learning algorithms have been used to classify ECG signals. The two main branches are beat classification and rhythm classification. This paper will explore beat classification primarily. In this section, three different algorithms will be explained and used with different versions of the same data set. Details about each version of the data will also be given.

6.1 DATA

The data used comes from the MIT-BIH Arrhythmia database. This open-source data set contains ECG recordings from 47 different individuals displaying various heart conditions. Approximately half of the samples are representative of signals found in routine clinical use. The other half were selected because they contained less common but clinically important examples of arrhythmias. Cardiologists independently classified each beat and annotated changes in rhythm for each of the 30 minute signals. They then consulted together to resolve any discrepancies. Through the years, minor corrections have also been made to this data set of over 100,000 beats. Due to its careful creation and a lack of similar available data sets, the MIT-BIH Arrhythmia database is commonly used in machine learning problems.

In order to prepare the data for beat classification, each single lead (MLII) signal was cut into segments centered around the R peak. Each sample was then interpolated to standardize the length of each signal. The corresponding label for each sample was taken directly from the 16 distinct MIT-BIH beat annotations. These beats include:

- Normal
- Left Bundle Branch block
- Right Bundle Branch block
- Atrial premature
- Aberrated atrial premature
- Junctional premature
- Supraventricular premature
- Premature ventricular contraction
- Fusion of ventricular and normal
- Ventricular flutter
- Atrial escape
- Junctional escape
- Ventricular escape
- Paced

- Fusion of paced and normal
- Non-conducted P-wave (blocked PAB)

Due to the very high volume of normal beats, only 10,000 normal samples were retained in the data. The rest of the classes had between 2 and 9,000 samples, with an average of 2,000 samples. The data was then divided into a training and validation set. For well represented classes, 20% of samples were reserved for validation. For the three classes with less than 100 samples, 30-50% were used for validation to more accurately measure the model's performance for those classes.

6.2 EXPERIMENTAL DESIGN

As previously mentioned, this data set is very imbalanced. The following experiments were conducted to compare various forms of synthetic data:

- (i) Baseline: No modification to the training data
- (ii) Naive: Naive random over-sampling to achieve class balance
- (iii) Perturb: Using the perturbation described in 5.0.1 to achieve class balance
- (iv) SMOTE: Using SMOTE to achieve class balance
- (v) Synthetic: Using synthetic data from the dynamical model on a circle to achieve class balance
- (vi) Synthetic/SMOTE: Using a 50/50 combination of synthetic/SMOTE to achieve class balance
- (vii) Pure Synthetic: Using a data set of synthetic data only

After splitting data into training and test groups, the normal beats were the majority class with 7,981 unique samples. Thus, when the methods above were used to achieve class balance, each class ended up with 7,981 samples, for a total of 127,696 training samples.

To evaluate the performance of these seven different scenarios, we looked at the accuracy and F1 score within each class. The F1 score for each class, was calculated with the formula: $2(\textit{precision} * \textit{recall}) / (\textit{precision} + \textit{recall})$. The weighted mean of these values then makes up the overall F1 score, whereas the unweighted mean is called the average class F1 score. For each model type, the confusion matrix of both the baseline experiment and a high performing additional experiment will be shown to visualize the performance of each class. Each will show the predicted labels versus the

true labels; hence, a perfect model would have values only along the main diagonal. The results in the tables and the matrices were averaged over three different data splits using different random seeds.

The three models used to test these experiments were XGBoost, a fully-connected neural network and a convolutional network. The structure and training method were held constant within each model type. The details of each model are given in the following sections. Results from the various experiments and models are compared and discussed in Chapter 7.

6.3 CLASSIFICATION

The goal of a classification problem is to take an input x and map it to one of c classes. In our scenario, each x is a vector containing the amplitude of a single EGG beat, which we hope to map to one of 16 possible classes. In this section, we explain three different classification methods and show their results.

6.3.1 Tree-based Methods. Tree-based methods, such as random forests have had success in a number of classification problems. A random forest is made up a group of decision trees. Each decision tree takes an input and follows a series of splits based on decision rules to end up in a final branch. Each tree in the forest votes as to the final classification of the input.

When training a random forest, each tree uses available parameters to see which split best separates the classes. The user decides how many splits each tree should have, how many trees should be in the forest, and how much to learn from each sample, (along with various other hyper-parameters).

Gradient boosted trees are trees that helps prevent over-fitting. They fit a model to the data, then fit a model to the residuals, and then combine them to create a better model. One particularly popular gradient boosted tree model is XGBoost. Its relative speed and performance make it a common model in research and application.

To train an XGBoost model to classify ECG heart beats, the model was simply given the training data set. The algorithm uses cross entropy loss to train. After some experimentation the classifier used here had a learning rate of 0.1, with 100 trees of depth 4. The results of the experiments can be seen below.

