UNIVERSITY OF ESSEX

Deep Learning Approaches for ECG-Based Detection and Diagnosis of Coronary Artery Disease

by

Atitaya Phoemsuk

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

School of Computer Science and Electronic Engineering

February 2026

Declaration of Authorship

I, Atitaya Phoemsuk, declare that this thesis titled, 'Deep Learning Approaches for ECG-Based Detection and Diagnosis of Coronary Artery Disease' and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed: อาทิตยา เพ็งสุข

Date: 17/02/2026

"PhD: Proof that one can survive on caffeine, panic, and stubbornness."

— Atitaya, the author

Abstract

Coronary artery disease (CAD) is among the most prevalent and life-threatening cardiovascular conditions worldwide. Early detection is essential for improving patient outcomes and improving the efficiency of healthcare systems. Electrocardiography (ECG) is widely used for assessing cardiac function, yet manual interpretation of ECG signals can be inconsistent and prone to error. Developing reliable automated methods is therefore of great importance for enabling earlier and more consistent CAD detection. Although recent advances in deep learning have achieved strong performance in ECG analysis, many existing methods remain limited in practice. Current state-of-the-art models are often computationally complex and thus unsuitable for deployment on resource-constrained platforms.

To address the limited attention given to CAD in current ECG research, various deep learning models were developed in this thesis. First, a one-dimensional convolutional neural network (1D-CNN) was proposed to detect CAD directly from raw ECG signals without manual feature extraction. The model achieved 97.3% accuracy and demonstrated strong generalisability when using 250-point ECG segments. Then, a feature engineering approach was applied to select high-quality signal segments using sample entropy and normalisation techniques, further improving both accuracy and robustness. Next, a lightweight neural network architecture (CADNet) was developed, outperforming existing lightweight models by offering lower complexity and smaller size without compromising accuracy. The fourth study focused on 12-lead ECG, introducing a depth-wise, squeeze-and-excitation-based model that captured both lead-specific and inter-lead patterns, and achieved efficient deployment on an STM32 microcontroller. Finally, an attention-driven model was proposed for the detection of multiple cardiovascular diseases from a single ECG recording, demonstrating high diagnostic capability. Our qualitative evaluations in this study demonstrate that lightweight deep learning models can provide reliable ECG-based CAD diagnosis while remaining suitable for real-time deployment. The proposed approaches are robust to ECG variability and show consistent performance across different experimental and diagnostic scenarios, supporting their practical use in resource-constrained environments. Overall, the results highlight the relevance of lightweight deep learning architectures in enabling ECG-based diagnostic methods to be used as potential pre-screening tools.

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Dr Vahid Abolghasemi, for his unwavering support, encouragement, and belief in me throughout every step of my PhD journey. His mentorship has guided my academic path and inspired me to grow as both a researcher and an individual. Thanks to his support, I have had the privilege of presenting my work on international stages and publishing with IEEE. These are opportunities I never imagined when I began. I will carry his wisdom and kindness with me always, in both professional and personal aspects of life.

To my mother, Miss Samlee Noomsri, words will never fully capture how grateful I am. Her constant love, strength, and sacrifice, both physical and financial, have been the foundation of everything I have achieved. Her faith in me, even during the most difficult times, gave me the courage to keep going.

To my father, Mr Direk Phoemsuk, whose strength in living with heart disease has inspired this research. Seeing your journey gave me the motivation to find ways technology could help others in a similar situation. I know it is not easy to face each day with this condition. Without your support and example, this work may not have happened. This thesis is for you, and I hope it can help others in the future.

I am also thankful to my partner, Mr Alfie Brown, and his family, whose care, patience, and presence have been a source of comfort and stability. Their support has meant the world to me, especially during challenging moments. I am endlessly grateful for their kindness and for standing by me in both life and learning.

To the rest of my family and friends, for your constant support in so many ways. Your love, patience, and encouragement gave me strength during the most difficult moments of this journey. I am truly grateful for everything you have done to help me reach this point.

Contents

List of Publications

The following publications contribute to the originality of the work presented in this study:

Journal Papers:

1. A. Phoemsuk and V. Abolghasemi, "CADNet: A lightweight Neural Network for Coronary Artery Disease Classification Using Electrocardiogram Signals," *IEEE Journal of Biomedical and Health Informatics*, doi: 10.1109/JBHI.2025.3582872.
2. A. Phoemsuk and V. Abolghasemi, "Enhanced Coronary Artery Disease Classification Through Feature Engineering and One-Dimensional Convolutional Neural Network," *IEEE Access*, vol. 13, pp. 114306-114317, 2025, doi: 10.1109/ACCESS.2025.3584735.
3. A. Phoemsuk and V. Abolghasemi, "Real-time Coronary Artery Disease Detection from 12-Lead ECG Using A Lightweight Deep Network," *Biomedical Signal Processing and Control*. (Under review)

Conference Papers:

1. A. Phoemsuk and V. Abolghasemi, "Coronary Artery Disease Classification Using One-dimensional Convolutional Neural Network," *2024 IEEE Conference on Artificial Intelligence (CAI)*, Singapore, Singapore, 2024, pp. 389-394,

   doi: 10.1109/CAI59869.2024.00078.
2. A. Phoemsuk and V. Abolghasemi, "Multi-Disease Cardiovascular Detection from ECG Signals Using an Attention-Driven Deep Network," *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, Copenhagen, Denmark, 2025.

List of Figures

List of Tables

Abbreviations

AF	**A**trial **F**ibrillation
AI	**A**rtificial **I**ntelligence
AMI	**A**cute **M**yocardial **I**nfarction
ANOVA	**A**nalysis **o**f **V**ariance
ASY	**A**systole
AUC	**A**rea **U**nder the **C**urve
BCE	**B**inary **C**ross-**E**ntropy
BS	**B**atch **S**ize
CAD	**C**oronary **A**rtery **D**isease
CAM	**C**lass **A**ctivation **M**ap
CCDD	**C**hina **C**ardiovascular **D**isease **D**atabase
CHF	**C**ongestive **H**eart **F**ailure
CNN	**C**onvolutional **N**eural **N**etwork
CSE	**C**ommon **S**tandards for **E**lectrocardiography
CVD	**C**ardiovascular **D**iseases
DBN	**D**eep **B**elief **N**etwork
DCNN	**D**eep **C**onvolutional **N**eural **N**etwork
DINN	**D**ual **I**nput **N**eural **N**etwork
DWT	**D**iscrete **W**avelet **T**ransform
ECHO	**E**chocardiography
ECG	**E**lectrocardiogram
FE	**F**eature **E**ngineering
FLOPs	**F**loating-**P**oint **O**perations
FP	**F**alse **P**ositive
FPGA	**F**ield-**P**rogrammable **G**ate **A**rray

FPU	**F**loating **P**oint **U**nit
FN	**F**alse **N**egative
FB	**F**usion **B**eat
GPU	**G**raphics **P**rocessing **U**nit
HGEA	**H**igh **G**rade **V**entricular **E**ctopic **A**ctivity
HRV	**H**eart **R**ate **V**ariability
HSMM	**H**idden **S**emi-**M**arkov **M**odel
IMI	**I**nferior **M**yocardial **I**nfarction
ICA	**I**ndependent **C**omponent **A**nalysis
KNN	**K**-**N**earest **N**eighbours
LBBB	**L**eft **B**undle **B**ranch **B**lock
LMI	**L**ateral **M**yocardial **I**nfarction
LSTM	**L**ong **S**hort-**T**erm **M**emory
LVH	**L**eft **V**entricular **H**ypertrophy
MAC	**M**ultiply-**A**ccumulate
MCU	**M**icro**c**ontroller **U**nit
MI	**M**yocardial **I**nfarction
MPI	**M**yocardial **P**erfusion **I**maging
MRI	**M**agnetic **R**esonance **I**maging
NLP	**N**atural **L**anguage **P**rocessing
PAC	**P**remature **A**trial **C**ontraction
PCA	**P**rincipal **C**omponent **A**nalysis
PR	**P**aced **R**hythm
PTB	**P**hysikalisch-**T**echnische **B**undesanstalt
PVC	**P**remature **V**entricular **C**ontraction
Q	**U**nknown **B**eat
RBBB	**R**ight **B**undle **B**ranch **B**lock
RF	**R**andom **F**orest
ReLU	**R**ectified **L**inear **U**nit
ResNet	**R**esidual **N**eural **N**etwork
RNN	**R**ecurrent **N**eural **N**etwork
SampEn	**S**ample **E**ntropy
SCP	**S**tandardised **C**ardiac **P**atient

SE	**S**queeze-**a**nd-**E**xcitation
SHAP	**SH**apley **A**dditive **E**xPlanations
SPECT	**S**ingle-**P**hoton **E**mission **C**omputed **T**omography
SVEB	**S**upraventricular **E**ctopic **B**eat
SVM	**S**upport **V**ector **M**achine
t-SNE	**t**-**D**istributed **S**tochastic **N**eighbour **E**mbedding
TN	**T**rue **N**egative
TP	**T**rue **P**ositive
VB	**V**entricular **B**igeminy
VEB	**V**entricular **E**ctopic **B**eat
VF	**V**entricular **F**ibrillation
VFL	**V**entricular **F**lutter
VGAE	**V**ariational **G**ated **A**uto**E**ncoder
VT	**V**entricular **T**achycardia
WHO	**W**orld **H**ealth **O**rganisation
WPW	**W**olff-**P**arkinson-**W**hite

"To my family, whose love and support have been a constant source of strength throughout my journey. I also dedicate this work to all individuals living with heart diseases. It is my sincere hope that the research presented here will contribute to improved understanding, early detection, and better care for those affected by heart conditions..."

Chapter 1

Introduction

A large number of people have been diagnosed with cardiovascular diseases (CVDs) in the past decade. CVDs are the world's majority cause of death, as reported by the World Health Organization (WHO). Coronary artery disease (CAD) is one type of the CVDs that caused a million deaths in the past few years. Therefore, it is essential to keep track of the patients' health and ensure that they will acquire medical treatment. It will also help eventually lower the number of CAD and other CVDs patients. Around 17.9 million people died from CVDs in 2023; most died due to heart attack and stroke [5]. Additionally, CVDs often present no symptoms in their early stages. As such, CAD is one of the most prevalent heart conditions, with an increasing number of occurrences each year. Therefore, early identification of CAD and its contributing factors is crucial in preventing mortality.

Electrocardiogram (ECG) recording is one of the most widely used and non-invasive method for diagnosing CAD and other CVDs. An ECG measures the heart's electrical activity over time, producing waveforms that can reveal irregularities in heart rhythm, and blood supply. In clinical practice, cardiologists typically interpret ECG graphs alongside other diagnostic tools, such as blood tests, requiring extensive expertise to detect abnormalities accurately. However, human error may lead to overlooked abnormalities [6]. To address these challenges, an automated CAD diagnosis system have been explored to assist clinicians by analysing ECG signals using computational methods. Such systems can automatically identify key features liked to heart disease and highlight potential concerns for further review. By integrating the results with other diagnostic factors such as blood test results, patient weight, and chest X-rays, the system can support clinicians in reducing the likelihood of misdiagnosis and in improving patient outcomes.

Coronary artery disease (CAD), also referred to as myocardial ischaemia, remains a leading cause of sudden death worldwide, and its occurrence has increased in recent years. This trend imposes a growing financial and operational burden on public healthcare systems due to the high cost of long-term patient management. Earlier detection of CAD facilitates early intervention, reducing the mortality rate. In recent years, numerous studies have investigated the use of ECG data for heart disease diagnosis. Researchers have increasingly employed machine learning and deep learning techniques to identify complex patterns in ECG signals that may not be readily apparent to the human eye [7–9]. Some studies have developed automated CAD detection methods [10], while others have explored deep learning architectures such as convolutional neural networks (CNN) and long short-term memory (LSTM) network [11, 12]. These approaches have shown potential for improving diagnostic speed and accuracy, making them increasingly relevant to modern clinical practice. Building on this growing body of work, this thesis aims to develop a novel deep learning approaches for the detection of CAD that prioritises both computational efficiency and clinical applicability. The main objective is to reduce mortality rates and to enhance the reliability of diagnostic decision-making in clinical practice.

Despite the promising results reported in recent studies, many existing CAD detection models suffer from limited generalisability, often due to small or imbalanced datasets, inconsistent signal quality, or lack of model interpretability. Moreover, few systems are optimised for deployment in real-time or resource-constrained environments, such as wearable devices or mobile applications. The outcomes of this research have the potential to support clinicians in improving diagnostic workflows, reduce misdiagnosis rates, and enable cost-effective, scalable screening for at-risk populations, particularly in low-resource or remote settings.

This thesis investigates the development of a deep learning-based framework for the automatic detection of CAD using ECG signals. The primary objectives are as follows: (1) to acquire and preprocess suitable ECG datasets from open-access repositories such as PhysioNet; (2) to apply signal processing techniques to improve data quality; (3) to design and implement deep learning models for the classification of CAD and NON-CAD cases; (4) to explore the integration of deep learning with other algorithmic approaches to enhance diagnostic performance; (5) to assess the generalisability of the models across datasets; and (6) to benchmark the proposed methods against existing techniques reported in the literature.

Through these potential contributions, the work aims to advance the state of the art in CAD detection using ECG, with a particular focus on clinical applicability, model robustness, and the potential for real-time or wearable deployment.

1.1 Outline of Thesis

Chapter 1 provides a comprehensive introduction to CVDs, highlighting their substantial effect on the global population, describing heart anatomy, ECG, and the difficulties encountered in diagnosing heart disease through ECG. Moreover, the important role of machine learning and deep learning models in clinical practice is explored. Additionally, the gaps and contributions in existing literature are elaborated upon.

Chapter 2 provides further insights into ECG signal processing and its significance in diagnosing CVDs. Furthermore, traditional machine learning, deep learning and data pre-processing methods used in ECG-based analysis are further elaborated.

Chapter 3 explores the use of one-dimensional convolutional neural networks for the early detection of CAD using ECG signals to enhance detection accuracy while minimising network complexity. The research investigates the effects of varying ECG sample lengths and layer reduction techniques, utilising data from the MIMIC III and Fantasia databases.

Chapter 4 examines further pre-processing techniques, where feature engineering is applied to ECG data within a one-dimensional convolutional neural network (1D-CNN) model to enhance detection accuracy. A smart feature engineering (FE) process is proposed to eliminate unwanted and noisy signals before the main model training cycle. Furthermore, a novel 1D-CNN architecture integrated with feature engineering is introduced for diagnosing CAD.

Chapter 5 focuses on the development of a lightweight deep learning-based model, where a one-dimensional convolutional neural network is designed with consideration for both model performance and resource-constrained environments. State-of-the-art algorithms, including traditional machine learning, deep learning-based approaches, and existing lightweight models, are compared to validate its effectiveness.

Chapter 6 focuses on 12-lead ECG-based CAD diagnosis using a lightweight deep learning model, specifically designed for deployment in resource-constrained environments. The proposed model emphasises a balance between classification accuracy and computational efficiency. Furthermore, the model is deployed on the STM32F469I-DISCO development board, showcasing its feasibility on a cost-effective microcontroller unit. Detailed performance evaluation, including inference time, memory usage, and classification metrics, is presented to demonstrate the model's suitability for real-time healthcare applications.

Chapter 7 introduces a novel attention-driven deep network model for detecting multiple CVDs. Data augmentation is employed to improve dataset robustness, generalisability, and real-world applicability. Additionally, the impact of noise power on ECG signal classification is analysed to identify optimal augmentation settings. The model's effectiveness in diagnosing various heart conditions is further validated, demonstrating its reliability and efficiency.

Chapter 8 provides a detailed discussion of key findings, significance, and contributions of this research, evaluating the strengths and limitations of the proposed approaches in improving CAD detection accuracy and efficiency. Additionally, future research is discussed.

1.2 Background Overview

The biology of the heart is explained in this section. The components of the ECG signal will be introduced as well as functional representations of each wave. In addition, the changes in ECG signals caused by CVDs, including CAD, are briefly explained. Furthermore, the challenges in ECG-based diagnosis are discussed, followed by an introduction to traditional machine learning and deep learning models used in automated ECG diagnosis. Finally, the role of AI in the detection of CVDs using ECG signals is examined.

1.2.1 Heart

The human heart lies at the centre of the circulatory system, working continuously to pump blood that delivers oxygen and nutrients throughout the body as shown in 1.1. It consists of four chambers: the right atrium and right ventricle, which receive deoxygenated blood from the body and send it to the lungs for oxygenation, and the left atrium and left ventricle, which receive oxygen-rich blood from the lungs and pump it to the rest of the body. To ensure blood flows in the correct direction, the heart contains four valves: tricuspid, pulmonary, mitral, and aortic. These valves open and close with each heartbeat to prevent backflow during the cardiac cycle. The heartbeat itself is controlled by the heart's intrinsic electrical conduction system, which begins at the sinoatrial node, often referred to as the natural pacemaker. The electrical signal then passes through the atrioventricular node and a specialised network of fibres, coordinating the contraction of the heart chambers in a regular and efficient rhythm.

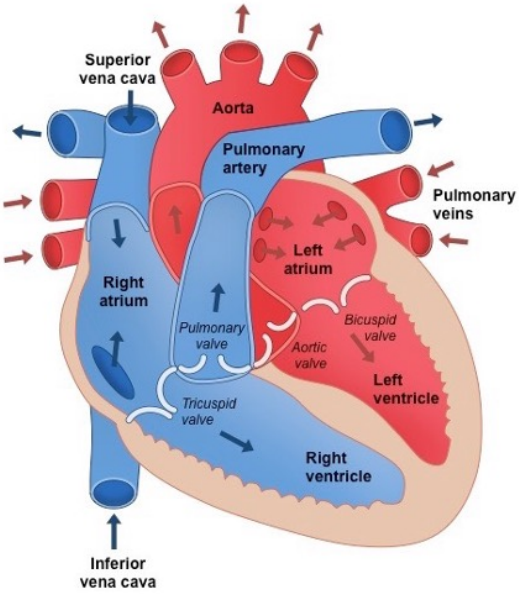

FIGURE 1.1: An anatomical diagram of the human heart showing the four chambers, major valves, and principal blood vessels, with arrows indicating the direction of blood flow [1]

CVDs are among the most common and serious non-communicable health conditions worldwide, accounting for an estimated 18.6 million deaths each year [13]. This substantial rise underscores the critical need for enhanced preventive strategies and timely therapeutic interventions in healthcare systems worldwide. In the United Kingdom, more than 7.6 million people are currently living with heart and circulatory diseases, which place a significant burden on the healthcare system and economy, with estimated costs totalling £28.385 billion annually [14]. These include CAD, heart failure, cardiomyopathies, and arrhythmias. Although many of these conditions are life-threatening if not identified and treated promptly, they often progress gradually and may initially cause only subtle or non-specific symptoms, such as persistent fatigue, chest discomfort, palpitations, or brief episodes of fainting. Atrial fibrillation (AF) is the most common sustained cardiac arrhythmia and is associated with a significantly increased risk of thromboembolic events. Arrhythmias can vary from benign to life-threatening, disrupting normal heart rhythms and significantly impairing cardiac function. The early detection of these conditions is crucial for preventing sudden cardiac death and often necessitates specific treatments.

Early diagnosis and prompt treatment play a vital role in reducing deaths and long-term health impacts linked to CVDs. However, many people delay seeking medical attention, often because they are unaware of the warning signs or misinterpret mild symptoms such as tiredness, chest discomfort, or shortness of breath. These subtle signs can easily be overlooked, especially in the early stages of disease, allowing the condition to progress

before it is recognised. Several factors influence a person's risk of developing CVDs. Some of these, such as smoking, lack of physical activity, unhealthy eating habits, and excessive alcohol consumption, are modifiable. These behaviours are closely linked to conditions such as high blood pressure, raised cholesterol, obesity, and type 2 diabetes, which significantly increase the risk of heart disease. Other factors, including age, sex, family history, and ethnicity, cannot be changed but still play an important role in overall cardiovascular risk.

1.2.2 Coronary artery disease

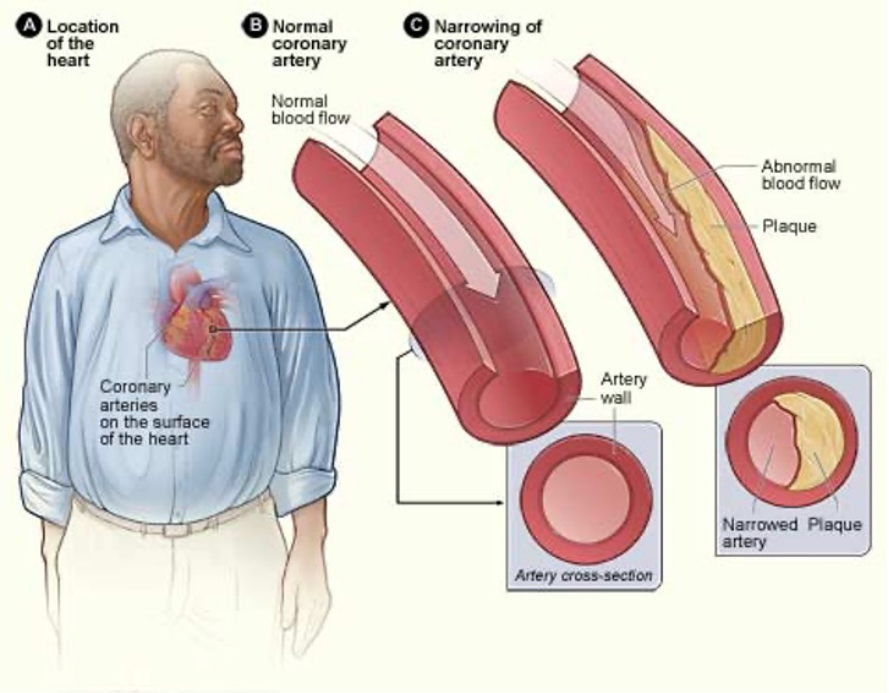

FIGURE 1.2: Illustration of CAD showing coronary artery anatomy, normal arterial blood flow, and arterial narrowing caused by plaque accumulation [2]

CAD is one of the most prevalent and life-threatening forms of CVDs, as highlighted by the WHO [5]. The WHO report highlights that in 2023, CVDs were responsible for approximately 17.9 million deaths globally, representing 32% of all worldwide fatalities [5]. Of these deaths, 85% were attributed to heart attacks and strokes, with a primary occurrence in low- and middle-income countries, which accounted for over three-quarters of such deaths. A study indicated that CVDs are anticipated to maintain its position as the foremost cause of mortality worldwide by the year 2030, with notable prevalence expected across both high and low-income nations [15]. This highlights the growing burden of cardiovascular disease and the urgent need for effective prevention [16]. CAD develops as a result of a progressive process known as atherosclerosis, wherein fatty deposits build up within the walls of the coronary arteries, restricting blood flow to the

heart as shown in Figure 1.2. This condition begins with endothelial damage, which may be caused by risk factors such as hypertension, smoking, elevated low-density lipoprotein cholesterol, diabetes, and chronic inflammation [17]. The build-up of plaques can cause the arteries to narrow, reducing the flow of oxygen to the heart. Studies have shown that advanced symptoms of CAD typically display during middle age, highlighting the correlation between aging and the likelihood of CAD development [18],[19]. CAD, involving the left anterior descending artery, is increasingly observed in young adults, driven by factors such as elevated cholesterol levels, hypertension, and tobacco use [20]. This condition typically presents as single-vessel disease rather than multi-vessel involvement. CAD continues to be a leading cause of mortality among individuals aged 35 and older across both developed and developing nations [21]. Typically, an individual coronary artery has to deliver oxygenated blood to the myocardium. Each artery is a muscular tube lined by smooth tissue and has three layers, namely intima, media and adventitia as shown in Figure 1.3. The intima is the inner layer lined by endothelium. The media is a middle layer of muscle that handles high pressure from the heart. Lastly, the adventitia is the outer layer which is a connective tissue that anchor arteries to other tissues. Myocardial repolarisation is a cause of CAD as the inner walls of coronary arteries are blocked and then cause the changes in the ST segments and T wave [22]. CAD is often asymptomatic in its early stages, which increases the risk of severe complications as the disease progresses. Many patients remain unaware of their condition until it reaches an advanced stage, where blood flow to the heart is significantly restricted. CAD may progress into severe and fatal conditions such as myocardial infarction (MI) if it is not identified and managed at an early stage. The most frequently reported symptom is chest pain, often characterised by sensations of pressure, burning, tightness or discomfort. CAD is associated with modifiable risk factors such as tobacco use, elevated cholesterol levels, hypertension, diabetes and excessive alcohol consumption. It is also influenced by non-modifiable factors such as age, sex and family history. CAD often progresses without noticeable symptoms for many years. However, early signs may include chest discomfort, shortness of breath, fatigue, and increased blood pressure or cholesterol levels. Therefore, early detection of CAD is crucial, as it can progress asymptomatically before leading to severe outcomes such as MI or congestive heart failure (CHF). Early diagnosis plays a vital role in preventing such events and mitigating behavioural risk factors associated with CVDs. A proactive diagnostic approach not only reduces mortality but also increases awareness among individuals who may be unaware of their underlying conditions [23]. To support this, the development and implementation of cost-effective diagnostic methods are essential for enhancing early detection and improving public health outcomes. However, despite advances in clinical practice, no single electrocardiographic feature can definitively diagnose CAD, as the condition may present with a variety of waveform and segment abnormalities [6, 9].

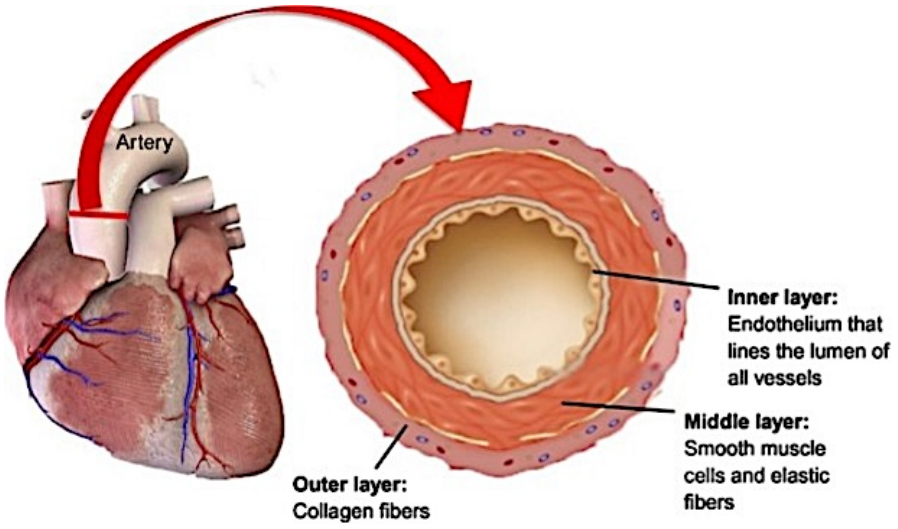

FIGURE 1.3: llustration of coronary artery anatomy showing the arterial wall layers, including the endothelium, smooth muscle layer, and outer connective tissue [3]

Maintaining good cardiovascular health through regular monitoring, such as routine heart rate measurements and scheduled medical assessments, is important for identifying risk factors at an early stage and preventing the progression of serious disease. As such, it is essential to adopt preventive strategies and implement pre-screening tools for those who are considered to be at increased risk. CAD is typically diagnosed through various diagnostic tests, including electrocardiography, treadmill ECG, echocardiography (ECHO), and angiography ECG is most commonly used for all CVDs initial screening in general practices due to being a cost-effective and widely accessible tool, capable of facilitating continuous monitoring, portability, and ability to provide real-time data. However, the current diagnosis using ECG is conducted manually, is time-consuming, and is subject to human error, posing significant challenges in clinical practices. Hence, to address these issues, we propose intelligent learning-based approaches to improve diagnostic efficiency. The proposed models serve as a cost-effective and user-friendly pre-screening tool for patients, requiring minimal specialised medical expertise.

1.2.3 Electrocardiogram

ECG is routinely used as a primary tool for initial screening of CVDs in general practice. It provides a non-invasive means of monitoring heart function and detecting abnormalities in cardiac activity. Recorded tracings are a key component of the diagnostic process, as general practitioners rely on them to support clinical decision-making and facilitate early identification of cardiac conditions. However, manual interpretation of ECG can be time-consuming, and variations in clinical expertise may result in occasional misdiagnosis of CAD. In addition, the absence of a clearly identifiable biomarker for CAD through

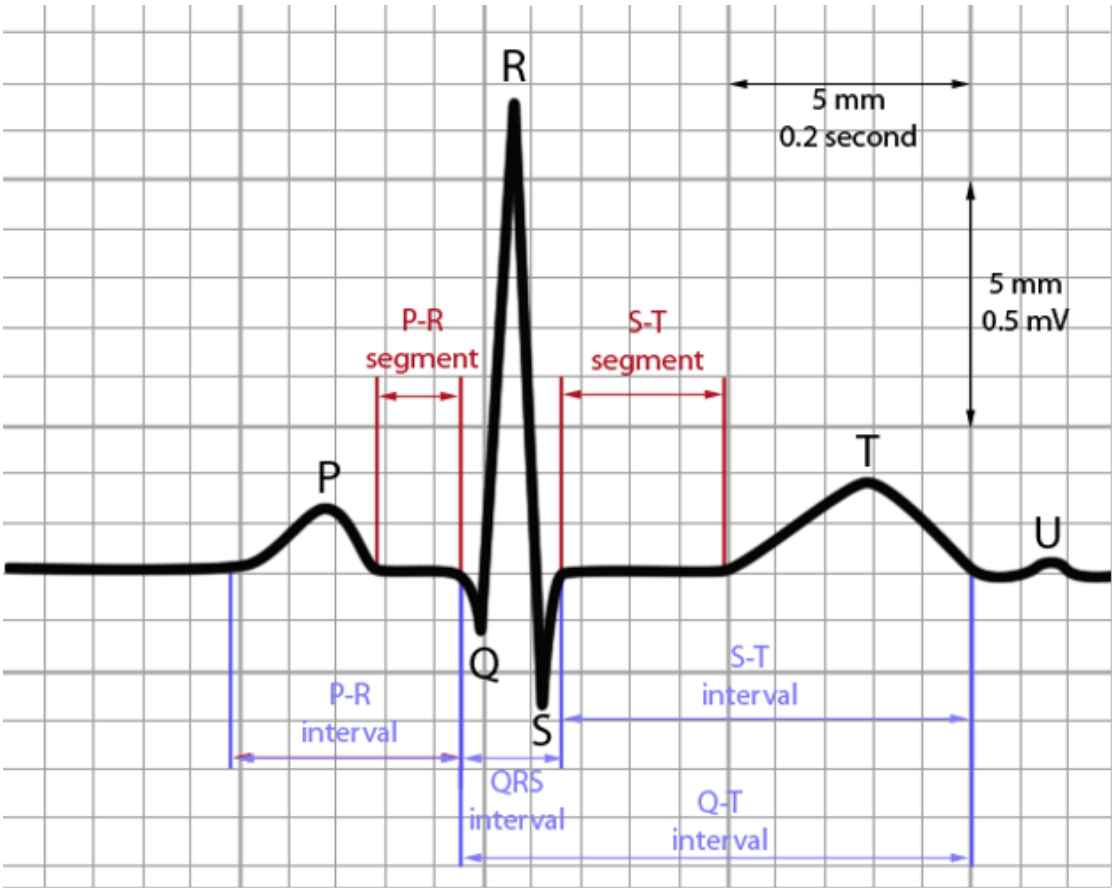

FIGURE 1.4: Representation of ECG waveform components]Representation of ECG waveform components, illustrating the P wave, QRS complex, and T wave associated with cardiac electrical activity [4]

non-invasive methods requires clinicians to consider multiple factors during assessment, which increases the complexity of diagnosis [9].

ECG represents the electrical activity of the heart and is composed of characteristic waveform patterns, including the P wave, QRS complex and T wave, as shown in Figure 1.4. These components reflect the processes of depolarisation and repolarisation in the atria and ventricles. Electrocardiography, the method of recording ECG signals, plays a fundamental role in the diagnosis of various heart conditions. Each segment and wave within the ECG trace corresponds to specific cardiac events, providing valuable information for identifying different types of heart disease.

The morphology of an ECG signal within a single cardiac cycle comprises distinct waveforms, namely the P, Q, R, S and T waves, as shown in Figure 1.4. The P wave reflects the spread of electrical impulses through the left and right atria and typically lasts between 80 and 100 milliseconds (ms). The PR interval, measured from the onset of the P wave to the beginning of the QRS complex, generally ranges from 120 to 200 ms. The QRS complex represents ventricular depolarisation and has a typical duration of

approximately 60 to 100 ms. The ST segment is the distance from the end of the QRS complex to the beginning of the T wave. It is used to study ischemic heart disease as ST depression and elevation need to be observed for CAD diagnosis [24]. ST segment deviation is a vital marker employed in the diagnosis of ischaemic conditions, such as CAD, MI, and others. ST depression suggests the presence of significant coronary lesions and highlights the need for early invasive intervention in unstable CAD. Conversely, ST elevation typically indicates complete coronary artery occlusion and is a hallmark of myocardial infarction. The T wave indicates the repolarisation of ventricles. The U wave sometimes appears after the T wave as it represents the last part of the ventricular repolarisation. QT interval starts from the onset of the QRS complex to the end of the T wave, around 200 to 400 milliseconds depending on heart rate. Abnormalities in the P wave, QRS complex, T wave, and ST segment are commonly used as diagnostic markers of various cardiac pathologies. variations in the QRS complex, particularly the R peak, are widely studied and frequently used in the diagnosis of cardiac abnormalities such as arrhythmia and AF [25–27].

ECG signals are acquired using surface electrodes placed at specific anatomical locations on the patient's body [28]. A standard clinical ECG is typically recorded using a 12-lead configuration, which includes the limb leads I, II, III, aVR, aVL and aVF, along with the chest leads V1 to V6. These leads record the heart's electrical activity from different angles, allowing detailed analysis in both the frontal and horizontal planes. Figure 1.5 illustrates the standard anatomical placement of electrodes used in 12-lead ECG acquisition. The electrodes measure voltage differences generated by myocardial depolarisation and repolarisation, which are then amplified and digitised for clinical interpretation and computational analysis. The signal acquisition process depends on factors such as electrode placement, skin preparation, and patient movement. Inadequate contact or motion during recording can introduce noise, such as baseline drift or muscle artefacts, which may compromise diagnostic reliability. Clinical ECG systems typically sample signals at rates ranging from 250 Hz to 1000 Hz to ensure adequate temporal resolution for capturing rapidly changing waveforms such as the QRS complex. The acquired signals are often preprocessed using filters to remove powerline interference, high-frequency noise, and baseline wander prior to feature extraction or classification. Accurate and reliable acquisition is essential, as errors at this stage may lead to incorrect diagnoses or misinterpretation by automated systems.

With advancements in state-of-the-art technology, ECG has become widely employed for the classification of CVDs. Numerous studies have demonstrated its effectiveness in diagnosing conditions such as arrhythmia [29, 30], CAD [6, 31], and AF [32–34]. Given the increasing prevalence of these conditions, improving ECG classification methods is vital

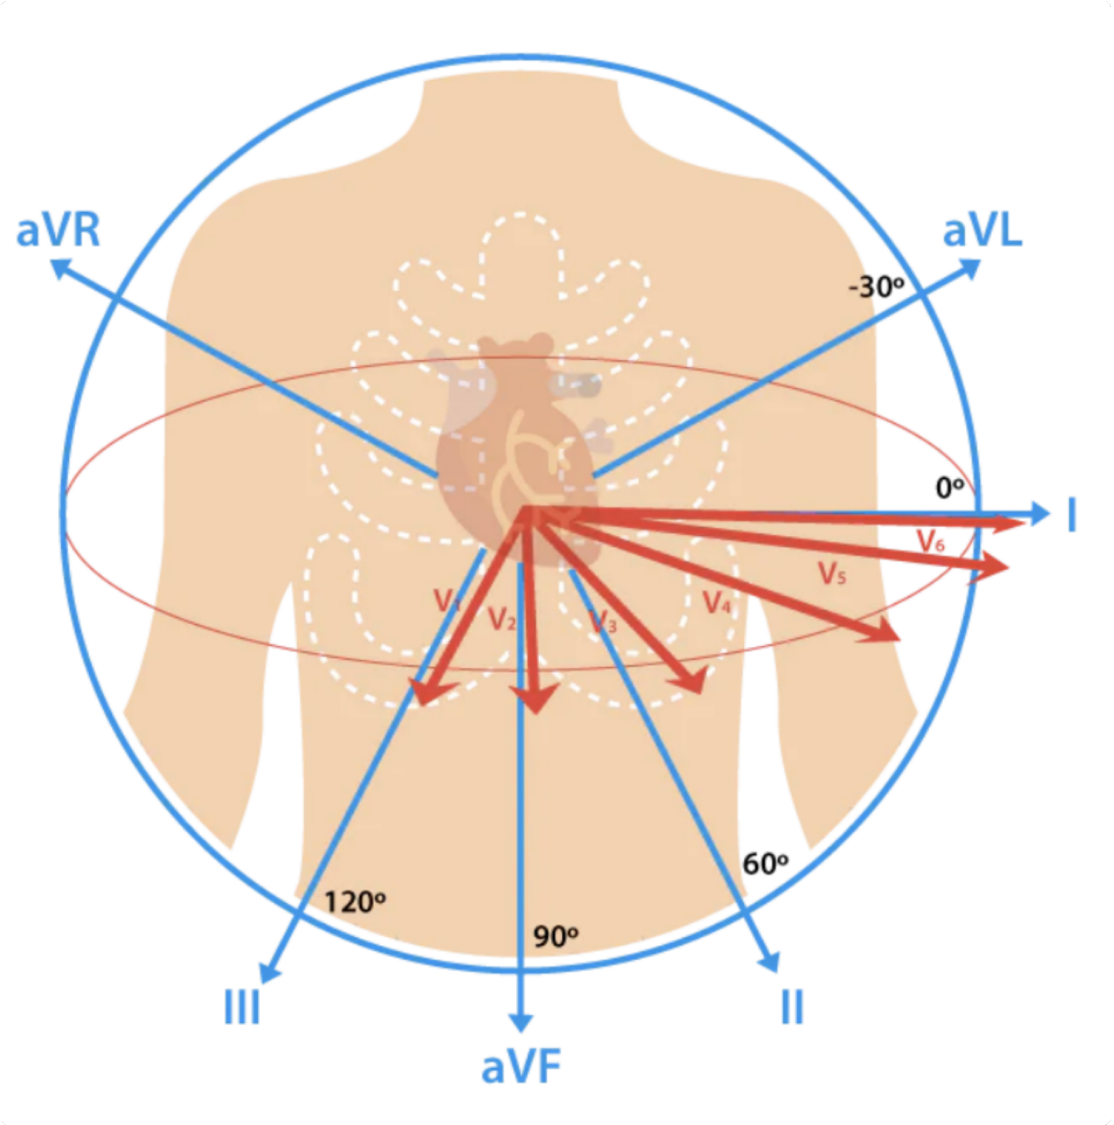

FIGURE 1.5: 12-lead ECG vector diagram showing lead angles and chest electrode positions around the heart

for enhancing patient outcomes and optimising healthcare resources. ECG signals contain detailed information in both the time and amplitude domains, which is critical for accurate classification. However, developing deep learning models capable of effectively capturing and interpreting these features across a range of cardiac conditions presents significant challenges. Furthermore, achieving generalisability across multiple diseases using a single model remains difficult due to considerable variation in ECG characteristics among patients. Although a range of diagnostic tools is used in cardiovascular assessment, including imaging and biochemical tests, ECG is often preferred due to its accessibility, non-invasive nature, and diagnostic value. In contrast, angiography, while also widely used for screening CVDs, has limited application in patient follow-up and treatment because of its invasive nature, high cost, and the requirement for specialised technical expertise. With the development of artificial intelligence (AI) technologies,

machine learning and deep learning techniques are being increasingly employed to analyse medical data, including physiological signals, X-rays, magnetic resonance imaging (MRI), and other modalities. In clinical practice, ECG continues to serve as a primary tool for the preliminary screening of various CVDs. Although a recorded ECG can support the initial diagnosis of conditions such as CHF and angina, further investigations such as echocardiography and exercise testing are typically required to confirm the diagnosis [35].

1.2.4 Challenges in ECG-based Diagnosis

ECG serves as a primary non-invasive technique for the detection of CVDs, including CAD. However, despite its clinical value, several challenges persist that limit both its diagnostic reliability and the development of automated detection systems. One key limitation is the need for human interpretation of ECG signals. Recognising features such as ST-segment changes or T-wave abnormalities often depends on clinical expertise, and even experienced practitioners may disagree when the findings are unclear. This variation in clinical judgement can lead to inconsistent diagnoses, particularly in busy or under-resourced healthcare settings [36]. Interpretation becomes challenging when abnormalities are not easily distinguishable. Moreover, similar ECG patterns may reflect different cardiac conditions, making accurate diagnosis difficult without additional clinical information. These challenges can result in delayed, missed diagnoses, unnecessary investigations, or inappropriate treatment, all of which may negatively affect patient outcomes. In high-demand clinical environments, manual ECG interpretation also adds to clinical workload and increases the risk of inconsistency.

In addition, ECG signals are frequently affected by various forms of noise and artefacts, such as baseline drift, movement-related noise, incorrect electrode placement, and power-line interference [37]. These disturbances can interfere with important components of the signal, within the P wave, QRS complex, and ST segment regions essential for identifying conditions such as arrhythmias or myocardial ischaemia. Poor signal quality is often caused by patient movement, inadequate skin preparation, or poor electrode contact. Therefore, ECG signals typically require preprocessing to ensure they are suitable for accurate interpretation or automated analysis. Common preprocessing techniques include filtering, baseline correction, and artefact removal. These steps help preserve diagnostically relevant features while reducing interference. Preprocessing is essential not only for improving clinical interpretation but also for enhancing the performance of automated systems that extract features and classify cardiac conditions.

Many cardiac conditions produce similar ECG patterns, which can make accurate diagnosis more difficult. These overlapping features often affect the same regions of the ECG waveform, making it challenging to distinguish between conditions using ECG data alone [38]. This similarity increases the risk of diagnostic error, particularly when interpretation relies solely on ECG signals. In clinical practice, additional information such as blood test results, symptom history and imaging is often used to support diagnosis. However, most automated systems are trained exclusively on ECG data, which limits their ability to differentiate between physiologically similar but clinically distinct conditions.

Data availability remains a key limitation in developing accurate CAD detection models. Many open-access ECG datasets lack specific annotations for CAD or contain only a small number of CAD cases. The limited size and diversity of these datasets also restrict model generalisability, as ECG morphology can vary with factors such as age, sex, and other health conditions [39]. Furthermore, ECG recordings used for CAD detection may also differ in important ways, including the type of acquisition equipment, lead configuration, sampling frequency, and labelling approaches. These variations can result in distribution shifts that reduce model performance when applied to external datasets. Addressing these challenges is essential to improve the reliability, fairness, and clinical utility of automated CAD detection.

1.2.5 Machine Learning Models

Machine learning approaches, such as support vector machines (SVM), decision trees, k-nearest neighbours (KNN), and random forests (RF), have been widely applied to ECG classification tasks. These methods typically rely on handcrafted features extracted from the ECG signal, including heart rate variability (HRV), QRS duration, and other signal characteristics. y identifying patterns and relationships within these features, the models can associate them with diagnostic categories. Their performance, however, is strongly influenced by the quality, relevance, and completeness of the extracted features. Moreover, the feature extraction process can be time-consuming and often requires domain-specific expertise, which limits scalability across varied clinical settings.

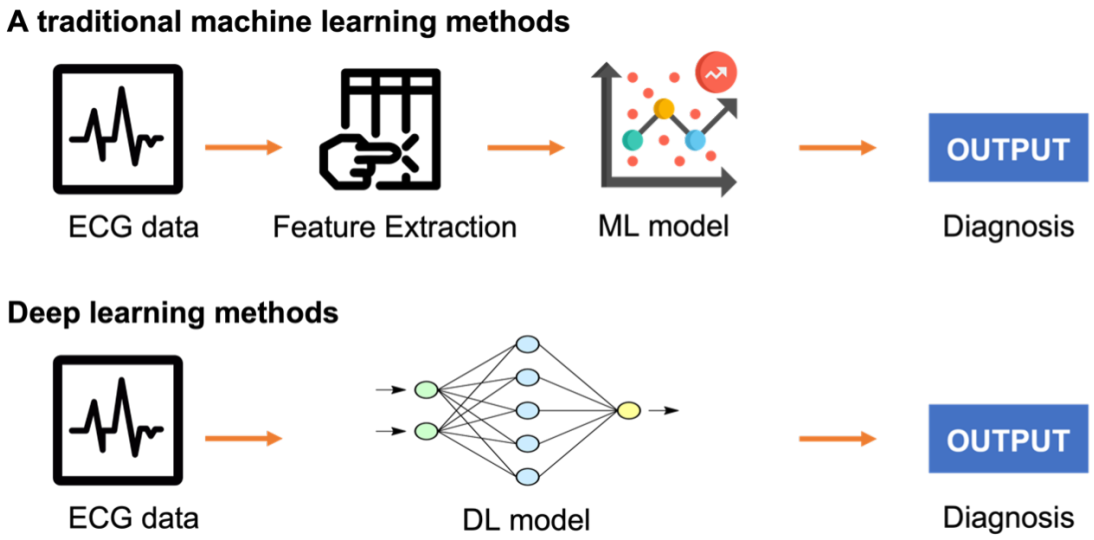

FIGURE 1.6: A comparison of ML and DL methods in ECG analysis

As illustrated in Figure 1.6, machine learning workflows for ECG analysis generally involve a series of preprocessing, feature extraction, and classification steps. Studies have shown that, when supported by well-selected features, algorithms such as SVM and RF can achieve competitive performance in detecting arrhythmias, myocardial infarction, and other cardiac abnormalities [40]. Ongoing research aims to optimise feature selection methods and strengthen model robustness to enable effective adaptation across diverse patient populations and recording conditions. In contrast to deep learning, which is capable of learning directly from raw ECG signals within a single integrated model, machine learning requires explicit feature engineering and separate processing stages. This key difference will be discussed further in Chapter 2

1.2.6 Role of AI in CVDs detection using ECG

With advancements in state-of-the-art technologies, low-cost diagnostic tools are increasingly being developed by leveraging AI. ECG remains the primary method for initial screening of CVDs in general medical practice. General practitioners rely on recorded ECGs as essential diagnostic tools during patient evaluation. Accurate diagnosis of CVDs requires a comprehensive understanding of risk factors alongside clinical expertise. Additionally, ECG is frequently employed as a preliminary screening tool due to its capacity for continuous, non-invasive monitoring and real-time data provision [41]. Although ECG is an effective diagnostic tool for recording the heart's electrical activity, current diagnostic procedures involving ECG are predominantly performed manually by practitioners and clinicians, which can be time-consuming and prone to error. Advances in state-of-the-art technologies have facilitated the development of various automated

systems for diagnosing CVDs, aiming to address these challenges [9, 42, 43]. However, such efforts have been underrepresented in the diagnosis of CAD. Therefore, developing automated ECG-based diagnostic methods is becoming increasingly important in modern healthcare, primarily due to the limitations of manual diagnosis. Automated diagnosis can be achieved by applying AI algorithms that learn from large datasets to provide accurate results without requiring specialist expertise. Furthermore, automated systems could offer greater convenience to both patients and healthcare professionals.

AI is increasingly being utilised in the diagnosis of CVDs through the analysis of ECG data. Deep learning is a branch of AI that has gained increasing attention for its ability to learn patterns directly from raw data. Its effectiveness has been demonstrated in a range of domains, including natural language processing (NLP), image analysis, and speech recognition, and this success has encouraged its application to ECG-based CVD detection [11, 44, 45]. Although these models typically require large datasets, extended training times, and powerful computing resources such as graphics processing units (GPUs), they offer advantages in scalability and accuracy. Deep learning algorithms have demonstrated considerable potential in detecting subtle and complex patterns within ECG waveforms that may be difficult for human observers to identify. This capability facilitates earlier and more accurate detection of conditions such as arrhythmia, AF, and CAD. Recent studies have shown that AI-based ECG analysis enhances diagnostic accuracy [46–48]. Furthermore, AI models can learn directly from raw or processed ECG signals, recognising key diagnostic features such as ST-segment deviations, T-wave abnormalities, and variations in QRS morphology. Consequently, AI presents an opportunity to reduce diagnostic errors and improve efficiency in both acute care and preventative settings.

Despite its potential, the use of AI in the diagnosis of CVDs still faces several challenges. One important limitation is the lack of large and diverse ECG datasets. Many publicly available datasets contain complexities such as inconsistent labelling, signal noise, or class imbalance, which can be difficult for AI models to handle effectively. In addition, differences in how data are collected, including variations in lead placement, device type, or sampling frequency, can alter the underlying characteristics of the data and reduce model performance when applied in real clinical settings. The interpretability of AI models is also a significant concern. Deep learning systems often offer limited insight into how decisions are made, which may lead to reduced trust among clinicians and create difficulties with regulatory approval. To address this, statistical methods such as SHapley Additive exPlanations (SHAP) value visualisations have been introduced to explain the predictions made by these models [49, 50]. However, their clinical usefulness remains limited, particularly when the features they highlight do not align with established physiological markers or the interpretive approaches commonly used in cardiology.

For AI to be successfully integrated into clinical practice, it must not only achieve high levels of predictive accuracy but also provide explanations that are clear, meaningful, and grounded in clinical understanding. AI holds considerable potential to improve the diagnosis of CVDs using ECG signals. It can support earlier detection, improve diagnostic accuracy, and increase access to timely care. However, these benefits can only be fully realised if the technical, clinical, and ethical challenges are addressed thoughtfully and responsibly within real-world healthcare settings.

1.3 Summary

This chapter has provided an overview of the foundational concepts relevant to this thesis. The anatomy and physiological function of the heart were outlined, followed by a discussion of CAD as a major cause of CVDs mortality. The ECG was introduced as a key diagnostic tool, and the principal challenges associated with ECG-based diagnosis were examined, including variability in signal quality and the potential for human error in interpretation. The review of machine learning approaches highlighted their dependence on handcrafted features. Although these features offer interpretability for clinical application, they may also restrict scalability across diverse healthcare settings. Finally, the emerging role of AI in the detection of CVDs from ECG data was considered, demonstrating its potential to support earlier and more accurate automated diagnosis.

Chapter 2

ECG Signal Processing and Its Role in Cardiovascular Disease Detection

2.1 Overview

CVDs remain one of the leading causes of death worldwide, and timely diagnosis plays a crucial role in improving patient outcomes. In light of this, the application of signal processing techniques is essential, as it enables the mitigation of artefacts and enhancement of diagnostically relevant features in ECG recordings. With recent advancements in signal processing techniques, automated ECG-based diagnosis has achieved considerably higher levels of accuracy in the detection of cardiac abnormalities, including CAD.

This chapter introduces an overview of ECG signal processing approaches and their role in the diagnosis of CVDs. It begins by examining how ECG analysis is applied in clinical settings, outlining its clinical utility as well as its limitations. The discussion then moves on to traditional machine learning methods and more recent deep learning-based techniques used for ECG diagnosis, with focus on models developed for CAD detection. The review also distinguishes between single-lead and multi-lead ECG classification strategies, highlighting model architectures such as CNN, attention mechanisms, and lightweight neural networks suitable for deployment in real-time or resource-constrained environments. The chapter concludes by highlighting the ongoing challenges and the increasing need for ECG analysis systems that are accurate, efficient, and suitable for clinical deployment.

2.2 An Examination of ECG Signal Analysis in Clinical Practice

Electrocardiography is a method used to capture an ECG from a patient's body. ECG data is fundamental in diagnosing CAD within clinical settings. ECG assists in measuring the heart's electrical activity through surface electrodes placed on the skin, providing essential insights into cardiac rhythm, conduction disorders, ischaemic conditions, and structural heart diseases. The analysis of ECG signals is crucial for detecting arrhythmias, such as AF, ventricular tachycardia, and heart blocks, by identifying abnormalities in P wave morphology, PR interval, QRS complex, and QT interval. The standard 12-lead ECG records electrical activity from different regions of the heart, providing a comprehensive view that aids physicians in diagnosing various cardiac conditions. These include CAD, Acute Myocardial Infarction (AMI), Inferior Myocardial Infarction (IMI), Lateral Myocardial Infarction (LMI), and other abnormalities, where AMI refers to the acute necrosis of myocardial tissue caused by a sudden reduction in coronary blood flow, and IMI and LMI represent regional variants that occur when the blockage affects specific branches of the coronary arteries.

In everyday clinical settings, the ECG is often one of the first investigations carried out when a patient presents with symptoms such as chest pain, palpitations, or episodes of fainting. It offers a fast, cost-effective, and reliable method for assessing cardiac function and helps guide further diagnostic and therapeutic decisions. ECG is used not only in emergency departments but also widely across outpatient and primary care environments, where it remains an essential part of routine cardiovascular assessment [51]. Interpretation is generally performed by cardiologists or other trained healthcare professionals who rely on their clinical judgement and experience. However, inter-observer variability is a well-recognised limitation, particularly in cases where abnormalities are subtle or borderline. Diagnostic accuracy can also differ depending on the clinician's level of training, with junior doctors and generalists often demonstrating lower performance than experienced cardiologists. This inconsistency has been linked to both delayed diagnoses and incorrect clinical decisions in the management of cardiac diseases.

The interpretation of ECG is further complicated by patient-specific factors such as age, sex, comorbidities, genetic variation, and medication use, all of which can influence waveform morphology. Numerous studies have highlighted the challenges and potential inaccuracies associated with manual ECG interpretation [52]. Misclassification may lead to missed diagnoses of serious conditions or unnecessary investigations that contribute to patient anxiety. To address these limitations, automated diagnostic systems are

being developed with the aim of improving reliability, reducing clinician workload, and enhancing efficiency.

A further challenge in ECG analysis arises from the presence of noise and signal artefacts, which can result from muscle activity, respiration, or incorrect electrode placement. Such artefacts may alter critical waveform features and affect clinical interpretation. While conventional filtering techniques such as high-pass and low-pass filters can help suppress noise, they may also eliminate subtle signal components that are clinically significant. Moreover, manual ECG analysis is time-consuming and not well suited to high-throughput clinical environments, large-scale screening programmes, or continuous real-time monitoring. These limitations have led to growing interest in automated ECG classification methods based on machine learning and deep learning, which show considerable promise in improving diagnostic accuracy, reducing observer variability, and enabling real-time analysis across both clinical settings and wearable devices.

2.3 Advanced signal processing for ECG-based cardiac abnormality detection

Accurate detection of CVDs from ECG signals depends not only on the use of advanced classification methods but also on the quality of the input signal. Data preprocessing is a critical step that involves noise reduction, signal standardisation, and enhancement of morphological features such as the QRS complex [53]. These processes are essential to ensure that clinically relevant information is preserved and made suitable for automated ECG classification [25, 37]. Clinical ECG recordings are often affected by various types of noise and artefacts, which can make it harder to interpret important waveform features accurately. One of the most common issues is baseline wander, a slow drift in the signal typically caused by breathing or slight movements by the patient. This can interfere with the detection of smaller waveform components, such as the P and T waves. Muscle artefacts are another frequent problem, introducing high-frequency noise when a patient moves or tenses their muscles, for example during exercise or stress testing. Electrical interference from nearby equipment can also cause power-line noise, usually centred around 50 or 60 Hz depending on the local mains frequency. In addition, motion artefacts from loose electrodes or sudden body movements can introduce sharp, irregular disturbances into the signal. Together, these artefacts can alter the shape and timing of key parts of the ECG, such as the QRS complex, making it more challenging for automated systems to detect cardiac abnormalities [54].

Effective denoising plays an important role in ECG signal processing, as it helps ensure that meaningful features can be accurately extracted. Without proper noise removal,

the performance of classification models can be significantly affected [37]. The aim of denoising is to reduce unwanted artefacts while preserving the key components of the ECG signal, such as the P wave, QRS complex and T wave, which are essential for ECG classification and clinical interpretation. Over the years, various techniques have been developed to address this challenge, each offering distinct advantages and limitations depending on the quality of the ECG signal and its intended application. Traditional filtering methods, such as band-pass filters operating within the 0.5 to 40 Hz range, are widely used to suppress baseline wander, power-line interference, and high-frequency noise in ECG signals [55]. These filtering methods are relatively straightforward to implement and are effective in controlled clinical settings. However, their performance may decline in more variable conditions, such as when patients are moving or undergoing long-term monitoring outside hospital environments. In such cases, unpredictable noise can interfere with the ECG signal, and traditional filters may unintentionally remove or alter clinically important features, such as the P wave, T wave, or ST segment, particularly during periods of instability or when noise overlaps with low-amplitude components [56].

To address the limitations of traditional filtering techniques, the wavelet transform has emerged as a widely used method for ECG denoising [27, 53]. Its principal advantage lies in its ability to analyse signal characteristics in both the time and frequency domains. Discrete Wavelet Transform (DWT) decomposes the ECG signal into multiple levels of detail, enabling more precise separation of frequency components. This layered approach allows for the selective reduction of noise, while maintaining important morphological features. Among the various wavelet functions, the Daubechies 6 (db6) wavelet is often used for ECG analysis because its shape closely matches that of the QRS complex [57]. Many studies have shown that wavelet-based denoising not only improves the clarity of the signal but also increases the accuracy of later classification steps, often performing better than traditional filtering methods [58]. Following signal denoising, Sample Entropy (SampEn) has been applied in ECG preprocessing as a nonlinear measure to assess the complexity and regularity of HRV. It provides a means to capture the unpredictability within ECG components. Clinically, SampEn has been demonstrated to aid in the identification of cardiac conditions such as AF and CHF, which are often characterised by reduced signal complexity. Incorporating SampEn into the preprocessing stage also enhances the reliability of subsequent analyses, by providing a quantitative measure of signal regularity and the effectiveness of filtering techniques [59].

Although ECG signal processing and the detection of heart conditions have improved a great deal, several important challenges are still present. One key difficulty is the natural variation between patients. Factors such as age, sex, medical history, and medication can all influence the ECG signal, making it challenging to develop algorithms

that perform reliably across diverse patient populations. ECG recordings obtained in real-world settings frequently contain multiple types of noise that can overlap, making it difficult to remove interference without compromising important ECG features. Achieving effective noise reduction while preserving subtle yet clinically significant features remains a challenge in ECG signal processing. As ECG monitoring becomes increasingly common in everyday life through wearable and portable devices, there is a growing demand for methods that are both fast and energy-efficient and capable of operating in real time on devices with limited power and processing capabilities. Looking forward, integrating advanced signal processing with explainable AI could help develop systems that clinicians find more transparent and trustworthy. Additionally, incorporating data from other biosignals alongside models tailored to individual patient characteristics may enhance diagnostic accuracy and enable more personalised care. Ongoing research in these areas will be essential to translate current advanced techniques into practical tools that improve patient outcomes.

2.3.1 Pre-Processing

Pre-processing represents a fundamental stage in ECG analysis, as it aims to minimise artefacts and preserve physiologically meaningful characteristics prior to subsequent feature extraction and classification. Various techniques are applied to reduce noise in the signal, including baseline wander, power-line interference and motion artefacts.

Baseline wander is a low-frequency drift in the ECG signal that is commonly caused by respiration and electrode motion. Wavelet-based techniques have been widely investigated for ECG denoising and baseline wander removal. In [60], Daubechies-3 and Symlets-3 wavelets were shown to be effective in reducing baseline wander, while the stationary wavelet transform demonstrated superior performance in preserving important ECG components in [61]. The study in [62] combined the wavelet transform with particle swarm optimisation, an evolutionary search algorithm inspired by swarming behaviour, and obtained improved suppression of power-line interference. Adaptive filtering methods have also been investigated, and signal-piloted and adaptive-rate FIR filtering techniques demonstrated notable computational efficiency while maintaining signal quality [63]. A hybrid method using iterative filtering and lifting wavelet transform was proposed in [64], and this approach was able to remove noise whilst preserving clinically relevant waveform features.

