



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: VII Month of publication: July 2025

DOI: <https://doi.org/10.22214/ijraset.2025.73108>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

Cardiovascular Disease Detection in ECG Images: A Comprehensive Analysis of Machine Learning and Deep Learning Approaches

Mr. Nilesh Dagadu Navale¹, Dr. Tripti Arjariya²

¹Bhabha University, Bhopal

²Professor and Principal, BERI, Bhabha University

Abstract: Cardiovascular diseases (CVDs) represent a major global health burden, necessitating accurate and early diagnosis for effective treatment. Traditional methods of ECG analysis rely heavily on expert interpretation, which is often time-consuming and prone to variability. This study proposes a hybrid deep learning model that combines visual analysis of electrocardiogram (ECG) images with structured clinical data to enhance cardiovascular disease detection. Convolutional neural networks (CNNs) such as AlexNet and SqueezeNet are utilized to extract salient features from ECG images, while a fully connected neural network (FCNN) processes patient metadata, including age, blood pressure, cholesterol levels, and medical history. The fusion of image-based and tabular data within a unified model yields a more holistic diagnostic system. The proposed model demonstrates improved classification accuracy and robustness compared to traditional approaches. With its potential to support early diagnosis and personalized treatment planning, this work contributes significantly toward intelligent healthcare solutions. Cardiovascular diseases remain a leading cause of mortality worldwide, necessitating accurate and early diagnosis. This study proposes a hybrid deep learning framework that combines ECG image analysis with structured patient medical data to improve the classification of CVDs. Convolutional neural networks (CNNs) are used to extract features from ECG images, while a fully connected neural network (FCNN) processes clinical data such as age, blood pressure, and cholesterol. The fusion of visual and non-visual data enhances diagnostic accuracy, supporting early intervention and personalized treatment.

Keywords: Cardiovascular Disease (CVD), Electrocardiogram (ECG), Deep Learning, CNN, SqueezeNet, Hybrid Model, FCNN, Medical History, Disease Detection

I. INTRODUCTION

Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, accounting for an estimated 17.9 million deaths annually. Accurate and early detection plays a vital role in reducing morbidity and improving survival rates. The electrocardiogram (ECG), a widely adopted and non-invasive diagnostic tool, provides critical insights into cardiac electrical activity. However, traditional analysis of ECGs relies significantly on the expertise of clinicians, which introduces subjectivity and may delay intervention.

In many regions, particularly in resource-constrained environments, expert cardiologists may not always be available to interpret ECGs promptly. In such cases, automated systems can bridge the gap between patient need and expert diagnosis. Recent advancements in artificial intelligence (AI), particularly in deep learning, have revolutionized the way complex data like medical images and time-series signals are processed. CNNs, in particular, are well-suited for interpreting ECG images due to their ability to capture hierarchical spatial patterns.

However, real-world diagnosis involves more than just image analysis. Clinicians often consider a patient's demographic profile, blood test results, previous medical conditions, lifestyle, and genetic factors to reach a conclusive diagnosis. This broader context is often missed in image-only deep learning models. Hence, there is a growing need for hybrid models that combine visual ECG data with non-visual, structured clinical information to make accurate and personalized decisions.

This study addresses this gap by developing a deep learning-based diagnostic model that combines both image and clinical data to predict cardiovascular abnormalities. The goal is not only to increase the accuracy of diagnosis but also to improve interpretability and provide clinical context to the predictions.

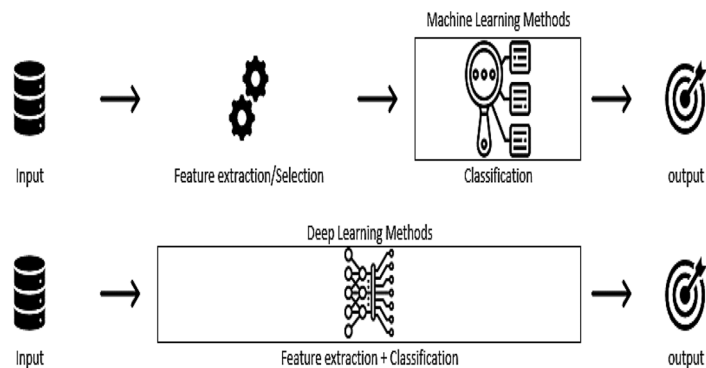


Fig. 1. Abstract concept of machine learning and deep learning.

Numerous machine learning (ML) techniques have been explored for the prediction and classification of cardiovascular diseases using benchmark datasets, such as the UCI Cleveland Heart Disease dataset. Among various algorithms applied—including Decision Tree (DT), Naïve Bayes (NB), K-Nearest Neighbors (K-NN), and Neural Networks (NN)—the DT algorithm demonstrated superior performance, achieving an accuracy of 89%. Dissanayake and Md Johar [16] investigated the influence of feature selection methodologies on the predictive capabilities of ML classifiers for heart disease diagnosis using the same dataset. Their study incorporated several feature selection techniques, including ANOVA, Chi-square testing, forward and backward selection, and Lasso regression. Subsequently, they implemented six classification models: DT, Random Forest (RF), Support Vector Machine (SVM), K-NN, Logistic Regression (LR), and Gaussian Naïve Bayes (GNB). Their results indicated that feature selection significantly enhanced classification accuracy, with the backward selection method in combination with DT achieving the highest accuracy of 88.52%.