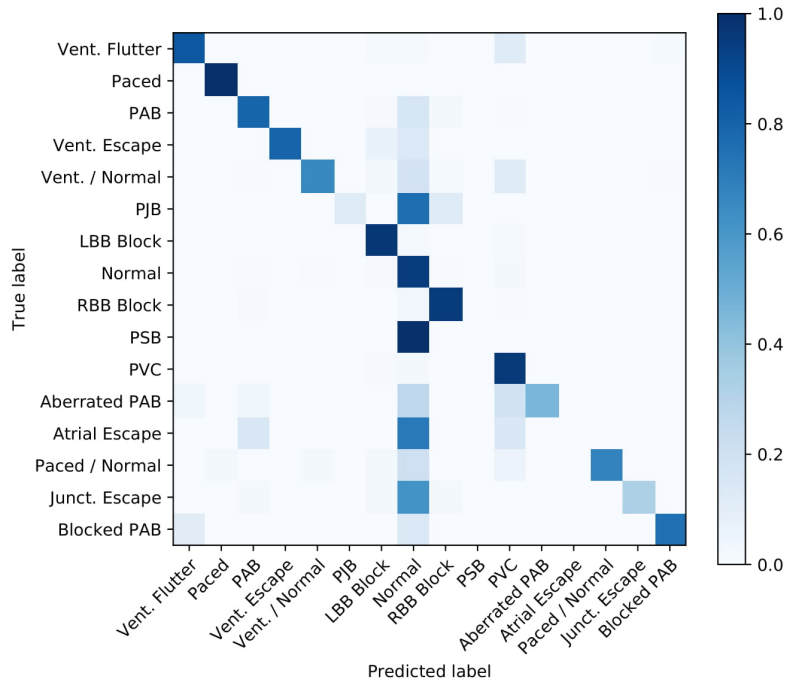


Figure 6.1: XGBoost baseline confusion matrix

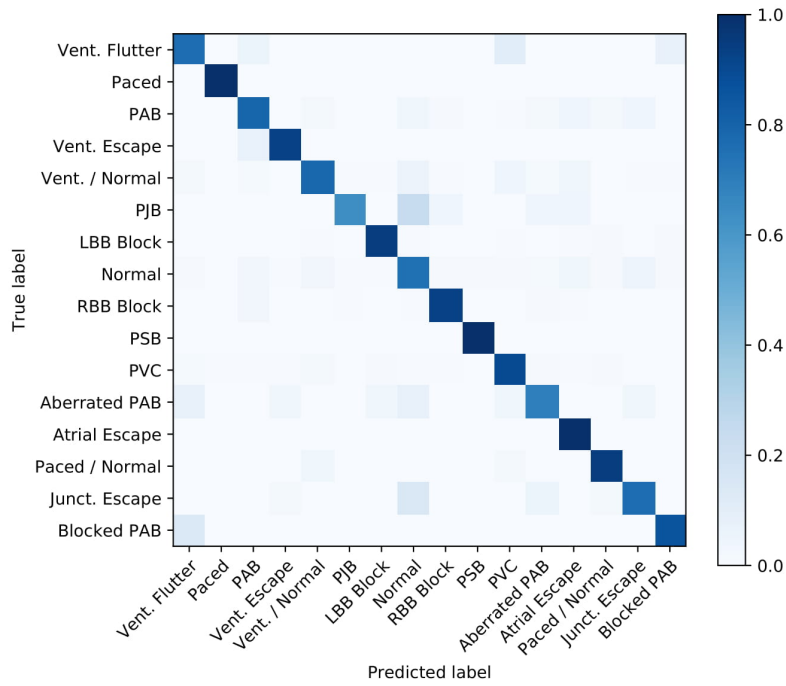


Figure 6.2: XGBoost synthetic/SMOTE confusion matrix

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM	Pure Syn
Vent. Flutter	0.845	0.893	0.833	0.893	0.690	0.762	0.714
Paced	0.999	0.997	0.994	0.999	0.993	0.994	0.982
PAB	0.793	0.853	0.774	0.822	0.787	0.793	0.795
Vent. Escape	0.800	0.867	0.933	0.867	0.933	0.933	0.933
Vent./Normal	0.656	0.800	0.838	0.806	0.794	0.788	0.800
PJB	0.120	0.760	0.880	0.800	0.360	0.640	0.560
LBB Block	0.970	0.962	0.952	0.962	0.938	0.942	0.934
Normal	0.951	0.831	0.758	0.819	0.764	0.752	0.672
RBB Block	0.956	0.942	0.921	0.933	0.929	0.930	0.912
PSB	0.000	0.000	0.000	0.000	1.000	1.000	1.000
PVC	0.958	0.928	0.899	0.919	0.894	0.897	0.850
Aberrated PAB	0.462	0.692	0.923	0.769	0.615	0.692	0.654
Atrial Escape	0.000	0.286	0.857	0.714	1.000	1.000	1.000
Paced/Normal	0.679	0.868	0.962	0.868	0.962	0.943	0.906
Junct. Escape	0.327	0.836	0.891	0.855	0.564	0.764	0.727
Blocked PAB	0.750	0.917	1.000	0.944	0.861	0.861	0.806
Avg. Class Accuracy	0.932	0.909	0.878	0.902	0.870	0.872	0.838
Overall Accuracy	0.642	0.777	0.838	0.811	0.818	0.856	0.828