Power-line interference refers to periodic noise introduced by mains electricity at 50 Hz, which can interfere with important ECG waveform components and thereby compromise diagnostic accuracy. Recent studies have therefore focused on developing effective

pre-processing techniques to reduce power-line interference in ECG recordings. Various techniques have been proposed to mitigate power-line interference in ECG recordings, including adaptive notch filters [65], signal-piloted filtering [66], local characteristic decomposition [67], and combined adaptive algorithms [68]. Adaptive notch filters with sharp spectral resolution have shown promise in suppressing power-line interference whilst preserving the morphology of the QRS complex [65]. In addition, several studies have explored hybrid approaches, such as iterative filtering combined with the lifting wavelet transform [64] and cascaded multistage adaptive noise cancellers [69], which achieve improved attenuation of power-line interference without compromising relevant ECG components.

Motion artefacts are high-frequency disturbances in ECG recordings that are caused by patient movement or muscle contraction and can distort the morphology of the ECG waveform. Motion artefacts represent a significant challenge in ECG and other cardiac monitoring signals, particularly in wearable and and in real-world monitoring conditions. Numerous methods have been proposed to detect and remove such artefacts, including machine learning-based techniques [70], wavelet transforms [71], Monte Carlo filtering [72] and adaptive filtering [73]. Independent component analysis (ICA) has also shown promise in separating noise from cardiac signals [74, 75]. While capacitive ECG systems offer certain benefits over conventional wet electrodes, they remain highly sensitive to motion artefacts [76]. Effective removal of motion artefacts is therefore essential for accurate diagnosis and monitoring, especially in intensive care settings where false alarms can affect clinical practice [77].

2.3.2 Classification

Traditional machine learning methods have been extensively used for automatic ECG interpretation. These methods generally involve multiple steps, including signal preprocessing, feature extraction, and classification using algorithms such as SVM and RF [40]. A variety of ML algorithms, including SVM, Gaussian naïve Bayes, k-means clustering, logistic regression and KNN have been applied to ECG data to detect patterns indicative of CAD [40, 78]. Commonly, these approaches utilise techniques such as R-peak detection and heart rate variability analysis to extract time-domain features from the ECG signals. However, manual feature engineering can struggle to capture the complex and subtle relationships present in ECG data [79]. To overcome these limitations, deep learning techniques have been introduced. Deep learning methods automatically learn intricate features directly from raw ECG signals, reducing reliance on handcrafted features. Nevertheless, feature extraction remains a fundamental step in ML-based ECG CAD detection as it significantly influences model performance, interpretability and diagnostic

reliability [80]. Typical feature extraction techniques include time-domain metrics such as RR intervals and QRS duration, as well as wavelet transform methods, which convert raw ECG data into more meaningful representations. Despite these advancements, several challenges persist in ECG-based CAD diagnosis. Variability in ECG signal acquisition caused by differences in devices, lead placement and environmental noise can adversely affect model performance. Hence, ML models require suitable preprocessing to remove noise and artefacts, ensuring data quality and consistency. Addressing these challenges demands effective data processing, sophisticated feature engineering and robust validation strategies to improve the reliability and generalisability of diagnostic models.

2.3.3 Traditional Machine Learning Approaches

Traditional machine learning methods usually require the extraction of meaningful features from ECG signals in order to classify CVDs effectively. These features are often based on prior knowledge and may include the shape of ECG waveforms, statistical values, and information from both the time and frequency domains. Once these features are prepared, a range of commonly used algorithms can be applied. These include SVM [81, 82], KNN [83, 84], k-means [85], logistic regression [26], and naive bayes classifiers [86]. These models depend heavily on hand-crafted features, which means that a strong understanding of ECG characteristics is needed, along with careful signal preprocessing.

Table 2.1 presents various studies that have applied traditional machine learning algorithms to the classification of heart conditions. A range of classifiers has been explored, including SVM, KNN, decision trees, RF, logistic regression, and k-means. These methods were often combined with feature extraction techniques such as wavelet transform, Principal Component Analysis (PCA), and time-frequency decomposition. These preprocessing steps were essential for reducing noise, enhancing relevant signal components, and improving classifier performance. SVM was among the most widely used classifiers [81, 86, 89]. It was frequently applied to both binary and multi-class classification tasks, including the detection of CAD, differentiation between normal and abnormal rhythms, and identification of various arrhythmia types. RF and other ensemble models were also widely adopted due to their ability to handle high-dimensional data and to combine the outputs of multiple weak learners for improved accuracy [87, 88]. KNN was commonly selected in studies focusing on wavelet-based or statistical features, particularly where simplicity and fast inference were prioritised [83]. Logistic regression was integrated with a hidden semi-Markov model (HSMM) to segment heart sounds into their physiological components, such as S1, S2, systole, and diastole [26]. K-means was

TABLE 2.1: Traditional Machine Learning Methods in Heart Disease Classification

Study	**Architecture(s)**	**Diagnosis**	**Dataset**
[26]	Logistic Regression HSMM	Heart sound segmentation (S1, S2, systole, diastole)	CinC Challenge 2011
[87]	SVM KNN RF Ensemble	Normal LBBB RBBB PAC PVC	MIT-BIH
[81]	PCA Optimized SVM	CAD NON-CAD	Long-Term ST 24-hour Holter
[83]	Wavelet Transform KNN	RBBB LBBB PR PAC VT VFL VB VF HGEA ASY	MIT-BIH: • Arrhythmia • Malignant ventricular
[85]	K-means	QRS detection	5 PhysioNet ECG
[86]	SVM Decision Trees Random Forest KNN	Normal Abnormal	UCI Machine Learning
[88]	Logistic Regression RF SVM	LVH	PTB-XL Georgia
[89]	SVM RF Ensemble	Normal Abnormal	Cleveland Clinic (UCI)
[40]	SVM RF	Normal SVEB VEB FB Q	MIT-BIH Arrhythmia

also employed for QRS complex detection [85]. These studies demonstrate that traditional machine learning algorithms continue to offer reliable performance, especially for ECG-based heart disease detection when applied alongside appropriate feature extraction techniques [40].

2.3.3.1 Feature Extraction

Accurate ECG classification in traditional machine learning depends heavily on effective feature extraction, which is essential for identifying critical signal characteristics associated with cardiac disorders. Given the complex and dynamic nature of ECG signals, various methods have been employed to derive meaningful diagnostic features. However, handcrafted features may not fully capture this complexity, potentially compromising model performance and diagnostic accuracy. Furthermore, feature extraction can be a time-intensive process, in contrast to deep learning techniques that automatically learn relevant features directly from raw data.

A commonly employed method for feature extraction in ECG signal analysis is the DWT [27]. When integrated with machine learning classifiers such as SVM, KNN, and RF, DWT-based approaches have achieved high classification accuracy in detecting various cardiovascular conditions. By decomposing the ECG signal into multiple resolution levels, DWT enables simultaneous analysis in both the time and frequency domains. This multi-scale representation is effective in capturing transient features such as the QRS complex, P-wave, and T-wave across various frequency bands. Its ability to localise signal features under changing heart rates and noisy conditions enhances its diagnostic utility [27, 83]. As a preprocessing step, DWT improves classification performance by isolating relevant features and suppressing noise and artefacts [90]. When integrated with advanced models, such as non-linear vector decomposed neural networks, DWT-based systems have achieved notable performance, with sensitivity, specificity, and accuracy reaching 92.0%, 89.33%, and 90.67%, respectively [90].

Another widely used technique is PCA, a statistical method that reduces the dimensionality of data while preserving the most significant variance. By transforming the original ECG signal into a set of orthogonal components, PCA simplifies the feature space, removes redundancy, and enhances classification performance. Studies have demonstrated that PCA can achieve high accuracy across a variety of segment lengths, highlighting its robustness and adaptability in ECG analysis [91]. PCA has been applied extensively in ECG classification tasks, particularly for the detection of arrhythmias. It is considered an effective approach for feature extraction, dimensionality reduction, and noise suppression in ECG signals. The PCA has been successfully implemented in a range of studies, including those focusing on sleep apnea [92], arrhythmia[93], sudden cardiac arrest diagnosis [94], and the detection of CHF [95]. In addition, the integration of PCA with other methods, such as those that analyse signals in both the time domain and the frequency domain, has demonstrated improved performance in the identification of cardiac abnormalities. Overall, PCA plays a significant role in ECG classification and

contributes to the development of accurate, interpretable, and computationally efficient diagnostic tools.

2.3.3.2 Support Vector Machines

SVM have demonstrated promising results in classifying ECG signals for arrhythmia detection. Numerous studies have combined SVM with feature extraction techniques to achieve high accuracy. SVM has been applied in ECG analysis not only for diagnosing CAD [81] but also for detecting other cardiac conditions such as AF [96–98], arrhythmia [99–101], and MI [102]. SVM has proven highly effective for ECG classification due to its ability to handle complex, high-dimensional data. Multiple studies have demonstrated the superior performance of SVM in detecting various arrhythmias, with accuracies reaching up to 99% [99–101]. The effectiveness of SVM is further enhanced when combined with advanced feature extraction methods, such as CNN or autoencoder [103]. These hybrid approaches tackle deep learning to perform automatic feature extraction while employing the robust classification capabilities of SVM.

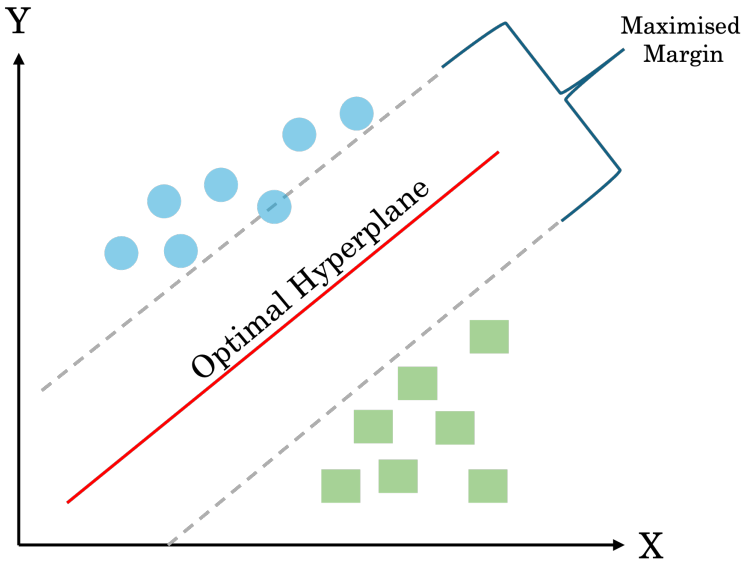

FIGURE 2.1: SVM classification showing the optimal hyperplane

In ECG binary classification, SVM is commonly applied due to their ability to create an optimal hyperplane that maximally separates two classes in a multidimensional feature space. Features such as RR intervals, QRS duration, and morphological characteristics of the ECG waveform are utilised as inputs. The support vectors closest to the maximised margin determine the decision boundary and play a crucial role in classification accuracy. Figure 2.1 illustrates how the SVM constructs the optimal hyperplane between the two classes by maximising the margin, thereby enhancing robustness and generalisability.

This method has proven effective in detecting cardiac abnormalities, providing reliable and interpretable outcomes for clinical diagnostics.

Although SVM has been widely employed for ECG classification, several limitations remain. The performance of SVM heavily depends on the quality of feature extraction, which often requires expert knowledge and extensive preprocessing. Furthermore, SVM can struggle with very large datasets due to computational complexity and may be less adaptable to the variability and noise inherent in ECG signals compared to deep learning models [97]. These challenges indicate that, while SVM remains a valuable tool, complementary approaches or hybrid models may be necessary to achieve optimal performance in real-world clinical applications.

2.3.3.3 K-Nearest Neighbours

KNN algorithm has attracted considerable attention from researchers for its efficacy in classifying ECG data, particularly in diagnosing arrhythmia, AF, and CAD [84, 104–106]. KNN assigns class labels based on the majority vote of the k nearest data points within the feature space, making it a straightforward yet effective method for ECG classification. Figure 2.2 highlights the importance of selecting an appropriate k value to optimise classification accuracy in ECG signal analysis. A new ECG segment (yellow square) is assigned to a class based on the majority label of its nearest neighbours in the feature space.

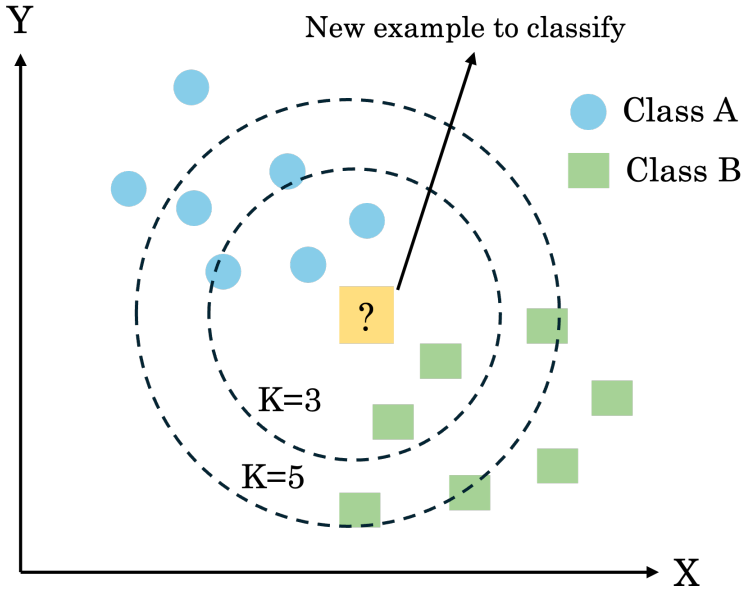

FIGURE 2.2: KNN classification example showing how the value of k affects the class assigned to a new data point based on its nearest neighbours.

Recent research has explored combining KNN with advanced techniques, such as deep learning-based feature extraction, to enhance its performance. Comparative studies have shown that KNN can outperform other classifiers in certain tasks. Nevertheless, challenges remain, including managing missing data within ECG recordings and improving the detection of minority classes [87]. In [107], the KNN algorithm with Euclidean distance and $k = 3$ reported QRS complex detection rates of 99.89% on the Common Standards for Electrocardiography (CSE) database and 99.81% on the MIT-BIH Arrhythmia database, demonstrating the model's robustness and reliability for accurate QRS detection. Furthermore, the integration of multirate ECG processing with KNN classification ($k = 5$) achieved an average classification accuracy of 91.87% across three cardiac classes, with the highest accuracy of 93.2% observed for Wolff-Parkinson-White (WPW) syndrome, thereby highlighting the potential of KNN-based methods for efficient and precise arrhythmia diagnosis [108].

2.3.3.4 K-means

K-means clustering is a widely used unsupervised learning technique for ECG analysis. Its ability to partition unlabeled data into distinct groups based on feature similarity makes it well suited for identifying patterns in complex ECG datasets. As shown in Figure 2.3, k-means iteratively assigns data points to clusters, refining these groupings until convergence is achieved. The algorithm has been effectively applied to CVDs disease classification [85].

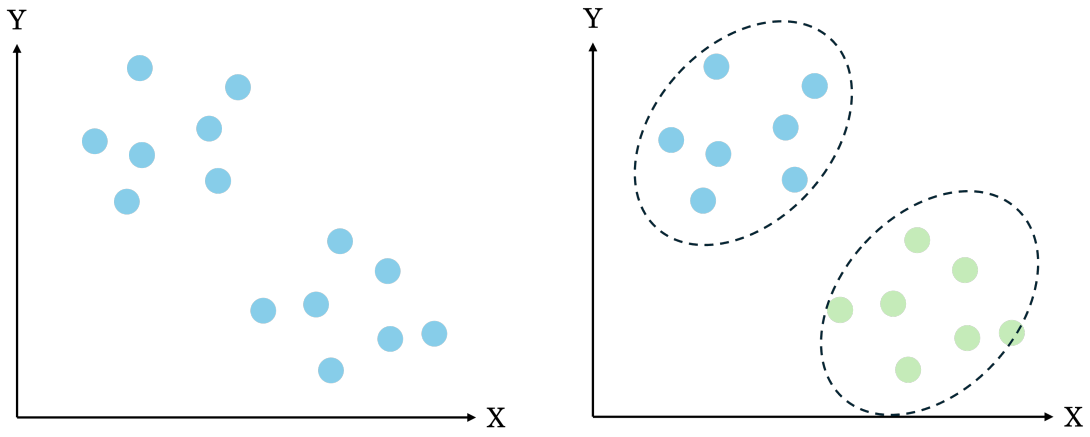

(a) Before K-means (b) After K-means

FIGURE 2.3: Visualisation of k-means cluster assignments

In [85], a robust algorithm was developed that combines an envelopment filter with K-means clustering for accurate QRS complex detection, achieving over 99.7% accuracy without the need for manual threshold tuning. This work underscores the effectiveness

of K-means in rapidly and reliably isolating key cardiac features. Similarly, in [109], K-means clustering was employed to segment ECG signals into regions with similar characteristics, allowing each segment to be processed according to its specific signal features. This facilitated more targeted analysis, supported downstream processing, and preserved clinically relevant information within the signal. Building upon these established strengths, recent research has extended the use of K-means clustering beyond classification tasks to support model interpretability. For instance, the ClusteredSHAP framework applies K-means clustering to group similar ECG samples, thereby reducing redundancy in explanation computations and substantially accelerating the generation of SHAP values for 12-lead ECG classification models [110].

2.3.4 Deep Learning Techniques

While traditional machine learning approaches rely on explicit feature extraction, deep learning techniques have emerged in recent years, including CNN [9, 31, 43, 111], LSTM [6, 112], hybrid CNN-LSTM models [9, 113], recurrent neural networks (RNN) [10, 114], Residual Neural Networks (ResNet) [115–117], and Autoencoders [118]. These methods have gained increasing attention in cardiovascular disease diagnostics, particularly for detecting MI and CAD. A key advantage of deep learning is its ability to perform feature extraction automatically within the learning process, thereby eliminating the need for manual and explicit feature engineering required by classical machine learning techniques. This automatic feature extraction capability arises from deep learning's robust pattern recognition and adaptable processing architectures. While these models have demonstrated exceptional performance in automatically extracting important ECG features, they require larger datasets and substantial computational resources, making deployment in real-time CAD diagnosis more challenging. Therefore, lightweight deep learning architectures have been introduced for the diagnosis of various heart diseases using ECG. In [119], a lightweight one-dimensional CNN model was proposed for cardiac arrhythmia diagnosis from ECG signals. Additionally, other DL-based models, such as SqueezeNet [120], EfficientNet [121], MobileNet [122], and ShuffleNet [123], have been explored to reduce computational complexity while maintaining high diagnostic accuracy, making them suitable for real-time and resource-constrained applications.

Researchers in [124] proposed a deep belief network (DBN) for the classification of CAD and NON-CAD cases. Features were extracted from short-term ECG signals using the Hilbert transform, with data sourced from the PhysioNet database. The model was evaluated using ten-fold cross-validation and achieved an accuracy of 98.1%. The DBN outperformed other approaches based on HRV, highlighting its potential for ECG-based

diagnosis. Although ECG signals are commonly used in deep learning-based classification of CVDs, they are not the only modality employed. Beyond the use of ECG, deep learning has been increasingly applied to medical imaging for CAD detection [125, 126]. These approaches aim to improve diagnostic performance through advanced image interpretation and feature extraction. In [125], a CNN was developed to classify myocardial perfusion imaging (MPI) data from single-photon emission computed tomography (SPECT), achieving high accuracy in identifying ischaemia. However, the study emphasised the need for larger imaging datasets to improve model generalisability. Similarly, in [126], a CNN was implemented for real-time detection of coronary artery lesions. Architectural enhancements included Rectified Linear Unit (ReLU) activation to reduce computational cost, residual learning to address training difficulties, and inception modules with skip connections for robust feature extraction. The model accurately detected stenotic lesions, although its applicability was limited by the restricted diversity of lesion types in the dataset. These studies collectively demonstrate the adaptability of deep learning approaches across different data modalities for CAD detection, while also highlighting common limitations related to dataset size, lesion diversity, and model generalisability. Recent studies have shown that CNN architecture is effective in various medical diagnostic tasks, such as detecting COVID-19 lesions in CT scans, diagnosing Alzheimer's disease from hippocampal MRI, classifying skin lesions in dermoscopy images, and identifying breast abnormalities in mammograms [127–130]. These examples demonstrate that CNN can successfully learn important medical features from complex data, supporting disease detection, localisation, and analysis across different types of medical images. CNN has also been applied to nonimaging biomedical signals, especially the ECG, which is widely used in the diagnosis of CVDs disease [131].

TABLE 2.2: Deep Learning Architectures in Heart Diseases Classification

Study	Architecture	Dataset(s)	Diagnosis
[9]	Hybrid CNN-LSTM	MIMIC II In-house	CAD
[6]	8-layer stacked CNN-LSTM	St Petersburg Fantasia	CAD
[132]	16-layer CNN-LSTM	St Petersburg Fantasia PTB Diagnostic BIDMC CHF	MI, CHF, CAD
[133]	DINN	Hospital	CAD
[134]	11-layer CNN	St Petersburg Fantasia	CAD
[124]	Deep Belief Network	St Petersburg Fantasia	CAD
[125]	CNN-RGB	In-house	MI
[126]	CALD-Net	In-house	CAD
[9]	Hybrid CNN-LSTM	MIMIC II In-house	CAD
[135]	1D-CNN	MIT-BIH	Arrhythmia
[131]	lightweight CNN	PTB-XL	CVDs
[34]	Gated CNN-Transformer	CSPC Challenge 2018	CVDs
[30]	patient-specific CNN	MIT-BIH	Arrhythmia
[7]	Hybrid CNN-LSTM	MIT-BIH	Arrhythmia
[136]	24-layer DCNN-BiLSTM	CINC challenge 2017	AF
[97]	Segment-based CNN with SVM	CINC challenge 2017	AF
[137]	DNN with residual blocks	MIT-BIH	Arrhythmia
[138]	1D-CNN	MIT-BIH	Arrhythmia
[139]	Hybrid CNN-LSTM	MIT-BIH	AF
[45]	CNN with CWT scalogram	MIT-BIH	Arrhythmia
[112]	LSTM	UCR Time Series	CVDs
[33]	Attention-based BiLSTM	In-house	AF
[140]	Residual Attention-CNN	MIT-BIH	Arrhythmia

Table 2.2 presents various deep learning models used for the classification of heart diseases. The table includes CNN, hybrid CNN-LSTMs, attention-based architectures, residual networks, and lightweight deep learning frameworks. These models have been evaluated using both public and private datasets, including MIT-BIH, Physikalisch-Technische Bundesanstalt (PTB) XL, the CinC Challenges, and in-house clinical data. The diagnostic targets include CAD, arrhythmia, AF, MI, and other cardiovascular conditions. The variety of approaches reflects the adaptability of deep learning methods to different ECG signal types and clinical objectives. This reflects a wider trend in recent research in which deep learning methods are often preferred over traditional machine learning techniques due to their ability to automatically extract relevant features, thereby reducing he need for manual feature engineering. Despite these advancements, many existing approaches rely on complex architectures or high computational resources, limiting their suitability for deployment in real-time or embedded clinical settings. To address this challenge, lightweight deep learning models are introduced later in the thesis. This model maintains high diagnostic accuracy while being optimised for efficient operation on resource-constrained device.

2.3.4.1 Convolutional Neural Network

CNN is a class of deep learning models well-suited to the analysis of data with spatial or temporal structure, including medical images and time-series signals such as ECGs. CNN architectures typically comprise convolutional layers for feature extraction, followed by fully connected layers for classification. In ECG classification tasks, the convolutional layers are designed to identify local patterns within the signal, such as QRS complexes or ST-segment deviations. The convolution operation may be expressed as:

$$y_i = \sigma\left(\sum_{j=0}^{k-1} w_j \cdot x_{i+j} + b\right) \tag{2.1}$$

where y_i denotes the output at position i, x is the input signal, w_j are the weights of the convolutional kernel of size k, b is the bias term, and σ represents a non-linear activation function, such as the ReLU. These layers enable the network to automatically learn clinically relevant temporal features from the ECG signal.

Following the convolutional layers, pooling layers are typically employed to reduce the spatial dimensionality of the feature maps while preserving the most relevant information. This downsampling operation reduces the computational complexity of the network and enhances generalisation by making the model less sensitive to minor shifts in the location of relevant features within the input signal. Pooling layers are widely

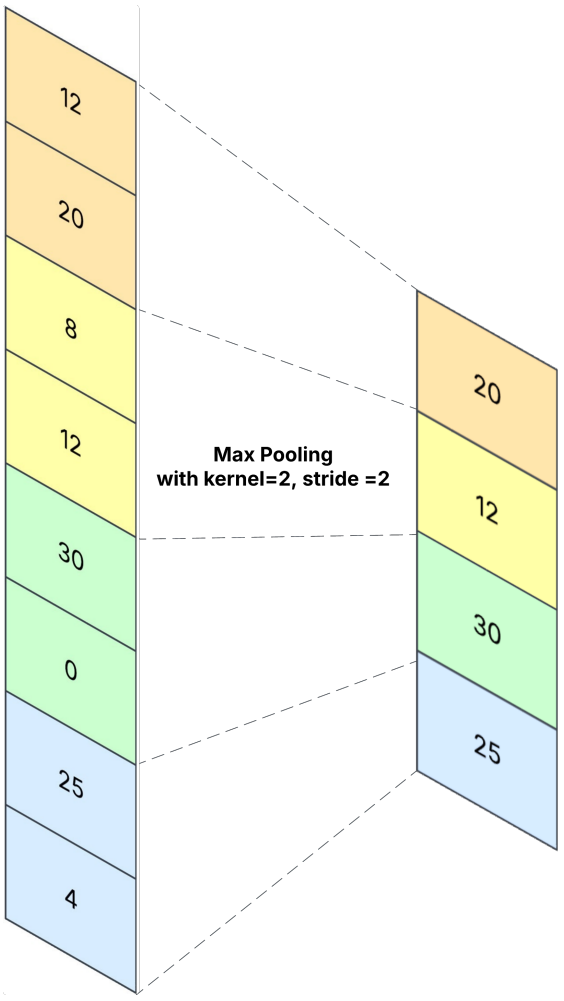

FIGURE 2.4: An example of max pooling operations

employed in studies on ECG classification. Among these, max pooling has consistently outperformed average pooling, offering improved feature extraction and contributing to higher classification accuracy [141]. Max pooling operates by extracting the maximum value from each local region of the feature map defined by the pooling kernel. In ECG classification, this enables the network to maintain the most important features of the signal while reducing spatial dimensionality as shown in Figure 2.4. Max pooling can mathematically be expressed as:

$$y_i = \max\{x_i, x_{i+1}, \ldots, x_{i+k-1}\} \tag{2.2}$$

where y_i denotes the max pooling output at index i, x represents the input signal, and k denotes the size of the pooling window.

Figure 2.5 illustrates the global average pooling technique, which is employed in CNN to summarise each feature map by computing the mean of all its activations. This

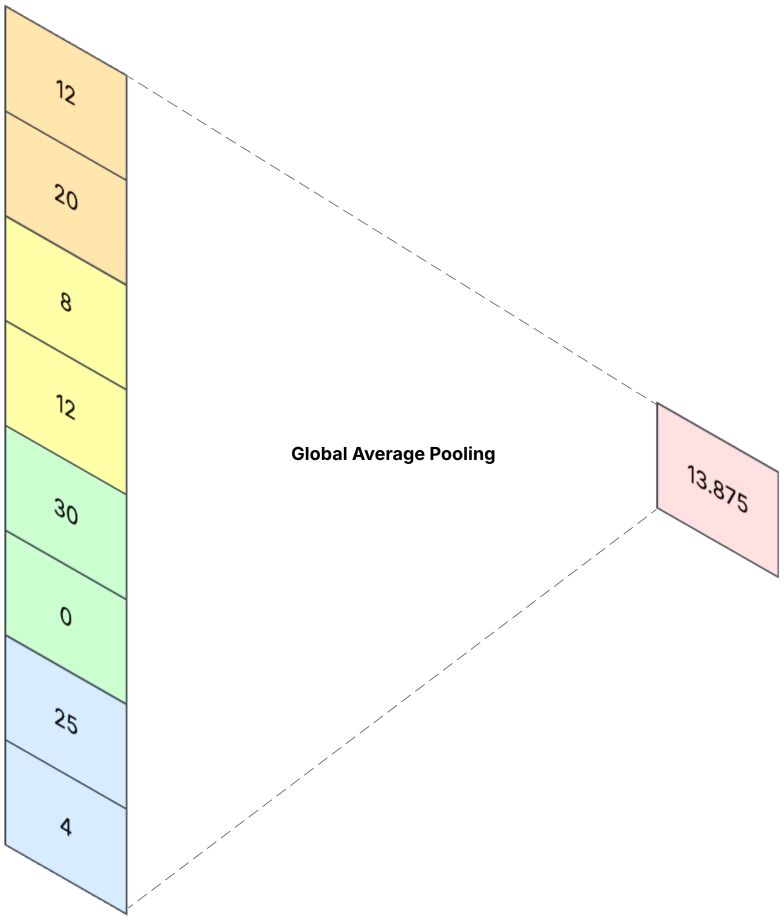

FIGURE 2.5: An example of global average pooling operations

operation reduces the spatial dimensions to a single scalar value per feature map and is commonly used as an alternative to fully connected layers in the final stages of the network. Furthermore, the integration of global average pooling has also been shown to enhance feature extraction in ECG classification [142]. Global average pooling can mathematically be expressed as:

$$y = \frac{1}{N} \sum_{i=1}^{N} x_i \tag{2.3}$$

where y denotes the global average pooling output, x_i represents the input signal, and N is the number of activations in the feature map.

Fully connected layers, also known as dense layers, are then employed at the final stages of CNN to map high-level features to the network's output. These layers establish connections between every neuron in the preceding layer and each neuron in the current layer, enabling the model to learn complex, non-linear relationships. In classification tasks, fully connected layers are often followed by activation functions (eg., sigmoid or softmax) to generate class probabilities.

The CNN has gained substantial attention in biomedical signal processing due to their ability to learn spatial hierarchies and extract features directly from raw data. Their effectiveness in identifying local patterns makes them well-suited for analysing ECG signals and detecting CVDs [30, 31]. Several studies have demonstrated the strong diagnostic capabilities of CNN-based models in classifying CAD and other cardiac conditions, often outperforming traditional machine learning approaches [6, 31, 134, 143–145]. CNN architectures have been employed for tasks such as detecting MI, arrhythmia, and left ventricular hypertrophy, achieving high accuracy, sensitivity, and specificity. For instance, Acharya et al. proposed a one-dimensional CNN for automated CAD detection, using ECG segments of two and five seconds to classify CAD and NON-CAD cases. Their approach integrated feature extraction and classification within a single framework, demonstrating promising results [134]. In [146], CNN-based models have significantly advanced ECG signal processing by enabling automatic and hierarchical feature extraction. These findings underscore CNNs' potential to improve diagnostic reliability in clinical applications. Although CNN is traditionally applied to image data, 1D-CNN has been successfully adapted for sequential ECG signals. These models are capable of learning relevant temporal patterns directly from raw input, reducing the need for manual feature engineering

While 1D-CNN is effective at capturing local spatial features in ECG signals, they may be limited in their ability to model long-term temporal dependencies. In addition, recent studies frequently employed CNN-LSTM architectures to capture both spatial and temporal characteristics of ECG signals. These hybrid models enable the joint learning of morphological features and temporal variations, such as HRV, thereby enhancing classification performance in the detection of CAD. Evidence suggests that combining spatial and sequential information improves the discriminative capacity of these models. Nevertheless, the need to incorporate additional clinical indicators, such as demographic data and patient history, to improve clinical relevance and generalisability remains [6, 9].

In addition, stacked CNN-LSTM models have enabled end-to-end classification of CVDs, significantly reducing the need for manual feature extraction. Although these models typically require large datasets and significant training time, they offer enhanced generalisability by capturing complex, non-linear relationships within the data [132]. This

is relevant in multi-class diagnostic settings, where deep learning models are employed to identify not only CAD but also related conditions, including MI and CHF. These models have demonstrated high classification accuracy, even in the presence of noisy signals, indicating a level of robustness. However, high computational demands and limited access to large, labelled datasets remain major challenges in applying these models to automated CVDs classification.

In addition to the traditional CNN, ResNets have been introduced for the classification of various CVDs. ResNet, a specific type of CNN, addresses several limitations associated with traditional CNN architectures. By incorporating shortcut connections, ResNets facilitate the construction of deeper network architectures, effectively mitigating the vanishing gradient problem often encountered in very deep networks [30, 32, 147]. Furtheremore, capsule networks have also been explored, with 1D-CADCapsNet achieving 99.44% accuracy on 2-second ECG segments [148]. Moreover, LSTM networks, a variant of RNN, have also been applied to cardiac diagnostics due to their ability to capture long-term dependencies in sequential data [149]. Given their complementary strengths, several studies have combined CNN with LSTM, where CNN extracts spatial features and LSTM model temporal dynamics over time [6, 136].

Beyond ResNet, attention mechanisms have recently been integrated into deep learning models to enhance their ability to focus on diagnostically important features in ECG data [33, 140, 150]. In addition, recent studies have explored CNN-LSTM-Attention frameworks that integrate the spatial feature extraction capabilities of convolutional neural networks with the temporal modelling strengths of long short-term memory networks. These architectures are further enhanced by self-attention mechanisms, which enable the model to focus more effectively on diagnostically significant segments of the ECG signal [151]. Furthermore, research has explored various input modalities for ECG classification, including raw time-series signals, spectrograms, and two-dimensional image representations of ECG waveforms [47, 152, 153]. In addition, multi-modality deep learning models that integrate both ECG waveforms and clinical features have demonstrated superior performance compared to traditional pre-test probability estimates [154]. These automated approaches have the potential to support clinicians in the early and accurate diagnosis of CAD, addressing some of the limitations associated with manual ECG interpretation [132, 155, 156]. Furthermore, this emphasis on automated feature extraction aligns with advancements in other biomedical domains. For example, a variational gated autoencoder (VGAE) has been applied to extract latent features from multiview biomedical data [157]. Such approaches illustrate the potential of deep learning architectures to handle complex, high-dimensional datasets, offering insights that may be extended to ECG analysis.

CNN has demonstrated significant promise in ECG classification, offering clear advantages over traditional machine learning methods. Their ability to automatically learn and extract relevant features directly from raw ECG signals eliminates the need for manual feature engineering, which is often time-consuming and prone to human error [158]. This automated feature extraction has contributed to high diagnostic accuracy across a range of cardiac conditions, including arrhythmias [159, 160], AF [98, 139], and CAD [6, 31]. In some applications, ECG signals are converted into two-dimensional representations, enabling CNN to more effectively capture both spatial and temporal features [161]. Furthermore, techniques such as data augmentation, transfer learning, and skip connections have been employed to enhance model robustness and training efficiency [162]. Recent studies have also applied CNN-based models in real-time ECG monitoring systems, demonstrating their potential for early detection of cardiac abnormalities in practical clinical settings [163, 164]. Overall, CNN provides a reliable and scalable solution for automated ECG analysis, improving both the accuracy and consistency of cardiac diagnosis. As research advances, their integration into clinical workflows may contribute to earlier intervention and better patient outcomes, particularly in environments where expert interpretation is limited.

Despite their demonstrated effectiveness in extracting spatial features from ECG signals, CNN exhibits several limitations in CVDs classification. A significant limitation of CNN is their reduced capacity to model long-range temporal dependencies, which are often necessary for identifying conditions that develop over extended time intervals. To overcome this, hybrid approaches integrating CNN with transformer-based architectures have been introduced, allowing for the extraction of both local and global features from ECG signals [165]. Additionally, class imbalance is a common issue in ECG datasets and may lead to biased model performance. Furthermore, although CNN models often achieve high classification accuracy, they are prone to overfitting when trained on small datasets. This issue can limit their ability to generalise to unseen data. Another important concern is the lack of interpretability in CNN models, which presents a barrier to clinical validation, where transparent and explainable decision-making is essential. Addressing these limitations requires improvements in model architecture, as well as the integration of interpretability tools and computational optimisation techniques. These developments are necessary to support reliable and real-time application of CNN models in environments with limited computational resources.

2.3.4.2 Attention Mechanism

Attention mechanisms are neural network components that allow models to dynamically focus on the most informative parts of an input sequence, rather than treating all elements as equally important. This is particularly advantageous in ECG analysis, where diagnostically relevant features such as abnormal QRS complexes, ST segment deviations, or other abnormality beats may appear infrequently within long and noisy recordings. The core operation involves computing a weighted sum of value vectors **V** based on the similarity between query vectors **Q** and key vectors **K**, as given by the scaled dot-product attention:

$$\textbf{Attention}(Q, K, V) = \textbf{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V, \tag{2.4}$$

where $Q = W_Q x$, $K = W_K x$, and $V = W_V x$ are the query, key, and value matrices derived from the ECG input x through linear projections, and W_Q, W_K, and W_V are learnable weight matrices. The term $\sqrt{d_k}$ scales the dot product to maintain gradient stability. The softmax function converts the similarity scores into attention weights, which are used to compute a weighted sum of the value vectors.

In ECG analysis, attention mechanisms are well-suited due to the characteristics of the signal. Pathological events such as AF, premature contractions, and ST-segment deviations tend to be localised, infrequent over time, and exhibit only slight or subtle changes in waveform shape. Traditional CNN and RNN architectures process signals either through fixed-length windows, which may lead to important features being overlooked or weakened during analysis. By contrast, attention mechanisms enable models to dynamically prioritise diagnostically important segments of the ECG waveform, allowing for more targeted feature selection during both training and inference. For example, in lengthy ECG recordings, an attention-based model can concentrate on arrhythmic beats or ischaemic episodes while disregarding normal segments. This not only improves classification performance but also enhances model interpretability, as attention weights can be visualised to highlight the regions of the signal that contributed most significantly to the prediction. In addition, when applied to multi-lead ECGs, attention mechanisms can operate across channels to capture inter-lead spatial relationships, which are essential for identifying conditions such as MI and CAD. Attention-based models have also shown strong performance in multilabel classification tasks [166], where the concurrent presence of multiple cardiac abnormalities must be taken into account. Recent research on ECG-based disease classification has explored the use of different attention mechanisms, including self-attention [34, 167, 168], channel attention [169, 170], and hybrid approaches [171, 172] that combine channel and spatial attention.

Self-Attention

Self-attention mechanisms have gained attention in ECG analysis due to their capacity to model global dependencies within the signal. A self-attention-based autoencoder architecture that integrates LSTM and 1D-CNN has achieved 99.71% accuracy and 99.86% sensitivity for arrhythmia detection on the MIT-BIH dataset, demonstrating the potential of temporal attention to enhance feature localisation and improve classification performance [167]. A dual-branch gated CNN-transformer architecture has been introduced to perform multi-disease ECG classification. This model integrates multi-head self-attention to extract both channel-wise and temporal features, thereby capturing spatial and sequential dependencies more effectively. Evaluation on the CPSC-2018 dataset obtained an accuracy of 92.2%, demonstrating the efficacy of self-attention mechanisms in handling complex ECG patterns [34]. In addition, a patient-adaptive beat-wise transformer has been proposed, which encodes symbolic ECG tokens to personalise classification for AF [168]. This approach has shown strong generalisation across long-term ECG recordings, supporting the potential of transformer-based models in real-world clinical applications.

Channel attention

Channel attention mechanisms enable neural networks to adaptively recalibrate the importance of individual feature channels, which is particularly useful in ECG classification involving multi-lead inputs. Each lead provides a distinct spatial perspective of cardiac electrical activity, yet not all leads are equally informative for the detection of specific conditions. Channel attention addresses this variability by allowing the model to focus on the most diagnostically relevant channels, thereby improving its discriminative capacity and overall classification performance. In [169], a channel attention mechanism based on the squeeze-and-excitation (SE) block was integrated into a convolutional neural network to enhance feature weighting at the lead level. The SE mechanism consists of two stages. The squeeze operation reduces global spatial information into a compact channel descriptor via global average pooling. This is followed by the excitation operation, which applies a learnable gating function to adaptively reweight each channel. When applied to the MIT-BIH arrhythmia dataset, the model achieved an accuracy of 99.2% for PVC classification and 97.4% for Supraventricular ectopic beats (SVEBs). Furthermore, a channel attention module was integrated with a spiking neural network and deployed on a field-programmable gate array (FPGA). The resulting system achieved real-time inference with an energy cost of 346.33 µJ per beat and a processing time of 1.37 ms per beat. These results demonstrate the feasibility of attention-based ECG classification in resource-constrained environments, such as wearable or ambulatory monitoring systems [170].

Hybrid approaches

Hybrid attention mechanisms combine channel and spatial attention to enable models to learn which leads are most informative and which regions of the signal require focus. ECG signals often exhibit spatial correlations across leads and temporal or morphological variability within waveforms. By incorporating both forms of attention, hybrid models can extract more meaningful and discriminative features than those using a single mechanism.A multiscale residual network incorporating both channel and spatial attention has been proposed for multilabel ECG classification. This model achieved F1 scores of 88.2% on the China Cardiovascular Disease Database (CCDD) dataset and 85.8% on the HF-challenge dataset, outperforming baseline architectures that did not incorporate attention mechanisms [171]. A further development introduced a guided spatial attention mechanism enhanced through the use of class activation maps (CAMs). Unlike approaches that rely solely on data-driven learning, this method incorporates clinical domain knowledge to guide the model's attention towards physiologically and pathologically significant regions. The use of CAMs during training enables the model to associate specific waveform segments, such as ST elevation or T-wave abnormalities, with the corresponding diagnostic labels. This hybrid attention strategy was evaluated across several ECG abnormality categories, including ST changes, atrial fibrillation, and WPW syndrome, and showed improved classification performance [172].

The integration of attention mechanisms into ECG classification offers several practical and methodological advantages. Firstly, attention-based models have demonstrated improved performance, achieving higher accuracy, sensitivity, and F1 scores across a range of CVDs classification [34, 167–172]. In addition to performance gains, attention mechanisms enhance model interpretability by enabling the visualisation of attention weights, which provides insight into the waveform segments that contributed to the decision-making process. This supports clinical validation and strengthens confidence in the model's outcomes. Attention mechanisms also enable the model to assign greater weight to ECG leads that contribute more significantly to the detection of specific cardiac conditions, while reducing the influence of less informative leads. This selective focus facilitates the learning of more discriminative representations, thereby improving classification performance and diagnostic reliability. Channel attention mechanisms further support the integration of information across multiple leads, which is beneficial in the interpretation of 12-lead ECG recordings [169]. Moreover, self-attention mechanisms can facilitate personalised modelling by adapting to patient-specific morphological variations through symbolic representations of ECG inputs [168]. Finally, several studies have demonstrated the suitability of deploying attention modules on energy-efficient hardware platforms, such as FPGAs, enabling real-time ECG monitoring in wearable or embedded systems [170, 172].

Despite their effectiveness, attention mechanisms present several technical limitations. Architectures that employ multi-head self-attention demand substantial computational resources, due to their memory complexity. This characteristic makes them less suitable for deployment in resource-constrained environments, such as wearable devices. Moreover, there is limited standardisation in the design and implementation of attention modules across existing studies. Furthermore, clinical validation of attention-based models also remains limited. Additionally, hybrid attention models that incorporate multiple attention mechanisms often introduce increased model's complexity. This can result in a larger number of parameters, increased training time, and reduced interpretability, which may further limiting their practicality for real-world clinical use.

2.4 Lightweight Neural Network Architecture

Lightweight neural networks achieve high predictive accuracy while substantially reducing computational complexity, memory requirements, and energy consumption. Compared to conventional deep networks, lightweight neural networks employ fewer parameters and streamlined architectures, thereby facilitating efficient inference on resource-constrained devices such as wearable devices.

Given the computational demands of traditional deep learning models that incorporate attention mechanisms or complex CNN architectures, there has been growing interest in lightweight neural network designs for ECG classification, particularly targeting deployment on resource-constrained devices. Techniques such as pruning, quantisation, and depthwise separable convolutions are commonly employed to reduce model complexity. These approaches enable real-time ECG signal processing and continuous patient monitoring by minimising computational and energy requirements. Lightweight CNN has gained significant attention due to their ability to offer faster inference and lower energy consumption, which are essential to facilitate rapid and reliable analysis in compact, low-power medical devices.

A study by Mewada [173], proposed a 2D-wavelet encoded deep CNN for ECG classification, transforming one-dimensional ECG signals into two-dimensional representations to incorporate spatial feature extraction capabilities. The author demonstrated that their approach improved arrhythmia detection accuracy without the need for extensive preprocessing. Similarly, another work examined lightweight CNN models applied to ECG datasets, emphasising the reduction of computational complexity while maintaining high classification performance for arrhythmias [174]. These approaches focus primarily on detecting abnormal heart rhythms, which are largely associated with electrical malfunctions in cardiac activity. Further, a study investigating lightweight CNN architectures

specifically for ECG signal classification showcased the effectiveness of optimised deep learning models in diagnosing arrhythmic events [175]. Additionally, researchers explored efficient CNN architectures to enhance real-time arrhythmia classification, optimising neural network layers to improve model interpretability and deployment in clinical environments [176]. In [177], a lightweight CNN for MI diagnosis was introduced, aiming to minimise computational and storage demands to facilitate deployment on portable devices. Utilising the PTB diagnostic database, various baseline CNN configurations were evaluated and compared to their proposed lightweight model. Results indicated that the proposed model maintained high accuracy while reducing model complexity. Several lightweight networks have been utilised for ECG analysis, including SqueezeNet [178], EfficientNet [179], MobileNet [180], and ShuffleNet [123]. These widely well-established models are specifically designed for deployment on resource-constrained devices, enabling efficient applications in such environments. In a notable advancement in novel technology, an ultra-lightweight end-to-end electrocardiogram classification neural network has been developed. The research employed advanced techniques aimed at reducing computational complexity while maintaining high-performance standards, making it suitable for integration into portable and wearable medical devices [181].

Designing deep learning models or ECG classification requires balancing model complexity with computational efficiency. While deeper architectures incorporating attention mechanisms often improve accuracy by capturing complex signal patterns, they increase parameter counts, memory usage, and inference time, thus limiting deployment on resource-constrained devices. Lightweight models address these challenges by employing techniques such as depthwise separable convolutions to reduce computational load, although this may sometimes compromise the model's ability to capture discriminative features critical for accurate diagnosis. Recent studies have examined these trade-offs, demonstrating that increased model complexity does not necessarily guarantee superior performance [182]. Nevertheless, lightweight architectures such as ShuffleNet and SqueezeNet achieve strong accuracy with fewer parameters [123, 175, 178], while novel models including CADNet and ArrhythmiNet maintain high accuracy alongside low computational demands [164, 183]. These advancements enable real-time, on-device ECG classification for arrhythmia, AF, CAD, and other cardiac abnormalities.

2.5 Implementation of Lightweight Models on Resource-Constrained Hardwares

Recent developments in lightweight neural network architectures have enabled efficient ECG classification on resource-constrained hardware where memory bandwidth, power consumption and computational throughput are critically limited. These models are engineered to optimise the trade-off between diagnostic accuracy and hardware resource utilisation, making them suitable for embedded systems and wearable platforms.

Table 2.3 summarises lightweight ECG models, their deployment platforms and targeted cardiac diagnostics. Quantised 1D-CNN, such as LMUEBCNet and SEmbedNet, have been implemented on STM32H7 microcontrollers, delivering high accuracy for ectopic beat classification [184]. Similarly, ULECGNet deployed on MSP432 microcontrollers achieves reliable arrhythmia detection, demonstrating feasibility on ultra-low-power processors [181]. Attention-based transformer architectures adapted for low-resource microcontrollers, such as the Tiny Transformer implemented on GAP9, utilise sparse self-attention mechanisms to sustain performance while lowering computational complexity [185]. FPGA implementations, including 1D-CNN deployed on the Xilinx Zynq ZC706 and spiking neural networks with channel-wise attention on Artix-7 devices, demonstrate hardware-accelerated inference with microsecond-scale latency and significantly reduced energy consumption [170, 186]. A combination of LSTM and RNN architectures with wavelet preprocessing has been effectively utilised for continuous ECG monitoring on consumer wearable devices such as the Moto360 Android Wear [187], reflecting the growing adoption of on-device processing for real-time cardiac monitoring. Furthermore, quantised 1D-CNN with residual connections deployed on Raspberry Pi platforms highlight the trade-off between reducing computational complexity and maintaining diagnostic performance in arrhythmia classification [188]. Additionally, lightweight ConvLSTM2D architectures implemented on STM32F746G microcontroller enhance ECG abnormality detection capabilities within resource-constrained environments [189], highlighting the adaptability of hybrid CNN-RNN models for time-series ECG data. These implementations illustrate consistent advancements in optimising advanced ECG classification models for deployment in resource-limited environments.

Table 2.3: Lightweight ECG Classification Models on Resource-Constrained Hardware

Study	Model Architecture	Hardware	Diagnosis
[184]	LMUEBCNet SEmbedNet	STM32H7	Ectopic beat
[181]	ULECGNet	MSP432	Arrhythmia
[185]	Tiny transformer	GAP9	Arrhythmia
[186]	1D-CNN	Xilinx Zynq ZC706	Arrhythmia
[170]	spiking neural network with channel-wise attention	Artix-7	Arrhythmia
[189]	ConvLSTM2D	STM32F746G	ECG abnormality
[187]	LTSM RNN Wavelet	Moto360 AndroidWear	ECG monitoring
[188]	Quantised 1D-CNN with residual connections	Raspberry Pi	Arrhythmia

Despite significant advances in lightweight ECG classification models for resource-constrained devices, research targeting CAD detection remains limited. Most existing studies focus on arrhythmia classification, creating a gap in the development and deployment of lightweight models for CAD diagnosis. This gap will be addressed in this thesis by developing lightweight deep learning models optimised for accurate and efficient CAD detection on resource-constrained hardware platforms.

Most lightweight architectures focus on detecting rhythm abnormalities, such as arrhythmia, which typically exhibit clear temporal irregularities. In contrast, CAD-related ECG waveforms are often morphology-based, making them more difficult to capture using models designed for arrhythmia detection. However, several reported architectures demonstrate the potential of lightweight 1D-CNN for ECG-based diagnosis, which could be adapted for CAD detection. While these approaches benefit from low computational complexity, the use of a single ECG lead may limit the ability to learn clinically relevant CAD-related features. To address this limitation, this study introduces a 12-lead ECG-based diagnosis model suitable for deployment in resource-constrained environments. In summary, although existing lightweight models can be adapted for CAD diagnosis, such adaptations often involve trade-offs between diagnostic performance, computational efficiency, and deployability, highlighting the need for CAD lightweight architectures.

2.6 Single lead and multi-lead ECG classification

ECG signals may be recorded using either a single lead or a multi lead configuration. A single lead ECG records electrical activity along one electrical axis and is commonly used in wearable or portable monitoring devices due to its simplicity and reduced data requirements. In contrast, a multi lead ECG, such as the standard 12-lead system, captures signals from multiple anatomical orientations and provides a more comprehensive representation of heart function. Single lead ECG data are well suited to lightweight models that extract key features from limited input. These models align with the resource-constrained devices and real-time applications. In contrast, multi-lead data offers more comprehensive data support for more complex architectures that capture inter-lead dependencies. While these models can improve diagnostic accuracy, they require greater computational resources and larger training datasets. The choice of lead configuration therefore, influences model design, performance, and suitability for deployment in both clinical and practical settings

Studies have investigated the application of CNN for CAD detection using ECG signals, with both single-lead and multi-lead approaches demonstrating encouraging results. Single-lead CNN has shown comparable performance to 12-lead systems, with only an average reduction of 8.7% in AUC [131]. Both single-lead and multi-lead CNN have exhibited strong performance in detecting CAD, with single-lead ECGs achieving accuracy close to that of multi-lead models [134, 190]. Advanced techniques such as multiscale analysis and lead-asymmetric pooling have further enhanced the performance of multi-lead CNN [191, 192]. Although multi-lead CNN generally offers higher accuracy, the performance gap between multi-lead and single-lead models is narrowing. Therefore, single-lead approaches are becoming increasingly practical for CAD screening, particularly in wearable device applications [31, 131]. The choice between single-lead and multi-lead ECG classification is largely determined by the intended application and the constraints of available resources. Multi-lead ECG remains the clinical standard due to its broad diagnostic capability [93, 169]. In contrast, single-lead ECG is increasingly employed in wearable and mobile health technologies due to its simplicity, ease of use, and lower hardware requirements [187].

2.7 Performance metrics

Performance metrics commonly used include accuracy, sensitivity, specificity, precision, F1-score, and the area under the receiver operating characteristic curve (AUC-ROC) [193, 194]. These measures are widely applied in the evaluation of machine learning and deep learning models for ECG-based classification of CVDs, including CAD. Each metric reflects a different aspect of model performance and is selected based on the clinical objective and the characteristics of the dataset.

Binary classification

To evaluate the performance of the proposed model, various performance metrics were introduced. Accuracy (Acc), Precision, Recall, F1 score, and Area Under the Curve (AUC) are used to assess performance.

To evaluate how the model performs on both CAD and Normal classes, Acc is measured as follows:

$$Acc = \frac{TP + TN}{TP + TN + FP + FN} \tag{2.5}$$

where True Positives (TP) refer to CAD cases correctly classified as CAD. True Negatives (TN) represent Normal cases that the model accurately identifies as Normal cases. False Positives (FP) occur when Normal cases are incorrectly classified as CAD, whereas False Negatives (FN) arise when CAD cases are mistakenly identified as Normal cases.

Recall (Sen) defines the proportion of all CAD cases that were classified correctly as CAD cases.

$$Recall = \frac{TP}{TP + FN} \tag{2.6}$$

Precision (Ppr) refers to the proportion of predicted CAD cases that are correctly identified as true CAD cases, assessing the reliability of positive classifications made by the proposed model.

$$Precision = \frac{TP}{TP + FP} \tag{2.7}$$

Specificity (Spr) refers to the proportion of NON-CAD cases correctly identified by the model, indicating its ability to reduce false positive predictions.

$$Specificity = \frac{TN}{TN + FP} \tag{2.8}$$

F1 score denotes a balance measure between precision and recall.

$$F1\ score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{2.9}$$

Misclassification Rate quantifies the proportion of incorrect predictions made by the model.

$$Misclassification\ Rate = 1 - Accuracy \tag{2.10}$$

Furthermore, AUC is utilised to quantify the model's ability to differentiate between CAD and NON-CAD cases as shown in (2.11).

$$\text{AUC} = \sum_{i=1}^{n-1} \frac{1}{2} \cdot (FPR_{i+1} - FPR_i) \cdot (TPR_i + TPR_{i+1}) \tag{2.11}$$

where:

$$FPR = \frac{FP}{FP + TN} \tag{2.12}$$

and FPR_i signifies the percentage of cases without CAD that are inaccurately classified as having CAD at the i^{th} threshold. Here, the threshold is defined as the probability cut-off used to determine whether a case is classified as CAD or NON-CAD. TPR_i (Sen) represents the proportion of cases with CAD that are correctly identified as having CAD at the i^{th} threshold. n is the total number of thresholds. $FPR_{i+1} - FPR_i$ is the difference in the proportion of cases without CAD that are incorrectly identified as having CAD between two consecutive thresholds. $TPR_{i+1} - TPR_i$ demonstrates the total proportion of individuals with CAD correctly identified as having CAD at the two consecutive thresholds.

Additionally, t-distributed Stochastic Neighbour Embedding (t-SNE) analysis is carried out to visualise high-dimensional feature representations in a two-dimensional space, offering insights into the ability of each key module within the models to differentiate between classes. The t-SNE equation is defined as follows:

$$KL(P \parallel Q) = \sum_i \sum_{j \neq i} p_{ij} \log \frac{p_{ij}}{q_{ij}} \tag{2.13}$$

where p_{ij} and q_{ij} denote the joint probabilities of a pair of ECG data points, i and j, in the high-dimensional feature space and the corresponding low-dimensional space, respectively.

Multi-class classification

In the case of multi-class classification, precision, recall, and F1 score are computed independently for each class. A macro-average is then calculated to provide an overall performance score, giving equal weight to each class regardless of class frequency.

2.8 Conclusion

In this chapter, an overview of ECG signal processing and its role in the diagnosis of heart diseases has been presented. Section 2.2 discussed the challenges and limitations of manual ECG interpretation in clinical practice, including the potential for diagnostic errors and the reliance on expert knowledge, which underscore the need for automated and reliable computational approaches. In Section 2.3, both traditional machine learning and modern deep learning techniques applied to CVDs diagnosis were explored. This included a detailed examination of feature extraction methods, SVM, KNN, and k-means clustering within traditional machine learning frameworks, alongside CNN and attention mechanisms within deep learning approaches. The discussion highlighted the strengths and limitations of these methods with regard to their diagnostic accuracy, interpretability, and computational demands. Furthermore, lightweight neural network architectures designed for efficient ECG classification were explored, with an emphasis on balancing model complexity and computational efficiency to enable real-time processing on resource-constrained devices. The implementation of such lightweight models on embedded hardware platforms was also examined, addressing practical challenges including memory limitations, power consumption, and inference latency, as discussed in Section 2.4 and 2.5. Section 2.6 reviewed the literature on both single-lead and multi-lead ECG classification methods, discussing their relative strengths and limitations in detecting various cardiac abnormalities. Section 2.7 then summarised the key performance metrics utilised in this thesis. In the next chapter, CAD diagnosis using 1D-CNN is presented. The aim is to used deep learning techniques rather than traditional machine learning methods, thereby removing the requirement for manual feature extraction and improving diagnostic accuracy.

Chapter 3

Coronary Artery Disease Classification Using One-Dimensional Convolutional Neural Network

3.1 Overview

To address early CAD detection, we propose utilising 1D-CNN to enhance detection accuracy and reduce network complexity. This study goes beyond traditional diagnostic methodologies, leveraging the remarkable ability of 1D-CNN to interpret complex patterns within ECG signals without depending on feature extraction techniques.

We investigate the impact of varying sample lengths on model performance and conduct experiments involving layers reduction. The ECG data employed were obtained from the PhysioNet databases, namely the MIMIC III and Fantasia datasets, with respective sampling frequencies of 125 Hz and 250 Hz. The highest accuracy for unseen data is obtained with a sample length of 250. These initial findings demonstrate the potential of 1D-CNNs in CAD diagnosis using ECG signals and highlight the sample size's role in achieving high accuracy. These preliminary findings demonstrate the potential of 1D-CNN for reliable CAD detection from ECG signals and support further exploration of lightweight architectures for clinical and real-time applications.

3.2 Related works

In light of recent technological advancements, a multitude of investigators has devised cutting-edge computational diagnosis systems to facilitate the diagnosis of diverse CVDs [25, 134, 195]. The analysis of ECG signals for diagnosing CVDs has gained significant attention and has been the focus of an increasing number of investigations. Numerous deep learning techniques have been employed to classify heart diseases using ECG signals; CNN, LSTM, RNN, and autoencoder [114, 118]. In recent years, there has been a significant surge in interest among researchers towards using deep learning techniques to diagnose CAD. The majority of researchers have primarily focused on employing CNN techniques for the diagnosis of AF [195], MI [196], and arrhythmia [7]. However, the current research in CAD diagnosis is not yet conclusive due to the limited availability of data and the complex nature of ECG signals in CAD diagnosis. Therefore, a relatively small group of researchers have conducted their work on CAD diagnosis using similar techniques [6, 9, 134]. In [6], the 1D-CNN was combined with LSTM for CAD diagnosis. It was implemented to extract relevant features from CAD ECG signals. Then, LSTM and a fully connected layer were utilised to conduct the classification. The model is fully automated and needs less feature engineering. However, the limitation of CAD data caused a lower diagnostic performance. Several studies have implemented 1D-CNN for the automated detection of CAD, aiming to enhance diagnostic accuracy and improve patient outcomes. For instance, Acharya et al. [134] proposed an automated CAD diagnosis system based on 1D-CNN, which demonstrated promising results in terms of both accuracy and computational efficiency. Feature extraction techniques were combined into a model structure, which later obtained reliable accuracy. However, the model training was time-consuming and needed a large amount of data. Feature extraction still plays an important role in the data pre-processing stage by identifying and selecting informative features in numerous ECG signal processing works [133, 134, 197, 198].