In another study [17], the performance of NB, SVM, and DT classifiers was assessed using ten-fold cross-validation on the South African Heart Disease dataset comprising 462 records. The NB classifier yielded the most favorable results, with an overall accuracy of 71.6%, sensitivity of 63%, and specificity of 76.16%.

Kim et al. [18] conducted a comparative evaluation of several ML algorithms—namely NN, SVM, Classification based on Multiple Association Rules (CMAR), DT, and NB—for the prediction of cardiovascular conditions. Their experiments utilized two types of biomedical data: ultrasound images of carotid arteries (CAs) and heart rate variability (HRV) extracted from ECG signals. The fusion of features from both CAs and HRV datasets led to improved classification performance. Among the tested models, SVM and CMAR exhibited the highest accuracies, achieving 89.51% and 89.46%, respectively.

Unlike traditional ML approaches that rely on manual feature engineering, deep learning (DL) models autonomously learn complex patterns and extract significant features from raw data. As a specialized branch of ML, DL employs architectures composed of multiple hidden layers to enhance learning capabilities. Convolutional Neural Networks (CNNs), a prominent DL model, have shown remarkable effectiveness in various image-based diagnostic applications, including medical imaging and ECG signal analysis.

II. LITERATURE REVIEW

Electrocardiography (ECG) remains one of the most reliable and widely used diagnostic tools for detecting and monitoring cardiovascular conditions. An electrocardiogram (ECG or EKG) provides valuable insights into the electrical and muscular activity of the heart. Despite its relatively straightforward acquisition, accurate interpretation of ECG waveforms demands significant clinical expertise. Historically, ECG records were maintained on paper, making the process of manual inspection labor-intensive and prone to human error. Digitizing such records offers a promising avenue for enabling automated diagnosis and faster clinical workflows. The primary objective of recent research has been to convert paper-based ECG charts into digital 1-D signal formats through machine learning techniques. This involves segmenting the original ECG images into 13 leads, applying image processing techniques such as smoothing, thresholding, and binary conversion, and extracting significant cardiac waveform components (P, QRS, and T waves). Post-extraction, dimensionality reduction techniques like Principal Component Analysis (PCA) are utilized to interpret the underlying data patterns effectively. Classification models—including K-Nearest Neighbors (KNN), Logistic Regression (LR), Support Vector Machine (SVM), and ensemble-based voting classifiers—are evaluated using performance metrics such as accuracy, precision, recall, and F1-score to finalize the most effective diagnostic model. This model aids in the classification of major heart conditions such as Myocardial Infarction (MI), Abnormal Heartbeat (AH), or Normal (NP) based on ECG images.

Complementary research efforts also underscore the growing utility of deep learning (DL) models. For instance, MobileNetV2—a lightweight and computationally efficient convolutional neural network—has been successfully applied in pneumonia diagnosis using chest X-ray (CXR) images. The use of image enhancement techniques such as white balance correction and Contrast Limited Adaptive Histogram Equalization (CLAHE) improves the quality of CXR images, which are then input into a pretrained MobileNetV2 model. The system achieved impressive classification accuracies of **99.76%** for binary and **91.17%** for three-class classification tasks, validating its effectiveness for medical image diagnosis [2].

Further, a novel lightweight CNN architecture has been proposed for ECG image classification, achieving an accuracy of **98.23%** using only CPU resources, thus making it feasible for resource-constrained environments. Additionally, when this CNN was used as a feature extractor for traditional ML models, the Naïve Bayes classifier attained an impressive **accuracy of 99.79%**. Such integration presents potential for deployment in Internet of Things (IoT)-based healthcare systems [3].

Moreover, multi-modal approaches have emerged to enhance prediction capabilities by integrating demographic information (e.g., age, gender, height, weight) with ECG signal data. One such study employed Fast Fourier Transform (FFT) for feature extraction and Random Forest for feature selection. The refined features were then classified using a Multi-Layer Perceptron (MLP) model trained on a combined dataset from the PTB-XL repository. This comprehensive method significantly improved the robustness and accuracy of cardiac disease prediction [4].

These advancements highlight the transformative potential of artificial intelligence in cardiovascular diagnostics. By integrating deep learning, signal processing, and multi-modal data fusion, researchers are moving closer to developing highly accurate, automated, and scalable tools for early detection and classification of heart diseases.

A. Convolutional Neural Networks (CNN)

In the domain of deep learning, Convolutional Neural Networks (CNNs) represent a specialized class of deep neural networks, primarily developed for image recognition and processing tasks [40]. Unlike traditional neural networks, CNNs utilize a three-dimensional architecture where neurons are organized along the dimensions of height, width, and depth (i.e., the number of channels). For instance, a typical input image of size $227 \times 227 \times 3$ denotes a width and height of 227 pixels, with three color channels (RGB). The fundamental purpose of CNNs is to automatically extract hierarchical and spatially significant features from input images. This is achieved through a combination of key architectural components, mainly convolutional layers and pooling layers. Convolutional layers perform the core operation of feature extraction by applying a set of learnable filters (or kernels) across the input data. As these filters slide over the image, they perform localized matrix multiplications (convolutions), generating feature maps that capture spatial patterns such as edges, textures, or shapes.