Table 6.1: XGBoost accuracy results

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM	Pure Syn
Vent. Flutter	0.855	0.732	0.686	0.732	0.607	0.621	0.580
Paced	0.997	0.993	0.994	0.995	0.994	0.991	0.987
PAB	0.853	0.804	0.774	0.792	0.759	0.770	0.724
Vent. Escape	0.857	0.743	0.519	0.743	0.437	0.459	0.373
Vent./Normal	0.755	0.661	0.672	0.670	0.630	0.622	0.597
PJB	0.214	0.437	0.250	0.339	0.273	0.471	0.368
LBB Block	0.970	0.969	0.966	0.971	0.958	0.961	0.940
Normal	0.908	0.886	0.850	0.882	0.846	0.842	0.789
RBB Block	0.966	0.961	0.950	0.957	0.953	0.953	0.943
PSB	0.000	0.000	0.000	0.000	0.067	0.087	0.051
PVC	0.943	0.937	0.927	0.933	0.924	0.925	0.894
Aberrated PAB	0.585	0.444	0.304	0.374	0.276	0.295	0.264
Atrial Escape	0.000	0.444	0.167	0.526	0.095	0.095	0.082
Paced/Normal	0.791	0.742	0.689	0.676	0.607	0.606	0.565
Junct. Escape	0.486	0.467	0.366	0.433	0.307	0.362	0.280
Blocked PAB	0.761	0.717	0.692	0.739	0.496	0.559	0.563
Avg. Class F1 Score	0.933	0.928	0.932	0.930	0.921	0.925	0.908
Overall F1 Score	0.684	0.684	0.623	0.673	0.577	0.601	0.562

Table 6.2: XGBoost F1 score results

6.3.2 Neural Networks. Neural networks have been around for many years but have recently gained attention due to advances in computational power and various architectural advances. Graphic processing units (GPUs) make training large networks feasible. With increased width and depth, various other techniques make that training effective. A few methods will be covered in the following sub sections, though the networks used in this paper are not very complex. Thus a brief overview of neural networks along with their applications to this specific problem will be given here.

Neural Network Overview. The basic idea behind a neural network is fairly straightforward. The goal of a neural network (just like other machine learning tools) is to take in some input x and return some desired output y . The neural network is simply a function f such that $f(x) = y$. We begin by considering a simple example where x is a vector of size 2. Let $x = [x_1, x_2]^T$. First, we multiply each input by a weight:

$$\begin{aligned}x_1 &\rightarrow x_1 * w_1 \\x_2 &\rightarrow x_2 * w_2\end{aligned}\tag{6.1}$$

These weighted inputs are then added together with a bias b :

$$(w_1 * x_1) + (w_2 * x_2) + b\tag{6.2}$$

These two steps can be represented in a condensed form by creating a weight matrix $W = [w_1, w_2]^T$ such that

$$W^T x + b = w_1 x_1 + w_2 x_2 + b\tag{6.3}$$

as above.

We can let our function $f(x) = W^T x + b$. Given weights and a bias, we can map any x to its corresponding y . This function, however, will mostly likely not be robust enough for our needs. Often we want to stack these transformations. Suppose we have two weight matrices W_1 and W_2 and two biases b_1 and b_2 . Then composing these two transformations would yield

$$\begin{aligned}
g(x) &= W_2^T(W_1^T x + b_1) + b_2 \\
&= W_2^T W_1^T x + W_2^T b_1 + b_2 \\
&= W^T x + b
\end{aligned}
\tag{6.4}$$

for some $W = W_1 W_2$ and $b = W_2^T b_1 + b_2$. So even though we added another layer, it collapsed back onto itself because the composition of linear functions is linear.

Thus, in order to truly add depth, we need to add a non-linearity (or activation function) σ . Common options for this function include the sigmoid function, hyperbolic tangent, ReLU, and others. These are all functions that take numbers from $(-\infty, \infty)$ and map them to a smaller region. The one that typically works the best for other steps within the neural network process is ReLU, which is short for rectified linear unit. This function is defined to be

$$\sigma(x) = \begin{cases} 0, & x \leq 0 \\ x, & x > 0, \end{cases}
\tag{6.5}$$

which can also be thought of as $f(x) = \max(0, x)$.

By adding this non-linearity to our function f , we get that $f(x) = \sigma(W^T x + b)$. When composing functions of this type, the non-linearity prevents the collapse seen in 6.4. We can now have as many layers as we want. Layers between the input layer and output layer are known as hidden layers. Networks with more layers are known as deep and, ones with long hidden layers (many neurons or nodes) are considered wide.




Figure 6.3: A simple neural network

The network displayed in Figure 6.3 has an input of size 2. This then gets mapped to a hidden layer with four neurons before getting mapped to a final output. This means our first weight matrix would be a 2x4 matrix so that $h = W^T x$ would have dimension 4. By varying the dimensions of the weight matrices, one can change the number of neurons in the subsequent layer.

Now, in order to have a useful network, we need the weights to help transform our input into our desired output. This requires us to train the neural network to learn this mapping. In order to evaluate how well our model is doing, we need an objective or loss function to compare our predicted output with the desired output. We train our network by minimizing this loss function over many iterations.

Categorical cross entropy loss was used as the function to compare our predicted probability distribution with the true class distribution. Categorical cross entropy loss is defined as:

$$L = - \sum_{c=0}^{15} y_{o,c} \log(p_{o,c}), \quad (6.6)$$

where $y_{o,c}$ is a binary indicator of whether or not class label c is correct for observation o and $p_{o,c}$ is the probability that observation o is of class c .