In this study, the primary focus is to investigate the potential applications of a novel and compact 1D-CNN architecture with reduced complexity, with a specific aim of early onset detection. Early detection of CAD is crucial, as it enables timely and suitable treatment, resulting in better health outcomes for patients. To achieve this goal, the proposed 1D-CNN architecture will be developed in order to maintain high performance and minimise computational resource usage. This enhances the model's applicability when performing real-time processing with limited resources. Additionally, the proposed model will be applied to CAD ECG signals obtained from the MIMIC database. The model's purpose is to capture patterns and distinctive waveform characteristics that serve as markers for the early stages of CAD. To ensure consistent input for our model, a

data normalisation technique is implemented to standardise and adjust the ECG signal data, mitigating the impact of noise, variations and artefacts in medical data. By elaborating on these novel aspects of our 1D-CNN architecture, this study presents a comprehensive and impactful contribution to the field of CAD detection. The model's innovation lies in its efficient architecture, optimising filter counts and kernel size while using dropout layers strategically to enhance early-onset CAD detection precision while conserving computational resources. Through its advancements, the proposed model has the potential to revolutionise early onset CAD detection, ultimately leading to improved patient care and outcomes.

3.3 Methodology

The proposed method consists of three main steps: data collection, pre-processing, and deep learning model. The model will be designed and implemented for CAD classification using the ECG signals through our extensive experiments. Each step will be explained in detail in the following sections.

The main portion of ECG data used for training and testing is obtained from the MIMIC III and Fantasia database from the PhysioNet website [199, 200]. A total of approximately 2,840 patients, constituting approximately 7.1% of all hospital admissions, are identified as having coronary atherosclerosis of the native coronary artery in the MIMIC database. The Fantasia database contains ECGs of 40 healthy patients, including 20 young and 20 adult patients. Three distinct subsets of data are generated for the experiments; D_1, D_2 and D_3. The first subset (D_1) is created by selecting a cohort of 5 individuals diagnosed with CAD from the MIMIC database, and 5 healthy individuals are chosen from the Fantasia database for the purpose of training and testing the model. The second subset (D_2) is specifically composed to examine the predictive capabilities of our model further. It comprises 20 CAD subjects from the MIMIC database, alongside 20 NON-CAD individuals from the Fantasia database. A third subset (D_3) was compiled by selecting patients diagnosed with CAD from the St. Petersburg database [201]. The St. Petersburg database comprises a total of 7 CAD subjects, with each subject's record spanning a duration of 30 minutes. Each record in the St. Petersburg database consists of 12 standard leads, sampled at a frequency of 257 Hz. The subset D_2 and D_3 are then utilised for prediction.

TABLE 3.1: Summary of ECG datasets used in the experiments.

Dataset	Source(s)	CAD subjects	NON-CAD subjects
D_1	MIMIC III and Fantasia	5	5
D_2	MIMIC III and Fantasia	20	20
D_3	St. Petersburg INCART	7	–

3.3.1 Data pre-processing

In this thesis, all experimental evaluations employ subject-independent data partitioning, ensuring that ECG recordings from the same individual are not included in both the training and testing sets. The ECG signals were obtained from patient records, each exhibiting different lengths of signal recordings spanning several minutes. To conduct our experiments effectively, specific segments of these ECG signals were chosen. Initially, the ECG signal data was retrieved from the records of each patient, as shown in Figure 3.1. Subsequently, each ECG signal was selected, ranging from 0 to 1000 samples. This segment corresponds to approximately 8 seconds of signal data. The selected data contains a complete cycle of the cardiac waveforms and is then stored in a dataframe. Prior to inputting into the classifier, the pre-processed data is subjected to labeling. A binary label was assigned to each ECG segment within subsets. Specifically, a label of 0 indicated NON-CAD subjects, while a label of 1 indicated CAD subjects. This crucial step is essential for building the basis of supervised learning. The classifier can then acquire valuable features and make informed predictions based on the provided labels. During experiments, the sample lengths were potentially segmented to accommodate the study of the impact of varying lengths on model performance.

Data normalisation was then employed to transform numerical data into a standardised range, typically between -1 and 1. This process is achieved by scaling the data based on its mean and standard deviation or by applying a linear transformation to shift and re-scale the data. The standard deviation formula was used for re-scaling, as shown in 3.1.

$$s = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \overline{x})^2} \tag{3.1}$$

where s is the normalised signal, N is the number of samples, $\overline{x}$ is the average of a given signal, and x_i is the signal value at the i^{th} position in ECG data. The standard deviation measures the spread or dispersion of the signals in the dataset. A smaller standard deviation indicates that the signals are clustered closely around the average, while a larger standard deviation indicates that the signals are more spread out. By

re-scaling the data using standard deviation, we obtained normalised data, which was then utilised in classification.

3.3.2 CNN model

The proposed CNN model was designed as shown in Figure 3.2 and comprises four one-dimensional convolutional layers, a max-pooling layer, four dropout layers, a flattened layer, and a fully connected dense layer. The architecture was selected to enable effective temporal feature extraction from ECG signals while maintaining a compact and computationally efficient design. The first layer of the network comprises 512 filters with a kernel size of 32, and the subsequent layers contain 256 filters with the same kernel size. By utilising 512 filters with a kernel size of 32 in the initial layer and subsequently reducing the number of filters to 256 while keeping the kernel size consistent, the reduction in parameters contributes to enhancing the model's compactness. The ReLU activation function was used in the convolutional layers to introduce non-linearity into the model. Three dropout layers with a rate of 0.2 were added after the convolutional layers to prevent overfitting. The max pooling layer with a pool size of 128 was then applied to reduce the spatial size of the feature maps and improve generalisation. The flattened output of the max pooling layer was then fed into a fully connected (dense) layer with 128 neurons and ReLU activation, which enabled the model to learn complex representations of the input data. To further prevent overfitting, another dropout layer with a rate of 0.5 was introduced prior to the final output layer. The final output layer consists of two neurons and softmax activation, which enables the model to classify the input data into one of two possible categories. Additionally, to optimise the model's performance, the Adam optimizer with a learning rate of 0.0001 was chosen for parameter optimization. Adam's adaptiveness in adjusting the learning rate for each parameter based on past gradients and magnitudes is particularly beneficial for training CNNs, especially in ECG, where the model must effectively navigate complex, high-dimensional parameter spaces to accurately classify ECG data. The model was compiled using the binary cross-entropy loss function, which is particularly effective for binary classification tasks, such as distinguishing between CAD and NON-CAD. The hyperparameters used during model training, including learning rate, batch size, and number of epochs, are summarised in Table 3.2.

Figure 3.1 illustrates the process of classifying ECG signals using the proposed 1D-CNN model. The process starts with raw ECG signals, which are normalised to reduce the impact of variations in amplitude and baseline. The normalised signals are then inputted into the 1D-CNN model, which processes the data and extracts relevant features. The model subsequently uses these features to classify the signals as either CAD

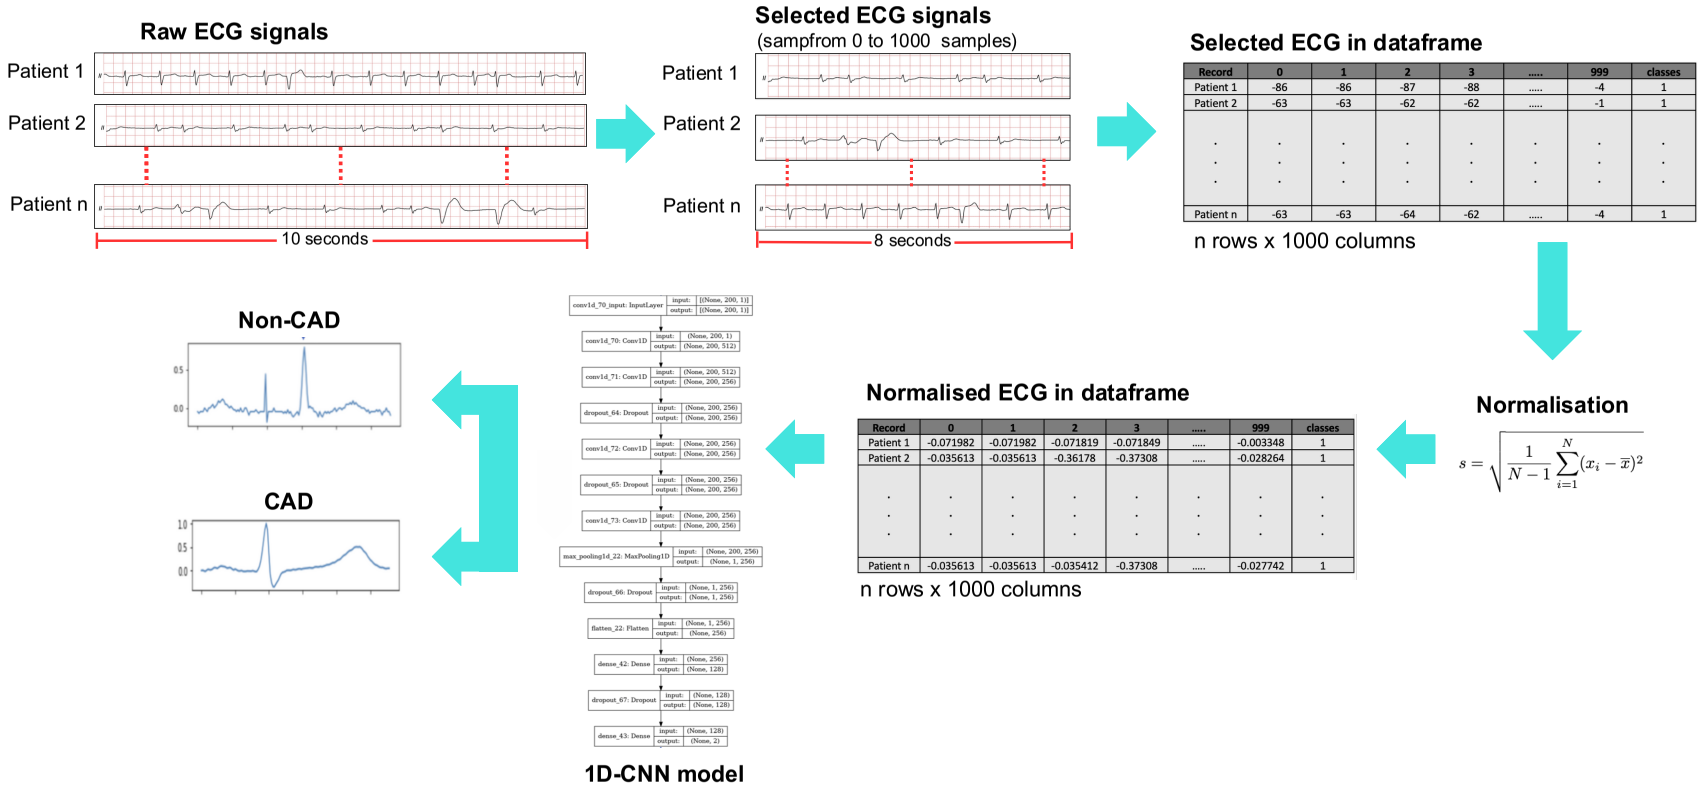

FIGURE 3.1: Flowchart of the proposed approach

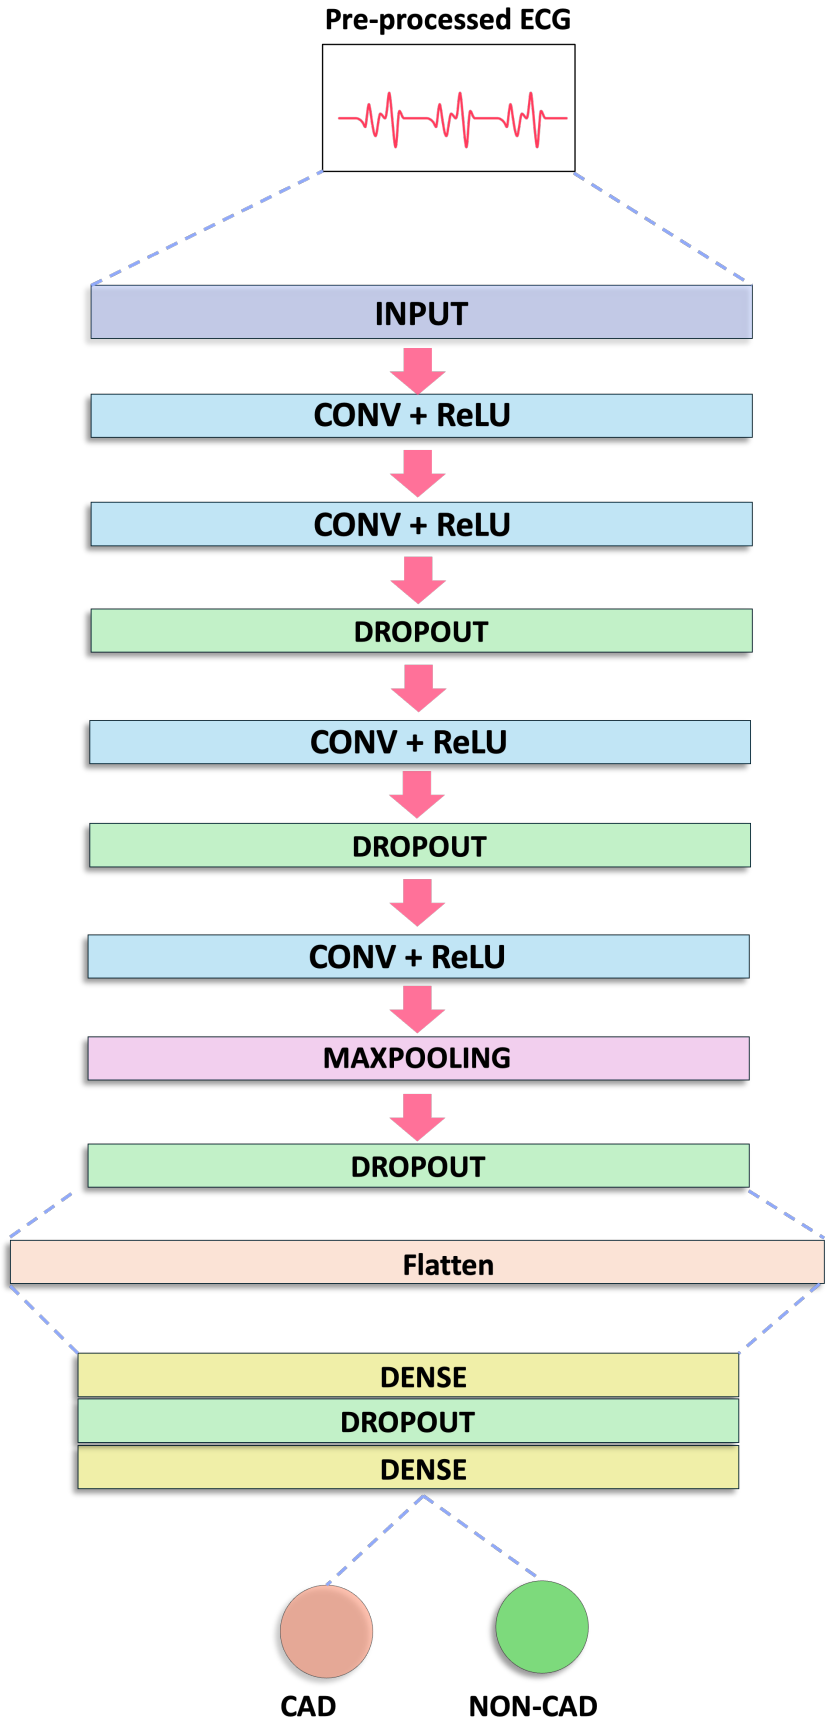

FIGURE 3.2: Overview of the proposed one-dimensional convolutional neural network model architecture for ECG-based CAD classification.

TABLE 3.2: Hyperparameters used for model training

Hyperparameter	**Value**
Learning Rate	0.0001
Batch Size	32
Number of Epochs	50
Optimiser	Adam
Loss Function	Binary Cross-Entropy
Hidden Layer Activation	ReLU
Output Layer Activation	Softmax

or NON-CAD. In our study, we proposed a modified 1D-CNN model that included some modifications, such as changes in filter size and the incorporation of dropout layers [202]. This model was employed throughout the CAD analysis process.

To evaluate the model, performance metrics including Accuracy, Misclassification Rate, Precision (Ppr), Sensitivity (Sen), and Specificity (Spr) were used to measure the performance of the classification model as shown in Equations (2.5), (2.10), (2.7), (2.6), and (2.8).

3.4 Experimental Results

In our experiments, we utilised the three subsets prepared during the data preparation stage. The subset D_1 was split into 70% for training and 30% for testing. This split is commonly used in machine learning because it allows for a sufficient amount of data to be used for training, while also providing enough data for testing the model's generalisation ability. Furthermore, the subset D_2 and D_3 were used for prediction where the trained model was put to test with these entirely new and unseen subsets.

3.4.1 Results and Discussion

Table 3.3 presents the overall performance of the 1D-CNN model in classifying ECG data into CAD and NON-CAD categories using different sample lengths, as we aim to determine the optimal sample length for accurate CAD diagnosis using the proposed 1D-CNN model. In the experiment, sample lengths of 1000, 500, 300, 250, 200, and 150 data points were manually selected from the lead II of each ECG subject. Varying the sample length of the input signal can reveal the impact of signal length on the classification model's accuracy. A longer sample length may provide more information about the ECG signal but may also require more sophisticated techniques and longer

processing times. On the other hand, a shorter sample length may not be as complex but may lead to lower accuracy due to the loss of critical information in the ECG signal. Hence, identifying the optimal length is essential.

Furthermore, the table shows the results of the experiment conducted on varied lengths of sample size on three different subsets. The model's accuracy in the subset D_1 was highest when the sample length was 300 data points, with training accuracy of 100% and testing accuracy of 96%, respectively. However, the model's accuracy remained relatively high across all sample lengths for all types of data. Moreover, it indicates that smaller sample lengths generally lead to slightly lower accuracy for train and test data in some sample lengths. The accuracy is significantly increased for unseen data in subset D_2 and D_3 when smaller sample sizes are employed. The results indicate that the model achieved the highest accuracy for unseen data when the sample length was 250, with an accuracy of 82.5% in D_2 and 85.7% in D_3.

Figure 3.3 illustrates the performance metrics for CAD detection in D_2 and provides valuable insights into the model's effectiveness. With an overall accuracy of 82.5%, the model demonstrates a solid capability in distinguishing between CAD and NON-CAD instances, although some misclassifications remain for both classes. The misclassification rate of 17.5% indicates room for improvement in accurately categorising cases. A precision of 85% represents that when the model identifies a positive case as CAD, it is correct approximately 85% of the time, showcasing its ability to minimise false positives. A recall of 80% reflects the model's success in capturing about 80% of actual CAD cases, which is crucial for avoiding missed diagnoses. Additionally, a specificity of 84% highlights the model's proficiency in accurately identifying negative cases, implying a satisfactory ability to distinguish NON-CAD instances.

Overall, the results suggest that a sample length of 250 data points might be optimal for achieving the highest accuracy in subsets D_2 and D_3 while still maintaining high accuracy for the train and test data in subset D_1. This finding could be due to the presence of key features in the ECG signals that indicate CAD, such as ST segments and other important ECG features that may be better represented in a sample length of 250. However, further research is needed to confirm this finding and to explore other factors that might impact the model's performance.

Table 3.4 provides a comparative overview of the performance of three distinct approaches for CAD detection using the MIMIC III dataset. The baseline 1D-CNN approach described in [9], when evaluated on our dataset, demonstrates a reasonable level of performance, achieving an accuracy of 83% on the training set. This result indicates its capacity to learn from the training data and identify patterns associated with CAD.

TABLE 3.3: An overall performance of 1D-CNN on CAD classification using different sample lengths on three subsets.

Sample length	Accuracy (%)			
	Subset 1 (D_1)		Subset 2 (D_2)	Subset 3 (D_3)
	Train	Test	Unseen	Unseen
150	95.5	85	75.5	81.3
200	94.6	90	73.8	85.6
250	97.3	89	82.5	85.7
300	100	96	71.9	82.4
500	98	82	66.7	63
1000	95.6	89	63.6	65

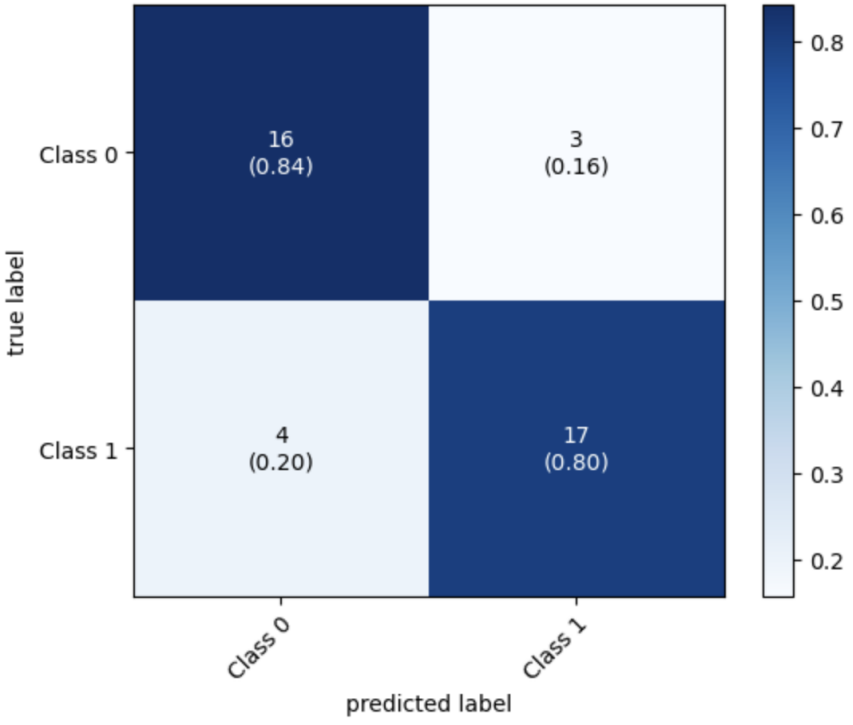

FIGURE 3.3: Confusion matrices on dataset D_2 with 250-sample length

However, a noteworthy observation is the decrease in accuracy to 74% on the test set. The complexity of the model could lead to the capturing of irrelevant features during training, resulting in a noticeable decrease in testing accuracy. The Hybrid CNN-LSTM approach enhances the ability to identify relevant CAD features by introducing LSTM layers [9]. This model achieved an accuracy of 94% in both the training and test sets, indicating effective generalisation and feature extraction capabilities. Lastly, our proposed model demonstrated remarkable accuracy, achieving 97.3% on the training set and

TABLE 3.4: Comparison of existing approaches performance on CAD applications using the MIMIC III dataset.

Architecture	Accuracy (%)	
	Train set	Test set
Baseline 1D-CNN [9]	83	74
Hybrid CNN-LSTM [9]	94	94
Proposed model	97.3	89

TABLE 3.5: Comparison of dropout layer configurations and probabilities in the proposed model

Architecture	Accuracy(%)	
	Unseen (D_2)	Unseen (D_3)
no dropout layers	62	60
one dropout (0.2)	68	71
two dropout (0.2)	65	57
three dropout (0.2)	65	62
three dropout (0.2) and a dropout of (0.5)	79	86

89% on the test set. However, the noticeable drop in accuracy on the test set warrants further exploration and investigation.

In addition, an ablation experiment encompassing diverse configurations of dropout layers was conducted to evaluate the efficacy of the integration of these layers on the optimal sample length, as shown in Table 3.5. The results of these experiments reveal a significant improvement in the performance of subset D_2, achieved through the incorporation of four dropout layers, each configured with dropout rates of 0.2 and 0.5. Notably, this configuration achieves the highest accuracy of 79% in D_2 and 86% in D_3. This improvement strongly suggests that adding dropout layers helps enhance the model's ability to generalise and effectively addresses concerns about overfitting. However, it is important to exercise caution when considering the inclusion of more dropout layers or higher dropout rates, as these adjustments may not necessarily lead to further performance gains. In fact, excessive dropout can potentially hinder the network's learning capacity. The experimental results demonstrate the effectiveness of the 1D-CNN model in accurately classifying ECG data into CAD and NON-CAD categories, regardless of the sample length. The model achieved balanced performance for both train and test data across all sample lengths, with the highest accuracy observed when the sample

length was set to 250 data points. This indicates that the model was able to learn and generalise well from various ECG samples, regardless of their length. Interestingly, the experimental findings suggest that reducing the sample length leads to a slight decrease in the accuracy of both train and test data. However, this is compensated by a significant improvement in the accuracy of unseen data, highlighting the potential for better generalisation of the 1D-CNN model with smaller sample sizes.

Additionally, we conducted an extensive examination of the model's complexity, as illustrated in Table 3.6. The Baseline 1D-CNN and the Hybrid CNN-LSTM models, each comprising 14 layers, exhibit significant differences in parameter usage [9]. The Baseline 1D-CNN employs 0.4 million parameters, while the Hybrid CNN-LSTM utilises 4 million parameters and achieves 83% and 94% accuracy, respectively. However, the proposed model features a relatively complex architecture with 12 layers and 8 million parameters. Remarkably, despite its relatively lower complexity and the smaller dataset size, it attains an impressive accuracy rate of 89%. Although the proposed model improves over the Baseline 1D-CNN, it is outperformed by the hybrid CNN–LSTM in accuracy.

Furthermore, the proposed model requires significantly more computational requirements due to the increased number of convolutional layers and parameters. However, significant temporal downsampling reduces the computational cost in terms of FLOPs compared with the Hybrid CNN-LSTM and Baseline 1D-CNN approaches. With only 65,792 FLOPs, our proposed model achieves outstanding performance, surpassing the accuracy of the Baseline 1D-CNN and Hybrid CNN-LSTM. This efficiency translates into a more cost-effective and energy-efficient deployment, making it an attractive option for real-world applications. However, it is important to note that a direct correlation between a number of parameters and FLOPs is not definitively established. Increasing these factors does not consistently lead to better performance.

To summarise, the proposed model emerges as a suitable choice, distinguished by its exceptional computational efficiency, high accuracy, and resource-efficient design in contrast to the remaining models in the table. It successfully balances complexity and performance, offering a practical and cost-effective solution for real-world applications.

3.5 Conclusion

Given the absence of precise CAD biomarkers, identifying robust classification features becomes crucial. Our study in this chapter highlighted the importance of extracting discriminative CAD features from ECG signals, while also emphasising the challenge of ECG-based CAD detection, as definitive ECG features are not consistently present

TABLE 3.6: Comparative analysis of model complexity with existing work

Metric	Proposed Model	Baseline 1D-CNN [9]	Hybrid CNN-LSTM [9]
Number of Layers	12	14	14
No. of Parameters	8 M	0.4 M	4 M
Activation Function	ReLU	ReLU	ReLU
Pooling Layers	Max Pooling	Max Pooling	Max Pooling
Dropout Rate	0.2 and 0.5	0.2	0.2
Learning Rate	0.001	0.003	0.003
Floating-Point Operations (FLOPs)	65,792	71,936	23,609,344
Performance (Accuracy)	89%	74%	94%
Computational Resources	Apple M2 Max	Intel® Xeon(R) 16-core	Intel® Xeon(R) 16-core

and some patients with CAD may present with normal ECG recordings. Exploring alternative ECG channels for CAD detection is also deemed essential. Extracting CAD-specific data from diverse patient records in the MIMIC III database proved challenging due to the varied ECG storage methods across channels. This study highlights the importance of feature extraction techniques and their influence on model performance, and also indicates that other factors, such as the number of ECG leads, sample size and underlying medical conditions, may have a significant impact and therefore require further investigation. One of the most pivotal aspects to further address is reducing network complexity, as it is directly correlated with model accuracy. This aspect will be addressed from Chapter 5 onwards, where the development of lightweight models and the minimisation of network complexity are explored in depth.

Chapter 4

Enhanced Coronary Artery Disease Classification through Feature Engineering and One-Dimensional Convolutional Neural Network

4.1 Overview

Based on the findings in Chapter 3, which highlighted limitations in both diagnostic accuracy and model's complexity, this chapter propose a novel learning-based model for CAD diagnosis using only ECG signals. The proposed method works based on a 1D-CNN, offering a cost-effective alternative for sophisticated cardiac health monitoring. This study is motivated by the urgent need for more accessible and cost-effective methods for diagnosing CAD, which remains a leading cause of mortality worldwide. Existing CAD diagnostic tools rely on costly and complex biomarkers and scanners.

Furthermore, we introduce the concept of feature engineering to improve the quality of the model training process and mitigate the challenge of ill-conditioned ECG data. Unlike existing approaches, which often overlook signal quality, our model applies a smart feature engineering, ensuring that only diagnostically reliable signals are used. This design improves robustness, generalisability, and suitability for real-world clinical settings.

Utilising one of the most complex publicly available datasets, i.e., MIMIC III, sourced from PhysioNet, the performance of the proposed model, along with existing ones in classifying potential cases of CAD and NON-CAD is investigated. Our findings confirm that the proposed model exhibits outstanding performance, highlighting the effectiveness of our integrated FE approach with the CNN model. Moreover, our investigation extends to a thorough comparative analysis between the proposed model and existing models, aiming to discern their respective strengths, weaknesses, and overall performance on CAD applications.

4.2 Related works

A hybrid CNN-LSTM architecture is introduced for the detection of CAD, utilising anomalous ECG morphology and irregular HRV to discern CAD and NON-CAD cases [9]. Additionally, the CNN-LSTM architecture is applied for CAD identification [6]. However, the diagnostic accuracy of CAD is impeded by constraints in the available data, as public resources provide only limited datasets and lack certain definitive biomarkers. Several studies have implemented 1D-CNN for the automated detection of CAD, aiming to enhance diagnostic accuracy and improve patient outcomes. In [134], an automated CAD diagnosis system utilising 1D-CNN is presented, showcasing notable outcomes in accuracy and computational efficiency. The 1D-CNN approach effectively distinguished between CAD and NON-CAD subjects. The integration of feature extraction techniques into their proposed model yielded promising accuracy levels. Nevertheless, the training process of the model proved to be time-intensive and demanded a substantial volume of data.

A deep neural network, with its layered structure and various activation functions, learns to recognise patterns at different levels of complexity within the ECG signals. Specifically, the CNN stands out for its ability to extract and identify features within ECG signals, making it highly effective for tasks such as diagnosing cardiac conditions and predicting patient outcomes based on ECG data. The CNN architecture has been devised to diagnose various CVDs applications, including arrhythmia conditions [160, 203–205], AF [139], and CAD [6, 9, 134, 206]. The performance evaluation of the CNN model relies on its capacity to accurately process input signals and produce the intended prediction. This evaluation focuses on how effectively the CNN interprets and manages the provided ECG signals to classify potential classes. An 11-layer CNN model, combined with DWT, achieved an accuracy of 94.95% in 2-second segments and 95.11% in 5-second segments. The DWT with the Daubechies 6 (db6) mother wavelet was employed on

the ECG signals to mitigate noise and baseline wander [134]. The CNN model possesses the capability to extract features from ECG signals, but its efficacy depends on the quality of the ECG data. When data is corrupted by noise and artefacts, it compromises the model's learning capacity. Thus, preprocessing methodologies, including DWT [133, 134], various entropy computations [207], and the Fourier transform [53], are applied to enhance the quality of the signal before its incorporation into the CNN model.

Due to a significant shortage of available CAD data, a relatively small group of researchers have conducted their work on CAD diagnosis [6, 134, 208]. Even among these works, many studies have focused on detecting arrhythmia [160, 205], congestive heart failure [209], AF [139] and MI [42].

A recent study proposed a CNN–LSTM–SE architecture for classifying heart failure severity using lead II ECG signals from the MIMIC III database [48]. While the focus was on heart failure, it demonstrates the effectiveness of applying deep learning methods to ECG data from MIMIC III. Other recent studies have also explored classification tasks using MIMIC III, including a multimodal contrastive learning approach that combines ECG signals with clinical text for arrhythmia detection [210], and a traditional k-NN method using handcrafted ECG features for atrial fibrillation identification [211]. These studies highlight the potential of MIMIC III for ECG-based classification. Yet, despite its widespread prevalence, very limited attention has been given to CAD [31]. The main focus of many existing studies relies on these heart diseases, primarily due to the availability of datasets as previously mentioned. Furthermore, these diseases are easily distinguishable into potential classes due to their specific and distinctive biomarkers. On the contrary, CAD, which lacks certain biomarkers, presents more formidable challenges. Recent deep learning approaches to CAD detection have primarily focused on deeper CNN architectures, often overlooking critical factors such as signal quality issues and real-time applicability [212]. This study addresses these limitations by proposing a streamlined 1D-CNN architecture, designed for CAD detection. By targeting this underexplored area, the study contributes to closing a notable gap in the current literature.

In this study, we aim to design a specific deep learning-based model for accurate coronary artery disease classification. The summary of the contributions is as follows:

- A smart FE is proposed to remove unwanted and noisy ECG segments prior to model training, improving data quality and overall model reliability.
- A novel and streamlined 1D-CNN architecture is developed, integrated with the FE module for CAD diagnosis. This integration enhances noise tolerance, improves

generalisability to unseen data, and supports efficient deployment in real-time settings.

- The MIMIC III dataset is prioritised for CAD diagnosis, with extensive preprocessing conducted to address the complexity and noise inherent in its ECG signals.
- To verify the effectiveness of our proposed model, we conduct a comparative analysis of classical machine learning models alongside existing CNN-based models on well-established datasets.
- A one-class evaluation is conducted to assess the model's ability to distinguish CAD from NON-CAD cases, reflecting realistic screening scenarios in clinical practice.
- A cost-effective and time-efficient diagnostic pipeline is designed to support scalable CAD screening, particularly in clinical and resource-constrained settings.

Moreover, while much of the existing literature focuses on arrhythmia classification using ECG signals, the proposed model is specifically developed for CAD detection, which is a clinically significant yet relatively underexplored application. Our primary aim is to develop a cost-effective and time-efficient diagnostic tool suitable for use in real-world healthcare environments. To facilitate clinical use, the proposed model is designed for fully automated operation, enabling rapid and accurate CAD diagnosis to support improved the standard of patient care. Additionally, the model's robustness is evaluated under varying conditions by comparing its performance on data with and without FE techniques.

4.3 Materials and Methods

The proposed method comprises three key phases: data preparation, feature engineering, and developing a deep learning model. We aim to design and implement a model to classify CAD using ECG signals, based on a comprehensive series of experiments. Subsequent sections will provide detailed explanations of each step.

4.3.1 Data preparation

The primary source for training and testing ECG data is derived from the MIMIC III and Fantasia databases, accessible on the PhysioNet website [199, 200]. Around 2,840 patients, comprising roughly 7.1% of the total hospital admissions, have been identified within the MIMIC database as individuals diagnosed with coronary atherosclerosis in the native coronary artery. On the other hand, the Fantasia database is characterised

by a cohort consisting of 40 individuals, with an equal distribution of 20 young and 20 adult patients. The acquisition of ECG signals is facilitated by employing ECG sensors. The ECG electrodes are strategically positioned on the patient's body, enabling the sensors to accurately capture the activity of the heart. The ECG signal consists of the P wave, representing atrial depolarisation; the QRS complex, indicating ventricular depolarisation and contraction; and the T wave, demonstrating ventricular repolarisation and relaxation. ST depression indicates significant coronary artery lesions, highlighting a critical need for an early invasive treatment. Moreover, ST elevation displays a complete blockage of the artery, indicating CAD, MI, or a heart attack. To create a balanced dataset for each class, forty patients were meticulously selected from each respective database and partitioned into 1-second segments of ECG data, resulting in approximately 500,000 seconds. Maintaining class balance is crucial for minimising model bias and ensuring reliable detection of minority class instances. In the presence of class imbalance, the model may exhibit reduced sensitivity to CAD cases, potentially compromising diagnostic performance. The segmentation process facilitates a thorough analysis of the dynamic changes and patterns within ECG signals over short time frames. This approach has been widely employed in most of the previous works to ensure consistency in the ECG signal [6, 148, 213]. Hence, in our study, the ECG data was segmented into small segments to maintain consistency. The aim is to guarantee that each segment contains a complete ECG cycle, without relying on QRS detection [148]. These segments will be employed to apply FE techniques, which include Sample Entropy and Standard Normalisation, to systematically eliminate irrelevant and noisy ECG data that may potentially affect the predictive accuracy of the model (see next section).

4.3.2 Feature Engineering

ECG signals are fundamental clinical tools for diagnosing cardiac diseases. However, the integrity of these signals is frequently compromised by prevailing challenges, including baseline drift, muscular interferences, powerline artifacts, and electrode motion disruptions. These challenges have a significant negative impact on the quality of ECG signals, rendering specific segments of the data inadequate for precise diagnosis [134, 214]. As a result, poor ECG signal quality may hinder reliable assessment of an individual's cardiac health status and delay the identification of underlying cardiac conditions. FE plays a significant role as it assists in removing missing or inconsistent ECG data resulting from human and equipment errors. It aids in transforming raw ECG data into informative features, which enable the model to better capture underlying patterns and improve its predictive accuracy. Sample entropy serves as a metric utilised to assess the quality of time series data. Its efficacy in mitigating noisy ECG channels has been demonstrated

in previous research [59]. Furthermore, standard normalisation is employed to eliminate flat time series data and to mitigate any potential data that could impact accuracy. The computation of sample entropy is presented as follows in Equation 4.1.

$$SampEn = -\ln\left(\frac{\sum_{i=1}^{N-m} Q_i^m(r)}{\sum_{i=1}^{N-m+1} P_i^{m+1}(r)}\right) \tag{4.1}$$

where $SampEn$ indicates the quality of the ECG signals. N represents the number of samples within each 1-second ECG segment. m refers to the embedded dimension, which represents the length of consecutive samples or data points analysed together. $Q_i^m(r)$ quantifies the instances of vector pairs of dimension m whose mutual distance falls below a predefined threshold r, suggesting a degree of similarity or regularity within the signal. $P_i^{m+1}(r)$ quantifies the number of vector pairs of dimension m+1 that are similar within a predefined threshold level r, thereby extending the comparison to sequences of increased length. A value of r =0.1 was employed in this chapter, consistent with prior literature suggesting that values within the range of 0.1 to 0.25 are effective for preserving the quality of ECG signals [215].

Standardisation is then applied to minimise flat regions and noisy artifacts in the ECG signals, thereby improving the representation of key morphological features. The computation of standard normalisation is detailed in Equation 4.2.

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(x_i - \overline{x})^2} \tag{4.2}$$

where σ also indicates the quality of the ECG signals, N is the number of signals, $\overline{x}$ is the average of a given signal, and x_i is the signal value at the i^{th} position.

To summarise, FE using $SampEn$ and standard normalisation was employed to remove ECG segments that exhibited flat lines or an absence of clear clinical features. These artefacts were caused by baseline wander, motion noise or poor electrode contact during signal acquisition. Although DL-based models perform automatic feature extraction, FE was applied to improve signal quality by removing low-information ECG segments. This step ensures that the input retains clinically meaningful waveform components, thereby enhancing model robustness and generalisability.

Importantly, FE is applied during both training and prediction to ensure that only clinically relevant and high-quality ECG segments are used. This consistency improves the model's robustness and generalisability in real-world clinical applications. In clinical

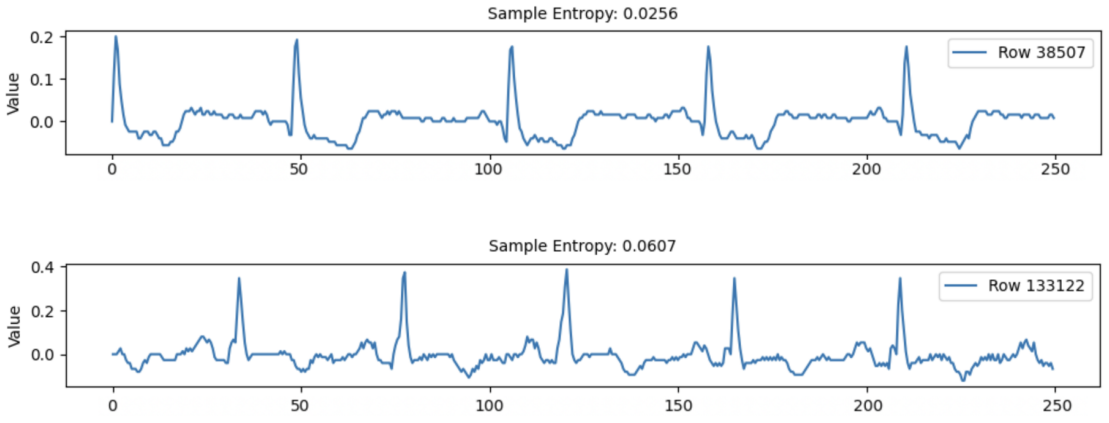

(a) Exemplary ECG Signal Demonstrating High Quality

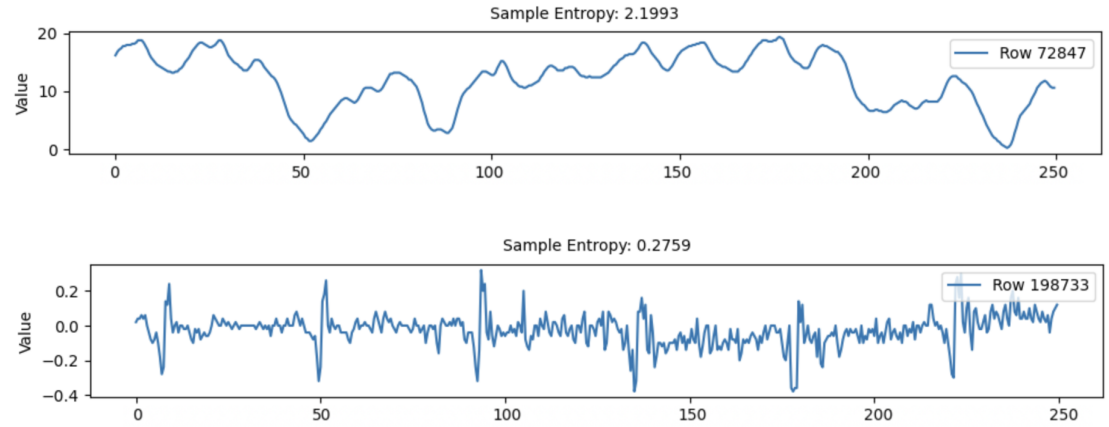

(b) Exemplary ECG Signal Demonstrating Poor Quality

FIGURE 4.1: Signal Quality Evaluation: ECG signals with sample entropy below the threshold ($r = 0.1$) were retained for further processing.

use, it also enables the system to focus on meaningful ECG input and reduces the risk of misclassification caused by noise or incomplete signals.

From a physiological perspective, ECG changes associated with CAD are often subtle and may appear as small variations in waveform shape or temporal behaviour. Retaining ECG segments with sufficient signal complexity therefore helps to preserve diagnostically relevant information that might otherwise be obscured by noise or low-information signals. In this study, *SampEn* is used as a signal-quality metric to assess temporal complexity, while σ is applied to identify and remove flat or low-variance ECG segments that do not contain meaningful cardiac information.

Figure 4.1 presents a comparison of exemplary ECG signals, showcasing examples of both high-quality and poor-quality signals. FE is applied to the 1-second ECG segments generated in the previous step. The threshold r is set at 0.1, which is used to determine the quality of the signal. Signals with a value less than r are classified as high-quality

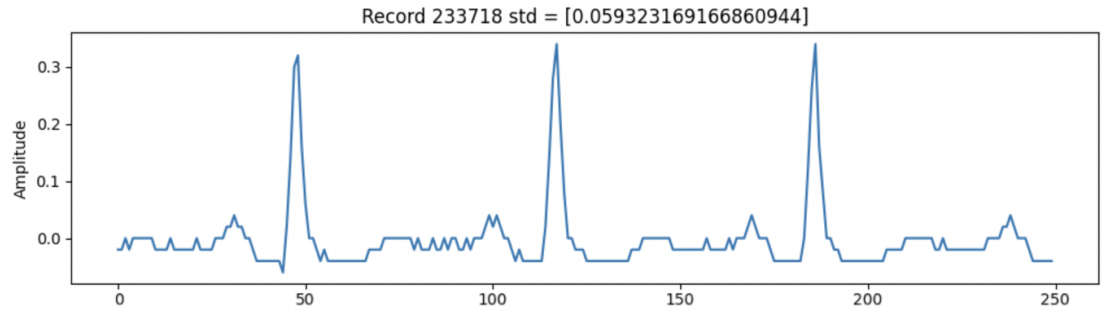

(a) Poor-Quality ECG Signal with flat line

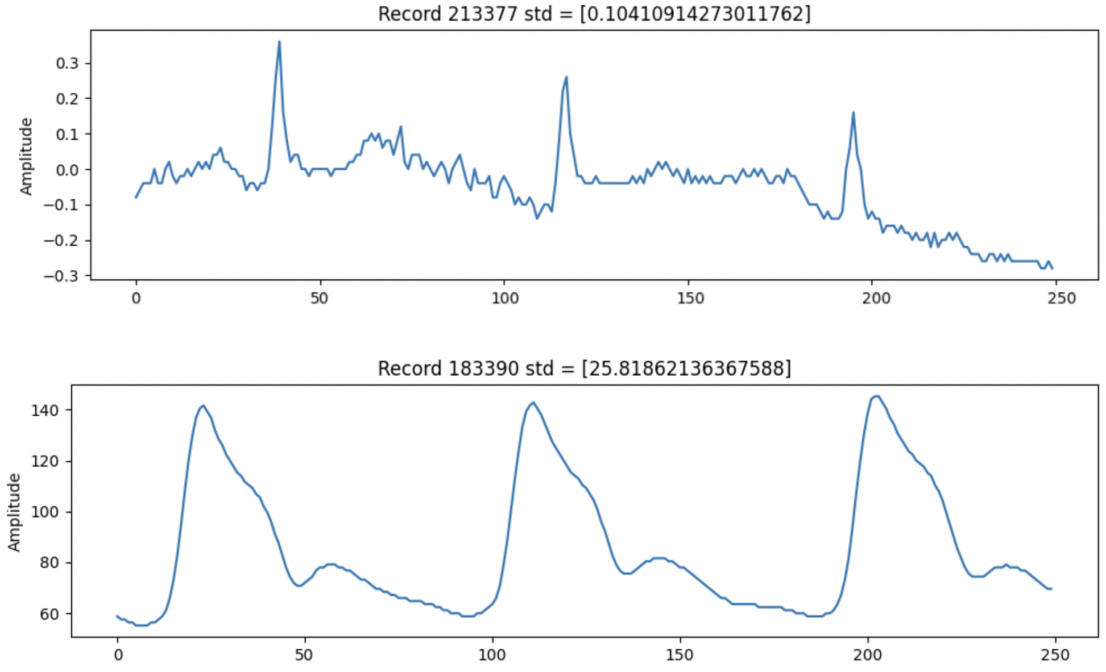

(b) Good-Quality ECG Signal without flat line

FIGURE 4.2: Signal Quality Evaluation: ECG signals with standard normalisation above threshold ($r = 0.1$) were retained for further processing.

signals. The segments in which ECG signals fail to meet the specified threshold r are excluded. In Figure 4.1(a), an exemplary ECG signal demonstrating high quality is depicted. This signal displays distinct ECG features, such as well-defined cycles containing QRS complexes and other waveforms. Conversely, Figure 4.1(b) illustrates an exemplary ECG signal demonstrating poor quality. This signal is characterised by the absence of clear ECG features, and it contains noise patterns indicative of potential recording or signal acquisition issues. Therefore, it is crucial to ensure the reliability of ECG signals in subsequent processes, as it helps mitigate the risk of generating inaccurate predictive outcomes.

Figure 4.2 illustrates a comparison of the quality of ECG signals after applying standard

normalisation to remove remaining noise and flatness. A signal is retained as high quality if its standard deviation exceeds the threshold r. Conversely, signals with the standard deviation below the threshold r are excluded. This FE technique is employed to eliminate any flatness present in the signal, thereby reducing the risk of low predictive accuracy. It serves as an optional approach to ensure signal cleanliness, aimed at preserving signal integrity and thereby enhancing the model's capacity for pattern recognition. After applying the FE techniques, the retained ECG signal is then inputted into the proposed model for the classification of CAD and NON-CAD cases.

4.3.3 CNN model

The CNN architecture employed in this chapter follows the same design as the model proposed in Chapter 3 and is evaluated here under different experimental conditions. A CNN model consisting of four convolutional layers, a max-pooling layer, four dropout layers, a flattened layer, and a fully connected dense layer was designed. ECG signals are taken as input to the first layer of the convolutional layer. The convolutional layer was used to learn and extract patterns of two possible classes (CAD and NON-CAD). The output of the convolutional layer is calculated as:

$$y(n) = \begin{cases} \sum_{i=0}^{k} x(n+i)h(i), & \text{if } n = 0 \\ \sum_{i=0}^{k} x(n+i+(s-1))h(i), & \text{otherwise} \end{cases} \tag{4.3}$$

where $y(n)$ represents the output signal at position n within the convolutional layer, $x(n)$ denotes the ECG input signal, and $h(i)$ signifies the convolutional kernel—a collection of learnable weights that the CNN acquires during training for feature extraction from the input signal $x(n)$. The parameter k corresponds to the size of the convolutional kernel, determining the receptive field, which defines the spatial area over which the filter operates on the input signal. s represents the stride length, indicating how the convolutional kernel moves through the ECG signal. s is set to 1 by default, indicating that the kernel progresses through the input signal with each step equivalent to the size of one ECG sample.

Figure 4.3 illustrates that the initial convolutional layer consists of 512 filters with a kernel size of 32, while the subsequent layers consist of 256 filters, each with the same kernel size. The ReLU activation function was used in the convolutional layers to introduce non-linearity into the model, thereby aiding in the comprehension of intricate patterns within the ECG signal. Three dropout layers with a rate of 0.2 were added after the convolutional layers to prevent overfitting. The max pooling layer with a pool

size of 128 was then applied to reduce the spatial size of the feature maps and improve generalisation by selecting the maximum value within a window size. The flattened output was then passed to a fully connected layer, consisting of 128 neurons and employing the ReLU activation function. To further prevent overfitting, an additional dropout layer with a rate of 0.5 was incorporated before the final output layer. The final layer comprises two neurons and utilises softmax activation, enabling the model to classify input data into one of two potential classes, thereby representing the probabilities of the input belonging to each class.

A binary cross-entropy loss function is used for CAD and NON-CAD classification due to its ability to handle binary classification tasks as shown in (4.4).

$$\text{BinaryCrossEn} = -\frac{1}{N}\sum_{i=1}^{N}\left(l_i \log(p_i) + (1 - l_i)\log(1 - p_i)\right) \tag{4.4}$$

where N is the number of samples in the corresponding ECG segment. l_i is the true label for ECG signals i, where $l_i \in \{0, 1\}$. $l_i = 0$ represents a NON-CAD case, indicating the absence of CAD features in the ECG signal i. Conversely, $l_i = 1$ exhibits a CAD case, representing the existence of relevant CAD features in the ECG signal i. p_i is a predicted probability that ECG signal i belong to class 1. It facilitates the model in determining its predictive certainty, thereby refining its prediction to achieve better alignment with the true labels.

4.4 Experimental results

In this section, the performance of the proposed model is evaluated, including classification accuracy and comparative analysis with other models.

4.4.1 Experiments

As outlined in Section 4.3.1, forty patients were deliberately sampled from both the MIMIC and Fantasia databases in identical proportions. Subsequently, the data was partitioned into 1-second segments (each with $N = 250$ samples) to enhance accuracy and focus on temporal aspects. After segmenting the ECG data into 1-second segments and implementing FE techniques, two distinct subsets of data are generated for the experiments: D_1 and D_2. The first subset (D_1) comprises 100 CAD samples and 100 NON-CAD samples, with each sample having a length of 1 second. To assess the model's performance robustly, we employed k-fold cross-validation on D_1 by partitioning it into

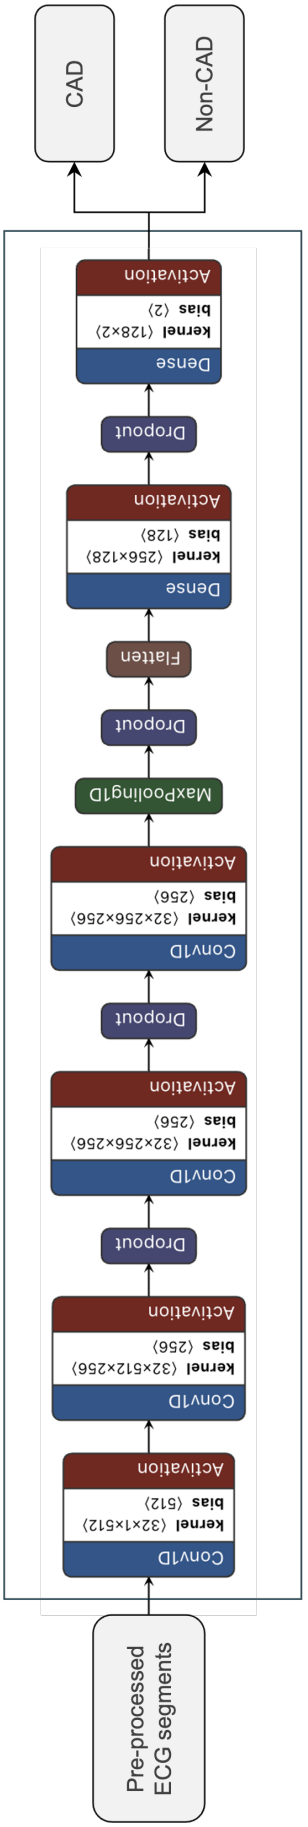

FIGURE 4.3: The proposed model architecture.

10 folds. In each iteration, 70% of the folds were used for training, and the remaining 30% for validation. The partitioning method is extensively utilised due to its dual capacity to provide substantial training data and allocate sufficient resources for assessing the model's ability to generalise. Moreover, subset D_2, containing an equivalent number of samples as D_1, was employed for predictive purposes to evaluate the model's performance against this unique and previously unencountered data subset. This additional dataset serves as an unseen dataset used to measure the model's performance. It has not been encountered during training or validation, making it crucial for evaluating the model's ability to generalise to unseen data.

The hyperparameters used during model training, including learning rate, batch size, and number of epochs, are summarised in Table 3.2. To evaluate the model, performance metrics including Accuracy, Error Rate, Precision (Ppr), Sensitivity (Sen), Specificity (Spr), and F1 Score ($F1$) were used to measure the performance of the classification model as shown in Equations (2.5) , (2.10) , (2.7), (2.6), (2.8), and (2.9).

4.4.2 Results and discussion

To assess the impact of FE, we conducted ablation experiments. Models trained without FE exhibited lower accuracy and generalisation, particularly on unseen data (D_2), confirming the role of FE in reducing irrelevant and noisy components. Table 4.1 demonstrates the investigation of the model's performance with and without FE, aimed at studying its impact on the proposed model. The effectiveness of the proposed model is notably high, achieving an accuracy of 99.3% on the training set and 98.5% on the testing (D_1). Moreover, it showcased robust performance on unseen data from D_2, achieving a noteworthy accuracy of 99.0%. In contrast, the proposed model yielded lower accuracy when FE was not applied to both datasets, with D_2 exhibiting an accuracy of 87.0%. To summarise this investigation, it proves that applying FE significantly improves the model's ability to accurately classify instances of CAD and NON-CAD cases. Particularly, FE facilitates the model in eliminating irrelevant data, thereby mitigating model potential issues, leading to higher accuracy rates.

Table 4.2 presents a comparative analysis of classical machine learning and deep learning algorithms and our proposed model on D_1 and D_2, with and without FE. In the absence of FE implementation, the proposed model performed admirably on D_1, achieving 85.1% accuracy on the training set and 87.9% on the testing set. It also achieved the highest accuracy on previously unseen data from D_2, reaching 87.0%. Although kNN reported a higher training accuracy of 96.3% on D_1, it demonstrated poor generalisability, with its accuracy dropping to 49.0% on D_2. SVM and K-means also showed reduced accuracy

TABLE 4.1: Accuracy Comparison with and without FE on Datasets D_1 and D_2.

Model	D_1		D_2
	Train (%)	Test (%)	Unseen (%)
Trained on D_1, Tested on D_2 without FE	85.1	87.9	87.0
Trained on D_1, Tested on D_2 with FE	99.3	98.5	99.0
Trained on D_1 with FE, Tested on D_2 without FE	99.3	98.5	69.0

** FE is a Feature Engineering*

without FE; SVM achieved 65.2% on the D_1 test set and 63.0% on D_2, while K-means attained 59.1% and 60.0%, respectively. The application of FE notably improved the performance of all classical models, underscoring their reliance on engineered features for effective classification.

In comparison to the classical machine learning methods, deep learning models exhibited stronger generalisation across both datasets. Without FE, the LSTM [113] model achieved a test accuracy of 78.3% on D_1 and 77.5% on D_2, while the CNN-LSTM [113] model yielded a comparable accuracy of 79.0% on D_2. With FE, these results improved further, with the LSTM reaching 90.0% and CNN-LSTM achieving 89.0% on D_2, demonstrating the ability of temporal models to learn sequential dependencies directly from raw ECG signals.

Among all classical machine learning and deep learning models, the proposed model demonstrates exceptional performance in binary classification on both datasets when FE is applied, achieving accuracy rates of 99.3% on the training set and 98.5% on the testing set of D_1. Additionally, the trained model demonstrated strong performance by achieving a classification accuracy of 98.5% on unseen data from D_2. SVM and kNN demonstrated strong performance on D_1, obtaining accuracy rates of 95.5% and 96.9%, respectively. However, kNN struggled to generalise to unseen data from D_2, while SVM maintained its effectiveness, achieving an accuracy of 96.4%. The performance of K-means is noteworthy, as its performance appears relatively subpar when evaluated on the training data D_1. However, it displays a notable improvement in classifying unseen data D_2, surpassing the performance of the kNN. The investigation reveals that the proposed model outperformed in the application of CAD, particularly in cases where no certain biomarker represented its signal. This underscores a significant aspect of classical machine learning models, as they heavily rely on the extraction of precise features for the model to effectively capture and learn patterns. Moreover, machine learning is adept at handling small datasets effectively. Nevertheless, as our dataset scales up considerably, deep learning emerges as a more appropriate approach owing to its capability to manage intricate and expansive data structures proficiently. As evidenced by the outcomes of

TABLE 4.2: Comparison of existing classical machine learning and deep learning algorithms performance on D_1 and D_2 with and without FE applied.

Model	Without FE			With FE		
	D_1		D_2	D_1		D_2
	Train (%)	Test (%)	Unseen (%)	Train (%)	Test (%)	Unseen (%)
SVM	64.2	65.2	63.0	95.5	96.1	96.4
K-means	60.4	59.1	60.0	86.6	93.9	89.0
kNN	96.3	82.5	43.5	96.9	97.5	49.0
LSTM [113]	76.2	78.3	77.5	92.6	86.7	90.0
CNN-LSTM [113]	79.3	78.3	79.0	92.9	86.7	89.0
Proposed Model	85.1	87.9	87.0	99.3	98.5	99.0

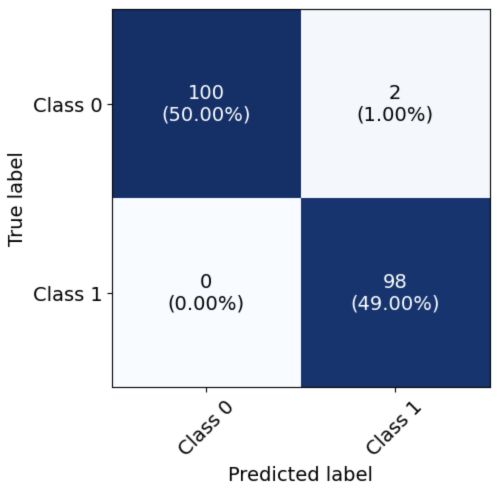

(a) on Training Data D_1

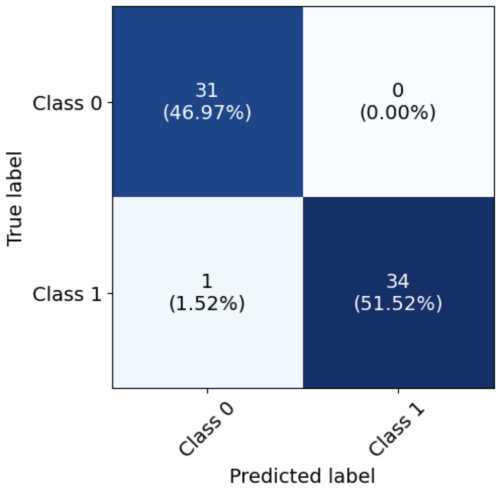

(b) on Test Data D_1

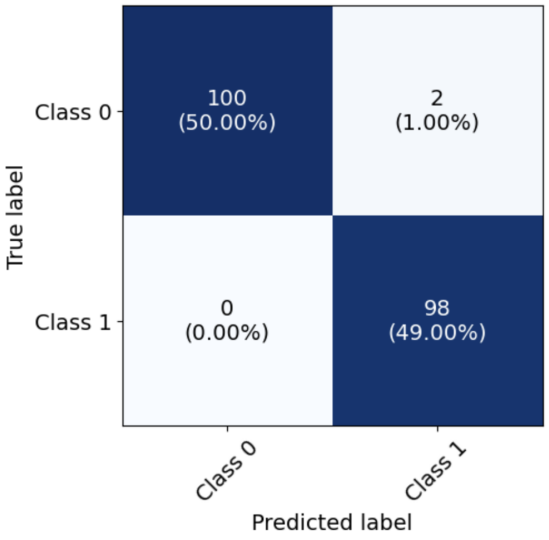

(c) on Test Data D_2

FIGURE 4.4: Confusion Matrix Analysis on D_1 and D_2

this investigation, it validates and clarifies our reasoning for selecting deep learning over classical machine learning algorithms in our application.

Figure 4.4(a) illustrates the performance metrics for CAD detection in D_1 and provides valuable insights into the model's effectiveness. With an accuracy of 99.3%, the model demonstrates its capability to correctly classify all CAD and NON-CAD instances, indicating solid overall performance with an error rate of 0.7%. The *Ppr* is 98.5%, indicating that nearly all positive predictions were correct. Similarly, *Sen* was observed to be 98.5%, suggesting that the model accurately identified the majority of actual CAD cases. Additionally, a *Spr* of 100% highlights the model's proficiency in accurately identifying negative cases, implying a satisfactory ability to distinguish NON-CAD instances.

Moreover, the *F1 score*, a metric that balances precision and recall, was calculated at 99.01%, indicating a strong overall performance of the model in accurately classifying both CAD and NON-CAD cases on the training data.

Figure 4.4(b) shows an overall accuracy of about 98.5% on testing data. The model correctly predicted 34 instances of Class 1. This indicates a strong ability of the model to identify CAD cases. Similarly, the model correctly predicted 31 instances of class 0, which demonstrates the model's effectiveness in identifying NON-CAD cases. There was 1 instance where the model incorrectly predicted class 1 as class 0. This error indicates a slight issue with the model's sensitivity to identifying CAD cases.

Figure 4.4(c) illustrates that the model has a high degree of proficiency, achieving an exceptional accuracy rate of 99.0%. The model accurately predicted CAD for 98 instances and NON-CAD for 100 instances. Additionally, there were 2 instances where the model incorrectly predicted NON-CAD as CAD cases, but there were no instances where CAD was incorrectly predicted as NON-CAD. Notably, the *Ppr* for class 1 demonstrates a commendable figure of 98%, indicating the model's accuracy in identifying CAD cases. Additionally, the *Sen* for class 1 is outstanding, signifying that the model successfully captures all instances of class 1 present in the dataset D_2. Furthermore, the *F1 score*, which combines precision and recall into a single metric, is approximately 98.99% for class 1, suggesting a well-balanced performance in this binary classification.

Figure 4.5 demonstrates the confusion matrices used to evaluate the accuracy of our proposed model in a one-class classification case. The model was trained on a balanced dataset comprising both CAD and NON-CAD samples. For evaluation purposes, we tested the classifier on two distinct datasets: one containing only CAD samples, with no NON-CAD samples, and the other containing only NON-CAD samples. This approach allows for a clear assessment of the model's capability to accurately identify each class in isolation. Figure 4.5(a) demonstrates that the classifier accurately is 99%, with only one sample misclassified as NON-CAD. Similarly, Figure 4.5(b) illustrates that the classifier correctly identified 98% of NON-CAD samples correctly. With a true positive rate of 99% for CAD samples and a true negative rate of 98% for NON-CAD samples, the model exhibits robust performance metrics. The notably low false negative rate in CAD detection is especially crucial, as it ensures that nearly all patients with CAD are accurately identified. Similarly, the low false positive rate in NON-CAD detection helps prevent misdiagnoses.

Figure 4.6 exhibits the performance comparison of classical machine learning algorithms and our proposed model, evaluated on CAD classification. The k-means model, depicted in blue, has an Area Under the Curve (AUC) of 0.10, showing weak performance due to its proximity to the random guess line, suggesting that the model's performance is not

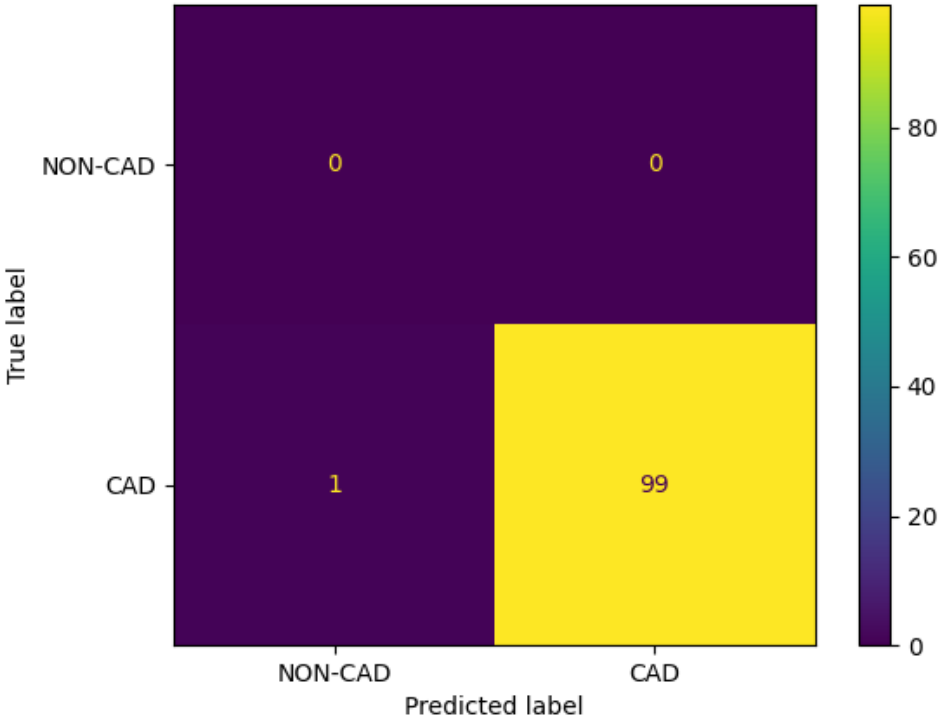

(a) on 100 CAD ECG samples

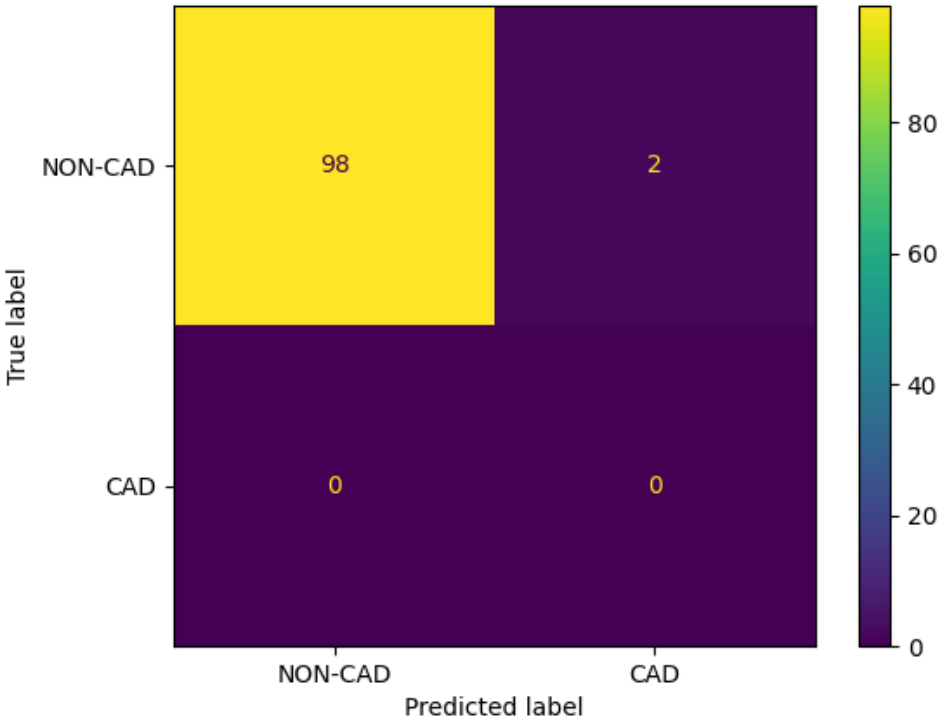

(b) on 100 NON-CAD ECG samples

FIGURE 4.5: Confusion Matrix: One-Class Testing Performance

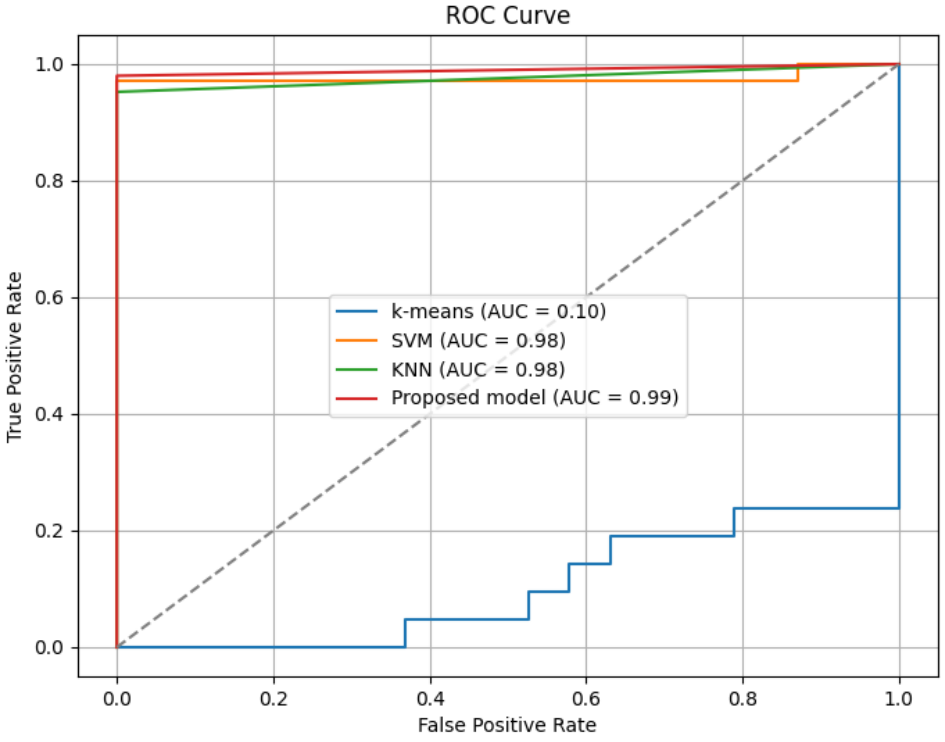

FIGURE 4.6: Analysis of Classifier Efficacy Through ROC Curve Metrics

much better than random chance. The SVM, depicted in orange, and the kNN, shown in green, both exhibit a high level of discriminative ability with an AUC of 0.98. The Proposed model is represented by the red line and has the highest AUC of 0.99, which suggests that it has the best performance among the models presented. The SVM and kNN models exhibit strong performance, as evidenced by their high AUC values, indicating their proficiency in effectively discriminating between CAD and NON-CAD cases. On the contrary, the proposed model exhibits slight enhancements when compared to SVM and kNN. This improvement stems from the CNN's ability to automatically extract relevant features from the dataset, identifying complex patterns and relationships that SVM and kNN may overlook. Consequently, the proposed model achieves a slightly higher level of discrimination, resulting in its slightly improved performance. This outcome is anticipated, as k-means clustering operates by grouping data points according to their similarity, rather than assigning them to predetermined categories, which is essential in classification tasks. As a result, among the other models, k-means exhibits the lowest performance in classifying CAD and NON-CAD cases.

Figure 4.7 illustrates the comparative performance of existing CNN models trained on the same dataset under identical experimental conditions as the proposed model. Following the findings presented by [216], which demonstrated a 93.33% accuracy rate for

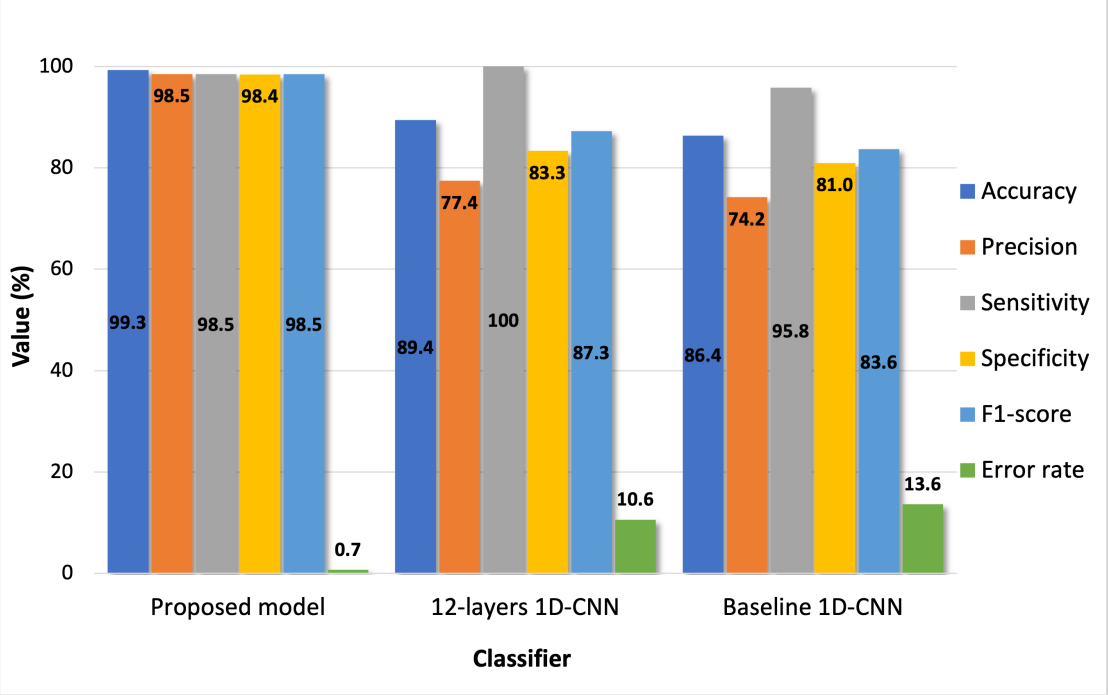

FIGURE 4.7: Leveraging Existing CNN Models with FE Applied

AF diagnosis using a 12-layer 1D-CNN architecture, our study aimed to expand upon this research. Employing the 12-layer CNN model on our datasets resulted in an accuracy of 89.4%. Furthermore, a baseline 1D-CNN model as described in [9] is employed for diagnosing cases of CAD, achieving an accuracy of 87%. In our study, our datasets, we achieved an accuracy of 86.4%. The proposed model exhibits strong predictive capabilities, achieving the highest accuracy of 99.3% compared to the other models. It also boasts a precision of 98.5%, sensitivity of 98.5%, specificity of 98.4%, and an F1 score of 98.5%. Impressively, it maintains a notably low error rate of 0.7%. Comparatively, the 12-layer 1D-CNN model also demonstrates strong performance, although with a significantly higher error rate of 10.6% compared to the proposed model. Furthermore, the 12-layer 1D-CNN appears to perform better compared to the baseline model. The baseline 1D-CNN model, on the other hand, falls short with a higher error rate of 13.6%, suggesting its inferior performance in comparison to the other two models. Therefore, the observed results indicate that the FE module contributes positively to the performance of both the existing CNN models and the proposed model by improving their ability to discriminate between CAD and NON-CAD cases, leading to enhanced overall diagnostic performance.

Alongside classification accuracy, computational efficiency is an important factor in ensuring a model is suitable for real-world use, particularly in clinical environments where time and resources may be limited. As shown in Table 4.3, we compared the average inference time per subject across a range of traditional and deep learning models. All

TABLE 4.3: Model inference time comparison for CAD classification

Model	Avg. Inference Time per Subject (ms)
SVM	0.051
K-Means	0.132
KNN	1.142
LSTM [113]	28.746
CNN-LSTM [113]	22.194
Proposed model	18.731

experiments were conducted on a machine equipped with an Apple M2 Max processor and 32 GB of unified memory. While traditional methods offer very fast inference times, they typically underperform in classification accuracy compared to deep learning approaches, as highlighted in Table 4.2. The proposed model completes inference in 18.731 ms, which is faster than more complex models such as LSTM at 28.746 ms and CNN-LSTM at 22.194 ms. These results suggest that the model is not only accurate but also efficient enough for real-time use, including applications in bedside monitoring or portable ECG devices where quick and reliable decisions are essential.

Several recent studies have applied machine learning and deep learning techniques to the detection of CAD [31, 133, 134]. For instance, hybrid CNN-LSTM models have been used to classify ECG signals, while traditional approaches based on handcrafted features have also shown promise in early diagnosis [9]. These benchmarking studies reflect the growing interest in automated CAD detection. In contrast to existing methods, the proposed model combines FE with a compact CNN architecture, achieving strong classification performance alongside improved computational efficiency. This design enables practical deployment in both clinical and portable settings. By filtering low-quality ECG segments prior to classification, the method enhances signal relevance while preserving key diagnostic features. Furthermore, the reduced model complexity supports real-time application, representing a meaningful advancement over previous CNN-based CAD approaches.

4.5 Discussion

Table 4.4 presents a comparison of recent ECG-based classification models evaluated on the MIMIC III dataset. While previous studies have addressed various cardiovascular conditions such as arrhythmia [210], heart failure [48], and atrial fibrillation [211], there is a notable absence of research focusing specifically on the detection of CAD using MIMIC III ECG signals. This gap highlights the limited exploration of CAD in large publicly available clinical ECG datasets, despite its high prevalence and clinical importance. To

TABLE 4.4: Comparison of ECG-based classification methods on the MIMIC-III dataset

Paper (Year)	**Disease/Task**	**Methods**	**Data Split**	**Accuracy (%)**
Vanitha et al. (2025) [210]	Arrhythmia	Multi-modal Contrastive Learning	Not specified	97.8
Zhang et al. (2024) [48]	Heart Failure	CNN-LSTM-SE	70% train, 10% val, 20% test	99.1
Bashar et al. (2020) [211]	Atrial Fibrillation	k-NN	10-fold + 2 independent test sets	99.3
Proposed Model	CAD	1D-CNN	10-fold + 2 independent test sets	99.0

address this, the proposed 1D-CNN model is developed and evaluated for CAD detection, achieving a competitive accuracy of 99.0%. The inclusion of FE contributes to improved model generalisability by enhancing discrimination between CAD and NON-CAD cases, thereby demonstrating the potential for automated CAD classification using MIMIC-III ECG data in this study.

Future studies should consider evaluating model performance across diverse patient subgroups, including those defined by age, sex, and ethnicity, to enhance both representativeness and clinical applicability of the proposed approach. Thus, Chapter 5 will consider the effect of age and sex on model performance to investigate the model's representativeness and clinical applicability. Although this study focuses on binary classification of CAD, examining the model's performance across different CAD subtypes could provide deeper insight into its ability to distinguish between them. Extending the model to support multi-class classification of these CAD subtypes would be a valuable step towards enhancing its clinical usefulness. Therefore, multiple CAD subtypes will be included in Chapter 6 for further investigation. While this study focuses on CAD, the approach could also be adapted to detect other forms of heart disease, such as arrhythmias or heart failure. This potential will be explored further in Chapter 7. Further studies could investigate alternative segmentation techniques that offer significant insights into ECG CAD signals. Furthermore, focusing on reducing complexity while maintaining accuracy could enhance real-time health monitoring capabilities for practical settings, and this will be addressed from Chapter 5 onwards. Following the reduction of network complexity, the model will be optimised and deployed the model on a practical sensing-and-processing device such as STM32F469I-DISCO. With its compact structure and ability to handle noisy signals, the proposed model is well-suited for use as a pre-screening tool in clinical settings, where it can assist with the review of ECG signals before cardiologist assessment. The deployability of models on STM32-based hardware platforms will be explored in Chapter 6.

4.6 Conclusion

Our study developed a CNN model to classify potential cases of CAD using ECG signals. By utilising data sourced from PhysioNet, it was revealed that the CNN model could independently classify binary classes. However, we observed a significant improvement in its performance when it was preceded by FE and pre-processing of the ECG data. The performance of the proposed model exceeded that of other CNN models investigated in our study, highlighting the importance of FEin increasing the model's ability to learn and make accurate predictions. Furthermore, our examination of three distinct classical

machine learning algorithms revealed that the CNN model surpassed these methods in predicting CAD. These findings indicate the significance of FE in enhancing the CNN model's performance, highlighting the CNN model's superiority over conventional methods in CAD diagnosis.

Chapter 5

CADNet: A lightweight Neural Network for Coronary Artery Disease Classification Using Electrocardiogram Signals

5.1 Overview

In Chapters 3 and 4, novel 1D-CNN models were introduced with a focus on reducing model complexity and incorporating FE technique to improve generalisability. However, further reductions in model complexity are still required, without compromising diagnostic accuracy. To address these limitations, this chapter introduces a lightweight one-dimensional convolutional neural network for CAD classification using non-invasive ECG signals. The proposed model, CADNet, consists of two key components: Feature Encoding and Compact Pooling. The feature encoding block extracts key temporal characteristics from ECG data using a convolutional layer, while the compact pooling block reduces temporal resolution, preserving essential ECG features for CAD diagnosis. CADNet comes with a novel feature engineering (FE) process to optimise computational efficiency and maintain high diagnostic accuracy. This approach aids convergence, significantly reduces the model parameters, and improves the model's ability to detect CAD patterns. Our extensive experiments with four diverse datasets show that CADNet achieves an average 99.3% accuracy, with 2,586 trainable parameters, surpassing state-of-the-art models' performance. This work highlights the potential of CADNet for real-time CAD detection in low-resource and clinical settings.

5.2 Introduction

Lightweight CNN models have gained attention for ECG classification due to their suitability for resource-constrained devices, offering fast inference and low energy consumption, which are important for real-time diagnosis. While many studies focus on arrhythmia detection using efficient architectures [173–176], Cao et al. [177] introduced a lightweight CNN for myocardial infarction detection with reduced complexity and high accuracy. Models such as SqueezeNet, EfficientNet, MobileNet, and ShuffleNet have also demonstrated strong potential for portable ECG applications [123, 178–181].

While these studies contribute valuable insights into the detection of arrhythmias, they do not directly address the challenges associated with coronary artery disease diagnosis. Unlike arrhythmias, which often present distinct electrical anomalies, CAD is characterised by subtle ischemic patterns and morphological changes in ECG signals. The detection of CAD requires domain-specific adaptations that extend beyond the capabilities of conventional arrhythmia-focused CNN models. In this study, we propose a lightweight one-dimensional convolutional neural network (CADNet) for automated CAD diagnosis. Our research distinguishes itself by targeting CAD classification, which presents unique challenges in early detection due to the absence of overt electrical abnormalities. We propose an advanced lightweight CNN framework that integrates specialised feature extraction techniques to identify ischemic patterns in ECG signals. By tailoring the model architecture to CAD-specific markers, we aim to improve diagnostic accuracy while ensuring computational efficiency for real-time deployment.

Compared with other existing lightweight models, deep learning-based models and traditional machine learning algorithms, our model demonstrates advantages in terms of reduced model size while maintaining high accuracy. This important capability facilitates the deployment of CADNet in clinical settings. This is particularly significant because traditional deep learning models require substantial computational resources, making them unsuitable and costly for real-time diagnosis and deployment on portable devices.

CADNet introduces several key innovations specifically tailored for CAD detection, addressing challenges that conventional arrhythmia-focused CNN models often overlook.

1. Innovative Feature Engineering (FE) Module: We introduce a novel data refinement process that combines sample entropy and standard normalisation, improving signal quality and computational efficiency before classification. This significantly enhances the reliability of CAD diagnosis.

2. Specialised for CAD Detection: Unlike arrhythmia classification, which often relies on clear electrical anomalies, CAD detection requires identifying subtle ischemic patterns in ECG signals. Our model incorporates domain-specific adaptations to improve sensitivity to these patterns, which are often overlooked by traditional CNN models.

3. Lightweight Yet High-Performing Architecture: CADNet achieves 99.3% accuracy with only 2,586 trainable parameters, significantly reducing computational overhead. This ensures its feasibility for real-time deployment in resource-constrained environments, an area where many existing deep learning models struggle due to high computational demands.

4. Novel Architectural Enhancements: Feature Encoding Block: Optimised convolutional layers to capture temporal characteristics relevant to CAD.

5. Compact Pooling Block: A new pooling strategy designed to retain essential ECG features while improving computational efficiency.

6. Comprehensive Benchmarking and Generalisation: We rigorously evaluated CADNet across four diverse datasets (PTB-XL, MIMIC-III, St. Petersburg, Fantasia), demonstrating superior accuracy and robustness. Additionally, we assessed its performance across different age groups to confirm its generalisability.

CADNet enables real-time CAD diagnosis, which is important given the often asymptomatic nature of CAD. CAD can develop over extended periods without clear clinical symptoms, resulting in many patients remaining undiagnosed until the occurrence of a CAD-related event. Therefore, real-time diagnosis and long-term health monitoring, when deployed on portable devices, are essential for supporting early detection and for reducing the workload on healthcare professionals.

The remainder of this chapter is structured as follows. Section 5.3 details the proposed methodology, including data preparation, purification, model architecture and optimisation. Section 5.4 describes the experimental setup, performance metrics, datasets, and results and discussions. Section 5.5 is devoted to drawing limitations of this research. Finally, Section 5.6 concludes the study.

5.3 Materials and Methods

This section outlines different steps for comprehensive CAD classification using our proposed CADNet model. Figure 5.1 illustrates the process of classifying binary classes,

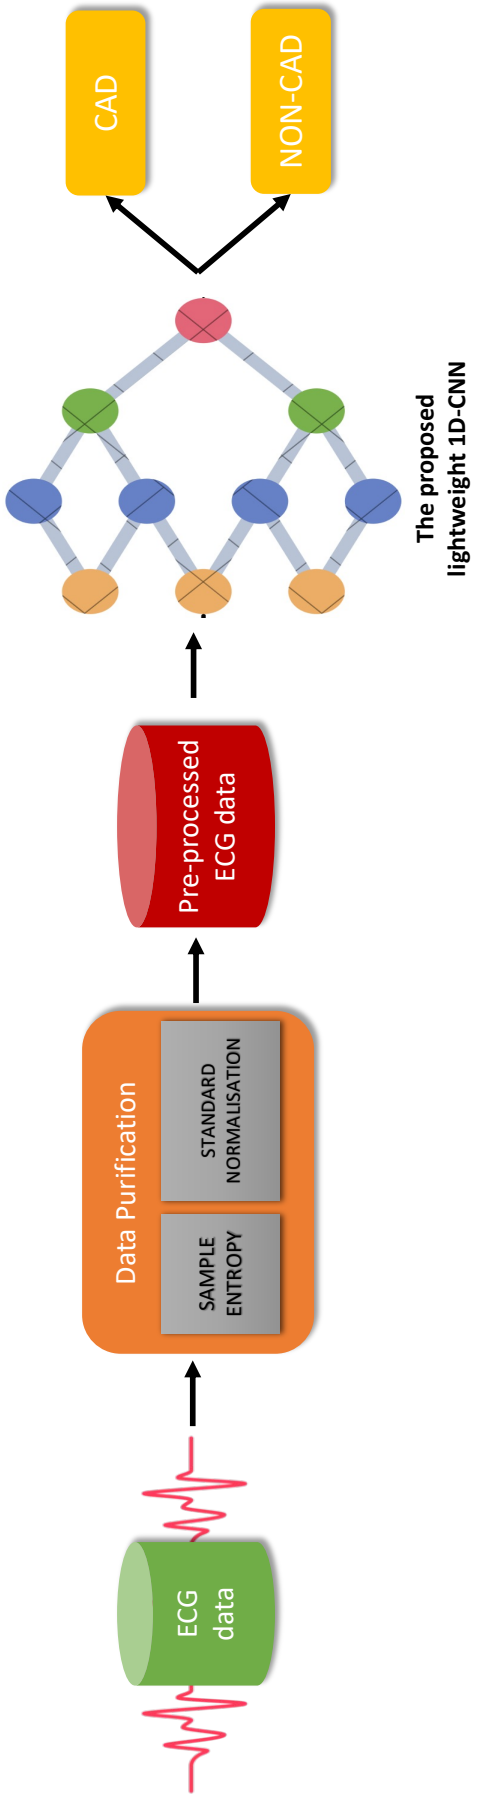

FIGURE 5.1: Flowchart outlining the operational framework of the proposed model

i.e. CAD and NON-CAD, from ECG data. The process initiates with raw ECG signals, obtained from the MIMIC III, Fantasia, St Petersburg and PTB-XL databases. For training, the ECG signals are meticulously chosen from MIMIC III and Fantasia databases, partitioned into 1-second segments and categorised into subsets, as outlined in Section 4.3.1. For testing, the same procedure is followed, but using St Peterburg and PTB-XL databases. A novel feature engineering technique is then applied to the subsets to eliminate noise and irrelevant data (Section 5.3.2). The proposed CADNet model (described in Sections 5.3.4 and 5.3.3) is subsequently employed to classify the ECG signals as either indicative of CAD or NON-CAD.

5.3.1 Data preparation

5.3.1.1 Data source

The ECG data used in this study are sourced from four publicly available databases accessible through PhysioNet: MIMIC-III [199], St. Petersburg [217], PTB-XL [218] and Fantasia [200]. The details of each dataset are outlined below:

MIMIC-III The MIMIC-III dataset includes a total of 2,840 patients diagnosed with coronary atherosclerosis in the native coronary artery, accounting for approximately 7.1% of total hospital admissions. ECG recordings from these patients were extracted and utilised for analysis.

St. Petersburg This dataset consists of 30-minute ECG recordings from 75 subjects, among whom 7 patients have been diagnosed with CAD. The ECG signals from these patients were selected for this study.

PTB-XL PTB-XL is a large-scale ECG dataset containing over 20,000 10-second recordings from approximately 18,000 patients. This dataset includes a variety of cardiac conditions, providing essential subclasses that support CAD diagnosis.

Fantasia The Fantasia dataset consists of ECG recordings from a cohort of 40 individuals, evenly distributed between 20 young and 20 adult subjects. This dataset was used to represent NON-CAD patients in our study.

5.3.1.2 Data preprocessing and segmentation

ECG recordings from CAD and NON-CAD patients were selected from each database and segmented into one-second ECG data segments, each comprising $N = 250$ samples as shown in Figure 5.2. A similar data selection strategy was used in [6], where ECG recordings from the Fantasia and St. Petersburg databases were used to represent normal and CAD subjects, respectively. This segmentation method enables detailed analysis of the dynamic variations and patterns in ECG signals within short temporal windows. Such an approach has been widely used to maintain consistency in ECG signal analysis [6, 148, 213], as it captures the ECG cycle without relying on waveform detection methods. Following segmentation, a feature engineering process (explained in the next subsection) was applied to remove irrelevant noise within the ECG segments. To facilitate classification, a binary label was assigned to each sample:

- Label 0: Non-CAD subjects from the Fantasia database.
- Label 1: CAD-diagnosed subjects from the MIMIC-III, St. Petersburg, and PTB-XL datasets.

5.3.1.3 Dataset composition for training and evaluation

A primary subset comprising 200 ECG signals was created for training and testing. The dataset maintains an equal distribution of CAD and NON-CAD subjects, with 100 ECG signals sourced from the MIMIC-III and Fantasia databases. To ensure robust evaluation, we employed ten-fold cross-validation, a widely used technique in both traditional machine learning and deep learning. For each fold, the dataset was split 70% for training and 30% for testing. To mitigate overfitting, dropout and early stopping strategies were implemented. Training was terminated if no improvement in validation loss was observed over eight consecutive epochs.

5.3.2 Feature Engineering

Since biomedical data are normally recorded with many unwanted noises and artifacts, proper feature engineering would play a pivotal role in ensuring the quality and consistency of input data (ECG here). And, the signal quality notably influences the model's performance [134, 214]. Consequently, researchers consider implementing various preprocessing techniques to ensure optimal accuracy.

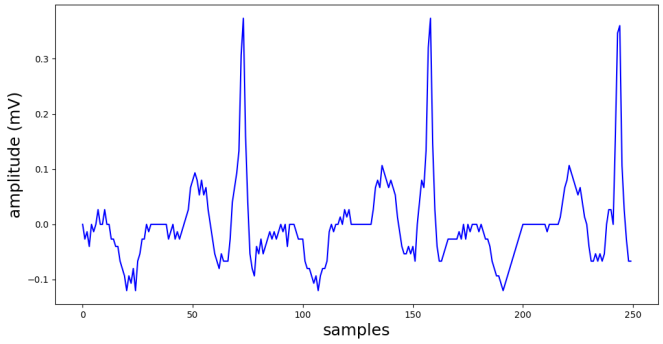

(a) ECG signal from MIMIC III.

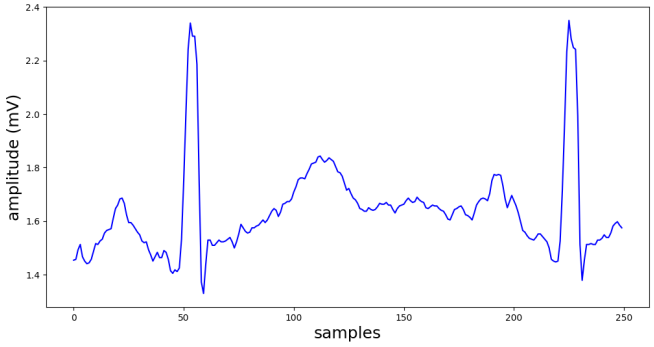

(b) ECG signal from St Petersburg.

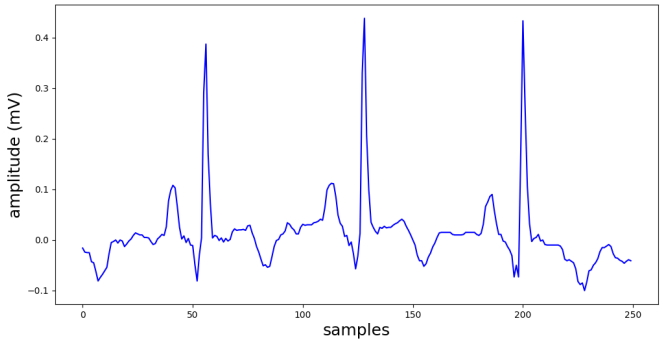

(c) ECG signal from PTB-XL.

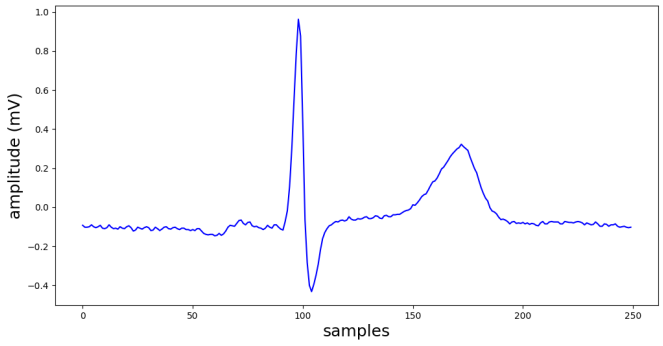

(d) Healthy ECG from MIMIC III.

FIGURE 5.2: Examples of ECG signals from different datasets used in this study.

Sample entropy emerges as a prominent method within temporal data applied in previous studies to enhance signal quality [59, 219]. Unlike standard entropy or approximate entropy, sample entropy provides a more consistent and robust estimate of signal complexity, especially in shorter and noisier time-series signals, which are common in real-world ECG recordings. Here, we propose a feature engineering technique, based on sample entropy, to facilitate CAD classification, combining the sample entropy concept and a standard normalisation step. In CADNet, this combination is applied to the raw ECG data to mitigate artifacts present in the signal. Sample entropy functions as a metric to evaluate the time series data quality as shown in Equation 4.1. Additionally, standard normalisation is applied to eliminate flat time series data and mitigate any potential influences that could affect the model's accuracy as shown in Equation 4.2.

Figure 5.3 illustrates examples of ECG signals evaluated during the feature engineering process. Figure 5.3(a) shows a high-quality signal exhibiting appropriate morphological complexity and variance, with a $SampEn$ below 0.1 and a standard deviation σ above 0.1. Figure 5.3(b) presents a signal with $SampEn > 0.1$, likely due to noise or random fluctuations. Figure 5.3(c) shows a signal exhibiting flatness, with a $\sigma < 0.1$. Only signals that met both criteria were used for further analysis.

To further quantitatively evaluate the impact of feature engineering, the average Root Mean Squared Error (RMSE) was calculated between the original ECG signals and the processed ECG signals after the feature engineering process, for both CAD and NON-CAD cases. RMSE represents the average difference between the original and processed ECG signals over time and provides a quantitative measure of how much the waveform changes as a result of feature engineering. The RMSE was 3.66 for CAD and 4.26 for NON-CAD signals. These relatively low values suggest that the feature engineering step effectively removed inappropriate ECG segments while preserving the overall ECG waveform structure. The slightly higher RMSE in the NON-CAD cases may be due to the more regular and stable patterns typically found in healthy ECG signals, which can be more easily removed during the feature engineering. However, this does not compromise the model's ability to learn from the ECG signals. Importantly, the RMSE values remain low, indicating that any changes introduced by feature engineering are minor and unlikely to affect diagnostically relevant ECG features in both CAD and NON-CAD cases. Visual inspection confirmed that key clinical components, such as the P-wave, QRS complex, and T-wave, remained clearly visible.

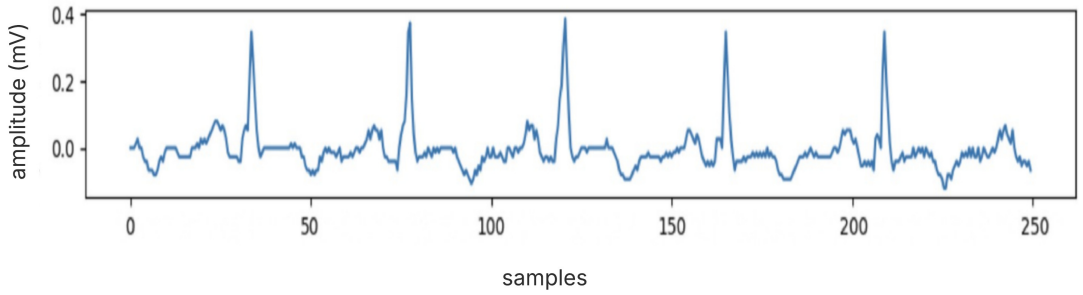

(a) Signal with $SampEn < 0.1$ and $\sigma > 0.1$

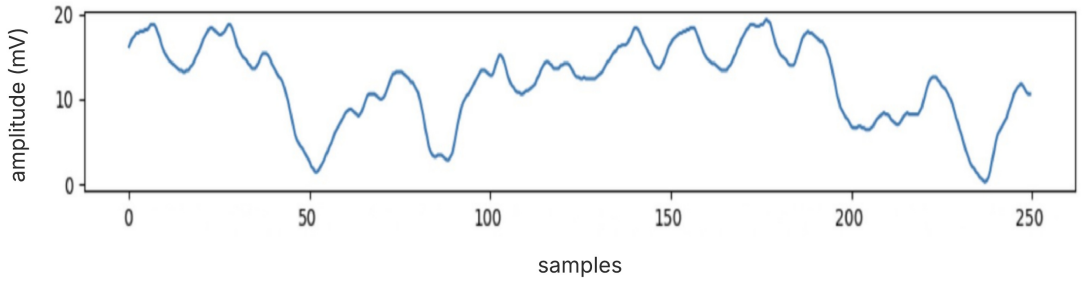

(b) Signal with $SampEn > 0.1$

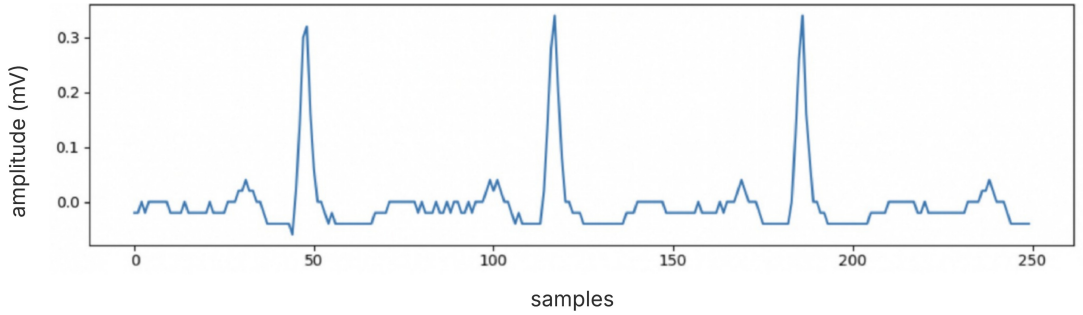

(c) Signal with $\sigma < 0.1$

FIGURE 5.3: Examples of ECG signals during the feature engineering process. Signals with SampEn < 0.1 were considered to exhibit appropriate complexity, while a standard normalisation threshold of $\sigma > 0.1$ was used to exclude flat signals. Only ECG signals satisfying both criteria were retained for further analysis.

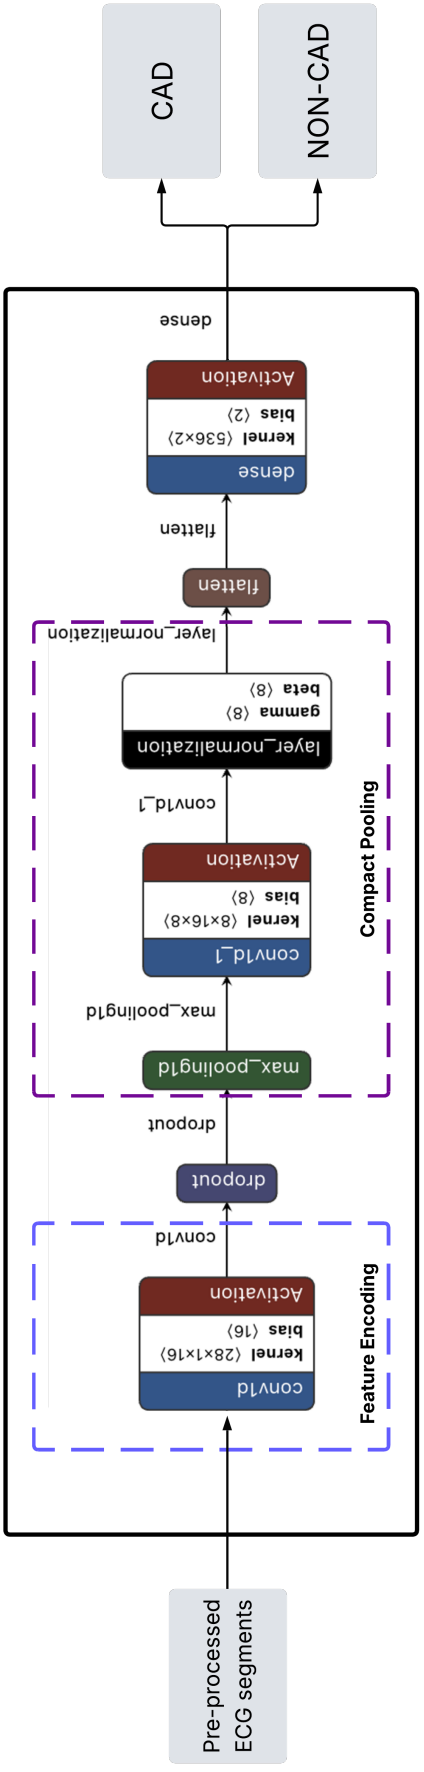

FIGURE 5.4: Detailed architecture of the proposed CADNet including feature encoding, compact pooling and classification blocks.

5.3.3 Lightweight deep learning architecture

Figure 5.4 illustrates the operational framework of our CADNet model. It is a lightweight 1D-CNN designed for efficient and effective CAD detection and optimises computational resources while delivering robust performance in a CAD classification task. The proposed CADNet model consists of two main blocks: Feature Encoding and Compact Pooling layers. These components enhance the model's robustness to noise and fluctuations in ECG signals, thereby improving its ability to classify CAD with precision and efficiency. The fully connected layer comprises two neurons and utilises Softmax activation, facilitating the classification of the input ECG signal into one of two potential classes, thereby representing the probabilities of the input belonging to each class. Furthermore, a dropout layer is incorporated between the primary blocks to augment the model's capacity for generalisation.

Feature Encoding block plays a crucial role in transforming raw ECG input data into a higher-dimensional feature space, efficiently capturing critical temporal patterns while minimising dimensionality. This layer preserves essential diagnostic characteristics, including P waves, QRS complexes, and T waves, facilitating the efficient extraction of significant ECG features with minimal computational cost. This block is a convolutional layer consisting of 16 filters. Each filter is responsible for identifying specific patterns or features within ECG data. The kernel size for each filter is set to 28. Consequently, during each convolution operation, the filter examines a window of 28 consecutive ECG samples from the ECG signal. This approach enables the filter to capture local patterns or features within the ECG signal that may indicate particular cardiac events or abnormalities of CAD. ReLU activation functions are employed within the block for their simplicity and computational efficiency. These functions introduce non-linearity to the model, facilitating the learning of complex patterns in the data while managing computational capacity effectively. The output of the feature encoding layer is calculated as:

$$y_i = \sum_{h=1}^{K} x_{i+h-1} \cdot w_h + b \tag{5.1}$$

where y_i represents the output signal at i^{th} position within the feature encoding layer. K denotes the number of weights within the filter, determining the size of the window that slides the input ECG signal. The index h represents the position within the filter. x_{i+h-1} represents each element of the ECG input aligned with a specific weight in the filter, denoted by w_h, during the computation of the output signal at i position. The bias term, represented by b, is initialised to zero. This initialisation enables the model to achieve greater flexibility in fitting the ECG data by adjusting the activation function.

The compact pooling layer is employed to reduce the temporal resolution of the ECG signals, refine the extracted features, and ensure that only the most relevant information for classification is passed to subsequent layers. This layer comprises max pooling, convolutional, and layer normalisation layers. The max pooling layer is introduced to operate within sliding windows of size 3, effectively reducing the temporal resolution of the signal while preserving essential temporal ECG features. Subsequently, a one-dimensional convolutional layer with a filter size and kernel size of 8 each is employed. The ReLU activation function is then applied to facilitate the detection of complex patterns indicative of CAD. Finally, layer normalisation is used to ensure promoting convergence and enhancing the model's ability to discern relevant CAD patterns in the ECG data. The layer normalisation can be computed via:

$$y_i = \gamma \left(\frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} \right) + \beta \tag{5.2}$$

where y_i represents the output signal of the normalisation layer for the i^{th} position. Learnable parameters are denoted by γ and β, respectively. x_i indicates the input signal at i^{th} position. μ represents the average of the input signal. σ^2 denotes the variability or dispersion of ECG features, where ϵ would be a small constant added to the denominator for numerical stability, ensuring that the denominator is never zero or too close to zero during computations.

5.3.4 Optimisation

The model is trained using binary cross-entropy loss, the Adam optimiser, ReLU and Softmax activation functions. The hyperparameters employed during model training, including learning rate, batch size, and number of epochs, are provided in Table 5.1. A binary cross-entropy (BCE) loss function is employed for CAD and NON-CAD classification purposes due to its effectiveness in handling binary classification tasks:

$$\text{BCE} = -\frac{1}{N} \sum_{i=1}^{N} \left(l_i \log(p_i) + (1 - l_i) \log(1 - p_i) \right) \tag{5.3}$$

where N represents the total number of samples contained within the respective ECG segment. l_i denotes the actual label assigned to ECG signals i, where $l_i \in \{0, 1\}$. $l_i = 0$ corresponds to a NON-CAD case, indicating the absence of CAD features within ECG signal i. Conversely, $l_i = 1$ indicates a CAD case, signifying the presence of relevant CAD characteristics in ECG signal i. p_i denotes the anticipated probability that ECG signal i is associated with class 1. It assists the model in determining the accuracy of

TABLE 5.1: Hyperparameters used for model training

Hyperparameter	**Value**
Learning Rate	0.0001
Batch Size	32
Number of Epochs	50
Optimiser	Adam
Loss Function	Binary Cross-Entropy
Hidden Layer Activation	ReLU
Output Layer Activation	Softmax

its predictions, thereby facilitating adjustments to enhance its predictive capability to align more closely with the actual labels.

Adam optimiser was utilised to enhance the efficacy of training our CADNet, employing a learning rate of 0.0001. Its capability to dynamically adjust learning rates for individual parameters ensures proficient optimisation, accommodating the nuanced gradients inherent within ECG data.

ReLU activation function, a key component of convolutional layers due to its simplicity and effectiveness, is presented via:

$$\text{ReLU}(x) = \max(0, x) \tag{5.4}$$

where an ECG input signal represented by a vector $x = [x_1, x_2, \ldots, x_n]$. ReLU transforms negative inputs to zero while preserving positive inputs, rendering it computationally efficient and ensuring differentiability across its domain, except at zero. This capability facilitates the network in capturing intricate patterns and features inherent in ECG data.

5.4 Experimental results

Extensive experiments were conducted to evaluate the performance of the proposed model. A machine equipped with an Apple M2 Max processor and 32 GB of unified memory was used to run all the experiments. The implementation was carried out using Python 3.9.6. We also provide comparative analyses with classical machine learning algorithms, existing DL-based models, and well-known lightweight models. Table 5.1 summarises the hyperparameters used in our experiments.

In this study, standard classification metrics were employed, namely accuracy (Acc), precision (Ppr), Sensitivity (Sen), Specificity (Spr), and F1 Score ($F1$) for evaluating the classification performance as shown in Equations (2.5), (2.7) , (2.6), (2.8), and (2.9).

5.4.1 Ablation study

Table 5.2 provides an ablation study to systematically evaluate the effectiveness of different developments we applied to the baseline 1D-CNN [31], leading to the proposed CADNet model. As found from Table 5.2, the baseline 1D-CNN model achieved a high accuracy of 99.3%, requiring approximately 8 million trainable parameters, a size of 32,190 KB, and a runtime of 854.8449 seconds. This highlights its substantial complexity, runtime demands, and storage and computational requirements, making it computationally intensive and unsuitable for resource-constrained environments, such as wearable devices or real-time monitoring systems. To address this limitation, we developed a lightweight model through a feature encoding layer. This new architecture reduces the trainable parameters to 8,552, with a size of 33.41 KB, while reducing the accuracy to 97.8% (Table 5.2). This suggests that the feature encoding layer could significantly lower parameters at a price of a slight reduction in accuracy compared to the baseline 1D-CNN. To further improve the model performance, the feature encoding layer and compact pooling were integrated to attempt a further reduction in the number of trainable parameters while maintaining accuracy. The integration of these two layers successfully reduced the number of trainable parameters and size (Table 5.2). We have implemented several robust regularisation techniques in our proposed CADNet architecture to avoid overfitting. Specifically, we have employed:

- Dropout tuning: We have incorporated dropout layers within our model architecture. By randomly deactivating neurons during training, these layers reduce the model's dependency on any single feature, thereby enhancing its ability to generalise to unseen data.
- Early stopping: To further mitigate overfitting, we applied an early stopping mechanism. Training is terminated if no improvement in the validation loss is observed over eight consecutive epochs. This prevents the model from training excessively on the training data and helps in avoiding overfitting.
- Model efficiency: As detailed in Table 5.2, CADNet features a minimal number of trainable parameters and a compact storage size, which naturally limits the model's capacity to overfit. Despite these constraints, our model achieved an accuracy of 99.3% and maintained a runtime of 10.0856 seconds, making it highly suitable for real-time CAD diagnosis in devices with limited computational resources.

TABLE 5.2: Ablation analysis among the baseline 1D-CNN model and different layers of the proposed CADNet.

Model Architecture	**Trainable params**	**Size**	**Acc (%)**	**Run time (s)**
Baseline 1D-CNN [31]	8,439,426	32.19 MB	99.3	854.8449
Feature Encoding layer	8,552	33.41 KB	97.8	5.0911
Feature Encoding layer + Compact Pooling layer	2,586	10.10 KB	98.5	8.3467
CADNet: Feature Encoding layer + Compact Pooling layer + Dropout	2,586	10.10 KB	99.3	10.0856

As a result, the model becomes less sensitive to noise in the training set, enabling it to better generalise to unseen samples. The proposed CADNet (Table 5.2) ultimately featured a minimal number of trainable parameters and compact storage size, while maintaining a high level of accuracy at 99.3% with a shortest runtime of 10.0856 seconds, making it suitable for real-time CAD diagnosis. This significant improvement ensures its feasibility for devices with limited computational resources, bridging the gap between high performance and practical applicability.

5.4.2 Computational analysis

Figure 5.5 indicates the impact of each proposed architecture in reducing the number of trainable parameters. The baseline model was able to achieve an overall classification accuracy of 99.3%. However, the number of trainable parameters remains notably high, indicating significant computational demands. It can be observed that as the proposed feature encoding, compact pooling, and dropout layers are successively added, the number of trainable parameters decreases from approximately 8 million to 8,552, 2,586, and 2,586, respectively. Remarkably, even as the model becomes more streamlined with fewer parameters, its performance remains consistently high, maintaining high accuracy. This observation underscores the effectiveness of the proposed techniques in optimising the model's complexity without compromising its predictive capability.

5.4.3 Comparative study

Table 5.3 illustrates a comparison among different traditional ML-based, DL-based network and lightweight network architectures. Additionally, it highlights differences in model size and runtime across various architectures, providing insights into their computational efficiency and potential suitability for resource-constrained environments. The models are evaluated based on trainable parameters, file size, runtime and performance metrics including Acc, AUC, and $F1$ score. The classical ML-based methods applied to our subset include SVM, Gaussian Naive Bayes, K-Means, KNN, and Logistic Regression. The various DL-based models we compared include LSTM [113], CNN-LSTM [113], RNN [220], and a baseline 1D-CNN [31]. Moreover, well-known lightweight networks, including SqueezeNet [178], MobileNetV1 [180], EfficientNetB0 [179] and ShuffleNetV1 [123], are employed to evaluate performance on our subset, alongside our CADNet model. MobileNetV1 was modified to process one-dimensional ECG signals by replacing its two-dimensional convolutional layers with one-dimensional operations. Among the classical ML-based models, Gaussian Naive Bayes and Logistic Regression are highly efficient, with small sizes of 8.41 KB and 2.68 KB, minimal or no trainable

TABLE 5.3: Comparative analysis of traditional ML-based, DL-based and lightweight networks utilising ten-fold cross-validation.