Suppose we randomly initialize our weights and biases. Then we can feed our inputs into our network and get our predicted outputs and calculate the loss. By finding the gradient of the loss with respect to our various weights and biases, we can determine which weight we should change. In order to find the partial derivative of the loss with respect to a certain weight w_i , we use the chain rule to get

$$\frac{\partial L}{\partial w_i} = \frac{\partial L}{\partial \hat{y}} * \frac{\partial \hat{y}}{\partial h_i} * \frac{\partial h_i}{\partial w_i}, \quad (6.7)$$

where \hat{y} is the predicted output. This use of the chain rule to calculate derivatives through the layers of the network is known as back propagation. Then using stochastic gradient descent (SGD), we update our weights and biases

$$w_i \leftarrow w_i - \eta \frac{\partial L}{\partial w_i} \quad (6.8)$$

so that our loss will improve the next time we stick in that input or set of inputs (known as a

batch). The value η is a learning rate, which determines how large of a step to take in that descent direction. Once all of the samples have been passed through once, this is known as an epoch. Networks may be trained anywhere from five to hundreds of epochs, depending on the network structure and the data itself.

The Adam optimizer, an alternate to classic stochastic gradient descent, was used to minimize the loss over many training epochs. While still following the basic principles behind SGD, it makes use of individualized learning rates per parameter. To adapt these learning rates over time, it uses the exponential moving averages of both the first and second moments of the gradients with varying levels of decay over time to update each learning rate.

Fully Connected Network. The basic structure in the previous section is known as a fully connected or dense network (DNN). As depicted in Figure 6.3, each node is connected to every node of the following layer. In practice, a few of these connections may be randomly dropped to prevent the network from relying too heavily on any given node. This is a natural regularization technique known as dropout. Even with dropout this network architecture is still known as a dense network.

In order to use a DNN for classifying ECG beats, we used the architecture shown in Figure 6.4 starting with a hidden layer of size 256.

Note that we use the softmax activation at the end. This function is given by

$$\sigma(s)_i = \frac{e^{s_i}}{\sum_j e^{s_j}}, \quad (6.9)$$

and is used to make the final output a probability array, where each entry in the array is the probability of the sample belonging to that class. We may then choose our predicted label by picking the class with the highest probability. The results of the baseline model can be seen in Figure 6.5. Tables 6.3 and 6.4 show comparisons of the various experiments.

Convolutional Neural Network. Convolutional networks use a series of convolutional filters instead of fully connected layers. Consider a 32x32 black and white image. We can convolve a 5x5 filter with the image, by essentially sliding the filter across the image and computing dot products. This will give us a new 28x28 image. We can use another filter to give us another 28x28 image. If we use n 5x5 filters, we'll create a new "image" of size 28x28xn. To avoid the change in size, the




Figure 6.4: DNN architecture




Figure 6.5: DNN baseline confusion matrix

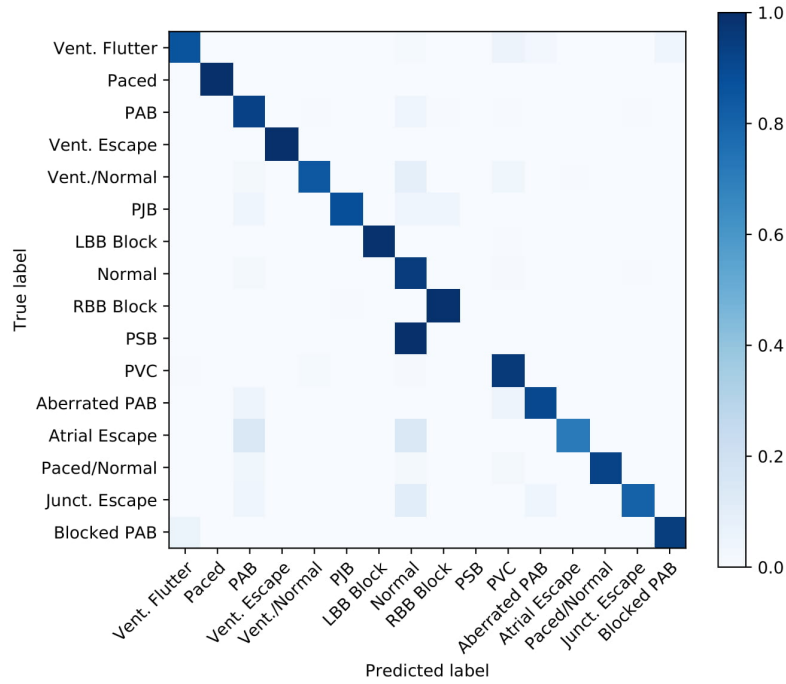


Figure 6.6: DNN SMOTE confusion matrix

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM.	Pure Syn
Vent. Flutter	0.833	0.889	0.903	0.861	0.847	0.861	0.806
Paced	0.999	0.999	0.999	0.999	0.999	0.999	0.997
PAB	0.903	0.927	0.933	0.931	0.893	0.915	0.881
Vent. Escape	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Vent./Normal	0.813	0.858	0.884	0.845	0.826	0.826	0.845
PJB	0.600	0.880	0.920	0.880	0.960	0.920	0.920
LBB Block	0.988	0.987	0.988	0.991	0.988	0.985	0.979
Normal	0.974	0.958	0.961	0.953	0.943	0.938	0.897
RBB Block	0.995	0.988	0.987	0.988	0.990	0.990	0.957
PSB	0.000	0.000	0.000	0.000	1.000	1.000	1.000
PVC	0.967	0.958	0.958	0.959	0.958	0.958	0.886
Aberrated PAB	0.762	0.905	0.905	0.905	0.810	0.905	0.810
Atrial Escape	0.000	0.571	0.714	0.714	1.000	1.000	1.000
Paced/Normal	0.868	0.925	0.981	0.925	1.000	0.981	1.000
Junct. Escape	0.660	0.723	0.787	0.809	0.787	0.766	0.787
Blocked PAB	0.946	0.946	0.973	0.946	0.946	0.946	0.946
Avg. Class Accuracy	0.968	0.966	0.969	0.966	0.961	0.961	0.929
Overall Accuracy	0.769	0.845	0.868	0.857	0.934	0.937	0.919