Model Architecture	**Trainable params**	**Size (KB)**	**Acc (%)**	**AUC (%)**	**F1 score (%)**	**Runtime (s)**	**Inference time(ms)**
SVM	0	38.72	83.5	82.5	83.1	0.7087	0.0527
Gaussian Naive Bayes	0	8.41	92.4	98.7	92.4	0.0039	0.0419
KNN	0	353.66	93.9	95.7	93.9	0.0030	1.0548
K-Means	0	5.24	75.0	79.1	74.5	1.1976	0.1046
Logistic Regression	251	2.68	84.8	82.1	84.7	0.0085	0.0272
LSTM[113]	51,102	199.62	95.5	96.0	95.4	85.7291	0.2998
CNN-LSTM [113]	63,630	248.55	97.0	97.0	96.9	49.6258	3.7010
RNN [220]	232,658	908.82	96.9	96.3	96.0	226.0266	31.4683
Baseline 1D CNN [31]	8,439,426	32,190	99.3	99.0	98.5	854.8449	12.8933
SqueezeNet [178]	354,370	1,401.85	54.5	55.0	39.7	191.7461	0.6807
EfficientNetB0 [179]	7,003,266	27,636.79	97.0	99.1	96.6	2,081.4195	186.4941
1D-MobileNetV1 [180]	3,167,554	12,392.10	49.0	51.9	32.8	1,475.6756	23.7306
ShuffleNetV1 [123]	753,578	3,062.25	48.5	44.8	44.5	106.2833	2.1184
CADNet	2,586	10.10	99.3	99.0	99.0	10.0856	0.3250

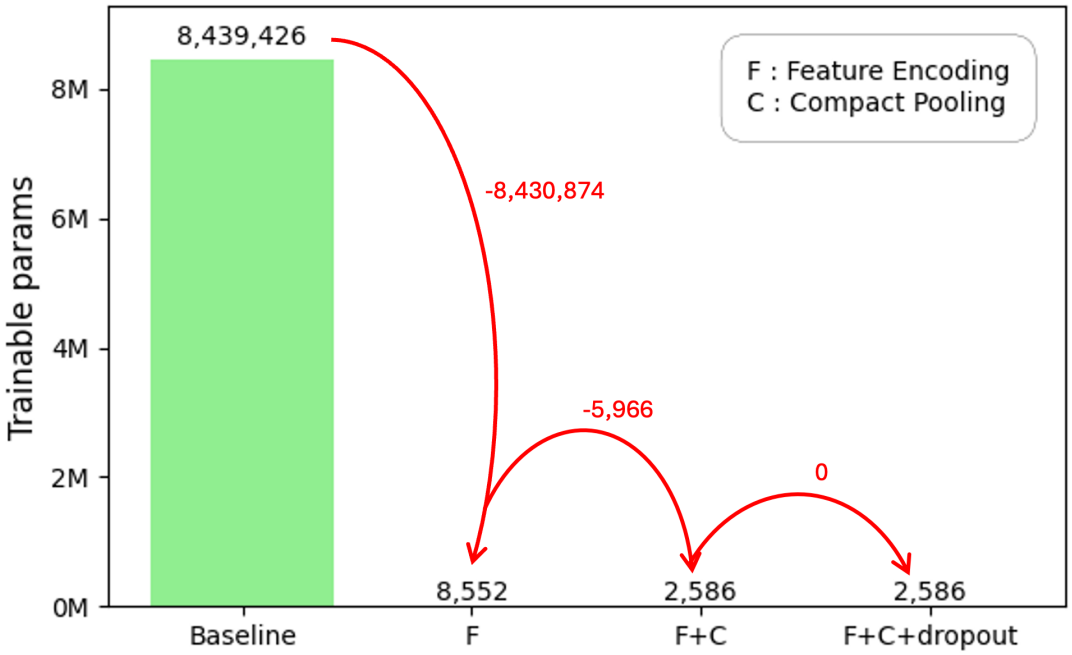

FIGURE 5.5: Comparison of parameter meter analysis between the established CNN model and our proposed lightweight model. The chart illustrates the key metrics derived from the analysis, showcasing the efficiency and effectiveness of our proposed model in terms of parameter utilisation.

parameters, and very low runtimes of 0.0039 and 0.0085 seconds, making them suitable for resource-constrained environments. Although KNN is relatively large at 353.66 KB and has no trainable parameters, it achieves a quick runtime of 0.0030 seconds. In contrast, K-Means, the smallest model at 5.24 KB, has a longer runtime of 1.1976 seconds, reflecting trade-offs between memory efficiency and computational speed across ML-based architectures. The KNN and LSTM models achieve the highest accuracy, at 93.9% and 95.5%, respectively. Gaussian Naive Bayes at 92.4% and Logistic Regression at 84.8% provide a favourable balance of accuracy and efficiency. Among the ML-based models, K-Means, with the lowest accuracy at 75.0%, demonstrates limited predictive capability. The baseline 1D-CNN exhibits the highest number of trainable parameters and the longest runtime among DL-based models, totalling 8,439,426 parameters and 854.844 seconds, which signifies the complexity of this model. Conversely, our CADNet stands out for its notably lighter parameter count, comprising only 2,586 parameters, thereby hinting at a more streamlined DL-based architectural design. Furthermore, LSTM, CNN-LSTM, and RNN also exhibit lower trainable parameter counts compared to the baseline 1D-CNN model. The sizes of the models display significant variability, with the baseline 1D-CNN being the largest, totalling 32,190 KB. In contrast, the CADNet is considerably smaller, occupying only 10.10 KB, while LSTM, CNN-LSTM, and RNN contain 199.62, 248.55, and 908.82 KB, respectively. The corresponding runtimes for these models are 85.7291 seconds for LSTM, 49.6258 seconds for CNN-LSTM,

226.0266 seconds for RNN, and 10.0856 seconds for the CADNet model, making our CADNet model particularly suitable for deployment in resource-constrained environments where memory limitations are a concern, as it is both significantly smaller in size and faster than other DL-based models. LSTM achieves an accuracy of 95.5%, which, while commendable, represents the lowest performance within the group. Conversely, the CNN-LSTM model, integrating convolutional layers with LSTM layers, demonstrates a notable enhancement in accuracy, achieving 97.0%. RNN showcases an accuracy closely approximating that of the CNN-LSTM, at 96.9%. Both the baseline 1D-CNN and the CADNet achieve the highest accuracy at 99.3%, significantly surpassing the other DL-based models. This heightened level of accuracy suggests that both models excel in feature detection and classification, likely attributed to the robust capability of CNN in extracting crucial features indicative of CAD from ECG data. Additionally, this suggests that despite its simplicity, the CADNet does not compromise on predictive performance. To further demonstrate the suitability of the CADNet model for deployment in resource-limited environments, comparisions are conducted with well-known lightweight models, including SqueezeNet, EfficientNetB0, 1D-MobileNetV1, and ShuffleNetV1. The CADNet model exhibits only 2,586 trainable parameters and a runtime of 10.0856 seconds, achieving the highest accuracy at 99.3%. In contrast, EfficientNetB0, with over 7 million trainable parameters, achieves a high accuracy of 97% but requires a considerably longer runtime of 2,081.4195 seconds, highlighting its substantial demand for computational resources. SqueezeNet, 1D-MobileNetV1, and ShuffleNetV1, with 354,370, 3,167,554, and 753,578 trainable parameters, respectively, exhibit lower accuracy and extended runtimes of 106.2833 seconds, 1,475.6756 seconds, and 106.2833 seconds. These results highlight the CADNet model's advantage in achieving high performance with minimal computational resources.

To further evaluate computational efficiency, inference time is assessed to evaluate the suitability of each model for deployment in real-time and resource-constrained environments. The results indicate that traditional machine learning methods, including SVM, Gaussian Naive Bayes, Logistic Regression, and KNN, achieve the lowest inference times. Among ML-based methods, Gaussian Naive Bayes is the most efficient, requiring only 0.0419 ms per ECG signal. KNN is comparatively slower, averaging 1.0548 ms per ECG signal. Among DL-based models, the CNN-LSTM shows a good balance between prediction accuracy and speed, with an average inference time of 3.7010 ms per ECG signal. In comparison, the RNN takes longer, averaging 31.4683 ms per ECG signal, due to the extra processing required for handling sequences and its larger number of trainable parameters. Lightweight DL-based architectures, such as SqueezeNet and ShuffleNetV1, show relatively fast inference times. In comparison, more complex models, including EfficientNetB0 and 1D-MobileNetV1, require much longer, with inference times of 186.4941

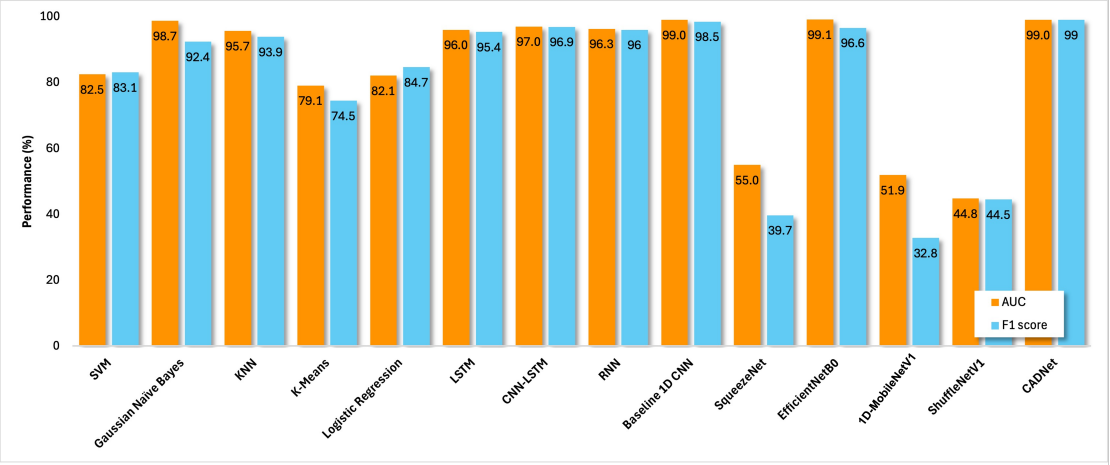

FIGURE 5.6: Performance comparison of ML-based, DL-based, lightweight models, and our CADNet model using AUC and F1 metrics to assess class distinction and predictive effectiveness.

ms and 23.7306 ms per ECG signal, respectively. The proposed CADNet model achieves an inference time of only 0.3250 ms per ECG signal, while using just 2,586 trainable parameters, and still maintains a high level of classification accuracy. These results suggest that CADNet is well-suited for use in real-time or resource-constrained environments.

Figure 5.6 illustrates the comparative performance analysis between ML-based, DL-based, lightweight algorithms and our proposed lightweight model on CAD classification presented in Table 5.3, including the CADNet model. Gaussian Naive Bayes, KNN, CNN-LSTM, RNN, and the CADNet model achieve high AUC and F1 values, demonstrating strong class separation and balanced precision-recall performance. With an AUC and F1 score of 99%, the CADNet model emerges as the top performer, making it highly suitable for CAD diagnosis requiring exceptional reliability and accuracy. SVM, K-Means, and Logistic Regression display moderate performance, with slightly lower AUC and F1 values, suggesting fair but less consistent precision-recall balance. In contrast, SqueezeNet, EfficientNetB0, 1D-MobileNetV1, and ShuffleNetV1 exhibit low AUC and F1 scores, indicating limited effectiveness in both class distinction and classification balance.

5.4.4 Generalisation and interpretability

Table 5.4 illustrates the CADNet model trained using the PTB-XL database with a subset of 400 ECG signals: 200 CAD and 200 NON-CAD. The model demonstrates high classification performance within the training dataset and robust generalisability across unseen subsets. Achieving a training accuracy of 95.90% and a test accuracy of 92.05% on the PTB-XL dataset, the model exhibits reliable learning and validation

Table 5.4: Performance of the CADNet Model Trained with PTB-XL database.

Metric(s)	**Train Accuracy**	**Test Accuracy**	**Precision**	**Recall**	**F1 Score**
Trained with PTB-XL	0.9590	0.9205	0.9635	0.9565	0.9600
Test on Fantasia and MIMIC III	-	0.8700	0.7400	1.0000	0.8506
Test on Fantasia and St Peterburg	-	0.8800	0.7600	1.0000	0.8636

performance, as indicated by a precision of 96.35%, recall of 95.65%, and an F1 score of 96%. These results demonstrate the model's effectiveness in accurately identifying CAD and NON-CAD cases. To further assess the model's generalisability, we conducted an initial validation on a subset of 200 ECG signals: 100 NON-CAD signals from the Fantasia database and 100 CAD signals from the MIMIC III database. The model achieved a test accuracy of 87%, with a perfect recall of 100%, though with a slight reduction in precision to 74% on this validation set. The resulting F1 score of 85.06% indicates stable performance. Subsequently, we conducted a second validation on a subset of 200 ECG signals: 100 NON-CAD ECG signals from the Fantasia database and 100 CAD ECG signals from the St. Petersburg database. The model achieved a test accuracy of 88%, recall of 100%, precision of 76%, and an F1 score of 86.36%. While the model demonstrates robust sensitivity for CAD detection across diverse datasets, the observed reduction in precision on validation subsets suggests a sensitivity-specificity trade-off that could benefit from further optimisation to enhance specificity and reduce false positives. This performance underscores the potential of the CADNet model for CAD diagnosis, with scope for refinement to improve adaptability across varied ECG data sources.

Table 5.5 presents the CADNet model trained on ECG signals from the PTB-XL, MIMIC-III, and Fantasia databases, using a subset comprising 100 ECG signals from each database. The model achieved a training accuracy of 99.5% and a testing accuracy of 95.96%, demonstrating reliable performance in classifying CAD and NON-CAD cases. The high precision of 99.29%, together with a perfect recall of 100%, suggests that the model accurately identifies CAD cases, with minimal risk of overlooking true positives, as reflected in an F1 score of 99.64%. This level of performance indicates that the model is highly proficient in distinguishing between CAD and NON-CAD ECG signals, ensuring a balanced approach between sensitivity and specificity. The first validation subset consisted of 100 NON-CAD ECG signals from the Fantasia database and 100 NON-CAD ECG signals from the MIMIC-III database. The model's robustness was demonstrated, achieving a test accuracy of 99%. Precision slightly decreased to 98%, while recall remained perfect at 100%, resulting in an F1 score of 98.99%. This slight reduction in precision, alongside perfect recall, suggests occasional misclassification of NON-CAD cases as CAD but consistent identification of all true CAD cases. Subsequently, the second validation subset was formed with 100 NON-CAD ECG signals from the Fantasia database and 100 CAD ECG signals from the St. Petersburg database. The model continued to demonstrate high performance, achieving a test accuracy of 98.49%, with precision reaching 100% and recall at 95.08%, resulting in an F1 score of 97.48%. This level of precision indicates that all identified CAD cases are true positives.

Table 5.5: Performance of the CADNet Model Trained with PTB-XL, MIMIC III and Fantasia databases.

Metric(s)	**Train Accuracy**	**Test Accuracy**	**Precision**	**Recall**	**F1 Score**
Trained with Fantasia, MIMIC III and PTB-XL	0.9950	0.9596	0.9929	1.0000	0.9964
Test on Fantasia and MIMIC III	-	0.9900	0.9800	1.0000	0.9899
Test on Fantasia and St Peterburg	-	0.9849	1.0000	0.9508	0.9748

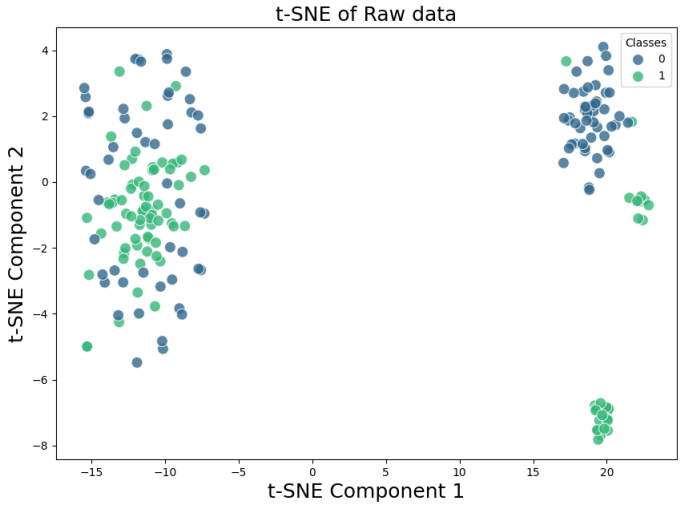

(a) T-SNE of original ECG signals

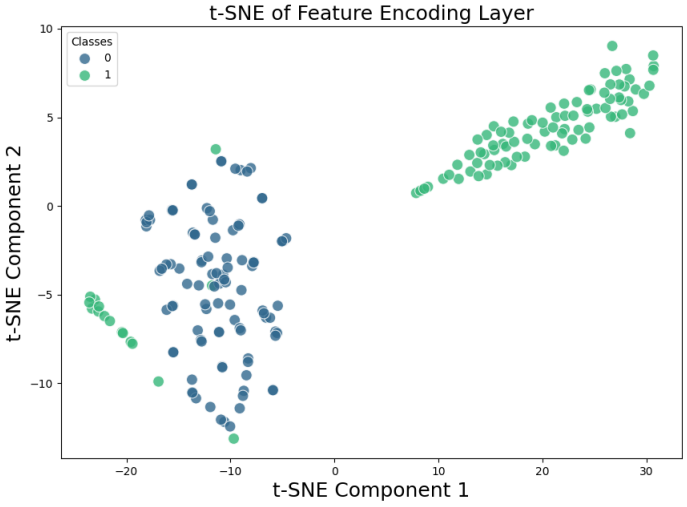

(b) T-SNE of Feature Encoding Layer

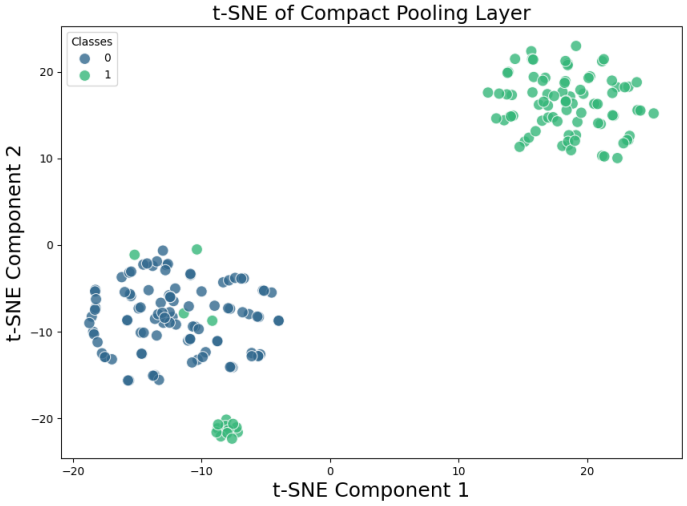

(c) T-SNE of Compact Pooling layer

FIGURE 5.7: T-SNE visualisation of the feature space, with class labels: 0 (Normal) and 1 (CAD).

TABLE 5.6: Test Accuracy Across Age Groups (Fantasia and PTB-XL Databases).

Age Group (Years)	**Accuracy (%)**
10 – 20	97.87
20 – 30	98.35
30 – 40	98.49
40 – 50	98.83
50 – 60	98.27
Over 60	97.54
Average	98.23

To enhance interpretability, t-SNE analysis was performed to visualise high-dimensional data and examine the role of two key modules: feature encoding and compact pooling layers in distinguishing between the two classes as shown in Figure 5.7. ECG signals exhibit significant overlap in Figure 5.7(a), indicating that the feature encoding layer is beginning to learn distinguishing features between the two classes. Although some overlap remains, a clearer boundary begins to emerge between NON-CAD and CAD ECG signals, suggesting that the model is progressively capturing patterns that facilitate class separation as shown in Figure 5.7(b). The compact pooling layer achieves near-complete separation of the two classes in the feature space, as shown in Figure 5.7(c). This indicates that the model has effectively learned to differentiate between NON-CAD and CAD samples. The distinct separation in this layer underscores that the features learned in this layer are highly effective at distinguishing between classes.

To evaluate the classification accuracy of the proposed model across diverse populations and assess its reliability for real-world clinical use, we validated the model trained in Table 5.5 across different age groups. ECG signals from the Fantasia and PTB-XL databases were used for this assessment, as presented in Table 5.6. As seen from this table, CADNet achieves the highest accuracy in the 40–50 age group at 98.83% and the lowest in the over 60 age group at 97.54%. Although the overall average accuracy is high at 98.23%, minor variations are evident across different age groups. Notably, performance declines slightly in both the 10–20 and over 60 age groups, suggesting that the model may be slightly less effective at generalising to these demographics. This trend could reflect underlying physiological variations in ECG signals associated with age.

5.5 Limitations

One of the main limitations in CAD diagnosis using ECG signals is the limited availability of datasets, primarily due to the absence of clear medical indicators of CAD,

such as single-vessel or multi-vessel disease. After a thorough exploration of publicly available datasets, we identified MIMIC-III, St. Petersburg, and PTB-XL as suitable and comprehensive sources for evaluation. In Chapter 6, we will expand our research by classifying different types of CAD within the PTB-XL dataset.

Although the CADNet achieved high accuracy on multiple public datasets, these may not fully reflect real-world clinical variability. Factors such as device differences, recording conditions, and patient diversity can affect performance. While we considered this by testing across different age groups, further validation in real clinical settings is needed to confirm generalisability.

While we focused on the superclass-level diagnostic categories to be able to use all four datasets simultaneously, analysing model performance at the subclass level could offer deeper insights into the model's discriminative ability across various CAD subtypes. Hence, subclass-level evaluation is a promising extension for future work and will be explored in Chapter 7 to refine the model's clinical applicability further.

5.6 Conclusion

Our study introduced the CADNet model to differentiate between cases of CAD and NON-CAD, to reduce the complexity of the model and facilitate its deployment on resource-constrained devices. Through the utilisation of data acquired from PhysioNet, our findings demonstrated the model's capability to independently classify these binary classes while maintaining its simplicity. The performance of the CADNet which averaged 99.3% accuracy with 2,586 trainable parameters, surpassed that of other classical machine learning, DL-based and lightweight models, highlighting the reduction in computational resources or complexity without compromising predictive performance. The code used to implement the proposed method is available from the corresponding author upon reasonable request.

Chapter 6

Real-time Coronary Artery Disease Detection from 12-Lead ECG Using A Lightweight Deep Network

6.1 Overview

While the previous chapters focused on CAD diagnosis using Lead II ECG only, this chapter expands the investigation to 12-lead ECG signals in order to capture inter-lead relationships and improve diagnostic robustness. In addition, specific CAD subclasses such as AMI, IMI and LMI are considered. Therefore, a lightweight deep learning-based classifier is purposed for the early diagnosis of CAD using 12-lead ECG signals. The model is designed to facilitate accessible clinical deployment by distinguishing between normal sinus rhythm and CAD cases, thereby supporting early, non-invasive detection aimed at reducing mortality. The proposed architecture comprises three key components: *i*) convolutional layers, to extract localised ECG signal patterns, *ii*) residual connections, to capture deeper and more informative ECG features, and *iii*) squeeze-and-excitation blocks, to apply adaptive lead-wise attention. We have deployed this unique architecture on STM32F469I-DISCO embedded board to examine its computational efficiency, inference time, energy consumption and memory usage. The results show that the proposed model, while demonstrating promising classification accuracy achieving 95.45% accuracy in classifying CAD and NON-CAD cases from 12-lead ECG, processes each subject's signals in an average of 0.1121 seconds and requires only 7.39 mJ, highlighting its suitability for real-time ECG-based CAD diagnosis on resource-constrained devices.

The study presents a novel ECG-based CAD classifier that achieves accurate discrimination between normal and CAD cases while maintaining low power consumption and memory usage when deployed on STM32F469I-DISCO. This study establishes a unique real-time, energy-efficient CAD diagnostic tool on embedded systems, aimed at enabling accessible clinical prescreening.

6.2 Related works

Several studies have investigated the deployment of lightweight ECG models on resource-constrained hardware platforms, with the aim of facilitating real-time ECG diagnosis. In [181], a lightweight deep learning model was implemented on an MSP432 microcontroller, achieving high classification accuracy while maintaining low power consumption. A verification mechanism incorporated into a 1D-CNN improved R-peak detection performance on low-quality Holter recordings [221], and an adaptive dual-algorithm framework has been proposed to dynamically switch between a lightweight detector and a more complex algorithm in order to optimise energy efficiency during high-intensity activity [222]. Hardware-accelerated implementations using FPGAs have also achieved fast and accurate ECG analysis with low energy requirements [223–225]. In addition, IoT-based platforms have been developed to enable continuous ECG monitoring and remote data transmission [226], and real-time signal-quality assessment techniques have been investigated for long-term wearable monitoring [227]. These studies highlight the growing interest in the practical deployment of lightweight ECG models for real-time and continuous monitoring.

In this chapter, we propose a novel lightweight model for automated CAD detection. The proposed model uses convolutional blocks to efficiently extract features from ECG signals, incorporating depthwise separable convolutions to capture both spatial and temporal patterns while reducing computational costs. Residual blocks help stabilise the learning process and improve efficiency by ensuring smooth feature propagation across network layers. The inclusion of a squeeze-and-excitation attention block further strengthens feature representation by adjusting the importance of different leads, leading to improved diagnostic accuracy. To preserve important temporal features, the model applies optimised pooling methods, such as max pooling and global average pooling, to balance data reduction with information preservation. Additionally, data augmentation using white Gaussian noise is introduced to ECG signals to improve generalisation, making the model more robust to variations in signal quality. The proposed model is designed to minimise the number of parameters by incorporating lightweight depthwise convolutions, residual connections, and attention mechanisms, resulting in a

more efficient structure. This model architecture ensures high diagnostic accuracy while maintaining suitability for deployment on resource-limited devices. Furthermore, SHAP analysis is applied to determine lead importance, offering insights into feature relevance for CAD diagnosis. The summary of the contributions is as follows:

- A novel lightweight DL-based model incorporating depthwise convolutions, residual connections, and squeeze-and-excitation blocks for CAD diagnosis using 12-lead ECG
- A novel data augmentation strategy is introduced, incorporating the addition of white noise to the ECG signals to enhance the model's generalisation capabilities. Standard normalisation is applied to ensure consistency and stability.
- SHAP analysis is applied to determine the importance of individual ECG leads, providing valuable insights into the contribution of each lead to the model's classification decisions.
- Bayesian optimisation is utilised for hyperparameter tuning, facilitating the selection of the optimal configuration of model parameters to enhance performance and generalisation.
- t-Distributed Stochastic Neighbor Embedding (t-SNE) is employed to analyse the model's generalisation capability, offering a means to investigate the separation between normal and CAD cases in high-dimensional feature space.
- The proposed model is successfully deployed on the STM32F469I-DISCO MCU, showcasing its computational efficiency and demonstrating its potential for real-time ECG classification.

The rest of this chapter is organised as follows. Section 6.3 presents the data preparation and data pre-processing, SHAP-based lead importance analysis, the proposed model architecture, hardware implementation, performance metrics, and the hyperparameter tuning process. Section 6.4 showcases and discusses the results, including performance evaluation, comparisons, and visualisations. Finally, Section 6.5 concludes the study and outlines directions for future research.

6.3 Materials and Methods

6.3.1 Data preparation

The primary ECG data used for model training and testing are derived from the PTB-XL database, available on PhysioNet [218]. PTB-XL is a comprehensive, publicly accessible ECG dataset comprising over 21,000 12-lead ECG signals from more than 18,000 patients. It is one of the most extensive datasets available, encompassing a wide range of clinical conditions related to cardiac diseases. In this chapter, we have utilised ECG data corresponding to several specific cardiac conditions to facilitate the diagnosis of CAD. These conditions include IMI, AMI, and LMI, each of which represents a distinct manifestation of myocardial injury and falls under the broader category of CAD. Additionally, Normal sinus rhythm (Norm) data are utilised to represent the baseline for patients without cardiac abnormalities. Figure 1.5 illustrates the spatial configuration of the standard 12-lead ECG signal, highlighting the orientation and angular projection of each lead for the heart. The limb leads I, II, III, aVR, aVL and aVF record electrical signals in the frontal plane, while the chest leads V1 to V6 provide a view across the body.

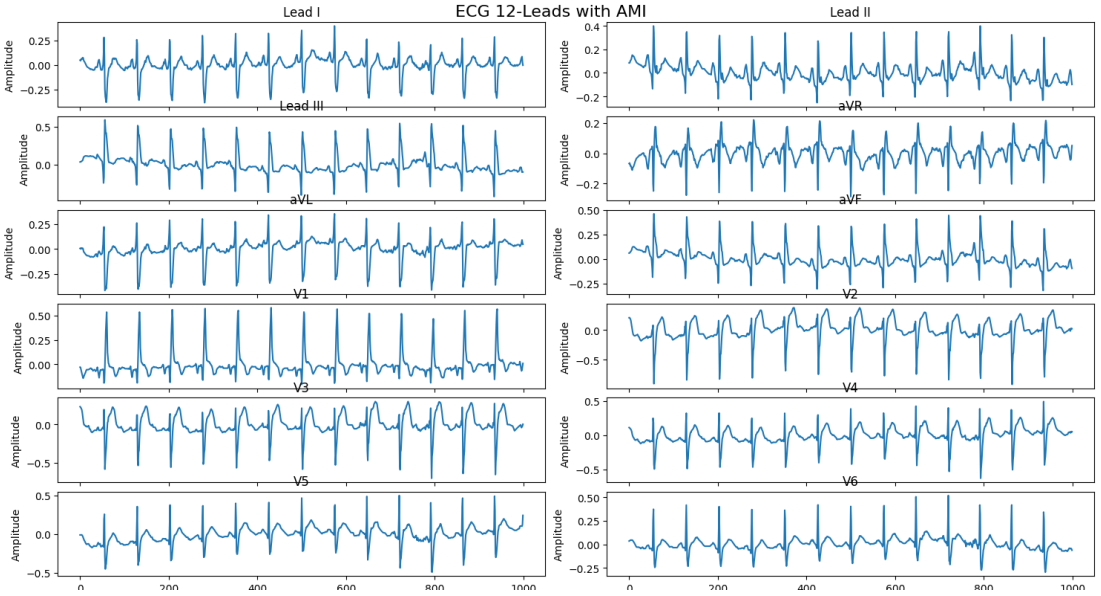

(a) An Example of 12-Lead ECG of CAD patients with AMI

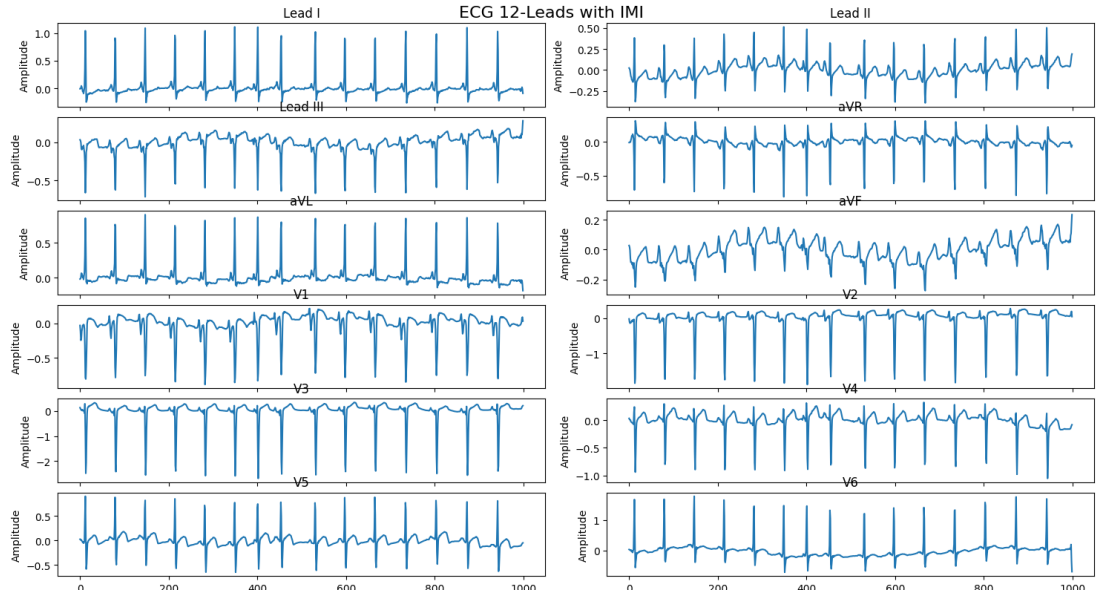

(b) An Example of 12-Lead ECG of CAD patients with IMI

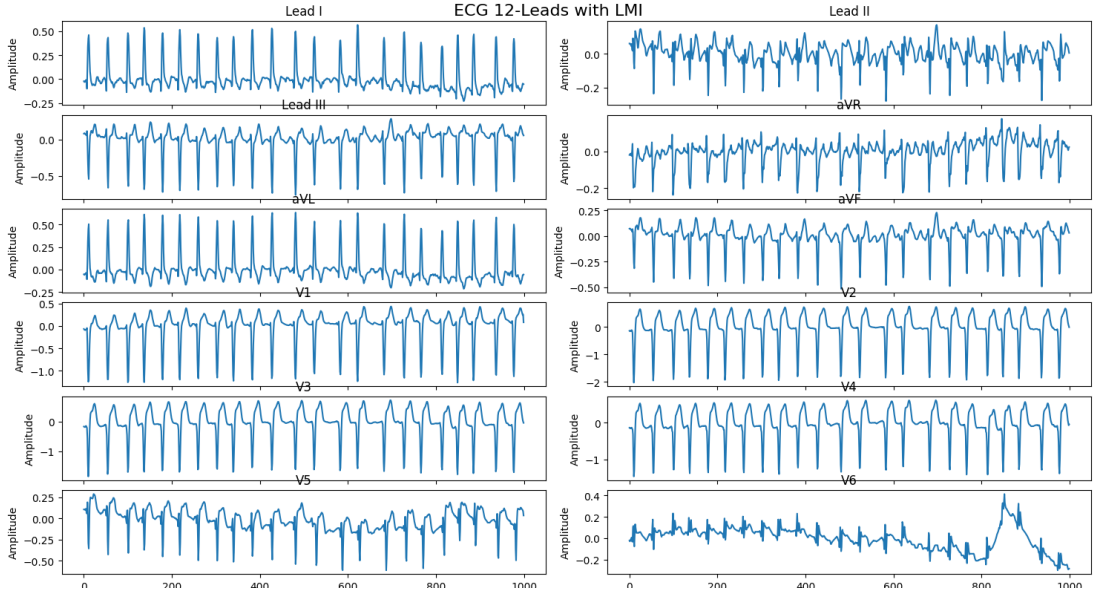

(c) An Example of 12-Lead ECG of CAD patients with LMI

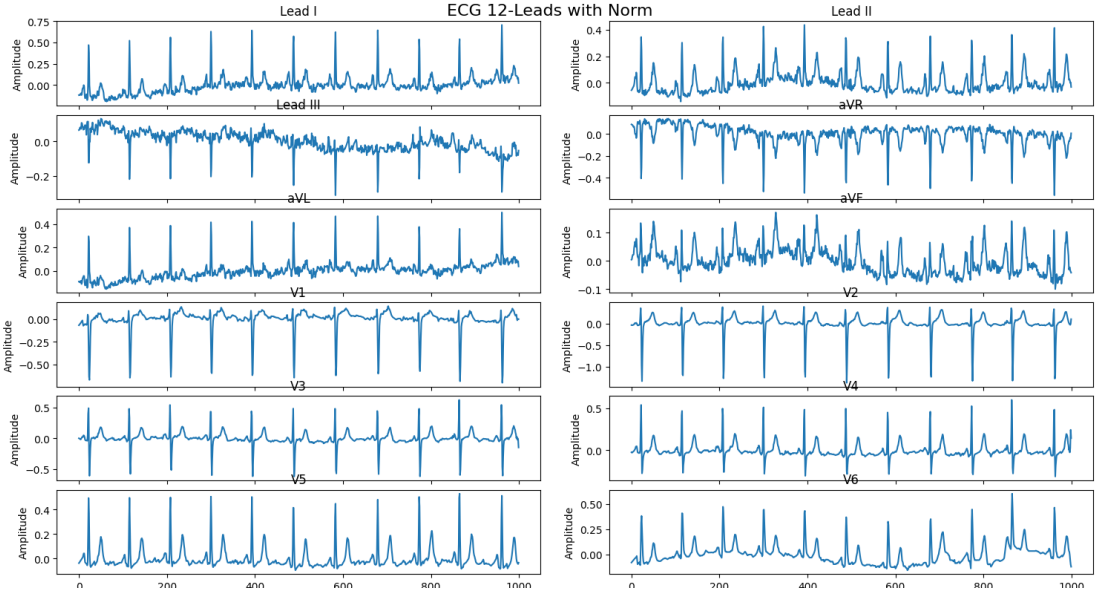

(d) An Example of 12-Lead ECG of patients with Normal sinus rhythm

FIGURE 6.1: Examples of 12-lead ECG recordings for CAD patients with different types of myocardial infarction (AMI, IMI, LMI) and a healthy patient with normal sinus rhythm.

To create a dataset ($\mathcal{D}$), 12-lead ECG signals with the following Standardised Cardiac Patient (SCP) codes: IMI, AMI, LMI, and Norm, were extracted from the PTB-XL database. Samples of these codes are illustrated in Figure 6.1. With the PTB-XL dataset sampled at 100 Hz, each ECG lead contains 10 seconds of data, corresponding to 1000 timesteps. After extracting the corresponding 12-lead ECG data based on SCP codes, the ECG data were labelled as 0 for normal (Norm) and 1 for CAD subclasses,

including AMI, LMI, and IMI. AMI, IMI and LMI are grouped under the broader category of CAD, reflecting their shared pathophysiological basis in acute ischaemia caused by coronary artery obstruction. This grouping addresses class imbalance and enhances model learning and generalisability. Although each subtype differs in anatomical location, all are critical manifestations of CAD that require timely recognition and intervention. The class distribution in the PTB dataset is as follows: Norm includes 9,514 subjects, IMI 1,022 subjects, LMI 619 subjects, and AMI 116 subjects. Additionally, to ensure balance among the different CAD subclasses, 100 subjects from each subclass and 300 Norm subjects were selected. The remaining ECG recordings were excluded from evaluation in order to preserve strict class balance across the training, validation, and test sets, thereby ensuring consistent and fair performance comparison between models. This strategy was used to maintain class balance and reduce variation during the early stages of model development and testing. It allows for focused analysis and consistent interpretation at the lead level. Each subject includes 12 leads, and each lead contains 1,000 timesteps, giving a total of 3.6 million timesteps. The following equation illustrates the structure of the data matrix created in this way:

$$\mathcal{D} = \begin{bmatrix} \text{Sample} & \text{Lead I} & \text{Lead II} & \text{Lead III} & \cdots & \text{V6} & \text{Classes} \\ 0 & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & \cdots & [\ldots,\ldots,\ldots] & 0 \\ 1 & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & \cdots & [\ldots,\ldots,\ldots] & 0 \\ 2 & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & \cdots & [\ldots,\ldots,\ldots] & 0 \\ \vdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ 599 & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & \cdots & [\ldots,\ldots,\ldots] & 1 \\ 600 & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & [\ldots,\ldots,\ldots] & \cdots & [\ldots,\ldots,\ldots] & 1 \end{bmatrix} \tag{6.1}$$

where $\mathcal{D}$, represents the data matrix containing 12-lead ECG samples with labels, which will be used in the next steps. The $\mathcal{D}$ comprise a total of 600 samples, resulting in a matrix of size 600×13. Each row in this matrix represents an individual patient, with the ECG signal amplitudes from the 12 leads distributed across 12 columns, each containing 1000 time points. The last column contains the corresponding class labels, where 0 represents normal and 1 represents CAD.

6.3.2 Data pre-processing

As seen in Figure 6.2, data augmentation is first applied to $\mathcal{D}$ to enhance generalisation performance, improve model robustness, and increase prediction reliability. In this work,

we propose an augmentation method by introducing zero-mean white Gaussian noise $\mathcal{N}(0, \sigma^2)$ to ECG signals. The augmented signal is defined as:

$$\tilde{x}_l = x_l + n_l, \quad n_l \sim \mathcal{N}(0, \lambda \cdot \sigma_{x_l}^2), \quad \lambda = 0.06, \tag{6.2}$$

where $\tilde{x}_l$ represents the augmented ECG signal for lead l, which white Gaussian noise n_l is added to the original ECG signal x_l . The noise n_l is sampled from a normal distribution with a mean of zero and a variance that is proportional to the variance of the original signal, expressed as $\mathcal{N}(0, \lambda \cdot \sigma_{x_l}^2)$. $\sigma_{x_l}^2$ indicates the variance of the original ECG lead x_l, ensuring that the noise level is appropriately scaled to the variability of the input ECG signal. The noise scaling factor λ is set at 0.06, which means the standard deviation of the noise is 6% of the standard deviation of the original ECG signal. $l \in \{0, ...12\}$ refers to the standard 12-lead ECG. This choice keeps the added noise within a controlled range, maintaining the accuracy of the original ECG signals while improving the model's ability to handle variations and small disturbances such as differences in electrode placement, patient movement and electrical interference.

Following this stage, the ECG signals undergo standard normalisation (Figure 6.2). Standardisation is applied to rescale the ECG signals, ensuring a mean of zero and a standard deviation of one. This transformation normalises the data distribution, making it more suitable for the proposed model by eliminating scale-related variations. It can be expressed mathematically as:

$$z_l = \frac{\tilde{x_l} - \mu_l}{\sigma_l} \tag{6.3}$$

where $\tilde{x_l}$ represents the ECG data after noise addition in the l^{th} lead, μ_l is the mean of the ECG signal in that lead, and σ_l represents the standard deviation of $\tilde{x_l}$ and is used for standardisation.

Standardising ECG signals preserves lead-specific diagnostic information, prevents bias from amplitude differences, and enhances model generalisation by improving robustness to inter-patient variations. Furthermore, adding noise to ECG signals enhances model robustness by simulating real-world signal variability, reducing overfitting, and improving the model's ability to generalise to unseen data.

6.3.3 SHapley Additive exPlanations analysis

SHAP analysis is widely used in quantifying the contribution of individual features to a model's predictions [49, 50]. In this chapter, SHAP was used to assess the relative

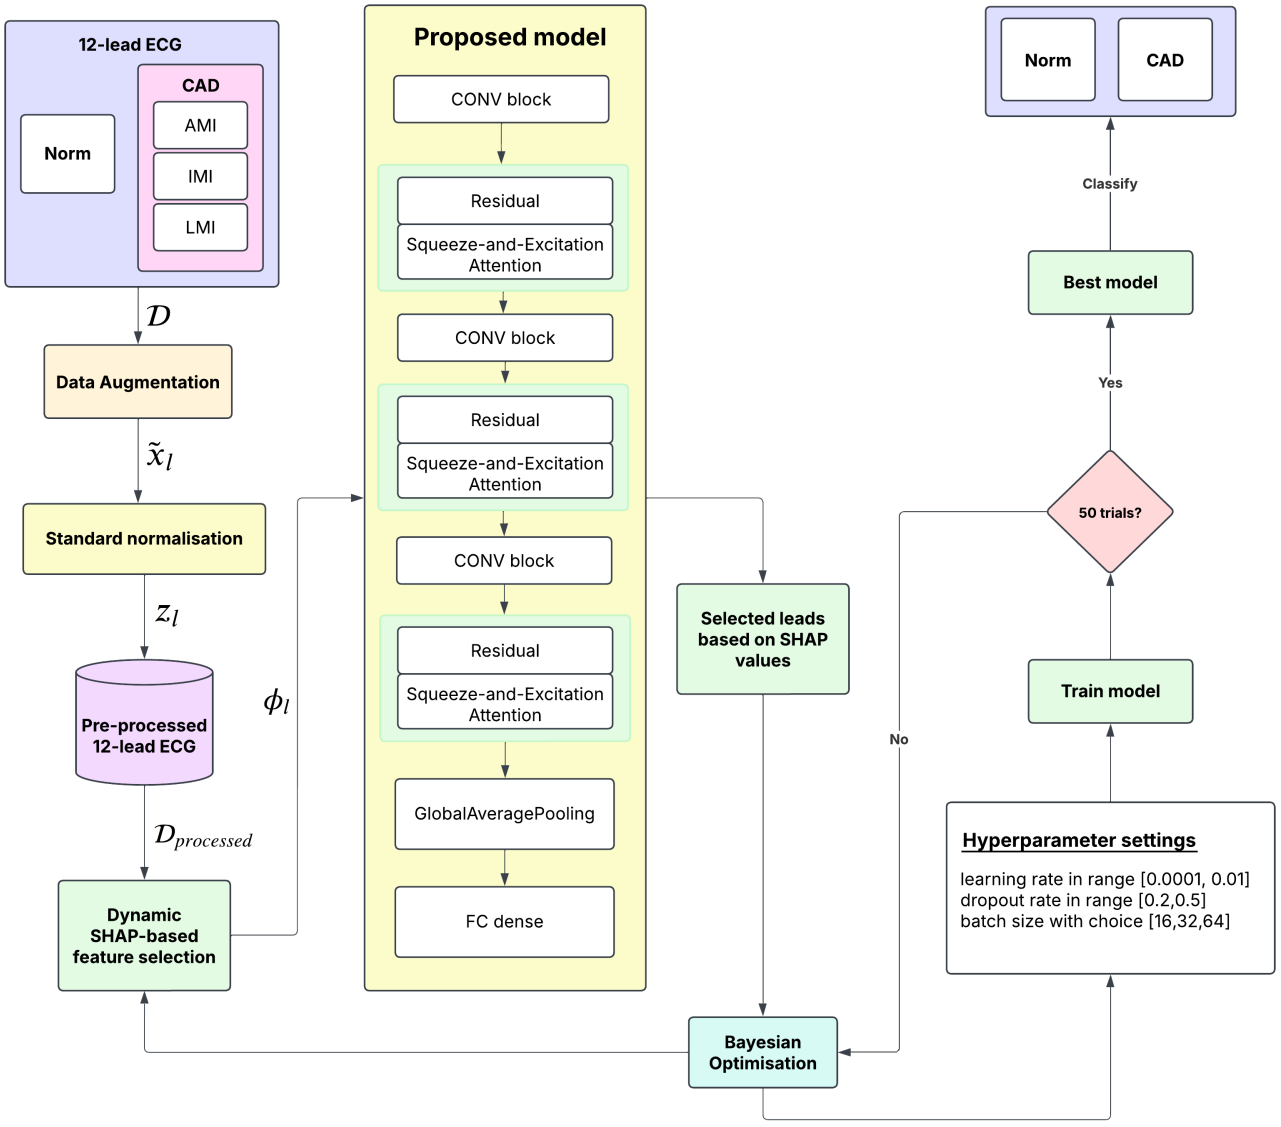

FIGURE 6.2: Flowchart outlining the operational framework of the proposed model

importance of each of the 12 ECG leads in distinguishing between normal and CAD cases. The SHAP Gradient Explainer was employed to compute feature importance scores, enabling the identification of the most influential leads (Figure 6.2). Models were retrained using different numbers of the highest-ranked leads, ranging from one lead to all twelve leads, to evaluate the effect of reduced-lead configurations on classification performance. This analysis was performed to explore potential lead reduction strategies for resource-constrained settings, while recognising that lead localisation is clinically important in myocardial infarction detection and that reduced-lead configurations may impact diagnostic performance. SHAP values can be mathematically computed as:

$$\phi_l = \sum_{S \subseteq L \setminus \{l\}} \frac{|S|!(|L| - |S| - 1)!}{|L|!} [f(S \cup \{l\}) - f(S)] \tag{6.4}$$

where L represents the total number of ECG leads, and l denotes a specific lead under evaluation for its importance. S is a subset of ECG leads, while $S \subseteq L \setminus \{l\}$ denotes all possible subsets that exclude lead l. The function $f(S)$ quantifies the performance of the model when using only the subset S, while $f(S \cup \{l\})$ measures the performance when the lead l is added to this subset. $\frac{|S|!(|L|-|S|-1)!}{|L|!}$ acts as a weighting factor, ensuring a fair distribution of the contribution of each lead in all possible S [228].

6.3.4 Lightweight deep learning architecture

We propose a novel, lightweight deep learning architecture for CAD classification that integrates three key computational building blocks: Convolutional (CONV), Residual, and SE blocks. The CONV block is a hierarchical feature extractor, capturing critical morphological and temporal patterns in ECG waveforms. The residual block is introduced to enhance gradient propagation and facilitate deeper feature abstraction while mitigating vanishing gradient issues. The SE block incorporates an adaptive attention mechanism that dynamically adjusts lead-wise feature significance, ensuring more effective feature representation.

6.3.4.1 Convolutional (CONV) block

In Figure 6.3, let $X \in \mathbb{R}^{T \times 12}$ denote the input ECG matrix, where T represents the total number of time steps and 12 corresponds to the number of ECG leads. X is first processed by a one-dimensional convolutional layer, which applies trainable filters to extract localised ECG patterns over time. The convolutional operation preserves the spatial arrangement of the leads, allowing the model to capture patterns within a

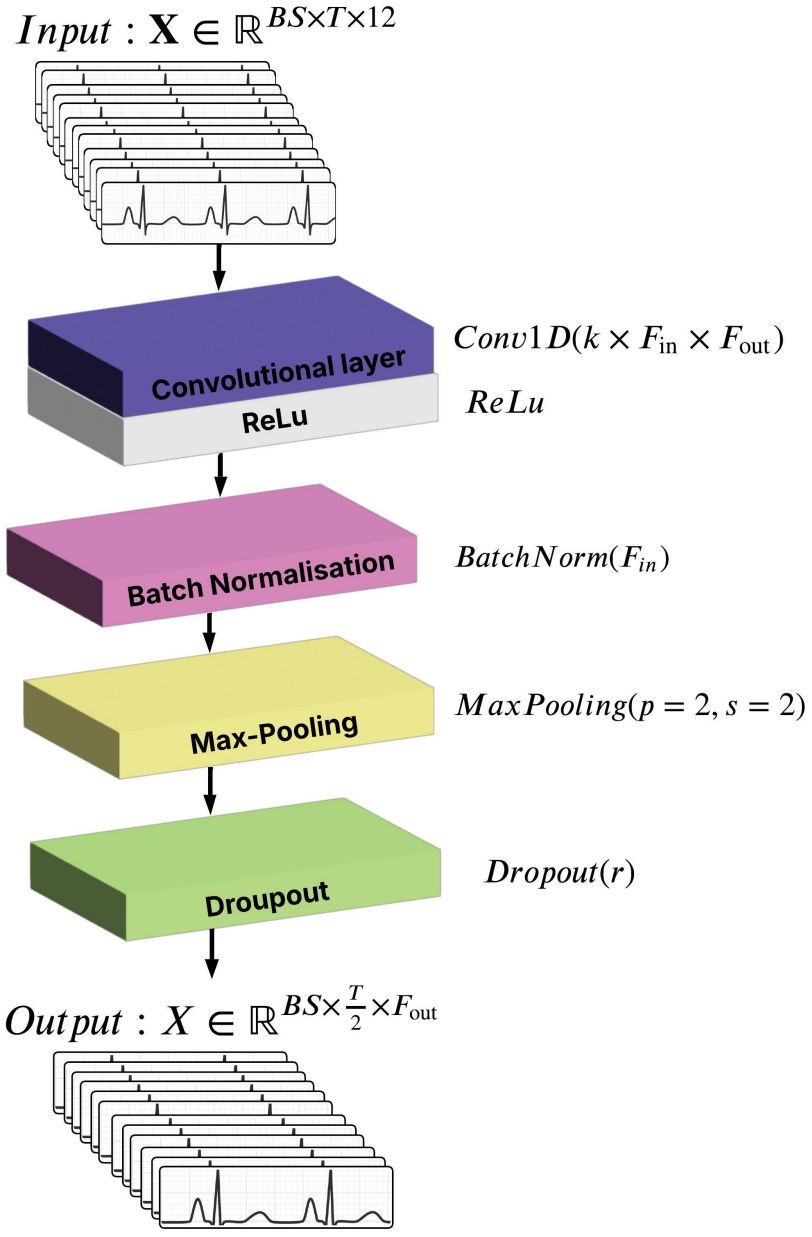

FIGURE 6.3: Convolutional (CONV) block

single ECG lead, learn how ECG signals change over time, and identify relationships between different ECG leads. As the model progresses through deeper network layers, it learns lead-to-lead correlations, enhancing its ability to recognise complex patterns. The convolutional transformation at time step i, for output filter $f \in \{1, \ldots, F_{\text{out}}\}$, is defined as:

$$y_i^{(f)} = \sum_{j=0}^{k-1} \sum_{l=1}^{F_{\text{in}}} x_{i+j,l} \cdot w_{j,l}^{(f)} + b^{(f)} \tag{6.5}$$

where $x_{i+j,l}$ denotes the input value at time step $i+j$ for lead l, $w_{j,l}^{(f)}$ is the learnable weight at offset j and lead l for the f^{th} filter, and $b^{(f)} \in \mathbb{R}$ is the associated bias term.

Batch normalisation is then applied to stabilise feature activations, addressing the issue of fluctuating activation values during training and improving overall training efficiency. For each lead l and timesteps T, batch normalisation is applied as:

$$y_{i,l} = \gamma \frac{x_{i,l} - \mu_{BS}}{\sqrt{\sigma_{BS}^2 + \epsilon}} + \beta \tag{6.6}$$

where $x_{i,l}$ is the activation of lead l at the i^{th} time step. This activation represents a learned feature extracted from the input matrix X after applying convolutional transformations. These features capture both morphological characteristics (e.g., QRS amplitude, ST-segment variations) and temporal dependencies within the ECG signal. μ_{BS} and σ_{BS}^2 are the mean and variance, respectively, computed over the batch. ϵ is a small constant added to prevent numerical instability when dividing by the standard deviation. The normalised activation is then scaled and shifted using the learnable parameters γ and β, where γ adjusts the variance, and β shifts the mean. This process ensures that activations remain within a consistent range, reducing internal covariate shift, including slower learning, gradient instability, and difficult optimisation, while improving the model's ability to learn robust features across different ECG leads.

Following batch normalisation, a max-pooling layer with a stride and pool size of 2 is applied to downsample the temporal resolution by a factor of two. This operation reduces computational complexity while retaining the most important features, improving the model's capacity to detect key waveform characteristics such as QRS complexes and ST-segment deviations. To further prevent overfitting and enhance the model's generalisation, a dropout layer is applied at the end of the block. The dropout probability is optimised through Bayesian hyperparameter tuning, effectively reducing overfitting and promoting robust feature learning. The output of the CONV block has the dimensions

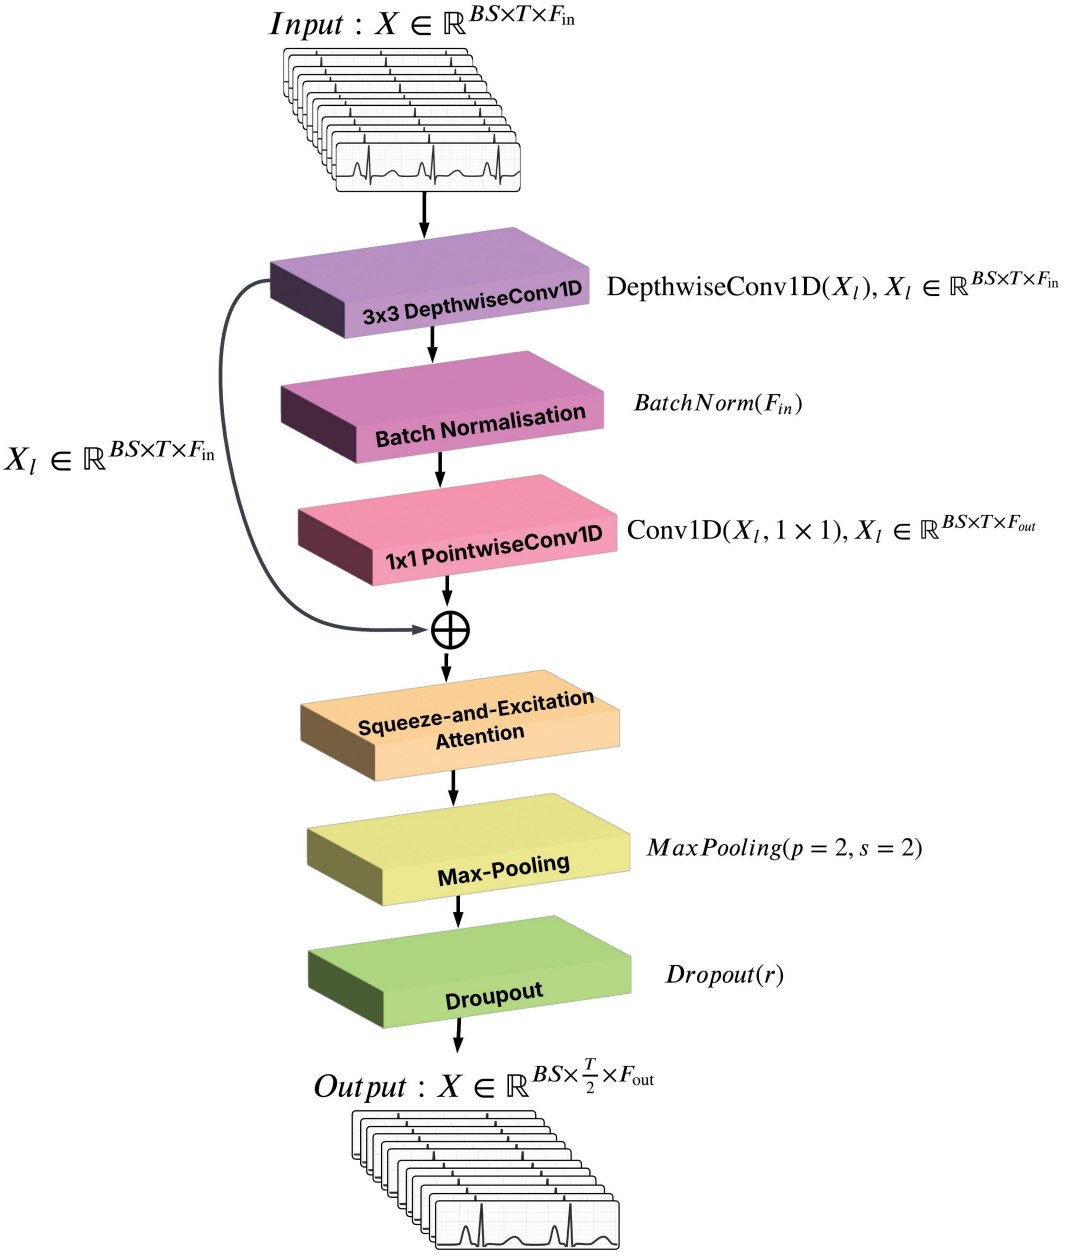

FIGURE 6.4: Residual block

$X \in \mathbb{R}^{\frac{T}{2} \times F_{\text{out}}}$, where the temporal resolution is reduced by half due to a max-pooling operation. This downsampling preserves the most important local patterns whilst reducing computational complexity.

6.3.4.2 Residual block

To extract deeper and more discriminative features, each residual block employs a 3×3 depthwise separable convolution, followed by batch normalisation and 1×1 a pointwise

convolution. A skip connection is employed to directly add the input ECG feature map to the output of the non-linear transformation, enabling the original signal characteristics to bypass intermediate operations and be integrated with the newly extracted features. The block further includes the SE block to adaptively reweight each lead, and concludes with max-pooling and dropout layers for temporal downsampling and regularisation, as shown in Figure 6.4.

Each residual block is formulated to learn a residual mapping, defined as:

$$F(x) = H(x) - x \tag{6.7}$$

where $x \in \mathbb{R}^{BS \times T \times F_{in}}$ denotes the input ECG feature map, with BS representing the batch size, T the number of time steps, and F_{in} the number of input feature representations. The function $H(x)$ corresponds to the complete nonlinear transformation applied to the input ECG feature map x, while $F(x)$ denotes the residual function, learned through a depthwise convolution, batch normalisation, and a pointwise convolution.

The output following the 1×1 pointwise convolution can be mathematically expressed as:

$$y = F(x) + x \tag{6.8}$$

By adding the input ECG feature map x to the output of the learned transformation $F(x)$, the skip connection allows the network to focus on learning the residual, instead of having to learn the full transformation all at once. This helps preserve important low-level features, such as QRS complexes and ST-segment variations, while enabling the model to learn more precise patterns relevant to CAD. The skip connection also improves training stability in deeper layers by facilitating more efficient gradient flow during training.

6.3.4.3 Squeeze-and-Excitation (SE) block

The SE block is introduced to enhance lead-wise attention by adjusting the importance of each ECG lead, helping the model focus on the most informative signals and improve representational efficiency. The SE block consists of a global average pooling layer followed by two fully connected or dense layers. A global average pooling is first applied to each feature map of each ECG lead to compute a single average value. This operation summarises the level of activation across the temporal dimension and provides input to

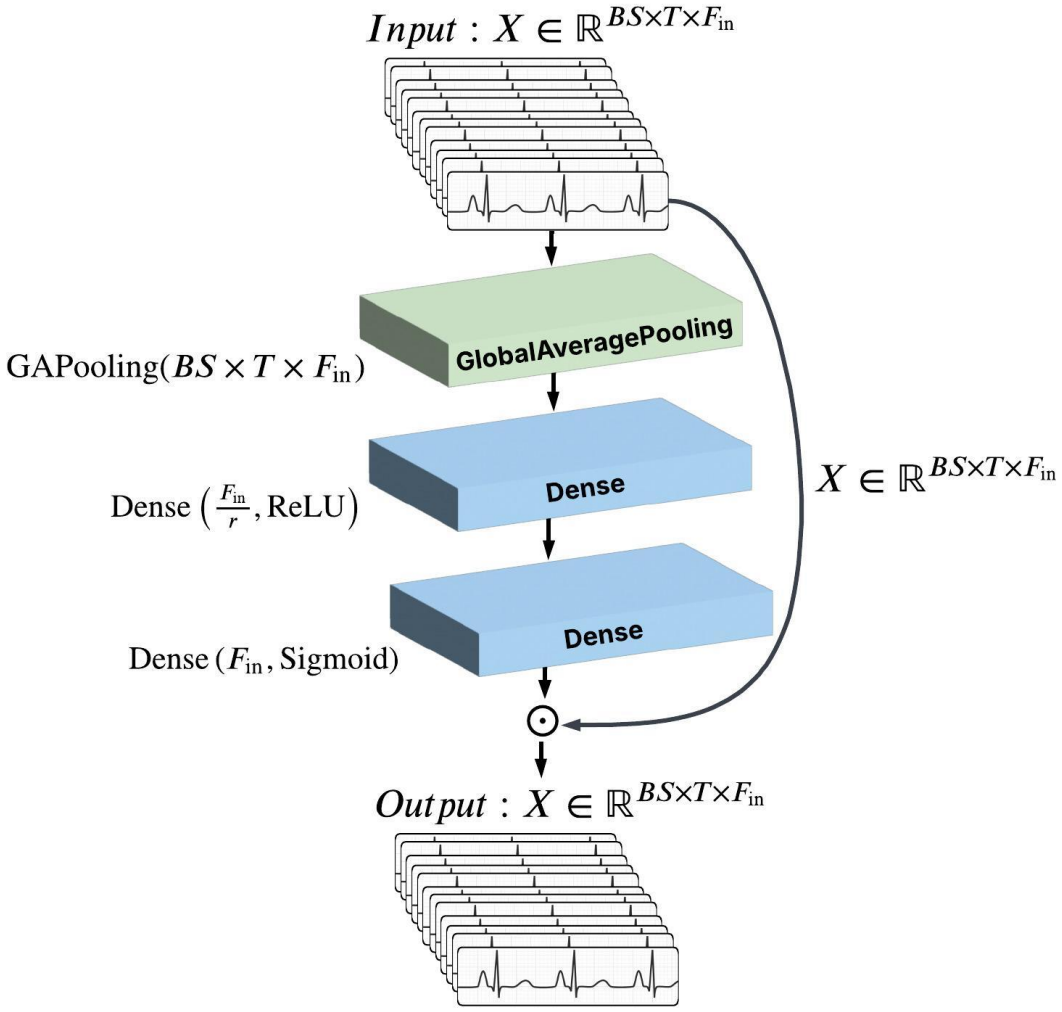

FIGURE 6.5: Squeeze-and-Excitation Attention (SE) block

the subsequent attention mechanism. The global average pooling can be mathematically expressed as:

$$\text{GAPooling}(x)_l = \frac{1}{T}\sum_{i=1}^{T} x_{i,l} \tag{6.9}$$

where T is the number of time steps and $x_{i,l}$ is the activation at time step i for lead l. Each lead is summarised by a single representative value that captures its mean activation across the temporal dimension.