Table 6.3: DNN accuracy results

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM.	Pure Syn
Vent. Flutter	0.816	0.853	0.839	0.849	0.813	0.844	0.686
Paced	0.999	0.999	0.999	0.999	0.998	0.999	0.998
PAB	0.919	0.915	0.915	0.916	0.917	0.915	0.870
Vent. Escape	0.978	0.957	0.957	0.978	0.786	0.863	0.698
Vent./Normal	0.860	0.821	0.825	0.826	0.828	0.823	0.710
PJB	0.732	0.698	0.697	0.772	0.828	0.821	0.582
LBB Block	0.991	0.991	0.993	0.992	0.992	0.990	0.983
Normal	0.960	0.961	0.968	0.958	0.955	0.953	0.931
RBB Block	0.992	0.991	0.992	0.991	0.990	0.991	0.970
PSB	0.000	0.000	0.000	0.000	0.500	0.667	0.500
PVC	0.968	0.965	0.965	0.962	0.964	0.961	0.925
Aberrated PAB	0.800	0.864	0.809	0.745	0.531	0.613	0.386
Atrial Escape	0.000	0.615	0.833	0.714	0.438	0.500	0.230
Paced/Normal	0.911	0.916	0.972	0.933	0.972	0.929	0.841
Junct. Escape	0.667	0.716	0.747	0.784	0.607	0.605	0.525
Blocked PAB	0.909	0.921	0.911	0.921	0.875	0.875	0.795
Avg. Class F1 Score	0.967	0.967	0.971	0.967	0.967	0.965	0.949
Overall F1 Score	0.781	0.824	0.839	0.834	0.812	0.834	0.727

Table 6.4: DNN F1 score results

original image may be padded on all sides, in this case by 2 pixels.

If we want to stack multiple convolution layers, we use non-linear activations like ReLU in between the layers. As we move throughout the layers, the model learns higher and higher level features. This becomes very useful in image classification problems, as the model can identify subfeatures to help classify the overall image. The advantage here with this approach is that the exact angle or location of the object within the image does not matter.

While some use a 2D image for classifying ECG signals, our signals are just one dimensional vectors. In this case, the convolutional filter is also one dimensional and slides across the signal horizontally. To classify ECG beats, each signal was fed into a series of one dimensional convolutional layers with kernel size 3, along with max-pooling and batch normalization. Max-pooling is a method for down-sizing the temporal dimension by taking the max value of every n samples. Batch normalization is used to normalize the intermediate values.

After the convolutional layers (starting with 256 filters and decreasing to 16), a final dense layer with a softmax activation was used to output a classification probability for each of the 16 classes. The details of the model architecture are shown in Figure 6.7.