The GAPooling value obtained from each lead l through global average pooling is passed through a two-layer fully connected network to capture inter-lead dependencies. The first dense layer reduces the size by a factor of r and applies a ReLU activation function, allowing for non-linear interactions between leads within a lower-dimensional space. The

second dense layer restores the original size and uses a sigmoid activation to produce a set of attention weights, one for each lead. These weights are then used to adjust the original input by increasing the influence of more important leads and reducing the impact of less relevant ones. A reduction ratio of $r = 8$ is used to compress the feature representation, providing a balance between computational efficiency and the capacity to capture meaningful inter-lead relationships [229].

6.3.5 Hypermeter tuning

Hyperparameter tuning is a critical stage in optimising deep learning models, as it ensures the selection of an optimal combination of parameters that enhances model performance. We employ **Bayesian Optimisation** to efficiently explore the hyperparameter space and identify the most effective configuration for the proposed model. The search space for key hyperparameters is defined in Table 6.1.

TABLE 6.1: Key hyperparameters search space configuration

Hyperparameter	**Search Space**
Learning Rate	range $[0.0001, 0.01]$
Dropout Rate	range $[0.2, 0.5]$
Batch Size	choice $\{16, 32, 64\}$

Bayesian Optimisation is employed due to its ability to efficiently explore high-dimensional search spaces by iteratively refining the search based on prior evaluations. This approach contrasts with traditional grid search, which is computationally expensive and inefficient for large parameter spaces. As depicted in Figure 6.2, the tuning process comprised 50 trials, each executed twice, with distinct hyperparameter configurations explored to ensure comprehensive coverage of the search space. This increases the possibility of identifying an optimal configuration that enhances model performance. Furthermore, each hyperparameter configuration is evaluated twice to mitigate variability introduced by factors such as stochastic weight initialisation, batch sampling, and optimiser behaviour. By averaging performance across multiple runs, this strategy enhances the stability and robustness of the optimisation process, reducing the impact of random fluctuations on model selection.

6.3.6 Hardware Implementation

The proposed lightweight CNN is designed for deployment on low-cost microcontroller units (MCUs) and is implemented in this work on the STM32F469I-DISCO development board. This platform is chosen for its optimal balance of processing power, energy

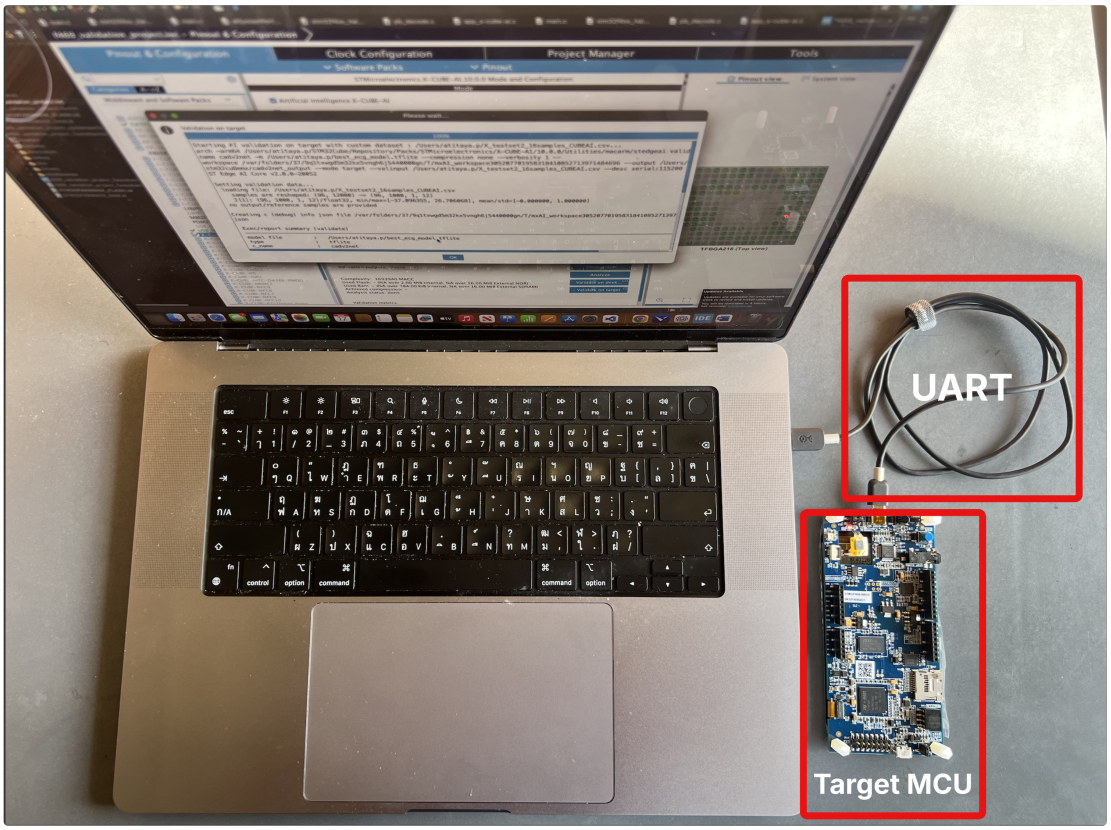

FIGURE 6.6: Implementation on STM32F469I-DISCO MCU

efficiency, and memory resources, making it suitable for deploying our proposed model for ECG CAD classification. This MCU features an ARM Cortex-M4 core, 2 MB of flash memory, and 324 KB of RAM. Its built-in 32-bit floating point unit (FPU) enables efficient execution of multiply-accumulate (MAC) operations required by the model. The STM32F4 platform is considered a suitable option for ECG classification, particularly for evaluating inference performance in resource-constrained environments [230]. Key deployment metrics include total flash usage, total RAM usage, average inference time per subject and energy consumption per subject.

The pre-trained model is run on a PC equipped with an Apple M2 Max processor and 32 GB of unified memory. The implementation is developed using Python 3.9.6. The STM32CubeMX and STM32CubeIDE tools are then used to convert the model and deploy it to the target development board. Figure 6.6 illustrates the experimental setup used to evaluate the total flash usage, RAM usage, average inference time per subject, and energy consumption of the MCU during deployment of the proposed model. The ECG signals are preloaded and transmitted to the MCU via a UART interface for validation. The MCU then executes the classification model and performs real-time inference.

To evaluate the performance of the proposed model, various performance metrics were introduced. *Acc* , Precision, Recall, F1 score, and *AUC* are used to assess performance as shown in Equations (2.5), (2.7), (2.6),(2.9), and 2.11.

6.4 Experimental results

We have conducted a thorough experimental study on the proposed approach with the real ECG data, described in Section 6.3.1. Once all the ECG signals were pre-processed, 70% of the data portion is allocated for training and the remaining for testing. Hyperparameter tuning is performed to determine the optimal model configuration using Bayesian optimisation with details provided in Section 6.3.5. The model is trained using the AdamW optimiser with an initial learning rate of 0.001 and a weight decay of 1×10^{-4}. Weight decay acts as a regularisation technique that helps prevent overfitting [231], thereby improving the model's ability to generalise to unseen ECG signals. Each trial is trained for up to 50 epochs with early stopping based on validation loss, and the search is conducted across up to 50 trials. The best model is then saved and used for prediction on the test set.

6.4.1 Statistical Analysis

Figure 6.7 illustrates a comparison of the distribution of ECG signal values between normal and CAD cases across all 12 standard ECG leads using kernel density estimation (KDE). As shown in Figure 6.7, the density curves for normal and CAD cases overlap significantly, suggesting that their signal distributions are quite similar. This indicates that distinguishing between normal and CAD cases based solely on raw ECG signal distributions may be challenging, as there is no clear separation between the two classes in most leads. From the KDE analysis, V2 and V3 stand out as the most informative leads. These leads exhibit broader distributions for CAD cases compared to normal cases, indicating greater variability in ECG signals affected by CAD. This variability could be due to ischemic changes that impact ventricular depolarisation, making these leads particularly valuable for detecting CAD. V5 also shows noticeable distributional differences, suggesting that it may capture additional diagnostic patterns. Similarly, aVR demonstrates a degree of separation between normal and CAD distributions, making it another potential lead for model training. In contrast, leads such as Lead I, Lead II, Lead III, aVL, aVF, and V6 exhibit highly overlapping distributions between normal and CAD cases. This suggests that these leads may contribute less discriminatory information to the model when used individually. However, this does not mean these leads should be discarded entirely, as the proposed model can extract hidden patterns when

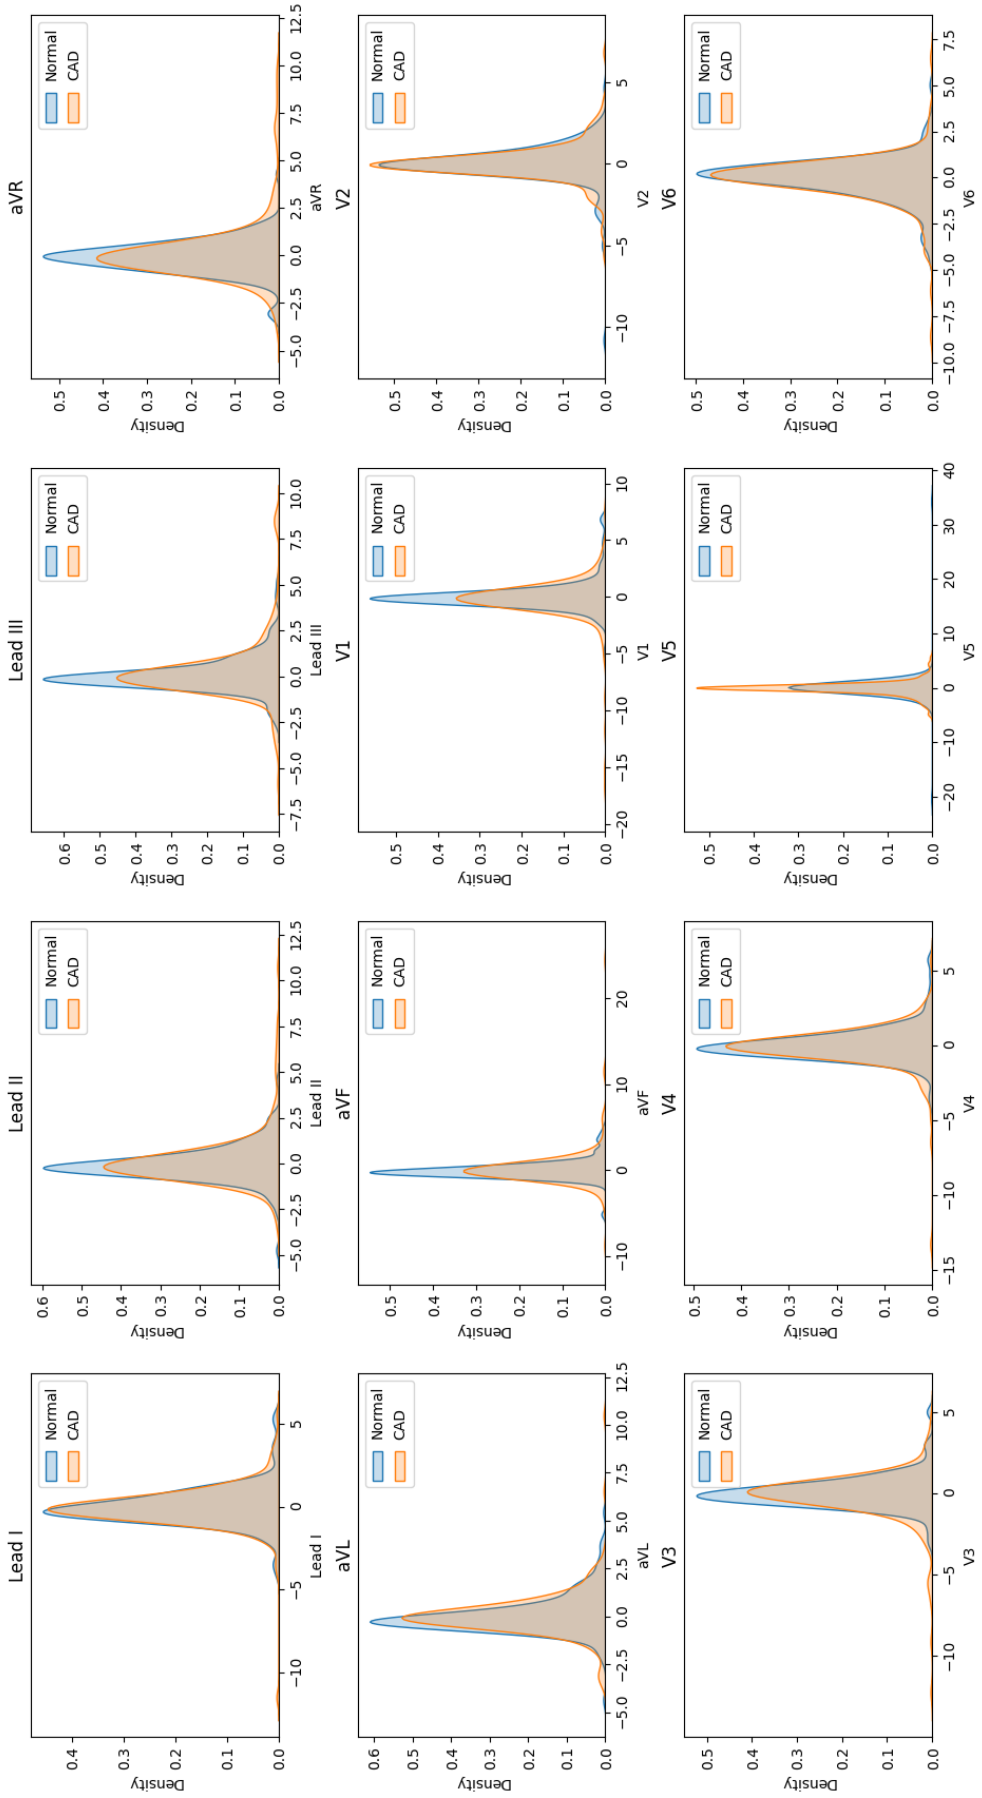

FIGURE 6.7: Kernel density estimation (KDE) of ECG signal distributions for 12-lead ECG, comparing Normal and CAD cases.

multiple leads are analysed together. The model is designed to benefit from combining high-variance leads with other leads to enhance overall robustness. These variations may be attributed to underlying pathological changes in cardiac electrical activity associated with CAD. To study the importance of features in different leads, we applied SHAP-based analysis to determine the contribution of each lead to the model's predictions.

TABLE 6.2: ANOVA Analysis of Feature Importance

Lead	**ANOVA p-value**
Lead I	0.010202
Lead II	0.331867
Lead III	0.055798
aVR	0.043652
aVL	0.008012
aVF	0.610952
V1	0.212404
V2	0.106015
V3	0.000027
V4	0.024050
V5	0.000820
V6	0.011649

Table 6.2 presents a statistical analysis of ECG lead importance using ANOVA (Analysis of Variance) p-values, offering insights into lead relevance for distinguishing between normal and CAD cases. ANOVA measures the extent to which the variance in ECG signal values between the two classes is significant, with lower p-values indicating stronger statistical differences. Conversely, higher p-values indicate that the lead does not exhibit significant variance between the two classes and may contribute less to the classification task.

Based on the ANOVA results, leads V3, V5, and aVL have the lowest p-values of 0.000027, 0.000820, and 0.008012, respectively, suggesting that they exhibit significant differences between normal and CAD groups. This indicates that these leads capture meaningful variations in cardiac electrical activity that could be associated with CAD. V4, Lead I, and aVR also demonstrate relatively low p-values of 0.024050, 0.010202, and 0.043652, respectively, indicating moderate statistical relevance. The significance of leads such as V3 and V5 aligns with clinical knowledge, as precordial leads often capture important ventricular activity changes linked to ischaemia or myocardial infarction, which are key indicators of CAD. On the other hand, leads with high p-values, such as aVF at 0.610952 and Lead II at 0.331867, suggest that these leads may not be as effective in distinguishing between normal and CAD cases. This suggests that their signal distributions are more similar across both groups, making them less informative in a purely statistical sense. V1 and V2 also exhibit relatively higher p-values

of 0.212404 and 0.106015, respectively, implying that while they may still contribute to classification, their individual statistical significance is weaker compared to leads such as V3 and V5. In addition to ANOVA being used to assess lead importance, SHAP analysis is also applied to provide further insights, as feature importance during training is often influenced by complex interactions between ECG leads that ANOVA alone cannot capture.

6.4.2 Model Performance Evaluation

Table 6.3 presents the classification performance of multiple ECG lead configurations using SHAP-based analysis to distinguish between normal and CAD cases. The table compares accuracy, precision, recall, and F1 scores, along with the optimised hyperparameters for each configuration. As shown in Table 6.3, the model's performance results demonstrate a clear relationship between the number of leads and diagnostic accuracy. The 12-lead configuration achieves the highest accuracy of 95.45% and an F1 score of 95.45%, reaffirming its superiority in ECG-based CAD detection. Among the reduced-lead configurations, the 8-lead setup, selected by the model from Lead I, Lead II, aVR, V5, V4, aVL, V1, and V2, achieves an accuracy of 94.95% and an F1 score of 94.94%, indicating its potential as an alternative in scenarios where a full 12-lead ECG is not available. The 5-lead configuration, which includes Lead I, Lead II, aVR, V2, and V4, follows closely, achieving an accuracy of 92.93% and an F1 score of 92.89%, demonstrating the applicability of deliberate lead selection in maintaining diagnostic reliability while simplifying hardware requirements. In contrast, the 1-lead configuration, where Lead II is selected, exhibits the lowest performance, with an accuracy of 80.81% and an F1 score of 80.35%, highlighting the limitations of single-lead ECG in CAD classification. Although it is commonly used in wearable monitoring devices, this configuration may not provide an accurate diagnosis. Notably, the accuracy improves to 89.39% with the 2-lead setup, in which Lead I and Lead II are selected, demonstrating that even small increments in the number of leads significantly enhance classification performance.

Beyond accuracy, the precision and recall metrics reveal further nuances in lead selection. In normal case classification, the 10-lead configuration achieves the highest precision of 97.73%, whereas the 12-lead setup yields a normal precision of 93.27%, suggesting that excluding two leads from a standard 12-lead ECG does not significantly impact overall precision. However, normal recall reaches 100% in both the 2-lead and 9-lead configurations, implying that certain lead combinations may improve detection capability while potentially reducing overall accuracy. For CAD classification, the 12-lead ECG exhibits

TABLE 6.3: Performance metrics for various ECG lead configurations identified through dynamic SHAP-based feature importance. The selected leads were determined by their relative importance in the proposed model, offering insights into their contribution to predictive performance.

Leads Configuration	**Accuracy (%)**	**Normal Precision (%)**	**Normal Recall (%)**	**CAD Precision (%)**	**CAD Recall (%)**	**Macro-Averaged F1 Score (%)**	**Selected Leads**	**Best Hyperparameters**
1 Lead	80.81	94.2	65.66	73.64	95.96	80.35	Lead II	LR: 0.0080, DR: 0.322, BS: 64
2 Leads	89.39	82.5	100.0	100.0	78.79	89.27	Lead I, Lead II	LR: 0.0041, DR: 0.224, BS: 16
3 Leads	91.92	95.6	87.88	88.79	95.96	91.91	Lead I, Lead II, V6	LR: 0.0047, DR: 0.266, BS: 64
4 Leads	91.41	87.27	96.97	96.59	85.86	91.38	Lead I, Lead II, Lead III, V1	LR: 0.0068, DR: 0.2, BS: 64
5 Leads	92.93	99.0	87.0	88.0	99.0	92.89	Lead I, Lead II, aVR, V2, V4	LR: 0.0084, DR: 0.2, BS: 16
6 Leads	92.42	94.68	89.9	90.38	94.95	92.42	Lead I, Lead II, V6, aVL, V3, Lead III	LR: 0.0047, DR: 0.266, BS: 64
7 Leads	88.38	81.15	100.0	100.0	76.77	88.23	Lead I, Lead II, V3, V5, V1, Lead III, aVL	LR: 0.0084, DR: 0.2, BS: 16
8 Leads	94.95	91.59	98.99	98.9	90.91	94.94	Lead II, Lead I, aVR, V5, V4, aVL, V1, V2	LR: 0.0079, DR: 0.229, BS: 16
9 Leads	89.39	82.5	100.0	100.0	78.79	89.27	Lead I, Lead II, V4, aVR, V2, Lead III, aVF, aVL, V1	LR: 0.0050, DR: 0.368, BS: 16
10 Leads	92.42	97.73	86.87	88.18	97.98	92.40	Lead I, Lead II, V5, V6, aVF, V3, aVR, Lead III, V1, V2	LR: 0.01, DR: 0.2, BS: 16
11 Leads	90.91	86.49	96.97	96.55	84.85	90.88	Lead I, Lead II, V5, aVF, V3, V4, V1, aVL, Lead III, aVR, V6	LR: 0.01, DR: 0.2552, BS: 16
12 Leads	**95.45**	**93.27**	**97.98**	**97.87**	**92.93**	**95.45**	**All standard 12 ECG leads**	**LR: 0.0047, DR: 0.266, BS: 64**

the highest precision of 97.87% and recall of 92.93%, reaffirming its robustness in diagnosing patients with CAD. Notably, the 8-lead configuration achieves a CAD precision of 98.9% and a recall of 90.91%, making it a practical option for reducing the number of monitored leads while maintaining strong performance. Conversely, the single-lead configuration exhibits the lowest normal recall of 65.66% and CAD precision of 73.64%, further emphasising its diagnostic limitations.

Model optimisation through hyperparameter tuning reveals that the 12-lead ECG achieves the highest performance with a learning rate of 0.0047, a dropout rate of 0.266, and a batch size of 64. Additionally, similar hyperparameter configurations applied to the 3-lead, 6-lead, and 8-lead setups contribute to stable classification performance. Differences in dropout rates and batch sizes across lead configurations have a minor impact on accuracy, further emphasizing the need for hyperparameter optimisation to ensure reliable model performance. These findings highlight the balance between classification performance and the number of leads used, offering insights into the practicality of reduced-lead ECG monitoring for clinical and telemedicine applications. While the 12-lead ECG remains the standard for CAD classification, the 5-lead and 8-lead configurations stand out as the most effective alternatives, maintaining diagnostic accuracy while improving ease of use.

These results explain the difference between the single-lead performance of the 12-lead proposed model proposed in this chapter and the strong performance of the CADNet model presented in Chapter 5. Although both approaches address binary CAD classification, the models are optimised under different assumptions. CADNet model was specifically designed and trained for single-lead ECG input, allowing it to learn lead-specific morphological features effectively. In contrast, the proposed model is trained to use information from multiple ECG leads, capturing differences in ECG waveform morphology across leads. When this multi-lead architecture is restricted to a single lead, a substantial amount of diagnostically relevant information is lost, which may lead to reduced performance.

As shown in Figure 6.8, aVR, Lead II, and V2 emerge as the most influential leads in the model's predictions. aVR ranks highest in SHAP importance, indicating its crucial role in distinguishing between normal and CAD cases. This finding aligns with the fact that aVR provides unique electrical information about the right atrium and the basal regions of the ventricles, which may be valuable for detecting CAD-related abnormalities. Lead II also ranks highly in SHAP importance, despite having a relatively high ANOVA p-value. This suggests that, although Lead II does not exhibit strong statistical differences in isolation, it interacts with other leads in a manner that enhances the model's predictive

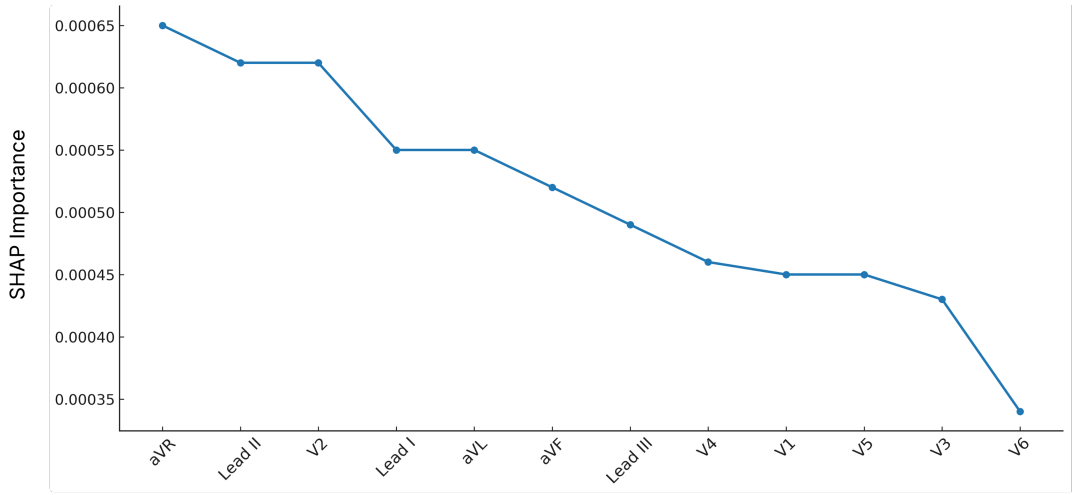

FIGURE 6.8: SHAP-based feature importance of 12-lead ECG signals

performance. Similarly, V2 demonstrates strong SHAP importance, reinforcing its role in capturing ventricular electrical activity changes associated with CAD.

On the other hand, the SHAP results indicate that Lead V6 has the lowest feature importance, suggesting that it contributes minimally to the model's classification decisions. This implies that V6 may not capture sufficient discriminatory information between normal and CAD cases compared to other leads. Notably, some leads with low ANOVA p-values, such as V3 and V5, do not rank as highly in SHAP importance. Hence, this suggests that while these leads demonstrate strong statistical separation between normal and CAD groups, the proposed model does not rely on them as significantly when making predictions, possibly due to their limited contribution to complex, non-linear interactions, where a lead's importance depends on the presence of other leads within the model during training, or due to their redundancy when replaced by other, more influential leads.

Table 6.4 presents a comparative analysis of various traditional machine learning, deep learning, and lightweight network architectures. It also examines differences in model size and inference time across these architectures, offering insights into their computational efficiency and suitability for resource-constrained environments. The evaluation considers trainable parameters, file size, and runtime, alongside performance metrics such as accuracy, AUC, and F1 score. Based on the results, traditional machine learning models such as SVM, Gaussian Naïve Bayes, K-Nearest Neighbours, K-Means, and Logistic Regression are computationally lightweight, with no or minimal trainable parameters. Among these models, Logistic Regression performs the best, achieving an accuracy of 86.87%, an AUC of 93.09%, and an F1 score of 86.86%. In contrast, K-Means demonstrates the lowest performance, with an accuracy of only 62.63% and an AUC of 62.63%, highlighting its limited ability to differentiate between classes. The

TABLE 6.4: Comparative analysis of traditional ML-based, DL-based and lightweight networks.

Model Architecture	**Trainable params**	**Size (MB)**	**Acc (%)**	**AUC (%)**	**F1 score (%)**	**Inference time per subject (s)**
SVM	0	0.1028	85.86	92.91	85.81	0.000083
Gaussian Naive Bayes	0	0.0021	69.70	86.48	68.41	0.000039
KNN	0	0.1510	84.34	92.26	84.05	0.002321
K-Means	0	0.0029	62.63	62.63	59.86	0.000125
Logistic Regression	49	0.0011	86.87	93.09	86.86	0.000028
LSTM[113]	55,502	0.6688	81.31	88.24	81.29	0.048771
CNN-LSTM [113]	64,686	0.7773	84.85	90.04	84.75	0.034116
Baseline 1D-CNN [31]	8,816,258	100.9503	93.94	98.41	93.93	0.051712
SqueezeNet[178]	361,762	1.3972	50.00	50.00	33.33	0.001610
EfficientNetB0[179]	7,004,322	26.9931	81.82	90.28	81.75	0.038880
1D-MobileNetV1[180]	3,168,610	12.1057	50.00	50.00	33.33	0.026323
ShuffleNetV1[123]	754,370	2.9935	74.24	81.10	74.24	0.003117
Proposed model	13,584	0.0530	95.45	98.96	95.45	0.019118

inference times for these traditional models are extremely fast, with Logistic Regression being the most efficient at just 0.000028 seconds per subject. Although Gaussian Naïve Bayes has the smallest file size of 0.0021 MB, its performance remains moderate, with an accuracy of 69.70% and an AUC of 86.48%.

In contrast, deep learning models, including LSTM [113], CNN-LSTM [113], and a baseline 1D-CNN [31], exhibit significantly higher model complexity, with a much larger number of trainable parameters. The CNN-LSTM model achieves slightly better accuracy of 84.85% than LSTM, which achieves 81.31%, but at the cost of increased inference time. Their AUC values, 90.04% and 88.24%, respectively, indicate strong classification performance. The baseline 1D-CNN, with over 8.8 million trainable parameters and a file size of 100.95 MB, delivers the highest performance among these models, achieving 93.94% accuracy and an AUC of 98.41%. However, with an inference time of 0.051712 seconds per subject, the baseline 1D-CNN is capable of near real-time processing in most clinical applications, although such latency may still be a consideration for highly resource-constrained or ultra–low-latency embedded implementations. In addition, lightweight network architectures such as SqueezeNet [178], EfficientNetB0 [179], 1D-MobileNetV1 [180], and ShuffleNetV1 [123] are optimised for efficiency, but their performance varies significantly. SqueezeNet, despite having a relatively small number of parameters at 361,762, demonstrates poor performance, achieving only 50% accuracy and an AUCof 50%, indicating its limited capability in distinguishing between classes. EfficientNetB0 performs better, reaching 81.82% accuracy with an AUC of 90.28%, though it remains inferior to traditional machine learning models. 1D-MobileNetV1, with an accuracy of 50% and an AUC of 50%, fails to provide competitive results. Among these architectures, ShuffleNetV1 offers a more balanced trade-off between efficiency and accuracy, achieving 74.24% accuracy with an AUC of 81.10% and an inference time of 0.003117 seconds per subject. This makes it a more practical choice than larger deep learning models in applications where computational efficiency is a priority while still providing moderate classification performance.

Hence, these results highlight that the proposed model significantly outperforms all other architectures, offering an optimal balance between accuracy, computational efficiency, and model size. Despite having only 13,584 trainable parameters and a compact file size of 0.0530 MB, it achieves the highest accuracy of 95.45%, an AUC of 98.96%, and an F1 score of 95.45%. These findings underscore its superior classification performance compared to other models. Furthermore, the proposed model maintains an inference time of 0.019118 seconds per subject, making it well-suited for real-time applications while delivering exceptional predictive capabilities. Additionally, when compared to other lightweight architectures such as SqueezeNet, EfficientNetB0, 1D-MobileNetV1, and ShuffleNetV1, the proposed model exhibits a significantly smaller model size while

achieving superior performance. While certain architectures, such as EfficientNetB0, attain reasonable accuracy levels, their inference times remain higher, rendering them less effective for real-time applications. The ability of the proposed model to balance high accuracy with computational efficiency makes it particularly well-suited for deployment in resource-constrained environments, where processing power and memory availability are limited, yet predictive performance remains a critical requirement. Furthermore, a key advantage of the proposed model is its ability to maintain a high *AUC* of 98.96%, demonstrating its strong capacity to differentiate between classes.

In order to study the performance of the proposed model across different CAD subclasses, as well as normal cases, Table 6.5 is depicted. First, 100 normal and 100 AMI cases (a subclass of CAD) were selected to evaluate the model's performance. The model achieved an accuracy of 99.19% on the training set and 93.94% on the test set, demonstrating reliable performance in distinguishing between normal cases and AMI. The high precision of 93.94%, combined with a perfect recall of 93.94% and an F1 score of 93.94%, suggests that the model not only correctly identifies AMI cases but also minimises false negatives. Additionally, the AUC of 96.51% further supports the model's strong capability to differentiate AMI cases from normal cases. Secondly, 100 normal cases and 100 IMI cases were selected to investigate how the model performs. The model maintained a high training accuracy of 99.04% and a test accuracy of 93.94%, suggesting consistent performance across different CAD subclasses. The precision of 94.10%, recall of 93.94%, and F1 score of 93.93% indicate that the model effectively distinguished IMI cases, with slightly improved precision compared to AMI classification. The AUC of 97.61% reflects a higher discriminatory ability, further highlighting the model's effectiveness in classifying IMI cases. Lastly, the model exhibited its strongest performance, achieving a training accuracy of 98.08% and a test accuracy of 96.97%, the highest among all CAD subclasses. The precision of 97.14%, recall of 96.97%, and F1 score of 96.97% demonstrate the model's ability to detect LMI cases with high reliability. The AUC of 99.63% indicates that the model has an exceptional capability to distinguish LMI cases from normal cases, suggesting a highly effective classification process.

The aforementioned analysis confirms that, in addition to binary classification, the model's performance was evaluated in a multi-class classification setting, where different CAD subclasses were treated as distinct categories. Compared to the binary classification results in Table 6.4, the model's accuracy decreased to 72.50%, suggesting challenges in classify between certain CAD subclasses. The model obtained a precision of 71.33%, a recall of 72.50%, and an F1 score of 72.74%. Despite the drop in an overall accuracy, the AUC remained high at 89.74%, indicating a strong distinguishing ability. However, some CAD subclasses were more difficult to classify correctly, as shown in Figure 6.9(a),

TABLE 6.5: Comparative performance of the 12-lead ECG lightweight model developed in this study across different datasets.

Classication	**Data used**	**Train Acc (%)**	**Test Acc (%)**	**Precision* (%)**	**Recall* (%)**	**F1 score* (%)**	**AUC(%)**
Binary	Norm:100 AMI:100	99.19	93.94	93.94	93.94	93.94	96.51
Binary	Norm:100 IMI:100	99.04	93.94	94.10	93.94	93.93	97.61
Binary	Norm:100 LMI:100	98.08	96.97	97.14	96.97	96.97	99.63
Multi-class	Norm:300 AMI :100 LMI:100 IMI:100	83.32	72.50	71.33	72.50	72.74	89.74

* Precision, recall, and F1-scores are reported as macro averages.

where the confusion matrix highlights misclassifications and potential overlap between CAD subclasses.

Furthermore, t-SNE analysis was conducted to visualise high-dimensional data and assess how well the model distinguishes between different CAD subclasses. In addition, a confusion matrix was used to evaluate the model's performance in a multiclass classification setting, where CAD subclasses were considered separately instead of using a binary classification approach as shown in Figure 6.9. In Figure 6.9(a), the model correctly identified 25 out of 30 normal cases, with one normal case misclassified as AMI, three as LMI, and one as IMI. This high accuracy suggests that the model effectively differentiates individuals with normal sinus rhythm from those with CAD. Among the 30 AMI cases, only 17 were correctly classified, while 3 were misclassified as normal, 3 as LMI, and 7 as IMI. Furthermore, the model demonstrated strong performance in classifying LMI cases, correctly identifying 25 out of 30 instances. However, five cases were misclassified as IMI, suggesting that while LMI features are well captured by the model, some instances share similarities with IMI. Lastly, the model's performance in distinguishing IMI from normal cases presented the greatest challenge, with 20 correct classifications out of 30. A significant number of 10 cases were misclassified as LMI. This highlights a high degree of feature similarity between IMI and other CAD subclasses, particularly LMI, which may explain the frequent misclassification. Based on these results, further refinement in feature extraction or the incorporation of more complex feature extraction layers may be necessary, as this could improve the model's ability to classify IMI cases more accurately.

To further understand the model's performance, the t-SNE technique was applied to visualise how well the model distinguishes between different CAD subclasses at two stages. Figure 6.9(b) illustrates the feature distributions in the first layer of the proposed model. As shown, the different CAD subclasses exhibit a significant degree of overlap, suggesting that their distinguishing features have not yet been fully captured. This indicates that the model has not yet learned meaningful features, making classification more challenging in the initial layers. On the other hand, the feature distribution in the final layer shows that the data points are more distinctly clustered. This suggests that the model has successfully extracted discriminative features for each CAD subclass, leading to improved class separation. Additionally, Normal and LMI cases appear well-separated, reflecting the model's strong ability to differentiate these classes. However, AMI and IMI still exhibit some overlap, which is consistent with the confusion matrix results.

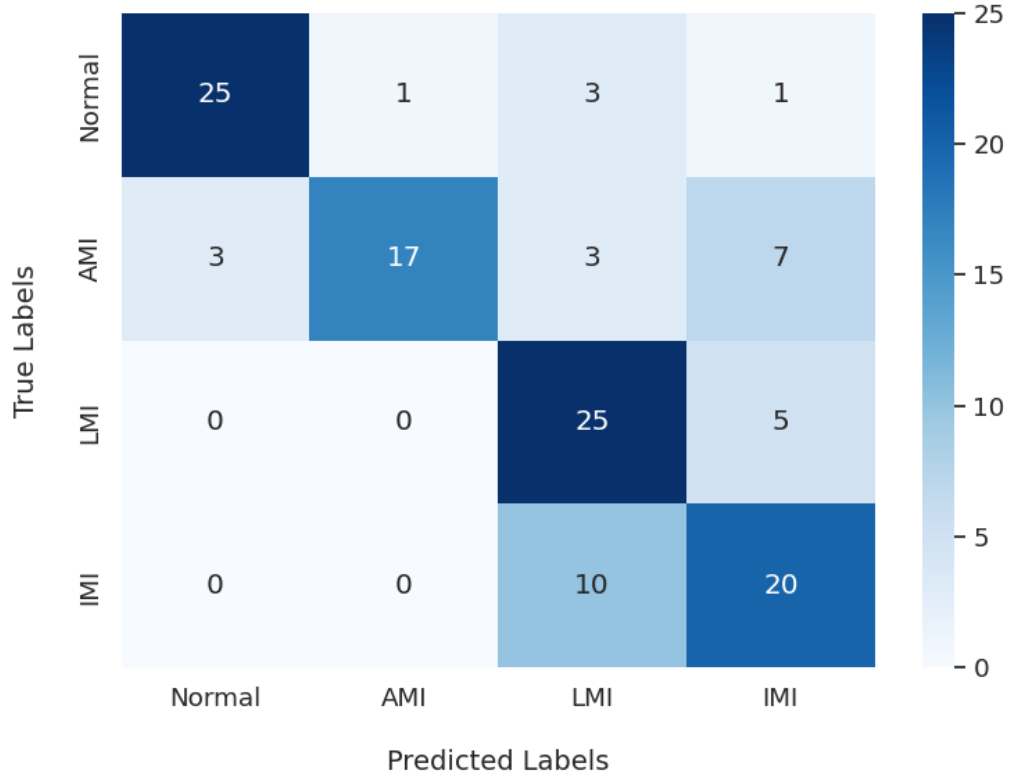

(a) Confusion matrix of the test data for the multiclass classification of CAD subclasses and normal sinus rhythm.

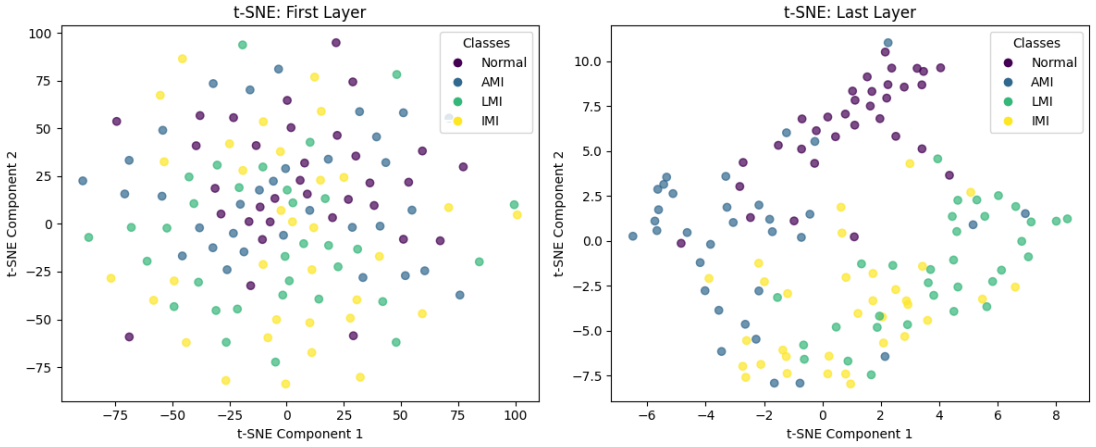

(b) t-SNE visualisation illustrating feature separability between normal sinus rhythm and CAD subclasses.

FIGURE 6.9: Comparison of classification performance and feature separability on the test set for multiclass classification: (a) Confusion matrix and (b) t-SNE visualization of learned feature representations.

6.4.3 Efficiency Analysis for Real-time Processing

Table 6.6 compares the performance of existing lightweight models with the proposed model, all deployed on the STM32F469I-DISCO microcontroller. This evaluation considers the model's computational complexity, memory requirements, inference time, and energy consumption, all of which are critical factors in the development of a real-time ECG CAD diagnostic tool. The proposed model demonstrated clear advantages in computational efficiency. It required only 1,693,940 MAC operations compared to 13,597,156 for 1D-MobileNetV1, representing a reduction of nearly 88%. This reduction in complexity was also reflected in the parameter count, with the proposed model using just 0.0553 MB of parameters, significantly less than the 0.4788 MB required by 1D-MobileNetV1. In comparison, the lightweight CNN [163] required 13,913,016 MAC operations and 0.2663 MB of parameters, which is lower than 1D-MobileNetV1 but still considerably higher than the proposed model. ResNetLite demonstrated improved computational efficiency, with MAC operations totalling 10,170,074. However, this remains greater than that of the proposed model. Regarding memory usage, the proposed model required only 0.0852 MB of flash memory, while 1D-MobileNetV1 needed 0.5261 MB, representing a reduction of approximately six times in storage size. Although the proposed model used slightly more RAM at 0.1117 MB, compared to 0.0671 MB for 1D-MobileNetV1, the increase was minor and remained well within the memory limits of resource-constrained devices. The lightweight CNN model consumed even more RAM at 0.1930 MB and required 0.2899 MB of flash memory, making it less suitable for low-memory environments. ResNetLite used 0.461 MB of flash memory and 0.1065 MB of RAM, both higher than those required by the proposed model, which makes it less ideal for resource-constrained environments. Most importantly, the proposed model showed clear improvements in both inference time and energy efficiency. The average inference time per subject was reduced to 0.1121 seconds, which is less than half of that required by 1D-MobileNetV1. This faster processing directly contributed to reduced energy consumption, decreasing from 16.51 mJ to 7.39 mJ per subject. In contrast, the lightweight CNN model had a much longer inference time of 0.7991 seconds and the highest energy consumption at 52.15 mJ per subject. These results demonstrate that the proposed model is more suitable for deployment in real-time and energy-efficient ECG-based CAD diagnosis operating on resource-constrained devices.

6.5 Conclusion

This chapter introduced a lightweight ECG CAD classification model designed to distinguish between CAD and normal sinus rhythm in real-time. To reduce the computational

requirements while preserving high classification accuracy, several optimisation strategies were implemented. Evaluated using the PTB-XL database, the proposed model achieved an accuracy of 95.45% and an AUC of 98.96%, indicating its strong performance. The model was composed of only 0.0553 MB of parameters and requires 1,693,940 MAC operations, significantly outperforming current state-of-the-art ECG classification models regarding resource efficiency. The model was successfully deployed on a cost-effective MCU, the STM32F469I-DISCO, where it consumed just 7.39 mJ per subject classification, ensuring real-time ECG CAD diagnostic capabilities.

TABLE 6.6: Comparison of the computational and memory efficiency of state-of-the-art existing lightweight CNN models and the proposed model on the STM32F469I-DISCO

Model	**MAC***	**Param# (MB)**	**Flash (MB)**	**RAM (MB)**	**Per Subject**	
					Time (s)	**Energy (mJ)**
1D-MobileNetV1[232]	13,597,156	0.4788	0.5261	0.0671	0.2501	16.51
Lightweight CNN [163]	13,913,016	0.2663	0.2899	0.1930	0.7991	52.15
ResNetLite [233]	10,170,074	0.4419	0.461	0.1065	0.5289	34.91
Proposed Model	1,693,940	0.0553	0.0852	0.1117	0.1121	7.39

* Multiply-Accumulate operations (MAC)

Chapter 7

Multi-Disease Cardiovascular Detection from ECG Signals Using an Attention-Driven Deep Network

7.1 Overview

While the previous chapters focused on developing lightweight models for CAD detection, in real-world clinical settings patients rarely present with a single cardiac condition. Patients commonly present with co-existing heart conditions, such as CAD and AF, and accurate diagnosis therefore requires distinguishing between multiple conditions within a single ECG recording. Building on the demonstration of efficiency and suitability in the earlier chapters, this chapter focuses on multi-disease classification by introducing a novel deep learning architecture for detecting multiple CVDs from ECG signals. The proposed model combines convolutional layers, residual connections, and attention mechanisms in a unified framework to improve both accuracy and computational efficiency. Convolutional layers enable the extraction of local signal features, while residual connections address the vanishing gradient problem, allowing the network to learn deeper and more complex patterns without degradation. Attention mechanisms enhance the model's focus on diagnostically relevant features, improving interpretability and precision. This model effectively captures both local and global features within ECG data, facilitating a comprehensive analysis of intricate cardiac patterns. Moreover, by eliminating the need for conventional feature extraction techniques, which often necessitate

substantial preprocessing in traditional machine learning methods, our approach streamlines the classification process, thereby enhancing both efficiency and accuracy in the detection of multiple CVDs. Our extensive experimental results demonstrate that the proposed model effectively achieves an average classification accuracy of 99.54%, which is superior to existing deep learning-based models, and enables the detection of multiple heart conditions from a single ECG reading, offering valuable insights for clinical decision-making.

7.2 Related works

CNN has demonstrated strong performance in capturing spatial features within ECG signals, making them highly suitable for the detection of CVDs [30, 31]. ResNets, a specific type of CNN, have been introduced to address the limitations of traditional architectures. By incorporating shortcut connections, ResNets enable the construction of deeper networks while mitigating the vanishing gradient problem often encountered in very deep models [30, 32, 147]. LSTM networks, a form of RNN, have also been utilised for the early detection of heart conditions due to their ability to model long-term temporal dependencies within sequential data [149]. Several studies have combined CNN with LSTM to benefit from both spatial feature extraction and temporal sequence learning [6, 136]. In addition, attention mechanisms have been increasingly integrated into deep learning models to improve classification performance by identifying and focusing on the most informative regions of the ECG input [33, 140, 150]. For example, [34] shows that attention can improve diagnostic accuracy by assigning greater weight to ECG leads that contribute more significantly to identifying specific heart conditions.

According to the literature, current models face significant challenges in multi-class disease classification from ECGs due to limitations in data availability, class imbalance, and inadequate feature extraction, which adversely affect their performance across various cardiovascular conditions [79]. These challenges underscore the need for a more effective and practical solution, prompting the development of our proposed model, which integrates advanced DL-based techniques to enhance both accuracy and robustness in the diagnosis of multiple cardiovascular diseases.

In this chapter, we propose a novel model for the classification of multiple CVDs, integrating convolutional layers, residual networks, and attention mechanisms. We propose to apply a data augmentation process to serve two purposes: 1) enriching the dataset to enhance its robustness, generalisability, and applicability to real-world scenarios, etc. and 2) examining the impact of noise power on ECG signal classification to determine

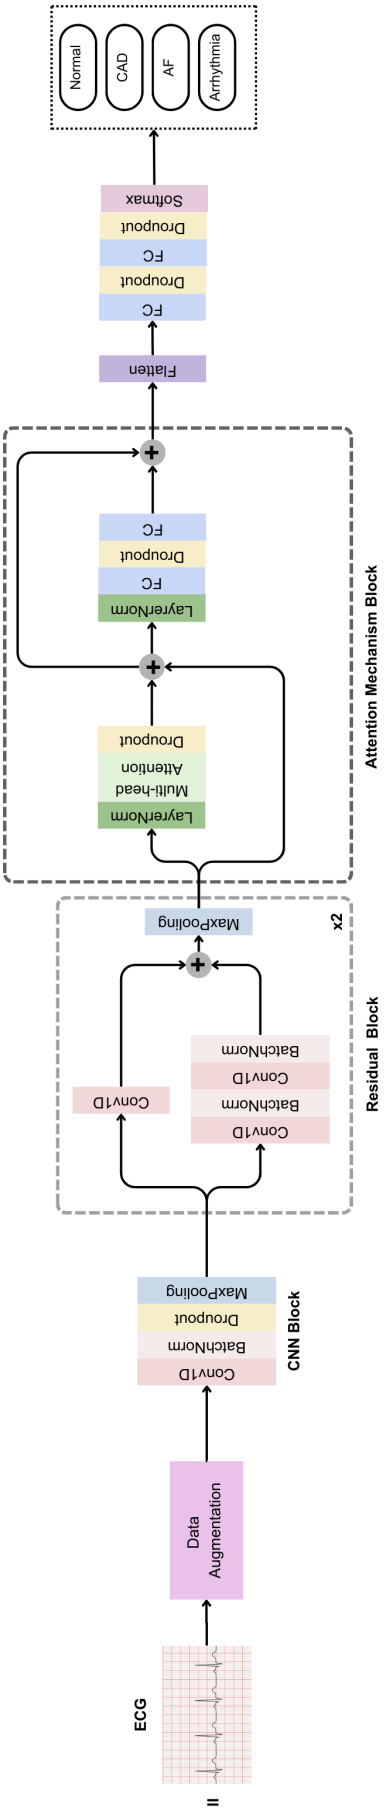

FIGURE 7.1: The proposed model architecture.

the optimal settings for augmentation. Our multi-label classification enables the detection of multiple conditions from a single ECG reading, potentially improving treatment decisions and revealing diseases that may have previously gone undetected. Furthermore, accurate multi-disease detection using only one ECG channel can potentially lead to feasible commercial health monitoring devices, such as smartwatches. The proposed model is designed to address the complexity of ECG signals, which often exhibit characteristics of multiple cardiac conditions simultaneously. This approach allows the model to capture subtle distinctions between various CVDs, thereby improving diagnostic accuracy and aligning more closely with real-world clinical practice. The proposed model can serve as a pre-screening tool, as ECG is non-invasive compared to other medical examination methods in hospitals. Furthermore, this approach can assist doctors in making informed decisions regarding the next steps in the diagnostic process.

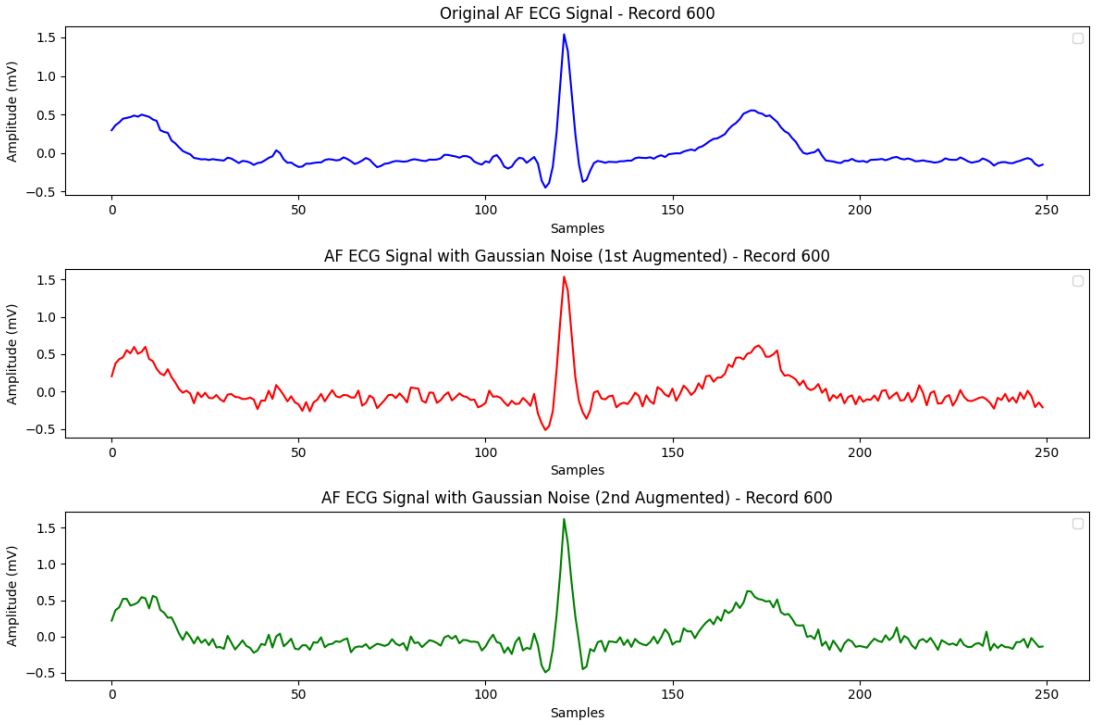

(a) ECG Signal demonstrating AF

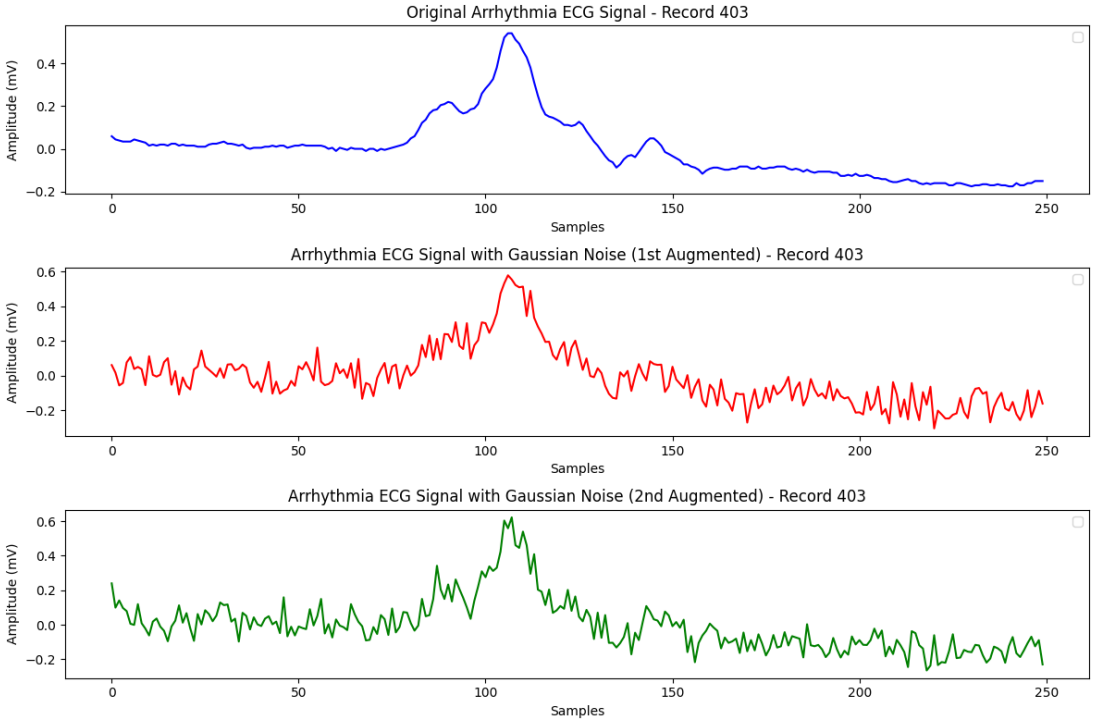

(b) ECG Signal demonstrating arrhythmia

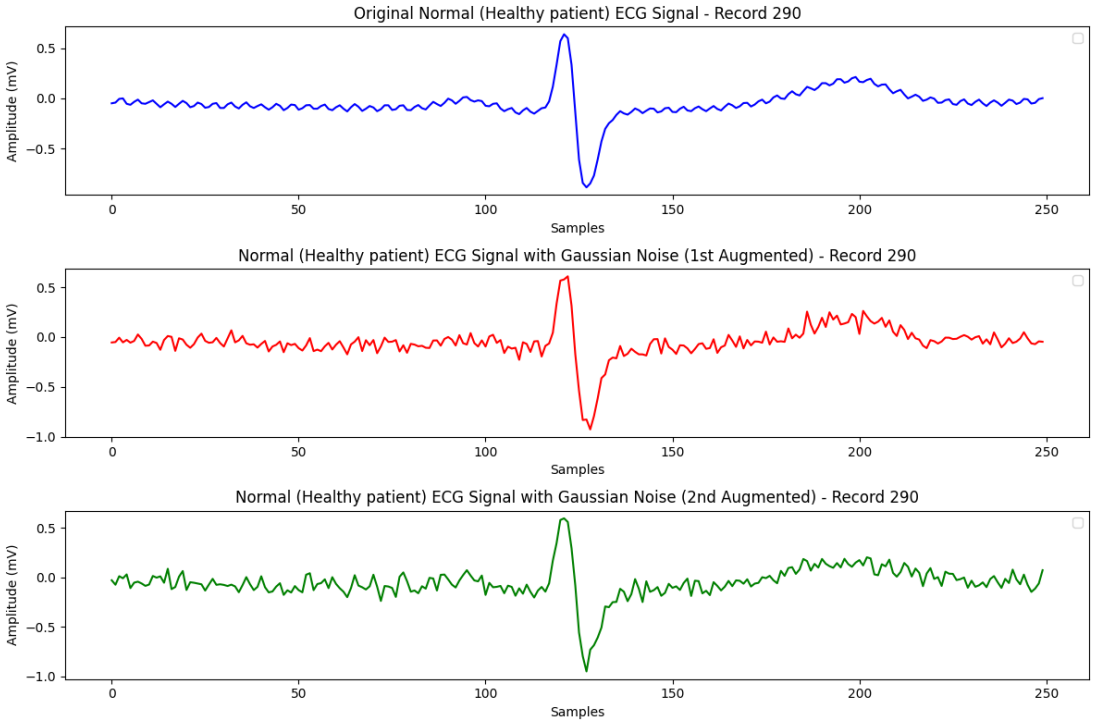

(c) ECG signal demonstrating normal sinus rhythm in a healthy patient

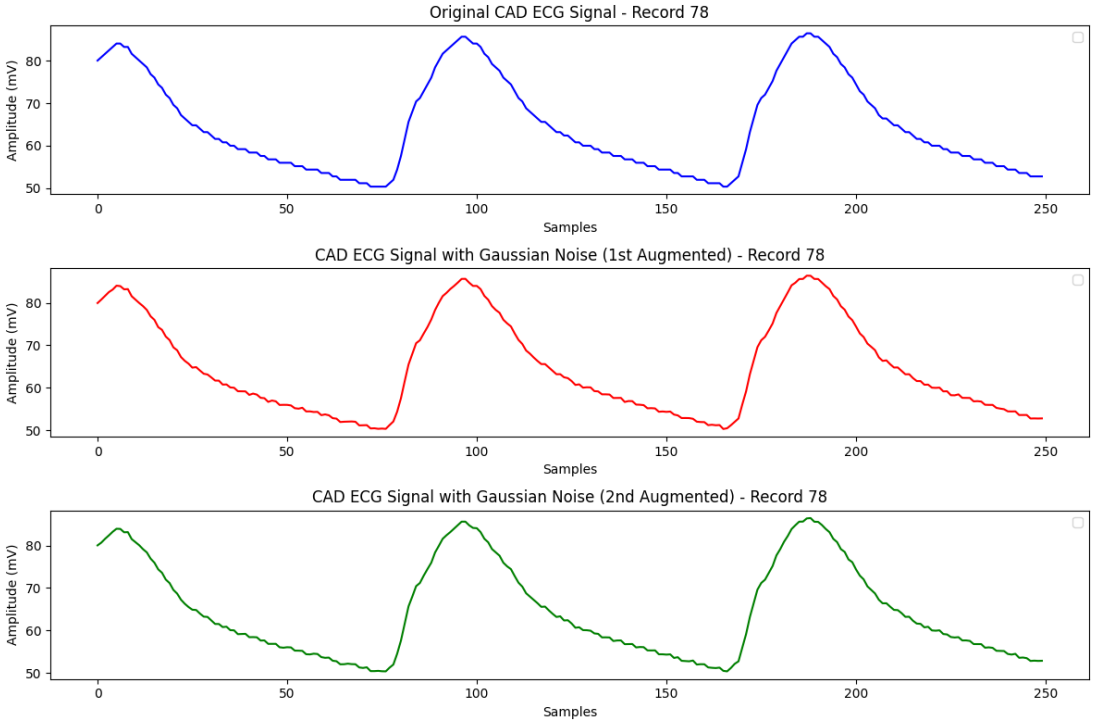

(d) ECG Signal demonstrating CAD

FIGURE 7.1: Example of ECG signals from four different classes (Normal, CAD, Arrhythmia, and AF) with applied data augmentation.