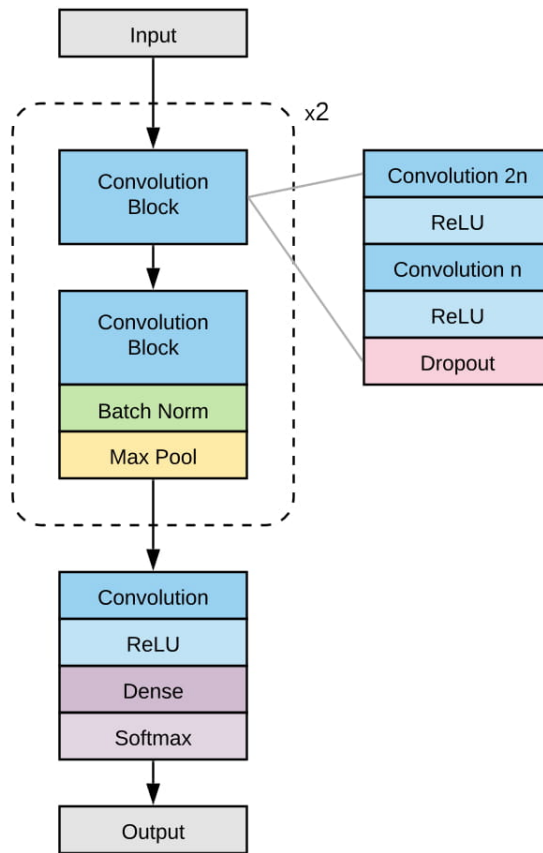


Figure 6.7: Convolutional neural network architecture




Figure 6.8: DNN baseline confusion matrix




Figure 6.9: DNN synthetic/SMOTE confusion matrix

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM	Pure Syn
Vent. Flutter	0.778	0.931	0.944	0.917	0.764	0.819	0.736
Paced	0.997	0.999	0.999	0.999	0.996	0.992	0.989
PAB	0.893	0.915	0.915	0.905	0.841	0.857	0.873
Vent. Escape	1.000	1.000	1.000	1.000	0.818	1.000	1.000
Vent./Normal	0.710	0.865	0.916	0.839	0.729	0.716	0.813
PJB	0.040	0.840	0.880	0.880	0.720	0.840	0.720
LBB Block	0.985	0.971	0.977	0.976	0.973	0.977	0.953
Normal	0.965	0.897	0.898	0.894	0.943	0.928	0.764
RBB Block	0.959	0.974	0.975	0.970	0.976	0.969	0.901
PSB	0.000	0.000	0.000	0.000	1.000	1.000	1.000
PVC	0.937	0.910	0.911	0.913	0.919	0.928	0.758
Aberrated PAB	0.429	0.810	0.905	0.857	0.619	0.714	0.619
Atrial Escape	0.000	0.571	0.429	0.429	0.143	0.286	1.000
Paced/Normal	0.887	0.981	0.943	0.981	0.925	0.925	0.943
Junct. Escape	0.596	0.809	0.809	0.809	0.340	0.745	0.617
Blocked PAB	0.703	0.946	0.946	0.919	0.811	0.838	0.919
Avg. Class Accuracy	0.946	0.937	0.940	0.936	0.937	0.939	0.854
Overall Accuracy	0.680	0.838	0.840	0.830	0.782	0.850	0.846

Table 6.5: CNN accuracy results

	Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM	Pure Syn
Vent. Flutter	0.718	0.766	0.795	0.759	0.651	0.690	0.444
Paced	0.999	0.999	0.999	0.998	0.998	0.996	0.989
PAB	0.856	0.842	0.840	0.847	0.877	0.862	0.765
Vent. Escape	0.978	0.898	0.846	0.786	0.766	0.800	0.389
Vent./Normal	0.797	0.786	0.807	0.772	0.774	0.782	0.687
PJB	0.069	0.677	0.677	0.611	0.720	0.824	0.414
LBB Block	0.986	0.979	0.981	0.983	0.977	0.982	0.966
Normal	0.936	0.927	0.930	0.924	0.921	0.926	0.849
RBB Block	0.975	0.983	0.983	0.981	0.982	0.980	0.944
PSB	0.000	0.000	0.000	0.000	0.667	0.667	0.074
PVC	0.951	0.928	0.933	0.935	0.927	0.936	0.837
Aberrated PAB	0.529	0.523	0.521	0.500	0.481	0.423	0.144
Atrial Escape	0.000	0.571	0.316	0.231	0.250	0.400	0.109
Paced/Normal	0.913	0.945	0.952	0.929	0.961	0.907	0.714
Junct. Escape	0.667	0.628	0.661	0.589	0.400	0.543	0.284
Blocked PAB	0.712	0.722	0.886	0.764	0.741	0.838	0.673
Avg. Class F1 Score	0.945	0.946	0.949	0.947	0.939	0.945	0.921
Overall F1 Score	0.693	0.761	0.758	0.726	0.756	0.785	0.580

Table 6.6: CNN F1 score results

6.3.3 Additional Benefits of Model Synthetic Data. The use of synthetic signals from a model is, in essence, another form of data augmentation, in that it adds the data available. Unlike the traditional data augmentation however, it does not rely on transforming real training data. Domain knowledge can be used to generate synthetic signals that are realistic enough to classify actual signals as seen in the results of the various experiments and in Figure 6.10.

Furthermore, the idea of synthetic modeling can be used for a more intelligent or controlled data augmentation. As previously mentioned, a synthetic signal may be fitted to a real training signal using a nonlinear least squares solver to obtain a parameterization of the real signal. Once those parameters are obtained, any subset of them may be modified to generate new signals. For example, a specific peak’s amplitude, width or location could be altered slightly in a way that doesn’t fundamentally change the classification but that does create a new signal. Varying amounts of noise may also be introduced into the signal. It provides a way to alter just a part of the signal in reasonable and useful ways.

Another use of the synthetic models comes from the parameters used to generate the signal. As mentioned previously, we may use a nonlinear least squares solver to find the Gaussian parameters than would replicate a given signal. Assuming we use seven Gaussians to construct our signal, this

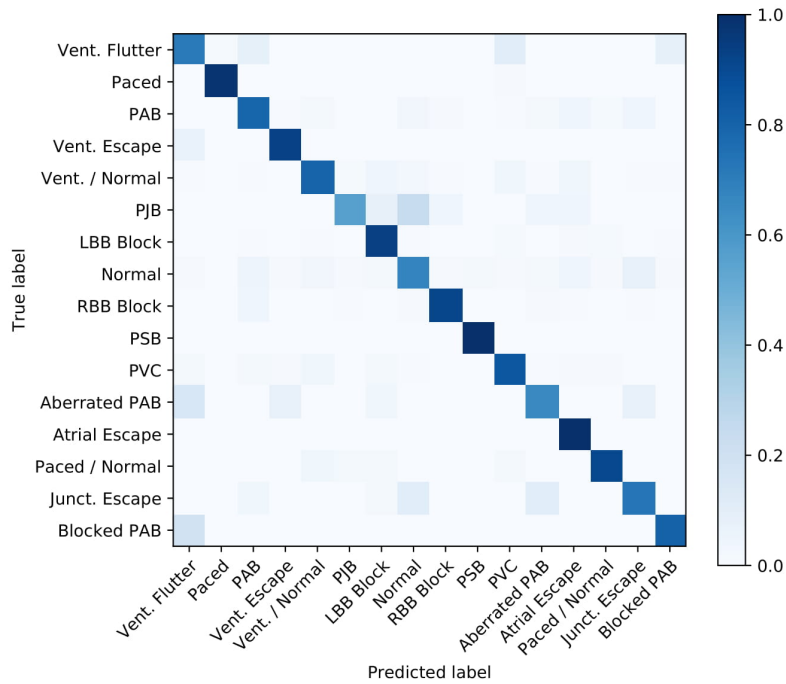


Figure 6.10: DNN Synthetic model confusion matrix

would convert our signal of length n into a vector of 22 parameters.