7.3 Methodology

Figure 7.1 illustrates an overview of our proposed approach. As seen from this figure, the process begins with raw ECG signals obtained from four classes (Normal, CAD, AF, and Arrhythmia) and only one ECG channel (lead II), which will be explained in detail in Section 7.3.1. Subsequently, the raw data will be randomly selected and used for data augmentation (see Section 7.3.2). The proposed model (with its structure described in Sections 7.3.3, 7.3.4, and 7.3.5) is then employed to classify the ECG signals. In our study, we propose a novel model that integrates CNN with residual blocks and multi-head attention mechanisms to enhance the processing and classification of ECG signals. This hybrid architecture effectively captures both local patterns and long-range dependencies, thereby improving feature extraction and overall model performance in identifying cardiac diseases.

7.3.1 Data preprocessing

The dataset used in this study consists of four distinct categories of ECG data: atrial AF, arrhythmia, CAD, and healthy patients. Arrhythmia data are obtained from a large-scale 12-lead electrocardiogram database [234]. AF data are sourced from the 4th China Physiological Signal Challenge 2020 [235], while CAD data are retrieved from the MIMIC III database [199]. Data from healthy patients are acquired from the Fantasia database [236]. All of these databases are publicly accessible via PhysioNet. Twenty patients were carefully selected from each respective database, and their ECG data were segmented into 1-second intervals. This segmentation allows for a detailed analysis of variations and trends in the signals over short time intervals, a methodology extensively used in previous research to ensure consistency in ECG signal evaluation [6]. Each segment comprises 250 samples, representing approximately one full cycle of the ECG signal. Class labels were assigned as follows: 0 for normal, 1 for CAD, 2 for arrhythmia, and 3 for AF. Subsequently, 200 segments from each class were randomly selected for data augmentation and further used in the experimental process.

7.3.2 Data augmentation

Data augmentation is employed by introducing random noise to ECG segments, significantly enhancing model generalisation and helping to prevent overfitting. Commonly applied in ECG signal analysis [79], this method effectively replicates the noisy conditions typically encountered in real-world scenarios. By enhancing the diversity of the training set, this approach enables the model to recognise and adapt to more complex

patterns, thereby improving its accuracy and performance on unseen data. Gaussian noise is added to generate synthetic examples, to improve the model's generalisation. A zero-mean white Gaussian noise, denoted as $\mathcal{N}(0, \sigma^2)$, with a standard deviation σ, is added to each sample. In this study, Gaussian noise is applied to each original ECG segment to generate two augmented segments as shown in Figure 7.1. The white Gaussian noise is introduced as follows:

$$X_{\text{augmented}} = X + \mathcal{N}(0, \sigma^2), \tag{7.1}$$

where $X = \{x_1, x_2, x_3, \ldots, x_n\}$ represents the original ECG segment composed of n samples, $\mathcal{N}(0, \sigma^2)$ represents Gaussian noise with mean 0 and variance σ^2.

7.3.3 CNN block

A CNN block utilises 50 convolutional filters, each with a kernel size of 28, to process the input ECG data. The convolutional operation involves sliding these filters across the temporal dimension of the input sequence to capture local patterns and features. The activation function applied is the ReLU, which introduces non-linearity into the model, enabling it to learn more complex patterns from the ECG signals. Furthermore, L2-norm regularisation is applied to the convolutional weights with a penalty term of 0.01. This regularisation technique helps to prevent overfitting by penalising large weights, thereby promoting model generalisation when interpreting ECG signals.

Batch normalisation is applied to adjust the activations from the previous layer by normalising them across the mini-batch. This process standardises the mean and variance of the inputs to each subsequent layer, which helps to reduce internal covariate shifts and stabilise the distribution of ECG signal features. The batch normalisation stabilises and accelerates the training process, ensuring that the data entering each layer is more consistent. This improved consistency facilitates more efficient gradient-based optimisation and enhances the overall training performance. To further address overfitting, the block incorporates a dropout layer with a dropout rate of 0.1, which randomly disables a fraction of neurons during training. This technique reduces the likelihood of overfitting by preventing the network from becoming excessively dependent on any individual neuron, thereby enhancing its ability to generalise to unseen ECG data. A max pooling layer with a pool size of 3 is used to decrease feature map dimensionality by selecting the maximum values from non-overlapping segments. This pooling operation reduces computational load and introduces translational invariance by emphasising the most significant features while discarding less relevant information from the ECG signal.

7.3.4 Residual block

A residual block is incorporated to enhance the model's ability to learn effectively as the network depth increases. The residual block contains two main convolutional layers; The first layer performs a one-dimensional convolution with 128 filters and 16 kernel size, using ReLU activation function. The filter size dictates the number of features extracted by each convolutional layer, whereas a kernel size of 16 specifies the window across which the convolution is applied, enabling the network to identify significant patterns within the temporal data. Each convolutional layer is followed by a batch normalisation layer, which standardises the outputs and improves the efficiency of convergence during training. A shortcut connection is introduced to allow the output from the CNN block to pass through directly if the dimensions of the input align with the dimensions of the processed output before addition. Mathematically, this can be expressed as:

$$\mathbf{y} = \mathcal{F}(\mathbf{x}, \{W_i\}) + \mathbf{x}, \tag{7.2}$$

where $\mathbf{x}$ denotes the input signal, $\mathcal{F}(\mathbf{x}, \{W_i\})$ represents the output of the convolutional layers with the learned weights $\{W_i\}$, and $\mathbf{y}$ corresponds to the final output of the residual block. If the dimensions of $\mathbf{x}$ and $\mathcal{F}(\mathbf{x}, \{W_i\})$ do not align, a one-dimensional convolution with a kernel size of 1 is applied to the shortcut connection to match the dimensions, thereby enabling the addition to be performed successfully. Following the two convolutional layers, a max pooling layer is applied with a pool size of 2 to downsample the feature map, reducing its dimensionality and summarising the most prominent features within each pooling window.

7.3.5 Attention block

An attention block integrates multi-head self-attention with feed-forward neural networks, using residual connections and layer normalisation to enhance the model's ability to capture complex dependencies in ECG data. This block is designed to maintain stable gradient flow, minimise overfitting, and address vanishing gradient issues, all of which are essential for effectively training deep neural networks. The block begins with layer normalisation, which standardises input activations across feature dimensions for each ECG instance. This is particularly beneficial for ECG data, where signal amplitudes and baselines vary. After normalisation, the block applies multi-head self-attention to capture relationships between different elements of the ECG sequence. The scaled dot-product attention, which forms the core of each attention head, is defined in Equation 2.4.

By employing 256 attention heads, each processing 512 dimensions of the input, the model captures a broader range of dependencies across the ECG sequence, enabling it to focus on different aspects or patterns within the ECG signal. After the multi-head operation, a residual connection is applied, combining the original ECG input x with the output from the multi-head operation. Following the multi-head attention mechanism, a feed-forward neural network is applied. This network comprises two fully connected layers. The first fully connected layer in the feed-forward network expands the input dimensionality to a size of 128, matching the size of the hidden layer. This expansion enables the model to learn richer and more abstract feature representations. By applying a ReLU activation, the layer introduces non-linearity, which enhances the model's ability to capture and process more complex patterns, a particularly important feature for ECG data. The second dense layer reduces the data back to its original input dimensionality. Regularisation is applied through both L1-norm and L2-norm penalties in this layer, while dropout is used between the layers to prevent overfitting. The feed-forward network operates independently on each time step of the ECG sequence, ensuring that the features extracted from each step are processed separately. Finally, the output of the feed-forward network is combined with the residual connection from the multi-head attention layer, preserving the original input information and facilitating effective gradient flow during backpropagation.

7.4 Experimental results

In this section, we evaluate the performance of the proposed model, with particular emphasis on classification accuracy, high-dimensional data visualisation, and a comparative study with other traditional ML and DL-based models.

7.4.1 Experiments

Four distinct classes were formed in our dataset, with each class containing 200 ECG segments, resulting in a total of 800 ECG segments, as mentioned in Section 7.3.1. After applying data augmentation, the dataset was expanded to 1,600 ECG segments, equivalent to 400,000 ECG samples, which were used for the subsequent experiments. To ensure a robust evaluation of the model's performance, we employed 10-fold cross-validation. In each fold, 90% of the data was used for training, while the remaining 10% was reserved for testing. This partitioning method is widely employed due to its ability to provide substantial training data while simultaneously allocating sufficient resources to assess the model's generalisation capability.

To evaluate the model, performance metrics including accuracy, precision, recall, and F1 score were used to measure the performance of the classification model as shown in Equations (2.5) , (2.7), (2.6), and (2.9).

Additionally, t-SNE analysis is conducted to visualise high-dimensional data and to investigate the impact of each important block in the proposed model on classifying between classes. This technique facilitates the observation of the distribution and clustering of data points in a lower-dimensional space. By applying t-SNE, insights can be gained into the relationships between different classes of ECG signals, aiding in the identification of patterns that may indicate various CVDs.

7.4.2 Results and discussion

Table 7.1 presents a comparative analysis of the existing traditional ML and DL-based model and our proposed model using 10-fold cross-validation. SVM achieves a training accuracy of 80.04%, but its test accuracy drops significantly to 63.25%, suggesting that the model may have overfitted the training data. Although SVM captures patterns in the training set, its generalisation capability is limited, as indicated by moderate precision of 65.52%, recall of 63.10%, and F1 score of 62.08%. KNN exhibits poor performance, with a training accuracy of 61.72% and a test accuracy of 50.50%, implying that it struggled to model the underlying relationships in the data. Its F1 score of 46.79% further reflects its difficulty in effectively handling this multi-class classification task. Logistic Regression shows the lowest overall performance among the traditional ML methods, with a training accuracy of 52.81% and a test accuracy of 48.45%. The result indicates that the linear nature of Logistic Regression is not suitable for the complexity and non-linearity of the data. In contrast, Naive Bayes outperforms both KNN and Logistic Regression, achieving a relatively balanced training accuracy of 67.95% and a test accuracy of 67.54%, indicating improved generalisation. Its higher precision of 79.07% suggests greater effectiveness in avoiding false positives; however, the moderate F1 score of 64.48% indicates that, like the other traditional models, it still struggles to fully capture the complexity of the ECG data

Among existing models, the Convolutional Attention-based model [150] exhibits strong performance, achieving a training accuracy of 98.11% and a testing accuracy of 97.50%. Its precision, recall, and F1 score are closely aligned, each around 97%, demonstrating a well-balanced and robust ability to handle both positive predictions and true positives effectively. Similarly, our baseline CNN model [31] also displays commendable results with a test accuracy of 97.13%, precision of 97.23%, and an F1 score of 97.11%, indicating reliable and consistent performance across key metrics. These results place it among the

TABLE 7.1: Comparative performance of existing Traditional ML and DL-based models and our proposed model using 10-fold cross-validation

Model	**Train Accuracy**	**Test Accuracy**	**Precision***	**Recall***	**F1 Score***
SVM	0.8004	0.6325	0.6552	0.6310	0.6208
KNN	0.6172	0.5050	0.6314	0.5053	0.4679
Logistic Regression	0.5281	0.4845	0.6085	0.4904	0.4800
Naive Bayes	0.6795	0.6754	0.7907	0.6753	0.6448
RNN-LSTM [149]	0.7977	0.7845	0.8146	0.7845	0.7839
BiLSTM [136]	0.4356	0.4195	0.4885	0.4195	0.3614
VGG19 [237]	0.3860	0.3866	0.4023	0.3866	0.2997
ResNet-16 [32]	0.7127	0.7104	0.7538	0.7104	0.7061
Convolutional Attention-based [150]	0.9811	0.9750	0.9767	0.9736	0.9742
Baseline CNN [31]	0.9449	0.9713	0.9723	0.9715	0.9711
Proposed model	**0.9979**	**0.9954**	**0.9953**	**0.9955**	**0.9953**

* Precision, recall, and F1-scores are reported as macro averages.

models, slightly behind the attention-based architecture but still demonstrating excellent generalisability.

However, several other architectures struggle to reach comparable levels of performance. For example, the BiLSTM model [136] and the VGG19 model [237] exhibit notably lower test accuracies of 41.95% and 38.66%, respectively. The VGG19 model, in particular, suffers from poor precision, with an F1 score as low as 29.97%, highlighting its imbalanced performance and difficulty in accurately predicting correct outputs. Similarly, the BiLSTM model produces only moderate results, with a precision of 48.85% and an F1 score of 36.14%, underscoring the limitations of these architectures for multi-class heart disease classification.

The RNN-LSTM [149] and ResNet-16 [32] models, although performing better than the BiLSTM and VGG19, still demonstrate only moderate performance. The RNN-LSTM achieves a test accuracy of 78.45%, precision of 81.46%, recall of 78.45% and an F1 score of 78.39%, while the ResNet-16 model achieves a test accuracy of 71.04%, precision of 75.38%, recall of 71.04% and an F1 score of 70.61%. In contrast, the proposed model clearly outperforms all existing models across the evaluated metrics. With a training accuracy of 99.79% and a test accuracy of 99.54%, it not only surpasses its competitors but also sets a new benchmark for performance in multiclass heart disease classification. Furthermore, its precision, recall, and F1 score are exceptionally high, reflecting the model's superior ability to balance true positive classifications with false positives, with values of 99.07%, 98.99%, and 98.97%, respectively. The close alignment of precision and recall further indicates that the model achieves a well-balanced trade-off between sensitivity and precision, which is crucial for ensuring reliable and accurate predictions in real-world applications.

The overall results demonstrate the superior performance of the proposed model, significantly outperforming other DL-based models, particularly in test accuracy and F1 score, both key indicators of model robustness and generalisability. These findings suggest that the architectural innovations integrated into the proposed model are highly effective, offering not only enhanced predictive accuracy but also more reliable and balanced performance across all key evaluation metrics. In contrast, traditional ML models exhibit limitations in their ability to generalise, especially when compared to the proposed model and other DL models. These limitations arise from the complexity of the ECG data and the non-linear relationships among the variables, which traditional models such as Logistic Regression, KNN, and SVM are unable to effectively handle. While Naive Bayes exhibits relatively satisfactory performance due to its simplicity and its premise of feature independence, it still fails to achieve the performance levels of more advanced DL models.

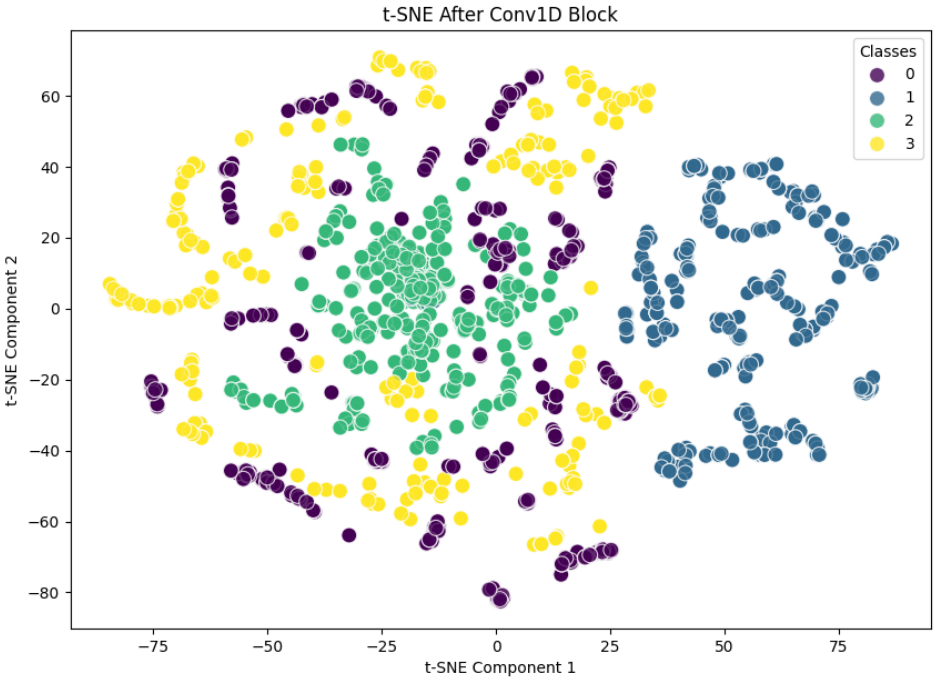

(a) t-SNE visualisation after CNN block

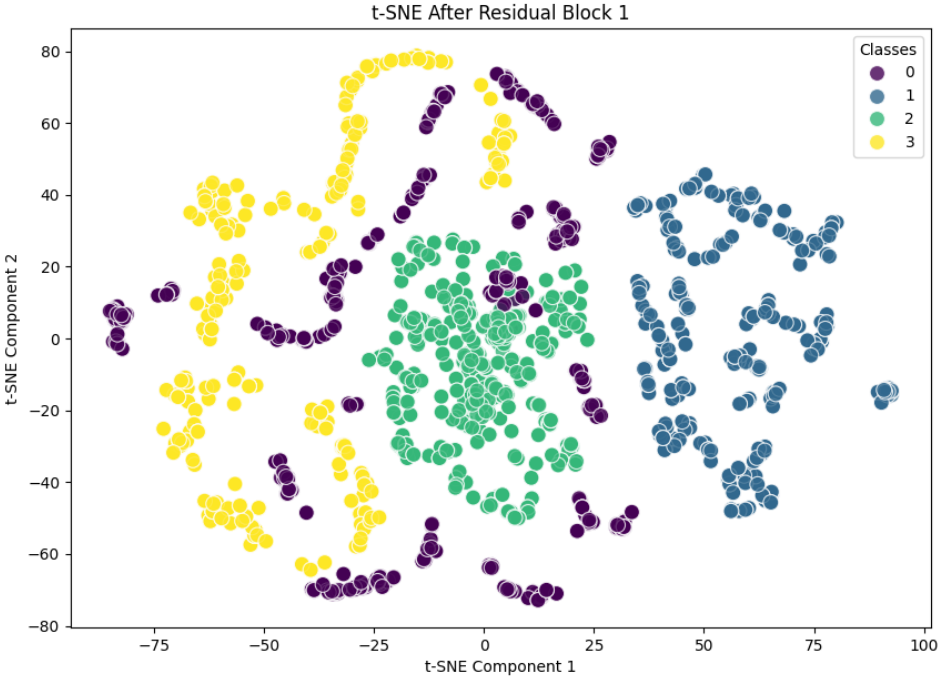

(b) t-SNE visualisation after Residual block 1

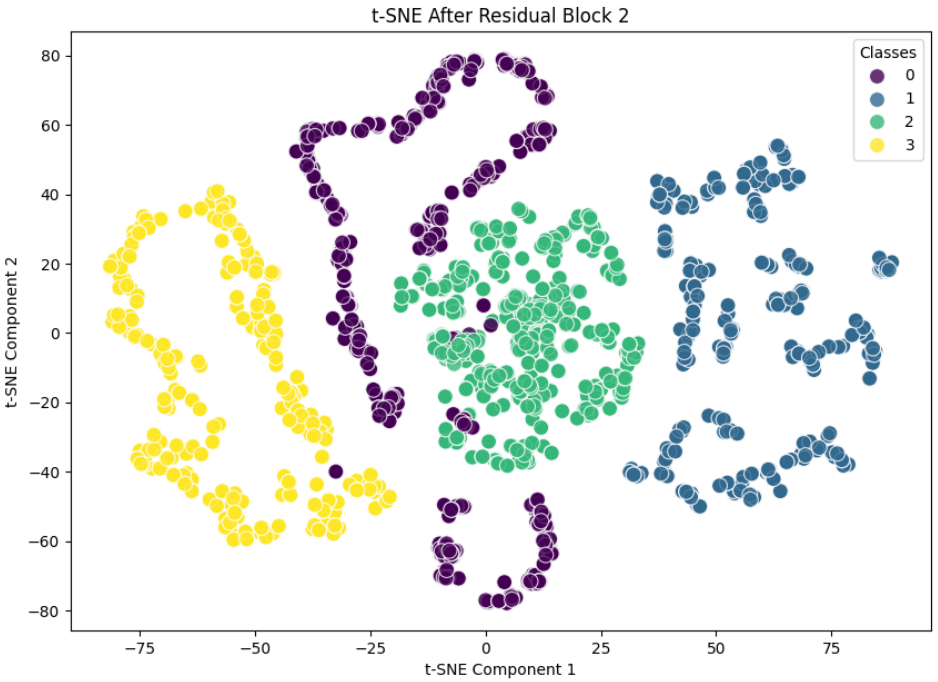

(c) t-SNE visualisation after Residual block 2

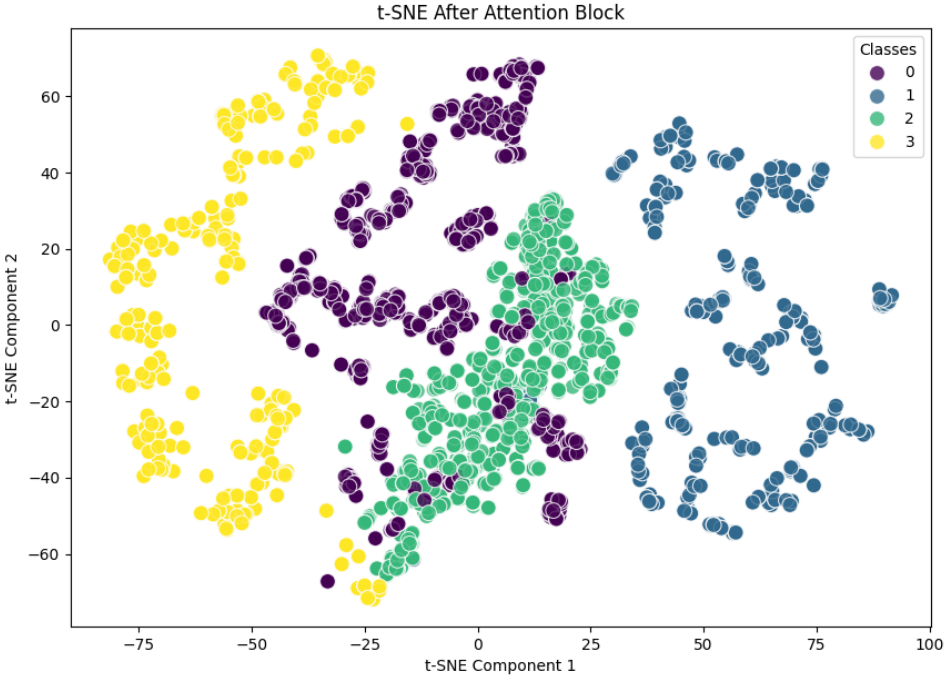

(d) t-SNE visualisation after attention block

FIGURE 7.2: t-SNE visualisations of the feature space.

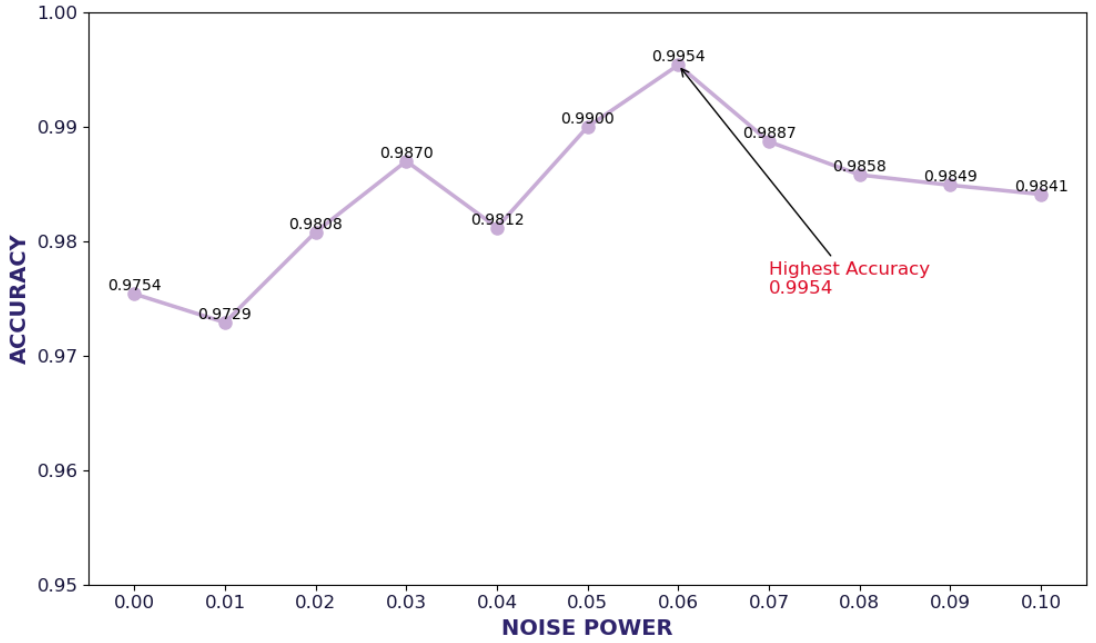

FIGURE 7.3: Impact of noise power variation on the accuracy of multi-class classification.

Figure 7.2 illustrates the effect of each block on feature extraction, visualised using t-SNE. Figure 7.2(a) shows that after the CNN block, the class separation is moderate, with some overlap between the clusters. Figure 7.2(b) demonstrates the effects after the first residual block, where distinct clusters begin to emerge, indicating clearer separation between classes. Following the second residual block, the clusters become more refined, demonstrating improved feature extraction, as shown in Figure 7.2(c). Figure 7.2(d) indicates that after the attention block, the class separation is more distinct, with clearly defined clusters for each class and only minor overlap between classes 0 and 2. This underscores the importance of residual and attention mechanism blocks in enhancing the model's ability to detect complex features across various classes.

Figure 7.3 illustrates the relationship between noise power and classification accuracy on the testing set for our proposed model. The noise power, a parameter in data augmentation that employs white Gaussian noise with a standard deviation, ranges from 0.01 to 0.1. At zero noise power, the baseline accuracy was 0.9754, indicating the model's performance in the absence of noise augmentation. With an increase in noise power to 0.01, a slight decrease in accuracy to 0.9729 was observed, suggesting that even minimal noise introduction may initially degrade performance. However, when noise power reached 0.02, the accuracy improved to 0.9808, surpassing the baseline. A further increase in noise power to 0.03 resulted in a substantial improvement, with accuracy rising to 0.987. This trend continued, with the highest accuracy of 0.9954 achieved at a noise power of 0.06. These results indicate that moderate noise augmentation can enhance

the model's robustness by mitigating overfitting and promoting improved generalisation. Therefore, a noise power of 0.06 was selected as the optimal level, as it provided the highest accuracy while also improving the model's generalisation and robustness. At noise power levels of 0.07 and 0.08, accuracy decreased slightly to 0.9887 and 0.9858, respectively. This decline continued as noise power increased to 0.1, leading to an accuracy of 0.9841. Despite this gradual decrease, the model maintained relatively high accuracy, demonstrating its resilience to higher noise levels.

TABLE 7.2: Comparison of Model Block Configurations in the Proposed Model

Block			Accuracy	Precision	Recall	F1 Score
CNN	**Residual**	**Attention**				
✓	X	X	0.9904	0.9912	0.9902	0.9905
✓	✓	X	0.9920	0.9925	0.9916	0.9919
✓	✓✓	X	0.9900	0.9908	0.9899	0.9898
✓	✓✓	✓	**0.9954**	**0.9953**	**0.9955**	**0.9953**

X: no block used, ✓: one block used, ✓✓: two blocks used.

In addition, an ablation experiment is conducted to evaluate various configurations for the integration of the convolutional block, residual block, and attention mechanism, as shown in Table 7.2. The results of this experiment underscore the significance of model components, demonstrating that the incorporation of residual and attention mechanisms can significantly enhance classification efficacy. The first configuration, which utilised only the CNN block, achieved an accuracy of 0.9904. This outcome illustrates that a baseline CNN architecture can effectively classify four distinct classes, indicating robust baseline performance. However, the absence of more advanced mechanisms highlights potential areas for further enhancement. Subsequently, the incorporation of a residual block alongside the CNN resulted in an improved accuracy of 0.9920. This enhancement highlights the effectiveness of residual connections in facilitating improved gradient flow during training, thereby enhancing the model's generalisation capabilities. In contrast, the incorporation of both the CNN and two residual blocks led to a slight decrease in accuracy to 0.9900. This finding suggests that increasing architectural complexity does not necessarily enhance performance. Nevertheless, the model continues to demonstrate robust performance while effectively mitigating the vanishing gradient problem. Lastly, the integration of all three blocks, CNN, two residual blocks, and Attention, resulted in the highest accuracy of 0.9954. This significant improvement underscores the critical

role of the attention mechanism in enabling the model to selectively focus on relevant features within the ECG data, thereby enhancing overall classification performance.

Figure 7.4 illustrates the corresponding confusion matrix for the proposed multi-class classification model in both training and testing phases. Figure 7.4(a) demonstrates a high level of training accuracy across all categories, correctly classifying 538 instances of Normal, 539 of CAD, 538 of Arrhythmia, and 540 of AF. Notably, only two instances from the Normal class were misclassified as Arrhythmia, and one instance of CAD was similarly misclassified as Arrhythmia. Furthermore, Figure 7.4(b) shows the correct classification of 58 instances of Normal, 62 of CAD, 57 of Arrhythmia, and 62 of AF in the testing phase. Notably, only one instance from the Normal class was misclassified as Arrhythmia, while no misclassifications were observed for CAD, Arrhythmia, or AF. These results highlight the robustness of the model in accurately differentiating between the classes, even when applied to unseen data. This performance suggests that the model maintains a high level of precision and effectiveness across both the training and testing datasets.

Despite the strong classification performance, the confusion matrix reveals occasional misclassification between CAD and arrhythmia classes, which may arise from overlapping ECG morphology between these conditions. This suggests that using a single-label multi-class framework may not fully represent real-world clinical scenarios, where patients can present with multiple CVD abnormalities. Multi-class classification enables differentiation between multiple cardiac conditions compared with binary diagnosis and provides more clinically informative predictions. Future studies will explore multi-label learning and investigate label correlation using multi-disease annotated datasets to model coexisting CVD conditions.

7.5 Conclusion

In this study, we propose a novel model for classifying multiple CVDs that integrates convolutional layers, residual networks, and attention mechanisms. This model effectively addresses noise in real-world ECG datasets through robust data augmentation, ensuring reliable performance across diverse cardiac conditions. Integrating noise into ECG signals before inputting the ECG data into the model is essential for simulating real-world clinical scenarios. This approach not only enhances the robustness of the model but also improves its accuracy in detecting cardiovascular conditions in the presence of typical environmental noise. Its multi-label classification capability has significant clinical implications, allowing for the simultaneous detection of multiple conditions from a single

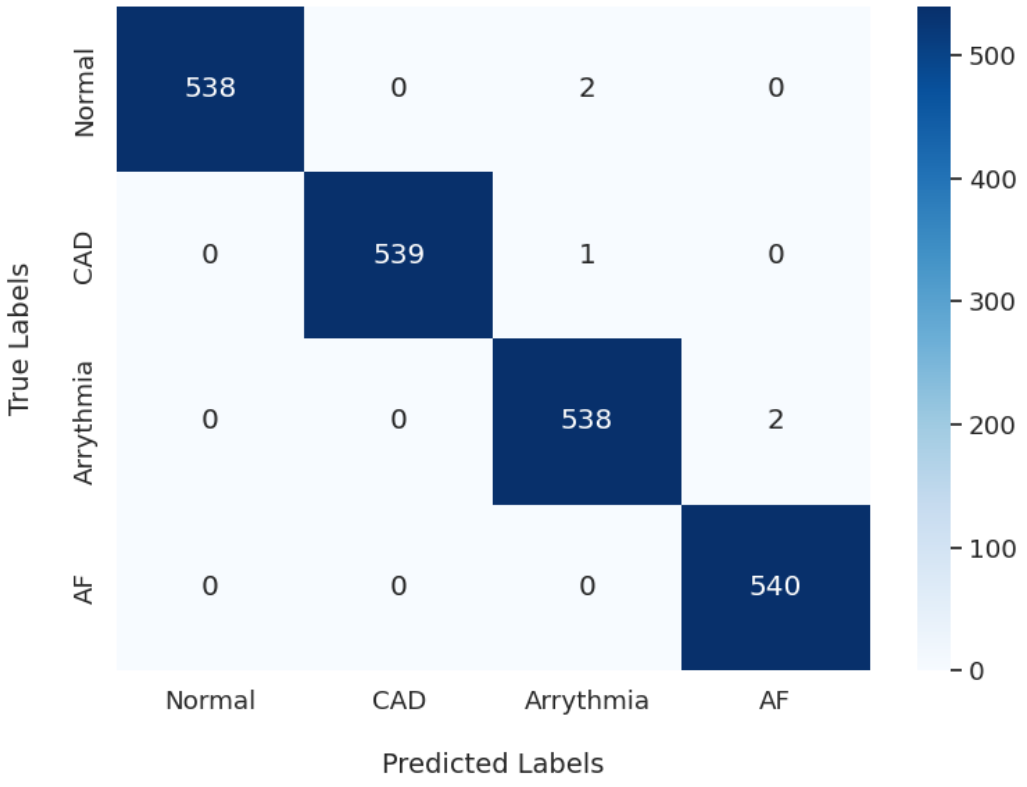

(a) On training set

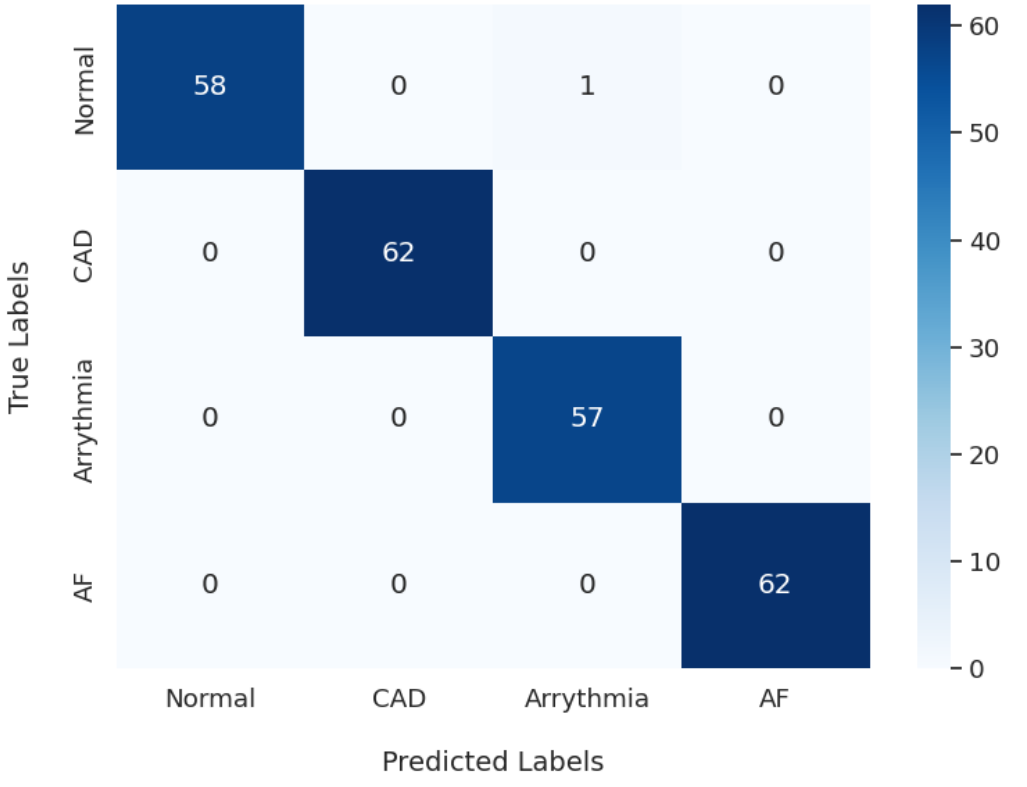

(b) On testing set

FIGURE 7.4: Confusion Matrix for the proposed model.

ECG reading. The proposed model can serve as a pre-screening tool, as ECG is non-invasive compared to other medical examination methods in hospitals. Furthermore, this approach can assist doctors in making informed decisions regarding the next steps in the diagnostic process. Future research may focus on refining the model for real-time monitoring and wearable devices, enhancing its clinical utility, and supporting proactive patient management.

Chapter 8

Conclusions and Future Work

This chapter provides a critical discussion and summary of the research contributions presented in this thesis, with a focus on their novelty, methodological significance, and potential clinical applicability. It begins by reflecting on the key innovations introduced across the proposed deep learning frameworks for ECG-based CAD diagnosis. A concise summary of the main findings from each chapter is then provided. This is followed by a discussion of the study's limitations, acknowledging the challenges and constraints encountered. The chapter concludes by outlining recommended directions for future research, with an emphasis on improving model generalisability and supporting real-world deployment.

8.1 Summary of Novel Contributions and Experimental Findings

This section outlines the major findings and considers their significance in relation to the aims of the study.

8.1.1 Chapter 3: Coronary Artery Disease Classification Using One-Dimensional Convolutional Neural Network

1. The proposed one-dimensional convolutional neural network is capable of accurately distinguishing between CAD and NON-CAD cases directly from raw ECG signals. Notably, it achieves this without the need for manual feature extraction, instead learning meaningful patterns automatically from the ECG data.

2. Experimental results indicate that shorter ECG samples, particularly those between 250 and 300 data points, improve generalisability to previously unseen data. This improvement may be attributed to reduced overfitting, as shorter segments contain fewer redundant or noise-prone signal components, enabling the model to focus more effectively on the most informative features. The highest test accuracy recorded was 96%, while the accuracy on unseen data ranged from 82.5% to 85.7% when using 250-point samples.

3. The proposed model achieved a classification accuracy of 97.3%, outperforming both the baseline one-dimensional convolutional neural network and the Hybrid CNN-LSTM model. Despite its strong performance and ability to generalise well from a relatively small dataset, the model's higher computational requirements may present challenges for deployment in resource-constrained environments.

4. CAD ECG data from the MIMIC III database exhibited significant variability in lead formats and storage formats, illustrating the complexity and practical challenges of working with real-world clinical datasets. This finding underscores the need for robust feature extraction methods that can accommodate such inconsistencies and help ensure dependable model performance in practical diagnostic settings.

8.1.2 Chapter 4: Enhanced Coronary Artery Disease Classification through Feature Engineering and One-Dimensional Convolutional Neural Network

1. Feature engineering was introduced to remove noisy and low-information ECG segments and to improve signal quality using sample entropy and standard normalisation. Applied prior to training the one-dimensional convolutional neural network, this approach significantly improved classification accuracy and robustness. The proposed model trained without feature engineering demonstrated poorer performance on unseen data, thereby highlighting the importance of signal quality and effective preprocessing.

2. A novel and streamlined one-dimensional convolutional neural network architecture has been developed for the diagnosis of CAD, integrated with a feature engineering module to enhance signal quality. This approach relies solely on ECG signals, offering a cost-effective alternative to existing diagnostic methods that depend on expensive biomarkers technologies.

3. The proposed one-dimensional CNN outperformed both classical and other deep learning models, achieving 98.5% accuracy on unseen data. Feature engineering

further enhanced generalisability, in temporal models such as LSTM and CNN-LSTM. These findings underscore the effectiveness of deep learning for CAD diagnosis using ECG, when applied to large and complex datasets.

4. The proposed model achieves fast inference times of 18.7 ms per subject, outperforming more complex models such as LSTM and CNN-LSTM. Combined with its high accuracy, this makes it well-suited for real-time applications such as portable ECG devices.

5. Although existing research on the MIMIC III dataset has primarily focused on conditions such as arrhythmia, CHF, and AF, relatively small attention has been given to the detection of CAD. This study addresses that gap by developing a one-dimensional convolutional neural network applied to the MIMIC III dataset.

8.1.3 Chapter 5: CADNet: A lightweight Neural Network for Coronary Artery Disease Classification Using Electrocardiogram Signals

1. A lightweight neural network architecture (CADNet) was developed for efficient classification of CAD using ECG signals, removing the need for extensive handcrafted features. The model combines low computational complexity with strong classification performance, making it ideal for resource-constrained CAD detection applications.

2. CADNet model significantly reduces computational complexity compared to the traditional baseline deep learning model, decreasing trainable parameters from approximately 8 million to just 8,552 and reducing storage size from 32.19 MB to 33.41 KB, while maintaining high accuracy and achieving faster runtime. Moreover, CADNet is considerably smaller than well-known lightweight deep learning models such as EfficientNetB0, SqueezeNet, 1D-MobileNetV1, and ShuffleNetV1. This effective balance of efficiency and performance renders CADNet highly suitable for real-time CAD diagnosis on resource-constrained devices.

3. A novel data purification step was applied to the ECG data subsets prior to model training. This technique removes noise and low-quality segments from the ECG signals. By filtering these components, the data purification process improves the overall quality of the ECG input data, thereby improving the model's ability to accurately classify CAD.

4. CADNet model comprises two key components: feature encoding block, which efficiently captures critical temporal ECG patterns such as P waves, QRS complexes,

and T waves, and compact pooling block, which reduces temporal resolution to focus on the most relevant CAD ECG features while enhancing robustness to noise.

5. CADNet was evaluated across four diverse datasets including PTB-XL, MIMIC-III, St. Petersburg and Fantasia, demonstrating superior accuracy and robustness. The model's consistent performance across different age groups, with highest accuracy in the 40 to 50 year range and slight declines in younger and older populations, further confirms its strong generalisability while highlighting the potential influence of age related physiological variations in ECG signals [238].

8.1.4 Chapter 6: Real-time Coronary Artery Disease Detection from 12-Lead ECG Using A Lightweight Deep Network

1. A novel lightweight deep learning model was developed for CAD diagnosis using 12-lead ECG, combining depthwise convolutions, residual connections and squeeze-and-excitation blocks. The model achieved strong discrimination performance with 95.45% accuracy, 98.96% AUC and a 95.45% F1 score. This result significantly outperforms the compared traditional machine learning, deep learning and other lightweight architectures.

2. The proposed lightweight model achieves an optimal balance between accuracy and efficiency, using only 13,584 trainable parameters, a 0.0530 MB storage size and 0.0191 seconds inference time per subject, while obtaining 95.45% accuracy and 98.96% AUC. The proposed ligtweight model outperforms traditional machine learning, conventional deep learning and other lightweight architectures, making it suitable for real-time deployment in resource-constrained environments.

3. The proposed model provides the optimal balance of efficiency and performance on the STM32F469I-DISCO. It requires far fewer MAC, 1.69 million, has the smallest parameter count at 0.0553 MB, operates within a flash size of 0.0852 MB and RAM of 0.1117 MB, and achieves the fastest inference time of 0.1121 seconds per subject with the lowest energy consumption of 7.39 mJ, outperforming MobileNetV1 [232], lightweight CNN [163], and ResNetLite [233].

4. t-SNE visualisation and confusion matrix demonstrate strong multiclass discrimination. Feature representations cluster distinctly for normal sinus rhythm and CAD subclasses, and the confusion matrix shows high true-positive rates with minimal misclassification, indicating that the proposed model reliably separates and classifies closely related CAD subclasses.

5. SHAP analysis indicated that aVR, Lead II and V2 contribute most strongly to distinguishing CAD from normal rhythm, whereas leads such as V3 and V5 have

comparatively less influence. This suggests the model leverages interactions between multiple leads and reduces redundancy rather than relying solely on individual leads. Previous studies have emphasised that Lead II is commonly preferred when a full 12-lead ECG is not available, as it provides a clear representation of the P wave, QRS complex and T wave, all of which are critical for detecting cardiac arrhythmias and ischaemic changes [131, 137].

6. 12-lead ECG data corresponding to several specific cardiac conditions were used to support the diagnosis of CAD. These conditions included IMI, AMI, and LMI, each representing a distinct form of myocardial injury within the subclasses of CAD. Additionally, data from patients with normal sinus rhythm provided a baseline representing the absence of cardiac abnormalities. The proposed model demonstrated strong performance when distinguishing normal cases from individual CAD subclasses, achieving accuracy of 93.94% and high AUC scores. However, accuracy decreased to 72.50% in multi-class CAD classification, highlighting the challenges and potential overlap in diagnostic features among different CAD subclasses.

8.1.5 Chapter 7: Multi-Disease Cardiovascular Detection from ECG Signals Using an Attention-Driven Deep Network

1. A novel model integrating convolutional layers, residual networks, and attention mechanisms was developed for multi-label classification of multiple CVDs from a single ECG reading. This integrated architecture captures both local and global features within ECG data effectively, enabling a comprehensive analysis of complex cardiac patterns. Furthermore, the proposed model achieves exceptional performance, with an average classification accuracy of 99.54% on the test set. This surpasses existing deep learning-based models, demonstrating the effectiveness of the proposed architecture.

2. A data augmentation process was applied to enhance dataset robustness and generalisability, as well as to investigate the impact of noise power on classification performance. The data augmentation with white Gaussian noise improves the model's robustness and generalisation. The study showed that moderate noise levels, with a noise power of 0.06, optimally enhance the model's ability to handle real-world ECG data and improve classification accuracy.

3. The proposed model detects multiple heart conditions from a single ECG reading and facilitates improved diagnosis and treatment. The study utilised publicly available ECG data from four categories, including AF, arrhythmia, CAD, and normal sinus rhythm, all sourced from PhysioNet databases. A key advantage

of the model is its ability to identify diseases even when overlapping features exist across different CVD conditions, thereby enhancing diagnostic accuracy and reducing the risk of undetected cases.

4. t-SNE visualisation was introduced to examine class separation and to assess the contribution of each major block in the proposed model. The results demonstrate the model's strong discriminative capability for different heart conditions, showing clearer class separation through the CNN, residual and attention blocks. The attention mechanism combined with residual blocks significantly enhances feature extraction from ECG signals, resulting in distinct clusters with minimal overlap.

5. The proposed model demonstrates strong potential as a non-invasive pre-screening tool for CVDs disease, achieving high accuracy while detecting multiple conditions from a single ECG recording. This capability could enhance early detection and improve patient outcomes in clinical settings.

8.2 Final Reflections

In this thesis, several novel methodologies were introduced to improve the efficiency and accuracy of ECG signal classification, with a focus on resource-constrained environments. Motivated by the advantages of structured and compact representations, this study investigated novel approaches within three key frameworks: lightweight deep learning architectures, signal preprocessing techniques, and resource-constrained deployment pipelines. In addition, the integration of these frameworks was examined to enhance model generalisability and enable real-time application. The principal contributions of the work are summarised as follows:

1. The design of compact convolutional neural networks tailored for one-dimensional ECG signals, enabling efficient training and inference without compromising performance

2. The implementation of feature engineering and augmentation techniques to enhance model resilience to noise and variability in ECG recordings.

3. The development of the lightweight CADNet model tailored for resource-constrained platforms, including inference performance evaluation and model size optimisation, thereby enabling its suitability for real-time ECG-based CAD diagnosis.

4. The development and deployment of a novel lightweight deep learning model for the classification of binary and multi-class CAD using 12-lead ECG data, designed

to operate efficiently on resource-constrained environments and support low-power tools for ECG-based screening and monitoring.

5. A novel multi-label electrocardiogram classification model was proposed, combining convolutional layers, residual blocks, and attention mechanisms to detect multiple cardiovascular conditions, including coronary artery disease, atrial fibrillation, and arrhythmia, from a single ECG reading.

8.3 Research Limitations

While this thesis contributes to the development of efficient and accurate deep learning models for CAD detection from ECG signals, it also identifies several potential limitations arising at different stages of the research.

In Chapter 3, the limited availability of CAD ECG data may have affected the diagnostic performance of the model. Extracting CAD ECG records from the MIMIC III database proved challenging due to the complexity and inconsistent structure of the dataset. The compact 1D-CNN model achieved strong accuracy, however, it requires substantial computational resources, which may limit its suitability in severely resource-constrained environments. Existing literature indicates a lack of clear ECG biomarkers or segments that reliably denote CAD, resulting in unavoidable uncertainty in the interpretation of ground truth. To address this, deep learning is employed to mitigate the absence of clearly defined ECG biomarkers by learning discriminative patterns directly from the ECG data. In this chapter, a sample length of 250 data points appeared to be a promising configuration for the proposed model. Nonetheless, further investigation is required to determine the optimal sample length and to assess the influence of other related factors on model performance.

Based on the findings in Chapter 3, a sample length of 250 data points was identified as optimal and then used in the investigations presented in Chapter 4. Feature engineering using Sample Entropy and standard normalisation was applied to extract appropriate ECG segments and improve signal quality prior to model training. Although this improved classification accuracy, it indicates that the model still relies on preprocessing, which may limit its applicability in situations where raw ECG data must be used directly. Despite this, the model maintained reasonable performance without feature engineering, suggesting a degree of robustness. The dataset used in Chapter 4 was also larger than that used in Chapter 3, which may have influenced the improved performance.

Despite promising results, several limitations are identified in Chapter 5. One of the key challenges lies in the limited availability of CAD datasets, particularly those containing

clinically confirmed indicators such as single- or multi-vessel disease. While four publicly available datasets were utilised, they may not fully reflect the diversity encountered in real-world clinical settings, where variations in ECG acquisition devices, recording conditions, and patient demographics can significantly influence model performance, as discussed in previous chapters. Although cross-dataset validation and age group analysis were conducted, further evaluation is required to confirm the generalisability of the model. Additionally, the current analysis focused on CAD superclass-level diagnoses, and the model's ability to distinguish between specific CAD subtypes was not examined in depth, potentially overlooking clinically meaningful distinctions. Moreover, the use of existing public datasets limits the opportunity to evaluate the model on newly acquired data. The model also lacks an integrated explainability framework, which is essential for promoting clinical trust in healthcare settings. Additionally, from Chapter 3 to Chapter 5, only Lead II ECG data were considered. The contribution of multiple leads was not explored, which may have limited the model's ability to capture distributed diagnostic information that could improve classification performance.

Following the encouraging performance results reported in Chapter 5, Chapter 6 focuses on the practical aspects of model deployment by examining the potentiality of implementation on resource-constrained hardware platforms. In Chapter 6, the evaluation was limited to the PTB-XL database, which may not capture the full diversity of patient populations, signal acquisition conditions or other real-world clinical variables, thereby constraining the external validity and generalisability of the results. Additionally, the proposed model demonstrated reduced reliability in the multi-class classification of CAD subclasses, due to the high degree of feature similarity between CAD subclasses. The current feature extraction strategy may be insufficient to resolve overlapping ECG characteristics, resulting in occasional misclassifications. In particular, improving the model's ability to distinguish between IMI and other CAD subclasses, especially LMI, remains a challenge due to their highly similar signal features. Furthremore, although per-inference energy consumption is low, the model's efficiency in real-time monitoring scenarios remains unaddressed, since power management and scheduling trade-offs may affect sustained operation. While deployment of the 12-lead ECG lightweight model was demonstrated on the STM32F469I-DISCO microcontroller, the proposed model was not evaluated across a wider range of resource-constrained devices. Its cross-platform performance and deployment potential therefore remain to be investigated. Furthermore, unlike the previous chapters, which were limited to Lead II ECG, Chapter 6 incorporated 12-lead ECG data, enabling a more comprehensive analysis of spatial features relevant to CAD classification.

Building on the CAD classification explored in previous chapters, where 12-lead ECG

data and CAD subclasses were considered, Chapter 7 introduces a novel deep learning-based model for the classification of multiple cardiovascular conditions from a single ECG reading. While the proposed approach demonstrates potential, it presents several important limitations. Reliable ECG datasets covering a wide range of cardiovascular diseases are limited in availability. Since most conditions are not contained within a single dataset, four publicly available datasets were combined to enable broader representation. This integration presents challenges related to inconsistencies in data structure, variability in annotation standards, and differences in ECG acquisition protocols. In addition, a notable class imbalance within the combined dataset may affect the model's ability to learn underrepresented conditions effectively. However, in Chapter 7, data augmentation was employed to help balance the class distribution among cardiovascular disease categories.

Across the chapters, the proposed models demonstrate promising performance. However, although subject-independent data partitioning was employed to mitigate data leakage, the data preparation strategy in Chapter 4, 5 and 7 may still introduce subtle dependencies within patient-specific signal patterns due to the segmentation of ECG recordings into fixed-length windows. While this reflects common practice in ECG-based model development, further validation using entirely independent, real-world clinical datasets would strengthen the assessment of model generalisability.

Chapters 5 and 6 prioritise the design of lightweight DL-based architectures to achieve computational efficiency and practical deployability. While knowledge distillation was not incorporated in the present work, it provides a promising way for lightweight networks to benefit from the knowledge learned by larger, more expressive models. In a recent study [239], knowledge distillation is employed to develop a deep learning model that is highly efficient in terms of resource utilisation and requires minimal computational power during deployment. Exploring such strategies in future studies may further enhance model performance without increasing computational complexity. In addition to knowledge distillation, prior studies indicate that advanced model optimisation strategies, including structured pruning, quantisation, and model compression, can reduce model size and inference latency while maintaining competitive performance [240–242]. These approaches represent promising directions for further extending the proposed lightweight frameworks.

8.4 Future Research Directions

Across all chapters, a recurring limitation relates to the availability and quality of data. Although many approaches report high performance on publicly available datasets, few

have been tested in clinical trials or real-world healthcare settings. In this thesis, several efforts were made to address the limitations related to data availability. ECG signals from multiple publicly available datasets were combined to capture greater variability, and clinical guidance was obtained from doctors in Thailand, who provided valuable input on the use of data and its clinical interpretation. An attempt was also made to access NHS data, although this was not successful due to the extensive approval processes required. These efforts highlighted challenges such as inconsistent labelling, class imbalance, and differences in recording methods, which made it difficult to compare datasets and limited generalisability. As external clinical data could not be accessed within the timeframe of this research, the use of independent datasets and closer integration with clinical trials has been postponed to future work. Future research will consider the use of externally acquired ECG datasets, rather than relying solely on publicly available sources. This would provide greater control over clinical variables within a more structured and consistent environment, thereby reducing the complexity often encountered in existing datasets due to variability in structure, validity, and completeness. In addition, collaboration with clinical professionals is essential to confirm accurate interpretation of ECG findings and to support model evaluation with the expectations and requirements of real-world clinical settings. Engaging clinicians during the development and validation stages can provide valuable insights into the clinical relevance of detected features, support the identification of diagnostically meaningful ECG patterns, and contribute to the optimisation of model outputs in a way that improves their clarity for clinical interpretation.

In Chapter 3 and 4, a sample length of 250 data points was identified as optimal. However, further investigation is required to confirm this finding by examining the influence of filter settings and other factors that may affect model performance beyond the choice of optimal sample length. Future work should consider to validate the optimal sample length across different datasets and patient groups to ensure robustness and generalisability. In addition, identifying clinically relevant features that contribute to accurate CAD detection remains a key area of investigation. While deep learning models are capable of automatically learning CAD-related features from raw ECG signals, identifying which features hold the greatest diagnostic value remains a key priority for enhancing clinical trust and transparency.

Chapters 3 to 5, and Chapter 7, focused only on a single ECG lead. This limitation was addressed in Chapter 6, where a 12-lead ECG configuration was investigated. The results indicated that incorporating multiple leads could improve model performance. However, the improvement was not consistent across all cases. Expanding the number of input leads can lead to redundancy and greater input complexity, which may diminish the model's capacity to extract meaningful and discriminative features. Future studies

should further investigate the individual diagnostic contribution of each lead in order to determine how the model responds to specific ECG inputs. Although Chapter 6 evaluated multiple lead configurations and employed SHAP-based analysis to investigate lead importance, it did not explore the specific contribution of each lead when assessed independently from the others.

Future studies should consider evaluating model performance across diverse patient subgroups, including those defined by age, sex, and ethnicity, in order to enhance both the representativeness and clinical applicability of the proposed approach. In this thesis, ECG signals are used as the primary data source for CAD diagnosis. In routine clinical practice, ECG is commonly employed as a pre-screening tool, while definitive diagnosis and subsequent treatment decisions rely on the integration of multiple sources of clinical information. Future work could therefore explore the fusion of ECG data with other modalities, including demographic information, photoplethysmography (PPG), wearable sensor data, and medical imaging techniques such as chest X-ray and echocardiography, to provide a more comprehensive assessment of cardiovascular risk. Binary and multi-class classification of CAD were both investigated in this thesis. In Chapter 6, a 12-lead lightweight deep learning model was developed to support the classification of CAD subtypes, representing a valuable step toward improving the clinical utility of the proposed approaches. However, further research should aim to address the challenge of overlapping ECG features among CAD subclasses by investigating more advanced architectural designs and more effective feature extraction strategies as previously mentioned in Chapter 6.

In Chapter 6, the proposed model demonstrated potential suitability for real-time and energy-efficient ECG-based CAD diagnosis on resource-constrained devices. Although the model was deployed on the STM32F469I-DISCO microcontroller, further research is required to evaluate its performance across a wider range of hardware platforms and to assess its scalability, power efficiency, and deployment feasibility in real-world clinical environments. In addition, evaluating the model's reliability during long-term operation in continuous or real-time monitoring scenarios is essential for validating its stability and practical utility. A promising direction for future work involves the direct integration of the model with embedded hardware connected to ECG recording devices, allowing seamless acquisition, processing, and diagnosis within a unified low-power system. Such integration would support the development of self-contained diagnostic tools suitable for both clinical and remote healthcare settings.