Essentially, we are breaking down the signal into component parts, like with principal component analysis. But by forcing this specific transformation, the parameterization is a standard transformation that could be used with any data set. These parameters may then be used instead for classification purposes and could even help to more clearly define the criteria for a certain cardiac condition.

To test this idea, SciPy’s nonlinear least squares solver was used to parameterize all the signals. The solver accepts a function to be optimized. In this case we minimized the difference between the original signal and a synthetic one created from the dynamical circle model. Other arguments include an initial estimate for the Gaussian parameters. These were found through domain knowledge as well as examining a few examples of each class. Bounds were also given to ensure that the Gaussians stay in order and to avoid extreme values from over-fitting. All the signals were put into the solver to get a 22 parameter representation of each. This was then used as the input for an XGBoost model with depth 3 and a DNN with structure shown in Figure 6.11.

The confusion matrix for the best model is shown in Figure 6.12. We can see that these smaller models perform relatively well, although not quite as well as the original models. This suggests

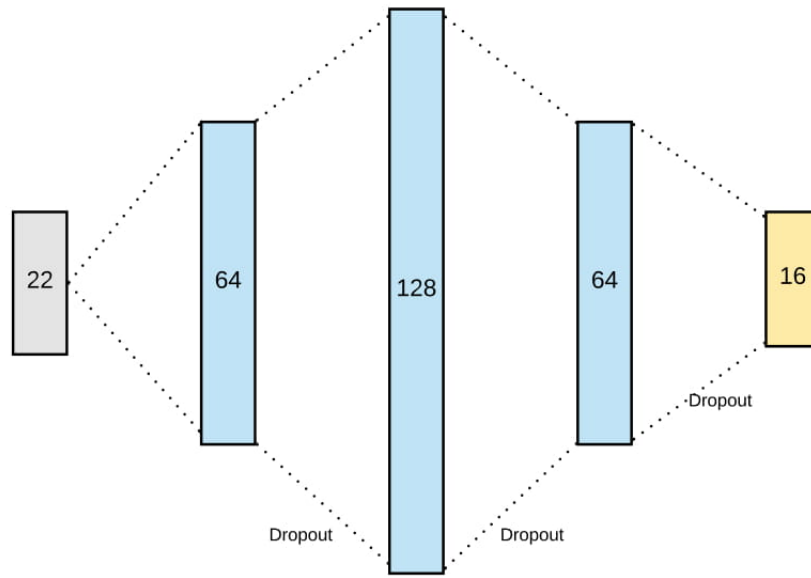


Figure 6.11: Parameterized DNN architecture




Figure 6.12: Parameterized DNN SMOTE confusion matrix

that this parameterization may have diagnostic value. To improve results, more care could be taken with the regularization, bounds, peak locations, etc. in order to improve and standardize the parameterization.

CHAPTER 7. CONCLUSION

In Table 7.1, we see that the best use of synthetic data varies from model to model and the metric in question. The baseline data seemed to work the best for XGBoost, whereas perturbation worked with the DNN. The synthetic/SMOTE combination worked well for the CNN and in the other models as well.

		Baseline	Naive	Perturb	SMOTE	Synthetic	Syn/SM	Pure Syn
XGB	Class Acc	0.932	0.909	0.878	0.902	0.870	0.872	0.838
	Class F1	0.933	0.928	0.932	0.930	0.921	0.925	0.908
	Overall Acc	0.642	0.777	0.838	0.811	0.818	0.856	0.828
	Overall F1	0.684	0.684	0.623	0.673	0.577	0.601	0.562
DNN	Class Acc	0.968	0.966	0.969	0.966	0.961	0.961	0.929
	Class F1	0.967	0.967	0.971	0.967	0.967	0.965	0.949
	Overall Acc	0.769	0.845	0.868	0.857	0.934	0.937	0.919
	Overall F1	0.781	0.824	0.839	0.834	0.812	0.834	0.727
CNN	Class Acc	0.946	0.937	0.940	0.936	0.937	0.939	0.854
	Class F1	0.945	0.946	0.949	0.947	0.939	0.945	0.921
	Overall Acc	0.680	0.838	0.840	0.830	0.782	0.850	0.846
	Overall F1	0.693	0.761	0.758	0.726	0.756	0.785	0.580

Table 7.1: Experiment Comparisons

As demonstrated by these various experiments, we conclude that the use of synthetic signals aids in the classification of various types of heart beats in many cases. Due to various challenges in collecting and augmenting ECG signals, this could be a very useful way to improve current classification models. Improvements were seen in the XGBoost model, as well as the fully-connected and convolutional neural networks, suggesting that the synthetic data may have impact on classification regardless of the model type.