A promising direction for future research is the integration of explainable AI techniques, such as ST-CNN-GAP-5 [243] to enhance the interpretability of the proposed model's predictions. These methods can help to visualise and clarify how the model processes

ECG signals and arrives at diagnostic decisions, offering insights into the specific regions or features that influence classification outcomes. This added transparency has the potential to improve clinician trust, support model validation, and enable the generation of automated interpretive reports to assist in clinical follow-up. Building on this, future work should also consider the development of AI-generated patient report outcomes, which could support clinical decision-making by providing structured summaries of diagnostic findings and potential treatment options. Such tools may ease the burden on healthcare professionals and contribute to more efficient patient management. Within the NHS, this could support earlier patient assessment and follow-up, thereby enabling rapid clinical intervention when necessary. As the role of AI in healthcare continues to expand, embedding interpretability within system design is expected to become essential for regulatory compliance, the development of clinician trust, and the safe and responsible adoption of AI in real-world clinical environments.

Bibliography

[1] T. H. Institute, "Heart information center: Heart anatomy," https://www.texasheart.org/heart-health/heart-informationcenter/topics/heart-anatomy.

[2] Phlbi.orge, "All about coronary artery disease," https://www.phlbi.org/divisions/heart-disease/coronary-artery-disease/.

[3] E. Kosari and K. Vafai, "Thermal stimulation of targeted neural circuits via remotely controlled nano-transducers: A therapy for neurodegenerative disorders," in *Advances in Heat Transfer*. Elsevier, 2020, vol. 52, pp. 543–581.

[4] T. M. P. V. Laboratory, "Cardiovascular lab: Electrocardiogram: Introduction'," http://www.medicine.mcgill.ca/physio/vlab/cardio/introecg.htm, July 2022.

[5] Who.int., "Cardiovascular diseases," https://www.who.int/health-topics/cardiovascular-diseases, 2019.

[6] J. H. Tan, Y. Hagiwara, W. Pang, I. Lim, S. L. Oh, M. Adam, R. S. Tan, M. Chen, and U. R. Acharya, "Application of stacked convolutional and long short-term memory network for accurate identification of CAD ECG signals," *Computers in Biology and Medicine*, vol. 94, pp. 19–26, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010482517304201

[7] S. L. Oh, E. Y. Ng, R. S. Tan, and U. R. Acharya, "Automated diagnosis of arrhythmia using combination of CNN and LSTM techniques with variable length heart beats," *Computers in Biology and Medicine*, vol. 102, pp. 278–287, 2018. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0010482518301446

[8] A. K. Bhoi, K. S. Sherpa, and B. Khandelwal, "Ischemia and arrhythmia classification using time-frequency domain features of QRS complex," *Procedia Computer Science*, vol. 132, pp. 606–613, 2018.

[9] R. Banerjee, A. Ghose, and K. Muthana Mandana, "A hybrid CNN-LSTM architecture for detection of coronary artery disease from ECG," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–8.

[10] L. Yao, C. Liu, P. Li, J. Wang, Y. Liu, W. Li, X. Wang, H. Li, and H. Zhang, "Enhanced automated diagnosis of coronary artery disease using features extracted from QT interval time series and ST–T waveform," *IEEE Access*, vol. 8, p. 129510–129524, 2020.

[11] Q. Nguyen, B. Nguyen, T. Nguyen, T. Do, J. Mbinta, and C. Simpson, "Stacking segment-based CNN with SVM for recognition of atrial fibrillation from single-lead ECG recordings," *Biomedical Signal Processing and Control*, vol. 68, May 2021.

[12] A. Saikhu, C. V. Hudiyanti, and A. Y. Wijaya, "Input feature selection in ECG signal data modelling using long short term memory," in *2021 4th International Seminar on Research of Information Technology and Intelligent Systems (ISRITI)*, 2021, pp. 229–234.

[13] World Health Organization, "United kingdom - who data," 2025, accessed: 28 February 2025. [Online]. Available: https://data.who.int/countries/826

[14] British Heart Foundation, "Cardiovascular disease statistics 2024," https://www.bhf.org.uk/what-we-do/our-research/heart-statistics/heart-statistics-publications/cardiovascular-disease-statistics-2024, 2024.

[15] C. D. Mathers and D. Loncar, "Projections of global mortality and burden of disease from 2002 to 2030," *PLoS Medicine*, vol. 3, no. 11, p. e442, 2006.

[16] B. Chong, J. Jayabaskaran, S. M. Jauhari, S. P. Chan, R. Goh, M. T. W. Kueh, H. Li, Y. H. Chin, G. Kong, V. V. Anand, J.-W. Wang, M. Muthiah, V. Jain, A. Mehta, S. L. Lim, R. Foo, G. A. Figtree, S. J. Nicholls, M. A. Mamas, J. L. Januzzi, N. W. S. Chew, A. M. Richards, and M. Y. Chan, "Global burden of cardiovascular diseases: projections from 2025 to 2050," *European Journal of Preventive Cardiology*, p. zwae281, 09 2024. [Online]. Available: https://doi.org/10.1093/eurjpc/zwae281

[17] C. Clinic, "Heart failure (congestive heart failure)," https://my.clevelandclinic.org/health/diseases/17069-heart-failure-understanding-heart-failure.

[18] M. J. Pencina, A. M. Navar, D. Wojdyla, R. J. Sanchez, I. Khan, J. Elassal, R. B. D'Agostino Sr, E. D. Peterson, and A. D. Sniderman, "Quantifying importance of major risk factors for coronary heart disease," *Circulation*, vol. 139, no. 13, pp. 1603–1611, 2019.

[19] R. Hajar, "Risk factors for coronary artery disease: historical perspectives," *Heart views: the official journal of the Gulf Heart Association*, vol. 18, no. 3, p. 109, 2017.

[20] N. Ramya, V. Prabakaran, A. Abbas, and S. P. Shankar, "Study on risk factors and angiographic pattern of coronary artery involvement in patients presenting with angina," *International Journal of Advances in Medicine*, vol. 6, p. 232, 2019.

[21] A. H. Abdelaziz, G. F. Gomaa, and A. M. Algamal, "Treatment strategies for patients with obstructive coronary artery disease (limited mansoura university experience) in 2021," *The Egyptian Journal of Hospital Medicine*, vol. 91, no. 1, pp. 5003–5008, 2023.

[22] NHS, "Overview - coronary heart disease," https://www.nhs.uk/conditions/coronary-heart-disease/, 2020.

[23] Who.int., "Cardiovascular diseases (cvds)," https://www.who.int/news-room/fact-sheets/detail/cardiovascular-diseases-(cvds), 2019.

[24] C. Physiology.com, "Cv physiology," https://cvphysiology.com/CAD/CAD012.

[25] J. Oster and G. D. Clifford, "Impact of the presence of noise on RR interval-based atrial fibrillation detection," vol. 48, no. 6, pp. 947–951, 2015.

[26] D. B. Springer, L. Tarassenko, and G. D. Clifford, "Logistic regression-HSMM-based heart sound segmentation," *IEEE transactions on biomedical engineering*, vol. 63, no. 4, pp. 822–832, 2015.

[27] F. Bouaziz, D. Boutana, and H. Oulhadj, "Diagnostic of ECG arrhythmia using wavelet analysis and k-nearest neighbor algorithm," in *2018 International conference on applied smart systems (ICASS)*. IEEE, 2018, pp. 1–6.

[28] K. Sanjo, K. Hebiguchi, C. Tang, E. A. Rashed, S. Kodera, H. Togo, and A. Hirata, "Sensitivity of electrocardiogram on electrode-pair locations for wearable devices: computational analysis of amplitude and waveform distortion," *Biosensors*, vol. 14, no. 3, p. 153, 2024.

[29] S. Nahak and G. Saha, "Ensembled feature based multi-label ECG arrhythmia classification," in *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–4.

[30] E. Merdjanovska and A. Rashkovska, "Patient-specific heartbeat classification in single-lead ECG using convolutional neural network," in *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2021, pp. 932–936.

[31] A. Phoemsuk and V. Abolghasemi, "Coronary artery disease classification using one-dimensional convolutional neural network," in *2024 IEEE Conference on Artificial Intelligence (CAI)*. IEEE, 2024, pp. 389–394.

[32] Z. Xiong, M. K. Stiles, and J. Zhao, "Robust ECG signal classification for detection of atrial fibrillation using a novel neural network," in *2017 Computing in Cardiology (CinC)*, 2017, pp. 1–4.

[33] W. Xie, C. Chen, R. Zhao, and Y. Lu, "Detection of atrial fibrillation based on feature fusion using attention-based BiLSTM," in *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–4.

[34] Y. Liu, H. Li, J. Lin, H. Li, H. Lei, C. Xia, C. Xiao, and B. Lei, "Gated cnn-transformer network for automatic cardiovascular diagnosis using 12-lead electrocardiogram," in *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2023, pp. 1–4.

[35] S. Compiet, R. Willemsen, K. Konings, and H. Stoffers, "Competence of general practitioners in requesting and interpreting ECGs - a case vignette study," *Netherlands Heart Journal*, vol. 26, 06 2018.

[36] K. Amini, A. Mirzaei, M. Hosseini, H. Zandian, I. Azizpour, and Y. Haghi, "Assessment of electrocardiogram interpretation competency among healthcare professionals and students of ardabil university of medical sciences: a multidisciplinary study," *BMC Medical Education*, vol. 22, no. 1, p. 448, 2022.

[37] Y. Jia, H. Pei, J. Liang, Y. Zhou, Y. Yang, Y. Cui, and M. Xiang, "Preprocessing and denoising techniques for electrocardiography and magnetocardiography: A review," *Bioengineering*, vol. 11, no. 11, p. 1109, 2024.

[38] M. U. R. Gondal, H. A. Mehdi, R. R. Khenhrani, N. Kumari, M. F. Ali, S. Kumar, M. Faraz, and J. Malik, "Role of machine learning and artificial intelligence in arrhythmias and electrophysiology," *Cardiology in review*, pp. 10–1097, 2024.

[39] A. Y. Hannun, P. Rajpurkar, M. Haghpanahi, G. H. Tison, C. Bourn, M. P. Turakhia, and A. Y. Ng, "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network," *Nature Medicine*, vol. 25, no. 1, pp. 65–69, 2019.

[40] T. Subba and T. Chingtham, "Comparative analysis of machine learning algorithms with advanced feature extraction for ECG signal classification," *IEEE Access*, vol. 12, pp. 57 727–57 740, 2024.

[41] I. Zeljković, H. Pintarić, M. Vrsalović, and I. Kruljac, "Effectiveness of cardiogoniometry compared with exercise-ECG test in diagnosing stable coronary artery disease in women," *QJM: An International Journal of Medicine*, vol. 110, no. 2, pp. 89–95, 2017.

[42] M. Dey, N. Omar, and M. A. Ullah, "Temporal feature-based classification into myocardial infarction and other CVDs merging CNN and Bi-LSTM from ECG signal," *IEEE Sensors Journal*, vol. 21, no. 19, pp. 21 688–21 695, 2021.

[43] V. Jahmunah, E. Y. K. Ng, T. R. San, and U. R. Acharya, "Automated detection of coronary artery disease, myocardial infarction and congestive heart failure using GaborCNN model with ECG signals," *Computers in Biology and Medicine*, vol. 134, p. 104457, 2021.

[44] E. Burns and K. Brush, "What is deep learning and how does it work?" https://www.techtarget.com/searchenterpriseai/definition/deep-learning-deep-neural-network., 2021.

[45] M. Rashed-Al-Mahfuz, M. A. Moni, P. Lio, S. M. S. Islam, S. Berkovsky, M. Khushi, and J. M. Quinn, "Deep convolutional neural networks based ECG beats classification to diagnose cardiovascular conditions," *Biomedical engineering letters*, vol. 11, no. 2, pp. 147–162, 2021.

[46] M. Liu and Y. Kim, "Classification of heart diseases based on ECG signals using long short-term memory," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2018, pp. 2707–2710.

[47] G. Liu, X. Han, L. Tian, W. Zhou, and H. Liu, "ECG quality assessment based on hand-crafted statistics and deep-learned s-transform spectrogram features," *Computer Methods and Programs in Biomedicine*, vol. 208, p. 106269, 2021.

[48] C.-J. Zhang, F.-Q. Tang, H.-P. Cai, and Y.-F. Qian, "Heart failure classification using deep learning to extract spatiotemporal features from ECG," *BMC Medical Informatics and Decision Making*, vol. 24, no. 1, p. 17, 2024.

[49] E. Kraev, B. Koseoglu, L. Traverso, and M. Topiwalla, "Shap-select: Lightweight feature selection using shap values and regression," *arXiv preprint arXiv:2410.06815*, 2024.

[50] K. Futagami, Y. Fukazawa, N. Kapoor, and T. Kito, "Pairwise acquisition prediction with shap value interpretation," *The Journal of Finance and Data Science*, vol. 7, pp. 22–44, 2021.

[51] R. E. Harskamp, "Electrocardiographic screening in primary care for cardiovascular disease risk and atrial fibrillation," *Primary Health Care Research & Development*, vol. 20, p. e101, 2019.

[52] H.-C. Chen and K. Shouryadhar, "Heart rate monitoring system using feature extraction in electrocardiogram signal by convolutional neural network," in *2021 12th International Conference on Ubiquitous and Future Networks (ICUFN)*. IEEE, 2021, pp. 384–388.

[53] B. Lal, D. R. Gopagoni, B. Barik, M. A. Khan, R. D. Kumar, and T. V. Lakshmi, "An efficient QRS detection and pre-processing by wavelet transform technique for classifying cardiac arrhythmia," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 11, no. 8s, pp. 490–498, 2023.

[54] M. Khalili, H. GholamHosseini, A. Lowe, and M. M. Kuo, "Motion artifacts in capacitive ECG monitoring systems: a review of existing models and reduction techniques," *Medical & Biological Engineering & Computing*, vol. 62, no. 12, pp. 3599–3622, 2024.

[55] P. M. Tripathi, A. Kumar, R. Komaragiri, and M. Kumar, "A review on computational methods for denoising and detecting ECG signals to detect cardiovascular diseases," *Archives of Computational Methods in Engineering*, vol. 29, no. 3, pp. 1875–1914, 2022.

[56] U. Satija, B. Ramkumar, and M. S. Manikandan, "A review of signal processing techniques for electrocardiogram signal quality assessment," *IEEE reviews in biomedical engineering*, vol. 11, pp. 36–52, 2018.

[57] S. Mahmoodabadi, A. Ahmadian, and M. Abolhasani, "ECG feature extraction using daubechies wavelets," in *the 5th IASTED International Conference on Visualization, Imaging, and Image Processing*, 2005, pp. 343–348.

[58] S. Lahmiri, "Comparative study of ECG signal denoising by wavelet thresholding in empirical and variational mode decomposition domains," *Healthcare Technology Letters*, vol. 1, no. 3, pp. 104–109, 2014.

[59] H. Ghonchi, S. Ferdowsi, and V. Abolghasemi, "Common spatial pattern with deep learning for fetal heart rate monitoring," in *2022 IEEE Workshop on Signal Processing Systems (SiPS)*, 2022, pp. 1–6.

[60] C.-C. Chen and F. R. Tsui, "Comparing different wavelet transforms on removing electrocardiogram baseline wanders and special trends," *BMC Medical Informatics and Decision Making*, vol. 20, no. Suppl 11, p. 343, 2020.

[61] A. Kumar, H. Tomar, V. K. Mehla, R. Komaragiri, and M. Kumar, "Stationary wavelet transform based ECG signal denoising method," *ISA Transactions*, vol. 114, pp. 251–262, 2021.

[62] A. Azzouz, B. Bengherbia, P. Wira, N. Alaoui, A. Souahlia, M. Maazouz, and H. Hentabeli, "An efficient ECG signals denoising technique based on the combination of particle swarm optimisation and wavelet transform," *Heliyon*, vol. 10, no. 5, 2024.

[63] S. M. Qaisar, "Cardiogram baseline wander and power line interference elimination by proficient adaptive-rate fir filtering," *Engineering Research Express*, vol. 2, no. 2, p. 025024, 2020.

[64] S. A. Malik, S. A. Parah, H. Aljuaid, and B. A. Malik, "An iterative filtering based ECG denoising using lifting wavelet transform technique," *Electronics*, vol. 12, no. 2, p. 387, 2023.

[65] B. Chen, Y. Li, X. Cao, W. Sun, and W. He, "Removal of power line interference from ECG signals using adaptive notch filters of sharp resolution," *IEEE access*, vol. 7, pp. 150 667–150 676, 2019.

[66] S. Mian Qaisar, "Baseline wander and power-line interference elimination of ECG signals using efficient signal-piloted filtering," *Healthcare Technology Letters*, vol. 7, no. 4, pp. 114–118, 2020.

[67] A. Dorostghol, A. Maghsoudpour, A. Ghaffari, and M. Nikkhah-Bahrami, "Power line interference and baseline wander removal from ECG signals using local characteristic decomposition," *Journal of Instrumentation*, vol. 17, no. 06, p. P06025, 2022.

[68] M. S. Ashraf and A. K. Ashraf, "Proposed network structures and combined adaptive algorithms for electrocardiogram signal denoising," *International Journal of Adaptive Control and Signal Processing*, vol. 34, no. 3, pp. 354–371, 2020.

[69] M. M. U. Faiz and I. Kale, "Removal of multiple artifacts from ECG signal using cascaded multistage adaptive noise cancellers," *Array*, vol. 14, p. 100133, 2022.

[70] S. Stalin, V. Roy, P. K. Shukla, A. Zaguia, M. M. Khan, P. K. Shukla, and A. Jain, "A machine learning-based big EEG data artifact detection and wavelet-based removal: An empirical approach," *Mathematical Problems in Engineering*, vol. 2021, no. 1, p. 2942808, 2021.

[71] D. Berwal, C. Vandana, S. Dewan, C. Jiji, and M. S. Baghini, "Motion artifact removal in ambulatory ECG signal for heart rate variability analysis," *IEEE Sensors Journal*, vol. 19, no. 24, pp. 12 432–12 442, 2019.

[72] S. Banerjee and G. K. Singh, "Monte carlo filter-based motion artifact removal from electrocardiogram signal for real-time telecardiology system," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–10, 2021.

[73] L. Xu, C. Rabotti, Y. Zhang, S. Ouzounov, P. J. Harpe, and M. Mischi, "Motion-artifact reduction in capacitive heart-rate measurements by adaptive filtering," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, no. 10, pp. 4085–4093, 2018.

[74] F. A. Castaño Usuga, C. Gissel, and A. M. Hernández, "Motion artifact reduction in electrocardiogram signals through a redundant denoising independent component analysis method for wearable health care monitoring systems: Algorithm development and validation," *JMIR medical informatics*, vol. 10, no. 11, p. e40826, 2022.

[75] C. Yang and N. Tavassolian, "An independent component analysis approach to motion noise cancelation of cardio-mechanical signals," *IEEE Transactions on Biomedical Engineering*, vol. 66, no. 3, pp. 784–793, 2018.

[76] V. G. Sirtoli, M. Liamini, L. T. Lins, M. Lessard-Tremblay, G. E. Cowan, R. J. Zednik, and G. Gagnon, "Removal of motion artifacts in capacitive electrocardiogram acquisition: A review," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 17, no. 3, pp. 394–412, 2023.

[77] S. K. Bashar, E. Ding, A. J. Walkey, D. D. McManus, and K. H. Chon, "Noise detection in electrocardiogram signals for intensive care unit patients," *IEEE Access*, vol. 7, pp. 88 357–88 368, 2019.

[78] R. Alizadehsani, M. Abdar, M. Roshanzamir, A. Khosravi, P. M. Kebria, F. Khozeimeh, S. Nahavandi, N. Sarrafzadegan, and U. R. Acharya, "Machine learning-based coronary artery disease diagnosis: A comprehensive review," *Computers in Biology and Medicine*, vol. 111, p. 103346, 2019.

[79] M. M. Rahman, M. W. Rivolta, F. Badilini, and R. Sassi, "A systematic survey of data augmentation of ECG signals for AI applications," *Sensors*, vol. 23, no. 11, p. 5237, 2023.

[80] A. Mincholé, J. Camps, A. Lyon, and B. Rodríguez, "Machine learning in the electrocardiogram," *Journal of Electrocardiology*, vol. 57, pp. S61–S64, 2019.

[81] A. D. Dolatabadi, S. E. Z. Khadem, and B. M. Asl, "Automated diagnosis of coronary artery disease (CAD) patients using optimized SVM," *Computer Methods and Programs in Biomedicine*, vol. 138, pp. 117–126, 2017.

[82] W. Zhu, X. Chen, Y. Wang, and L. Wang, "Arrhythmia recognition and classification using ECG morphology and segment feature analysis," *IEEE Transactions on Computational Biology and Bioinformatics*, vol. 16, no. 1, pp. 131–138, 2019.

[83] R. Saini, N. Bindal, and P. Bansal, "Classification of heart diseases from ECG signals using wavelet transform and kNN classifier," in *IEEE 2015 International Conference on Computing, Communication and Automation*, 2015, pp. 1208–1215.

[84] M. U. Khan, S. Aziz, S. Z. Hassan Naqvi, and A. Rehman, "Classification of coronary artery diseases using electrocardiogram signals," in *2020 International Conference on Emerging Trends in Smart Technologies (ICETST)*, 2020, pp. 1–5.

[85] M. Merino, I. M. Gomez, and A. J. Molina, "Envelopment filter and k-means for the detection of QRS waveforms in electrocardiogram," *Medical Engineering & Physics*, vol. 37, no. 6, pp. 605–609, 2015.

[86] S. R. Tithi, A. Aktar, F. Aleem, and A. Chakrabarty, "ECG data analysis and heart disease prediction using machine learning algorithms," in *2019 IEEE Region 10 Symposium (TENSYMP)*, 2019, pp. 819–824.

[87] M. Sraitih, Y. Jabrane, and A. Hajjam El Hassani, "An automated system for ECG arrhythmia detection using machine learning techniques," *Journal of Clinical Medicine*, vol. 10, no. 22, p. 5450, 2021.

[88] M. Taconné, V. D. Corino, and L. Mainardi, "An ECG-based model for left ventricular hypertrophy detection: A machine learning approach," *IEEE Open Journal of Engineering in Medicine and Biology*, 2024.

[89] U. Nagavelli, D. Samanta, and P. Chakraborty, "Machine learning technology-based heart disease detection models," *Journal of Healthcare Engineering*, vol. 2022, 2022.

[90] M. M. Suhail and T. A. Razak, "Cardiac disease detection from ECG signal using discrete wavelet transform with machine learning method," *Diabetes Research and Clinical Practice*, vol. 187, p. 109852, 2022.

[91] W. Yang, Y. Si, D. Wang, and G. Zhang, "A novel method for identifying electrocardiograms using an independent component analysis and principal component analysis network," *Measurement*, vol. 152, p. 107363, 2020.

[92] M. Zubair, R. K. Tripathy, M. Alhartomi, S. Alzahrani, S. R. Ahamed *et al.*, "Detection of sleep apnea from ECG signals using sliding singular spectrum based subpattern principal component analysis," *IEEE Transactions on Artificial Intelligence*, vol. 5, no. 6, pp. 2897–2906, 2023.

[93] X. Chen, Y. Wang, L. Wang *et al.*, "Arrhythmia recognition and classification using ECG morphology and segment feature analysis," *IEEE Transactions on Computational Biology and Bioinformatics*, vol. 16, no. 1, pp. 131–138, 2018.

[94] V.-S. Pham, A. Nguyen, H. B. Dang, H.-C. Le, and M. T. Nguyen, "Diagnosis of sudden cardiac arrest using principal component analysis in automated external defibrillators," *Scientific Reports*, vol. 13, no. 1, p. 8768, 2023.

[95] M. B. Selek, B. Yesilkaya, S. S. Egeli, and Y. Isler, "The effect of principal component analysis in the diagnosis of congestive heart failure via heart rate variability analysis," *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, vol. 235, no. 12, pp. 1479–1488, 2021.

[96] G. G. Geweid and J. D. Chen, "Automatic classification of atrial fibrillation from short single-lead ECG recordings using a hybrid approach of dual support vector machine," *Expert Systems with Applications*, vol. 198, p. 116848, 2022.

[97] Q. H. Nguyen, B. P. Nguyen, T. B. Nguyen, T. T. Do, J. F. Mbinta, and C. R. Simpson, "Stacking segment-based CNN with SVM for recognition of atrial fibrillation from single-lead ECG recordings," *Biomedical Signal Processing and Control*, vol. 68, p. 102672, 2021.

[98] M. Deng, L. Qiu, H. Wang, W. Shi, and L. Wang, "Atrial fibrillation classification using convolutional neural networks and time domain features of ECG sequence," in *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2020, pp. 1481–1485.

[99] E. H. Houssein, I. E. Ibrahim, N. Neggaz, M. Hassaballah, and Y. M. Wazery, "An efficient ECG arrhythmia classification method based on manta ray foraging optimization," *Expert Systems with Applications*, vol. 181, p. 115131, 2021.

[100] M. Naz, J. H. Shah, M. A. Khan, M. Sharif, M. Raza, and R. Damaševičius, "From ECG signals to images: a transformation based approach for deep learning," *PeerJ Computer Science*, vol. 7, p. e386, 2021.

[101] S. Celin and K. Vasanth, "A novel method for ECG classification using polynomial based curve fitting," in *2019 IEEE International Conference on Electrical, Computer and Communication Technologies (ICECCT)*, 2019, pp. 1–9.

[102] Z. Wang, L. Qian, C. Han, and L. Shi, "Application of multi-feature fusion and random forests to the automated detection of myocardial infarction," *Cognitive Systems Research*, vol. 59, pp. 15–26, 2020.

[103] M. K. Ojha, S. Wadhwani, A. K. Wadhwani, and A. Shukla, "Automatic detection of arrhythmias from an ECG signal using an autoencoder and SVM classifier," *Physical and Engineering Sciences in Medicine*, vol. 45, no. 2, pp. 665–674, 2022.

[104] M. Hammad, A. M. Iliyasu, A. Subasi, E. S. Ho, and A. A. Abd El-Latif, "A multitier deep learning model for arrhythmia detection," *IEEE Transactions on Instrumentation and Measurement*, vol. 70, pp. 1–9, 2020.

[105] V. Mazaheri and H. Khodadadi, "Heart arrhythmia diagnosis based on the combination of morphological, frequency and nonlinear features of ECG signals and metaheuristic feature selection algorithm," *Expert Systems with Applications*, vol. 161, p. 113697, 2020.

[106] S. K. Bashar, E. Ding, D. Albuquerque, M. Winter, S. Binici, A. J. Walkey, D. D. McManus, and K. H. Chon, "Atrial fibrillation detection in ICU patients: A pilot study on MIMIC III data," in *2019 41st Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2019, pp. 298–301.

[107] I. Saini, D. Singh, and A. Khosla, "QRS detection using k-nearest neighbor algorithm (KNN) and evaluation on standard ECG databases," *Journal of Advanced Research*, vol. 4, no. 4, pp. 331–344, 2013.

[108] S. M. Qaisar, M. Krichen, and F. Jallouli, "Multirate ECG processing and k-nearest neighbor classifier based efficient arrhythmia diagnosis," in *International Conference on Smart Homes and Health Telematics*. Springer, 2020, pp. 329–337.

[109] S. Yu, "Research on adaptive denoising method for ECG signal based on k-means algorithm," in *2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications (AIITA)*. IEEE, 2025, pp. 178–181.

[110] B.-Y. Mo, S. Nuannimnoi, A. Baskoro, A. Khan, J. Ariesta Dwi Pratiwi, and C.-Y. Huang, "Clusteredshap: Faster gradientexplainer based on k-means clustering and selections of gradients in explaining 12-lead ECG classification model," in *Proceedings of the 13th International Conference on Advances in Information Technology*, 2023, pp. 1–8.

[111] W. Yang, Y. Si, D. Wang, G. Zhang, X. Liu, and L. Li, "Automated intra-patient and inter-patient coronary artery disease and congestive heart failure detection using EFAP-Net," *Knowledge-Based Systems*, vol. 201, p. 106083, 2020.

[112] M. Liu and Y. Kim, "Classification of heart diseases based on ECG signals using long short-term memory," in *2018 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2018, pp. 2707–2710.

[113] Y. Obeidat and A. M. Alqudah, "A hybrid lightweight 1D CNN-LSTM architecture for automated ECG beat-wise classification." *Traitement du Signal*, vol. 38, no. 5, 2021.

[114] S. Singh, S. K. Pandey, U. Pawar, and R. R. Janghel, "Classification of ECG arrhythmia using recurrent neural networks," *Procedia Computer Science*, vol. 132, pp. 1290–1297, 2018.

[115] Y. Guan, Y. An, J. Xu, N. Liu, and J. Wang, "HA-ResNet: Residual neural network with hidden attention for ECG arrhythmia detection using two-dimensional signal," *IEEE Transactions on Computational Biology and Bioinformatics*, vol. 20, no. 6, pp. 3389–3398, 2023.

[116] S. Yang, C. Lian, Z. Zeng, B. Xu, J. Zang, and Z. Zhang, "A multi-view multi-scale neural network for multi-label ECG classification," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 7, no. 3, pp. 648–660, 2023.

[117] A. Ballas and C. Diou, "Towards domain generalization for ECG and EEG classification: Algorithms and benchmarks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 8, no. 1, pp. 44–54, 2024.

[118] S. Nurmaini, A. Darmawahyuni, A. N. Sakti Mukti, M. N. Rachmatullah, F. Firdaus, and B. Tutuko, "Deep learning-based stacked denoising and autoencoder for ECG heartbeat classification," *Electronics*, vol. 9, no. 1, 2020. [Online]. Available: https://www.mdpi.com/2079-9292/9/1/135

[119] B. Saha Tchinda and D. Tchiotsop, "A lightweight 1D convolutional neural network model for arrhythmia diagnosis from electrocardiogram signal," *Physical and Engineering Sciences in Medicine*, pp. 1–13, 2025.

[120] R. T. Nguimdo and A. Tiedeu, "Automated ECG analysis with dual-channel squeezenet using hilbert huang transform, fuzzy entropy, and hurst exponent," *Circuits, Systems, and Signal Processing*, pp. 1–26, 2025.

[121] N. Nonaka and J. Seita, "Electrocardiogram classification by modified efficientnet with data augmentation," in *2020 Computing in Cardiology*, 2020, pp. 1–4.

[122] H. Bechinia, D. Benmerzoug, and N. Khlifa, "Approach based lightweight custom convolutional neural network and fine-tuned mobilenet-v2 for ECG arrhythmia signals classification," *IEEE Access*, vol. 12, pp. 40 827–40 841, 2024.

[123] H. Tesfai, H. Saleh, M. Al-Qutayri, M. B. Mohammad, T. Tekeste, A. Khandoker, and B. Mohammad, "Lightweight shufflenet based CNN for arrhythmia classification," *IEEE Access*, vol. 10, pp. 111 842–111 854, 2022.

[124] G. Altan, "Diagnosis of coronary artery disease using deep belief networks," *European Journal of Engineering and Natural Sciences*, vol. 2, no. 1, pp. 29–36, 2017.

[125] N. Papandrianos and E. Papageorgiou, "Automatic diagnosis of coronary artery disease in spect myocardial perfusion imaging employing deep learning," *Applied Sciences*, vol. 11, no. 14, p. 6362, 2021.

[126] T. Du, X. Liu, H. Zhang, and B. Xu, "Real-time lesion detection of cardiac coronary artery using deep neural networks," in *2018 International Conference on Network Infrastructure and Digital Content (IC-NIDC)*. IEEE, 2018, pp. 150–154.

[127] I. Bakkouri and K. Afdel, "MLCA2F: Multi-level context attentional feature fusion for COVID-19 lesion segmentation from CT scans," *Signal, Image and Video Processing*, vol. 17, no. 4, pp. 1181–1188, 2023.

[128] I. Bakkouri, K. Afdel, J. Benois-Pineau, and G. C. F. t. A. D. N. Initiative, "BG-3DM2F: bidirectional gated 3D multi-scale feature fusion for alzheimer's disease diagnosis," *Multimedia Tools and Applications*, vol. 81, no. 8, pp. 10 743–10 776, 2022.

[129] I. Bakkouri and K. Afdel, "Computer-aided diagnosis (CAD) system based on multi-layer feature fusion network for skin lesion recognition in dermoscopy images," *Multimedia Tools and Applications*, vol. 79, no. 29, pp. 20 483–20 518, 2020.

[130] ——, "Multi-scale CNN based on region proposals for efficient breast abnormality recognition," *Multimedia Tools and Applications*, vol. 78, no. 10, pp. 12 939–12 960, 2019.

[131] A. Saglietto, D. Baccega, R. Esposito, M. Anselmino, V. Dusi, A. Fiandrotti, and G. M. De Ferrari, "Convolutional neural network (CNN)-enabled electrocardiogram (ECG) analysis: a comparison between standard twelve-lead and single-lead setups," *Frontiers in Cardiovascular Medicine*, vol. 11, p. 1327179, 2024.

[132] O. S. Lih, V. Jahmunah, T. R. San, E. J. Ciaccio, T. Yamakawa, M. Tanabe, M. Kobayashi, O. Faust, and U. R. Acharya, "Comprehensive electrocardiographic diagnosis based on deep learning," *Artificial Intelligence in Medicine*, vol. 103, p. 101789, 2020.

[133] H. Li, X. Wang, C. Liu, Y. Wang, P. Li, H. Tang, L. Yao, and H. Zhang, "Dual-input neural network integrating feature extraction and deep learning for coronary artery disease detection using electrocardiogram and phonocardiogram," *IEEE Access*, vol. 7, pp. 146 457–146 469, 2019.

[134] U. R. Acharya, H. Fujita, O. S. Lih, M. Adam, J. H. Tan, and C. K. Chua, "Automated detection of coronary artery disease using different durations of ECG

segments with convolutional neural network," *Knowledge-Based Systems*, vol. 132, pp. 62–71, 2017.

[135] D. Li, J. Zhang, Q. Zhang, and X. Wei, "Classification of ECG signals based on 1D convolution neural network," in *2017 IEEE 19th International Conference on e-Health Networking, Applications and Services (Healthcom)*, 2017, pp. 1–6.

[136] J. Cheng, Q. Zou, and Y. Zhao, "ECG signal classification based on deep CNN and BiLSTM," *BMC Medical Informatics and Decision Making*, vol. 21, pp. 1–12, 2021.

[137] M. F. Issa, A. Yousry, G. Tuboly, Z. Juhasz, A. H. AbuEl-Atta, and M. M. Selim, "Heartbeat classification based on single lead-II ECG using deep learning," *Heliyon*, vol. 9, no. 7, 2023.

[138] A. Rajkumar, M. Ganesan, and R. Lavanya, "Arrhythmia classification on ECG using deep learning," in *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*. IEEE, 2019, pp. 365–369.

[139] G. Petmezas, K. Haris, L. Stefanopoulos, V. Kilintzis, A. Tzavelis, J. A. Rogers, A. K. Katsaggelos, and N. Maglaveras, "Automated atrial fibrillation detection using a hybrid CNN-LSTM network on imbalanced ECG datasets," *Biomedical Signal Processing and Control*, vol. 63, p. 102194, 2021.

[140] P. Xu, H. Liu, X. Xie, S. Zhou, M. Shu, and Y. Wang, "Interpatient ECG arrhythmia detection by residual attention CNN," *Computational and Mathematical Methods in Medicine*, vol. 2022, no. 1, p. 2323625, 2022.

[141] J. Wang, J. Zang, Q. An, H. Wang, and Z. Zhang, "A pooling convolution model for multi-classification of ECG and PCG signals," *Computer Methods in Biomechanics and Biomedical Engineering*, pp. 1–14, 2023.

[142] M. Wasimuddin, K. Elleithy, A. Abuzneid, M. Faezipour, and O. Abuzaghleh, "Multiclass ECG signal analysis using global average-based 2-D convolutional neural network modeling," *Electronics*, vol. 10, no. 2, p. 170, 2021.

[143] P. Xiong, S. M.-Y. Lee, and G. Chan, "Deep learning for detecting and locating myocardial infarction by electrocardiogram: A literature review," *Frontiers in Cardiovascular Medicine*, vol. 9, p. 860032, 2022.

[144] G. Al Hinai, S. Jammoul, Z. Vajihi, and J. Afilalo, "Deep learning analysis of resting electrocardiograms for the detection of myocardial dysfunction, hypertrophy, and ischaemia: a systematic review," *European Heart Journal-Digital Health*, vol. 2, no. 3, pp. 416–423, 2021.

[145] J.-Y. Sun, H. Shen, Q. Qu, W. Sun, and X.-Q. Kong, "The application of deep learning in electrocardiogram: Where we came from and where we should go?" *International Journal of Cardiology*, vol. 337, pp. 71–78, 2021.

[146] A. K. Singh and S. Krishnan, "ECG signal feature extraction trends in methods and applications," *BioMedical Engineering OnLine*, vol. 22, no. 1, p. 22, 2023.

[147] E. Prabhakararao and S. Dandapat, "Congestive heart failure detection from ECG signals using deep residual neural network," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 53, no. 5, pp. 3008–3018, 2023.

[148] E. Butun, O. Yildirim, M. Talo, R.-S. Tan, and U. R. Acharya, "1D-CADCapsNet: One dimensional deep capsule networks for coronary artery disease detection using ECG signals," *Physica Medica*, vol. 70, pp. 39–48, 2020.

[149] F. Corradi, J. Buil, H. D. Cannière, W. Groenendaal, and P. Vandervoort, "Real time electrocardiogram annotation with a long short term memory neural network," in *2019 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, 2019, pp. 1–4.

[150] T. Zhu, W. Luo, and F. Yu, "Convolution-and attention-based neural network for automated sleep stage classification," *International Journal of Environmental Research and Public Health*, vol. 17, no. 11, p. 4152, 2020.

[151] A. Mukhtar, S. A. Malik, H. Asjad, S. Mufty, M. Tahir, and J. Sheikh, "Hybrid LSTM-Attention framework for enhanced classification of cardiac abnormalities from ECG signals," in *2025 International Conference on Emerging Technologies in Electronics, Computing, and Communication (ICETECC)*, 2025, pp. 1–6.

[152] K.-O. Cho and H.-J. Jang, "Comparison of different input modalities and network structures for deep learning-based seizure detection," *Scientific Reports*, vol. 10, no. 1, p. 122, 2020.

[153] K. C. Vidyasagar, K. R. Kumar, G. A. Sai, M. Ruchita, and M. J. Saikia, "Signal to image conversion and convolutional neural networks for physiological signal processing: A review," *IEEE Access*, 2024.

[154] R. Trivedi, I. Chiu, and D. Ouyang, "Multi-modality deep learning model for prediction of chronic obstructive coronary artery disease," *European Heart Journal*, vol. 45, no. Supplement_1, pp. ehae666–1364, 2024.

[155] Y. Yue and X. Zhu, "Automated coronary artery disease detection using deep learning on ECG datasets," in *Proceedings of the 2023 3rd International Conference on Bioinformatics and Intelligent Computing*, 2023, pp. 242–245.

[156] M. Leasure, U. Jain, A. Butchy, J. Otten, V. A. Covalesky, D. McCormick, and G. S. Mintz, "Deep learning algorithm predicts angiographic coronary artery disease in stable patients using only a standard 12-lead electrocardiogram," *Canadian Journal of Cardiology*, vol. 37, no. 11, pp. 1715–1724, 2021.

[157] Y. Guo, D. Zhou, X. Ruan, and J. Cao, "Variational gated autoencoder-based feature extraction model for inferring disease-mirna associations based on multiview features," *Neural Networks*, vol. 165, pp. 491–505, 2023.

[158] X. Xu and H. Liu, "ECG heartbeat classification using convolutional neural networks," *IEEE Access*, vol. 8, pp. 8614–8619, 2020.

[159] T. J. Jun, H. M. Nguyen, D. Kang, D. Kim, D. Kim, and Y.-H. Kim, "ECG arrhythmia classification using a 2-D convolutional neural network," *arXiv preprint arXiv:1804.06812*, 2018.

[160] A. Ullah, S. u. Rehman, S. Tu, R. M. Mehmood, and M. Ehatisham-Ul-Haq, "A hybrid deep CNN model for abnormal arrhythmia detection based on cardiac ECG signal," *Sensors*, vol. 21, no. 3, p. 951, 2021.

[161] J. Li, Y. Si, T. Xu, and S. Jiang, "Deep convolutional neural network based ECG classification system using information fusion and one-hot encoding techniques," *Mathematical Problems in Engineering*, vol. 2018, no. 1, p. 7354081, 2018.

[162] Z. Xiong, M. K. Stiles, and J. Zhao, "Robust ECG signal classification for detection of atrial fibrillation using a novel neural network," in *2017 Computing in Cardiology (CinC)*. IEEE, 2017, pp. 1–4.

[163] S. Kim, S. Chon, J.-K. Kim, J. Kim, Y. Gil, and S. Jung, "Lightweight convolutional neural network for real-time arrhythmia classification on low-power wearable electrocardiograph," in *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, 2022, pp. 1915–1918.

[164] A. Phoemsuk and V. Abolghasemi, "CADNet: A lightweight neural network for coronary artery disease classification using electrocardiogram signals," *IEEE Journal of Biomedical and Health Informatics*, 2025.

[165] A. J. Prakash and M. Atef, "Multi-scale feature extraction for ECG beat classification using a cnn-transformer network with imbalance mitigation," in *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2025, pp. 1–5.

[166] R. Hu, J. Chen, and L. Zhou, "A transformer-based deep neural network for arrhythmia detection using continuous ECG signals," *Computers in Biology and Medicine*, vol. 144, p. 105325, 2022.

[167] A. Shah, D. Singh, H. G. Mohamed, S. Bharany, A. U. Rehman, and S. Hussen, "Electrocardiogram analysis for cardiac arrhythmia classification and prediction through self attention based auto encoder," *Scientific Reports*, vol. 15, no. 1, p. 9230, 2025.

[168] S. Kim, J. Lim, and J. Jang, "Patient-adaptive beat-wise temporal transformer for atrial fibrillation classification in continuous long-term cardiac monitoring," *IEEE Access*, 2024.

[169] H. Tung, C. Zheng, X. Mao, and D. Qian, "Multi-lead ECG classification via an information-based attention convolutional neural network," *Journal of Shanghai Jiaotong University (Science)*, vol. 27, no. 1, pp. 55–69, 2022.

[170] Y. Xing, L. Zhang, Z. Hou, X. Li, Y. Shi, Y. Yuan, F. Zhang, S. Liang, Z. Li, and L. Yan, "Accurate ECG classification based on spiking neural network and attentional mechanism for real-time implementation on personal portable devices," *Electronics*, vol. 11, no. 12, p. 1889, 2022.

[171] S. Wang, R. Li, X. Wang, S. Shen, B. Zhou, and Z. Wang, "Multiscale residual network based on channel spatial attention mechanism for multilabel ECG classification," *Journal of Healthcare Engineering*, vol. 2021, no. 1, p. 6630643, 2021.

[172] Y. Huang, W. Liu, Z. Yin, S. Hu, M. Wang, and W. Cai, "ECG classification based on guided attention mechanism," *Computer Methods and Programs in Biomedicine*, vol. 257, p. 108454, 2024.

[173] H. Mewada, "2D-wavelet encoded deep CNN for image-based ECG classification," *Multimedia Tools and Applications*, vol. 82, no. 13, p. 20553–20569, Jan. 2023.

[174] F. DW, D. TG, A. YM, K. SR, and A. TF, "Lightweight multireceptive field CNN for 12-lead ECG signal classification," *Computational Intelligence and Neuroscience*, vol. 2022, p. 8413294, 2022.

[175] H. Mewada and I. M. Pires, "Electrocardiogram signal classification using lightweight DNN for mobile devices," *Procedia Computer Science*, vol. 224, pp. 558–564, 2023.

[176] K. Qin, W. Huang, T. Zhang, H. Zhang, and X. Cheng, "A lightweight SelfONN model for general ECG classification with pretraining," *Biomedical Signal Processing and Control*, vol. 89, p. 105780, 2024.

[177] Y. Cao, T. Wei, B. Zhang, N. Lin, J. J. Rodrigues, J. Li, and D. Zhang, "ML-Net: Multi-channel lightweight network for detecting myocardial infarction," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 10, pp. 3721–3731, 2021.

[178] N. Rahuja and S. K. Valluru, "A comparative analysis of deep neural network models using transfer learning for electrocardiogram signal classification," in *2021 International Conference on Recent Trends on Electronics, Information, Communication & Technology (RTEICT)*, 2021, pp. 285–290.

[179] N. Nonaka and J. Seita, "Electrocardiogram classification by modified efficientnet with data augmentation," in *2020 Computing in Cardiology*, 2020, pp. 1–4.

[180] M. Tang, Y. Fan, J. Huang, P. Gao, and Y. Zhan, "Classification of arrhythmia based on the fusion of mobilenet and BiLSTM models," in *2023 15th International Conference on Intelligent Human-Machine Systems and Cybernetics (IHMSC)*, 2023, pp. 45–48.

[181] J. Xiao, J. Liu, H. Yang, Q. Liu, N. Wang, Z. Zhu, Y. Chen, Y. Long, L. Chang, L. Zhou, and J. Zhou, "Ulecgnet: An ultra-lightweight end-to-end ECG classification neural network," *IEEE Journal of Biomedical and Health Informatics*, vol. 26, no. 1, pp. 206–217, 2022.

[182] T. Chen, E. Mazomenos, K. Maharatna, S. Dasmahapatra, and M. Niranjan, "On the trade-off of accuracy and computational complexity for classifying normal and abnormal ECG in remote cvd monitoring systems," in *2012 IEEE Workshop on Signal Processing Systems*, 2012, pp. 37–42.

[183] Z. Baig, S. Nasir, R. A. Khan, and M. Z. U. Haque, "ArrhythmiaVision: Resource-conscious deep learning models with visual explanations for ECG arrhythmia classification," *arXiv preprint arXiv:2505.03787*, 2025.

[184] Y.-L. Xie, X.-R. Lin, C.-Y. Lee, and C.-W. Lin, "Design and implementation of an ARM-based AI module for ectopic beat classification using custom and structural pruned lightweight CNN," *IEEE Sensors Journal*, vol. 24, no. 12, pp. 19 834–19 844, 2024.

[185] P. Busia, M. A. Scrugli, V. J.-B. Jung, L. Benini, and P. Meloni, "A tiny transformer for low-power arrhythmia classification on microcontrollers," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 19, no. 1, pp. 142–152, 2024.

[186] J. Lu, D. Liu, Z. Liu, X. Cheng, L. Wei, C. Zhang, X. Zou, and B. Liu, "Efficient hardware architecture of convolutional neural network for ECG classification in wearable healthcare device," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 68, no. 7, pp. 2976–2985, 2021.

[187] S. Saadatnejad, M. Oveisi, and M. Hashemi, "LSTM-based ECG classification for continuous monitoring on personal wearable devices," *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 2, pp. 515–523, 2019.

[188] H. D. M. Ribeiro, A. Arnold, J. P. Howard, M. J. Shun-Shin, Y. Zhang, D. P. Francis, P. B. Lim, Z. Whinnett, and M. Zolgharni, "ECG-based real-time arrhythmia monitoring using quantized deep neural networks: A feasibility study," *Computers in Biology and Medicine*, vol. 143, p. 105249, 2022.

[189] Z. Huang, L. F. Herbozo Contreras, W. H. Leung, L. Yu, N. D. Truong, A. Nikpour, and O. Kavehei, "Efficient edge-AI models for robust ECG abnormality detection on resource-constrained hardware," *Journal of Cardiovascular Translational Research*, vol. 17, no. 4, pp. 879–892, 2024.

[190] S. T. Vasudeva, S. S. Rao, N. K. Panambur, A. K. Shettigar, C. Mahabala, P. Kamath, M. P. Gowdru Chandrashekarappa, and E. Linul, "Development of a convolutional neural network model to predict coronary artery disease based on single-lead and twelve-lead ECG signals," *Applied Sciences*, vol. 12, no. 15, p. 7711, 2022.

[191] R. K. Tripathy, A. Bhattacharyya, and R. B. Pachori, "Localization of myocardial infarction from multi-lead ECG signals using multiscale analysis and convolutional neural network," *IEEE Sensors Journal*, vol. 19, no. 23, pp. 11 437–11 448, 2019.

[192] W. Liu, M. Zhang, Y. Zhang, Y. Liao, Q. Huang, S. Chang, H. Wang, and J. He, "Real-time multilead convolutional neural network for myocardial infarction detection," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 5, pp. 1434–1444, 2017.

[193] Q. Xiao, K. Lee, S. A. Mokhtar, I. Ismail, A. L. b. M. Pauzi, Q. Zhang, and P. Y. Lim, "Deep learning-based ECG arrhythmia classification: A systematic review," *Applied Sciences*, vol. 13, no. 8, p. 4964, 2023.

[194] D. K. McClish, "Analyzing a portion of the ROC curve," *Medical Decision Making*, vol. 9, no. 3, pp. 190–195, 1989.

[195] X. Fan, Q. Yao, Y. Cai, F. Miao, F. Sun, and Y. Li, "Multiscaled fusion of deep convolutional neural networks for screening atrial fibrillation from single lead short ECG recordings," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 6, pp. 1744–1753, 2018.

[196] U. B. Baloglu, M. Talo, O. Yildirim, R. S. Tan, and U. R. Acharya, "Classification of myocardial infarction with multi-lead ECG signals and deep cnn," *Pattern Recognition Letters*, vol. 122, pp. 23–30, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S016786551930056X

[197] V. Gupta, M. Mittal, V. Mittal, and N. K. Saxena, "A critical review of feature extraction techniques for ECG signal analysis," *Journal of The Institution of Engineers (India): Series B*, vol. 102, pp. 1049–1060, 2021.

[198] S. Karpagachelvi, M. Arthanari, and M. Sivakumar, "ECG feature extraction techniques-a survey approach," *arXiv preprint arXiv:1005.0957*, 2010.

[199] A. E. Johnson, T. J. Pollard, L. Shen, L.-w. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. Anthony Celi, and R. G. Mark, "Mimic-iii, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, pp. 1–9, 2016.

[200] N. Iyengar, C. Peng, R. Morin, A. L. Goldberger, and L. A. Lipsitz, "Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 271, no. 4, pp. R1078–R1084, 1996.

[201] V. Tihonenko, A. Khaustov, S. Ivanov, A. Rivin, and E. Yakushenko, "St petersburg incart 12-lead arrhythmia database," *PhysioBank PhysioToolkit and PhysioNet*, 2008.

[202] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, M. Adam, A. Gertych, and R. San Tan, "A deep convolutional neural network model to classify heartbeats," *Computers in Biology and Medicine*, vol. 89, pp. 389–396, 2017.

[203] J. Zhang, D. Liang, A. Liu, M. Gao, X. Chen, X. Zhang, and X. Chen, "MLBF-Net: A multi-lead-branch fusion network for multi-class arrhythmia classification using 12-lead ECG," *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 9, pp. 1–11, 2021.

[204] B. Wang, G. Chen, L. Rong, Y. Liu, A. Yu, X. He, T. Wen, Y. Zhang, and B. Hu, "Arrhythmia disease diagnosis based on ECG time–frequency domain fusion and convolutional neural network," *IEEE Journal of Translational Engineering in Health and Medicine*, vol. 11, pp. 116–125, 2023.

[205] A. A. Ahmed, W. Ali, T. A. Abdullah, and S. J. Malebary, "Classifying cardiac arrhythmia from ECG signal using 1d cnn deep learning model," *Mathematics*, vol. 11, no. 3, p. 562, 2023.

[206] A. Phoemsuk and V. Abolghasemi, "Multi-disease cardiovascular detection from ECG signals using an attention-driven deep network," in *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2025.

[207] U. R. Acharya, Y. Hagiwara, J. E. W. Koh, S. L. Oh, J. H. Tan, M. Adam, and R. San Tan, "Entropies for automated detection of coronary artery disease using ECG signals: A review," *Biocybernetics and Biomedical Engineering*, vol. 38, no. 2, pp. 373–384, 2018.

[208] P. Tang, Q. Wang, H. Ouyang, S. Yang, and P. Hua, "The feasibility of early detecting coronary artery disease using deep learning-based algorithm based on electrocardiography," *Aging (Albany NY)*, vol. 15, no. 9, p. 3524, 2023.

[209] K. Karboub, M. Tabaa, F. Monteiro, S. Dellagi, F. Moutaouakkil, and A. Dandache, "Automated diagnosis system for outpatients and inpatients with cardiovascular diseases," *IEEE Sensors Journal*, vol. 21, no. 2, pp. 1935–1946, 2021.

[210] R. Veerarajan and V. V, "Multimodal contrastive learning for optimized integration of ECG signals and clinical reports for arrhythmia classification," *International Journal of Basic and Applied Sciences*, vol. 14, pp. 453–463, 05 2025.

[211] S. K. Bashar, M. B. Hossain, E. Ding, A. J. Walkey, D. D. McManus, and K. H. Chon, "Atrial fibrillation detection during sepsis: study on MIMIC III ICU data," *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 11, pp. 3124–3135, 2020.

[212] S. Seoni, F. Molinari, U. R. Acharya, O. S. Lih, P. D. Barua, S. García, and M. Salvi, "Application of spatial uncertainty predictor in CNN-BiLSTM model using coronary artery disease ECG signals," *Information Sciences*, vol. 665, p. 120383, 2024.

[213] M. Sharma and U. R. Acharya, "A new method to identify coronary artery disease with ECG signals and time-frequency concentrated antisymmetric biorthogonal wavelet filter bank," *Pattern Recognition Letters*, vol. 125, pp. 235–240, 2019.

[214] R. Holgado-Cuadrado, C. Plaza-Seco, L. Lovisolo, and M. Blanco-Velasco, "Characterization of noise in long-term ECG monitoring with machine learning based on clinical criteria," *Medical & Biological Engineering & Computing*, pp. 1–14, 2023.

[215] R. Alcaraz, D. Abásolo, R. Hornero, and J. J. Rieta, "Optimal parameters study for sample entropy-based atrial fibrillation organization analysis," *Computer Methods and Programs in Biomedicine*, vol. 99, no. 1, pp. 124–132, 2010.

[216] H. Animikha, "ECG based atrial fibrillation classifier using 1D CNNs," 2021. [Online]. Available: https://github.com/animikhaich/ECG-Atrial-Fibrillation-Classification-Using-CNN

[217] A. L. Goldberger, L. A. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, J. E. Mietus, G. B. Moody, C.-K. Peng, and H. E. Stanley, "Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.

[218] P. Wagner, N. Strodthoff, R.-D. Bousseljot, D. Kreiseler, F. I. Lunze, W. Samek, and T. Schaeffter, "PTB-XL, a large publicly available electrocardiography dataset," *Scientific data*, vol. 7, no. 1, pp. 1–15, 2020.

[219] J. S. Richman and J. R. Moorman, "Physiological time-series analysis using approximate entropy and sample entropy," *American Journal of Physiology-Heart and Circulatory Physiology*, vol. 278, no. 6, pp. H2039–H2049, 2000.

[220] D. Fernandes, "ECG time-series classification." [Online]. Available: https://github.com/dave-fernandes/ECGClassifier

[221] M. U. Zahid, S. Kiranyaz, T. Ince, O. C. Devecioglu, M. E. H. Chowdhury, A. Khandakar, A. Tahir, and M. Gabbouj, "Robust R-peak detection in low-quality holter ECGs using 1D convolutional neural network," *IEEE Transactions on Biomedical Engineering*, vol. 69, no. 1, pp. 119–128, 2022.

[222] E. De Giovanni, T. Teijeiro, G. P. Millet, and D. Atienza, "Adaptive R-peak detection on wearable ECG sensors for high-intensity exercise," *IEEE Transactions on Biomedical Engineering*, vol. 70, no. 3, pp. 941–953, 2023.

[223] D. Lee, S. Lee, S. Oh, and D. Park, "Energy-efficient fpga accelerator with fidelity-controllable sliding-region signal processing unit for abnormal ECG diagnosis on iot edge devices," *IEEE Access*, vol. 9, pp. 122 789–122 800, 2021.

[224] S. Ran, X. Yang, M. Liu, Y. Zhang, C. Cheng, H. Zhu, and Y. Yuan, "Homecare-oriented ECG diagnosis with large-scale deep neural network for continuous monitoring on embedded devices," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–13, 2022.

[225] A. Giorgio, C. Guaragnella, and M. Rizzi, "FPGA-based decision support system for ECG analysis," *Journal of Low Power Electronics and Applications*, vol. 13, no. 1, p. 6, 2023.

[226] S. Raj, "An efficient IoT-based platform for remote real-time cardiac activity monitoring," *IEEE Transactions on Consumer Electronics*, vol. 66, no. 2, pp. 106–114, 2020.

[227] L. Smital, C. R. Haider, M. Vitek, P. Leinveber, P. Jurak, A. Nemcova, R. Smisek, L. Marsanova, I. Provaznik, C. L. Felton *et al.*, "Real-time quality assessment of

long-term ECG signals recorded by wearables in free-living conditions," *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 10, pp. 2721–2734, 2020.

[228] D. Zhang, S. Yang, X. Yuan, and P. Zhang, "Interpretable deep learning for automatic diagnosis of 12-lead electrocardiogram," *Iscience*, vol. 24, no. 4, 2021.

[229] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 7132–7141.

[230] Z. Jia, T. Zhou, Z. Yan, J. Hu, and Y. Shi, "Personalized meta-federated learning for IoT-enabled health monitoring," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 2024.

[231] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," *arXiv preprint arXiv:1711.05101*, 2017.

[232] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.

[233] S. S. Sumit, S. Anavatti, M. Tahtali, S. Mirjalili, and U. Turhan, "Resnet-lite: On improving image classification with a lightweight network," *Procedia Computer Science*, vol. 246, pp. 1488–1497, 2024.

[234] J. Zheng, H. Guo, and H. Chu, "A large scale 12-lead electrocardiogram database for arrhythmia study (version 1.0.0)," PhysioNet, 2022. [Online]. Available: https://doi.org/10.13026/wgex-er52

[235] X. Wang, C. Ma, X. Zhang, H. Gao, G. D. Clifford, and C. Liu, "Paroxysmal atrial fibrillation events detection from dynamic ECG recordings: The 4th china physiological signal challenge 2021 (version 1.0.0)," PhysioNet, 2021. [Online]. Available: https://physionet.org/content/challenge-2021/1.0.0/

[236] N. Iyengar, C. Peng, R. Morin, A. L. Goldberger, and L. A. Lipsitz, "Age-related alterations in the fractal scaling of cardiac interbeat interval dynamics," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 271, no. 4, pp. R1078–R1084, 1996.

[237] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

[238] A. Aggarwal, S. Srivastava, and M. Velmurugan, "Newer perspectives of coronary artery disease in young," *World Journal of Cardiology*, vol. 8, no. 12, p. 728, 2016.

[239] B. B. Borah, K. Das, U. Hazarika, and S. Roy, "A deep knowledge distillation heartcare framework for detection of multi-label myocardial infarction from multi-lead ecg signals," *SN Computer Science*, vol. 6, no. 1, p. 14, 2024.

[240] Y.-L. Xie and C.-W. Lin, "YOLO-ResTinyECG: ECG-based lightweight embedded AI arrhythmia small object detector with pruning methods," *Expert Systems with Applications*, vol. 263, p. 125691, 2025.

[241] X. Xu, Y. Suo, Q. Cai, Y. Zhang, Q. Chen, C. Gao, G. Wang, and Y. Zhao, "CardioLike-Net: An edge-end inter-patient arrhythmia classifier with quantization-aware-training for wearable ECG applications," *Biomedical Signal Processing and Control*, vol. 112, p. 108604, 2026.

[242] B. Neeraja, E. Satish, A. Alkwzahy, G. B. Priyanka *et al.*, "Lightweight Neural Compression with Quantization-Aware Attention for Real-Time Health Analytics in Internet-of-Things," in *2025 Third International Conference on Networks, Multimedia and Information Technology (NMITCON)*. IEEE, 2025, pp. 1–6.

[243] A. Anand, T. Kadian, M. K. Shetty, and A. Gupta, "Explainable AI decision model for ECG data of cardiac disorders," *Biomedical Signal Processing and Control*, vol. 75, p. 103584, 2022.