Though the dynamical model on a circle has little physiological backing, it can approximate signals well enough to improve accuracy and precision of classification. By essentially controlling the shape and location of each peak, synthetic signals can be tailored to specific needs, especially cases of class imbalance. We also saw that the model parameters themselves could be used for

classification.

Possible areas further areas of use include multi-lead beat classification, rhythm classification, using new and/or additional ECG data sources, and classification of other waveform signals altogether. However, the experiments performed do suggest that synthetic signals may be used to approximate and augment various types of real signals for use in training classification models.

BIBLIOGRAPHY

- [1] Dale Dubin. *Rapid Interpretation of EKG's*. Cover Pub Co, 6th edition, October 2000.
- [2] Clinical Skills Education LLC. Learn EKG interpretation. <http://ekg.academy>, Accessed 28 February 2019.
- [3] C. Vimal and B. Sathish. Random forest classifier based ecg arrhythmia classification. *International Journal of Healthcare Information Systems and Informatics*, 5(2), 2010.
- [4] Nahit Emanet. Ecg beat classification by using discrete wavelet transform and random forest algorithm. *2009 Fifth International Conference on Soft Computing, Computing with Words and Perceptions in System Analysis, Decision and Control*, September 2009.
- [5] Shweta H. Jambukia, Vipul K. Dabhi, and Harshadkumar B. Prajapati. Classification of ecg signals using machine learning techniques: A survey. *2015 International Conference on Advances in Computer Engineering and Applications*, March 2015.
- [6] Tae Joon Jun, Hoang Minh Nguyen, Daeyoun Kang, Dohyeun Kim, Daeyoung Kim, and Young-Hak Kim. Ecg arrhythmia classification using a 2-d convolutional neural network. *arxiv.org*, 2018.
- [7] Pranav Rajpurkar, Awni Y. Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y. Ng. Cardiologist-level arrhythmia detection with convolutional neural networks. *Nature Medicine*, 25(1), January 2019.
- [8] Patrick E. McSharry and Gari D. Clifford. Open-source software for generating electrocardiogram signals. *Biomedical Engineering*, 458, February 2005.
- [9] Ognian Boumbarov, Yuliyana Velchev, and Strahil Sokolov. Ecg personal identification in subspaces using radial basis neural networks. *2009 IEEE International Workshop on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications*, September 2009.
- [10] Elizabeth Ann Maharaj and Andrs M. Alonso. Discriminant analysis of multivariate time series: Application to diagnosis based on ecg signals. *Computational Statistics & Data Analysis*, 70:67–87, February 2014.
- [11] Reza Sameni, Gari D. Clifford, Christian Jutten, and Mohammad B. Shamsollahi. Multi-channel ECG and noise modeling: Application to maternal and fetal ECG signals. *EURASIP Journal on Advances in Signal Processing*, 2007(1):94–107, December 2007.
- [12] E. Ryzhii and M. Ryzhii. A heterogeneous coupled oscillator model for simulation of ecg signals. *Computer Methods and Programs in Biomedicine*, 117(1):40–49, October 2014.
- [13] M. Ryzhii E. Ryzhii. Modeling of heartbeat dynamics with a system of coupled nonlinear oscillators. *Communications in Computer and Information Science*, 404:157–165, 2013.
- [14] Patrick E. McSharry, Gari D. Clifford, Lionel Tarassenko, and Leonard A. Smith. A dynamical model for generating synthetic electrocardiogram signals. *IEEE Transactions on Biomedical Engineering*, 50(3):289–294, March 2003.

- [15] Reza Sameni, Mohammad B. Shamsollahi, Christian Jutten, and Gari D. Clifford. A non-linear Bayesian filtering framework for ECG denoising. *IEEE Transactions on Biomedical Engineering*, 54(12):2172–2185, November 2007.
- [16] R. Sameni, M. B. Shamsollahi, C. Jutten, and M. Babaie-Zade. Filtering noisy ECG signals using the extended Kalman filter based on a modified dynamic ECG model. *Computers in Cardiology*, pages 1017–1020, September 2005.
- [17] G. D. Clifford, A. Shoeb, P. E. McSharry, and B. A. Janz. Model-based filtering, compression and classification of the ECG. *International Journal of Bioelectricmagnetism*, 7(1), 2005.
- [18] J. Ben-Arie and K. R. Rao. Nonorthogonal signal representation by Gaussians and Gabor functions. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 42(6):402–413, June 1995.
- [19] Sandra R.F.S.M. Gois and Marcelo A.Savi. An analysis of heart rhythm dynamics using a three-coupled oscillator model. *Chaos, Solitons & Fractals*, 41(5):2553–2565, September 2009.
- [20] Mario A. Quiroz-Jurez, Omar Jimnez-Ramrez, Rubn Vzquez-Medina, Elena Ryzhii, Maxim Ryzhii, and Jos L. Aragn. Cardiac conduction model for generating 12 lead ecg signals with realistic heart rate dynamics. *IEEE Transactions on NanoBioscience*, 17(4):525–532, October 2018.
- [21] J. Nagumo, S. Arimoto, and S. Yoshizawa. An active pulse transmission line simulating nerve axon. *Proceedings of the IRE*, 50(10):2061–2070, October 1962.
- [22] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. Smote: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 26, June 2002.