Pooling layers, typically following the convolutional layers, are used to reduce the spatial dimensions of the feature maps, thereby decreasing computational complexity and aiding in generalization by retaining the most prominent features. At the deeper levels of the network, fully connected (dense) layers are used to integrate and interpret the abstracted features. The final output layer typically employs a sigmoid or softmax activation function to produce class probabilities, enabling the network to perform multi-class or binary classification tasks effectively. CNNs have proven to be highly efficient in a wide range of visual recognition tasks, including medical image analysis, where precise feature extraction from ECG signals or imaging data is critical for disease diagnosis.

1) CNN Algorithm

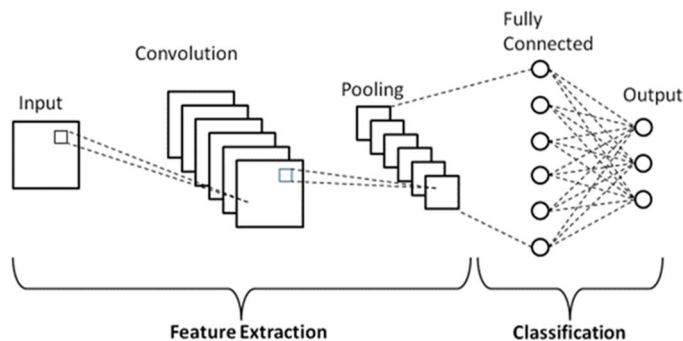


Figure 3: CNN Algorithm

Convolutional Neural Networks (CNNs) are a powerful tool for machine learning, especially in tasks related to computer vision. Convolutional Neural Network CNNs, are a specialized class of neural networks designed to effectively process grid-like data, such as images. A Convolutional Neural Network (CNN) is a type of deep learning algorithm that is particularly well-suited for image recognition and processing tasks. It is made up of multiple layers, including convolutional layers, pooling layers, and fully connected layers. The architecture of CNNs is inspired by the visual processing in the human brain, and they are well-suited for capturing hierarchical patterns and spatial dependencies within images.

Features of CNN are as follows:

- **Convolutional Layers:** These layers apply convolutional operations to input images, using filters (also known as kernels) to detect features such as edges, textures, and more complex patterns. Convolutional operations help preserve the spatial relationships between pixels.
- **Pooling Layers:** Pooling layers downsample the spatial dimensions of the input, reducing the computational complexity and the number of parameters in the network. Max pooling is a common pooling operation, selecting the maximum value from a group of neighboring pixels.
- **Activation Functions:** Non-linear activation functions, such as Rectified Linear Unit (ReLU), introduce non-linearity to the model, allowing it to learn more complex relationships in the data.
- **Fully Connected Layers:** These layers are responsible for making predictions based on the high-level features learned by the previous layers. They connect every neuron in one layer to every neuron in the next layer

2) Proposed CNN Architecture

The proposed Convolutional Neural Network (CNN) architecture is composed of a total of 38 layers, including input and output layers. Specifically, the architecture integrates six 2D convolutional layers, three max-pooling layers, three fully connected (dense) layers, eight LeakyReLU activation layers, eight batch normalization layers, five dropout layers, two depth concatenation layers, and a final softmax layer for classification. The complete architecture is illustrated in Figure 4. This model is designed as a dual-branch architecture to enhance the ability to extract and learn more representative and discriminative features. The input to the model is an image of size $227 \times 227 \times 3$, which is simultaneously fed into both the stack branch and the full branch.

3) Stack Branch

The stack branch comprises three sequential blocks of 2D convolutional layers, each with a 3×3 kernel size. Each convolutional layer is followed by a LeakyReLU activation function (with a negative slope coefficient of 0.1), a batch normalization layer, and a max-pooling layer. LeakyReLU is used in place of ReLU to mitigate the issue of inactive neurons ("dying ReLUs") [46]. Batch normalization accelerates training and enhances stability by normalizing activations for each mini-batch. Max-pooling reduces spatial dimensions and computational overhead by retaining only the most prominent features in each region. In this branch, max-pooling layers with a 6×6 filter and a stride of 3 are used. The number of filters employed in the convolutional layers is 64, 128, and 224, respectively, progressively increasing the model's capacity to extract deep hierarchical features. At the end of the stack branch, the output feature map has a spatial dimension of $2 \times 2 \times 224$.

4) Full Branch

The full branch begins with a fully connected layer containing 16 neurons, which distinguishes it from the convolution-based stack branch. In fully connected layers, each neuron is connected to every neuron in the preceding layer, in contrast to the localized connectivity in convolutional layers. This layer is followed by a LeakyReLU activation, a batch normalization layer, and a dropout layer to mitigate overfitting and improve generalization. Following this block, two additional convolutional layers—Conv04 and Conv05—are used in parallel to capture broader spatial features. Conv04 consists of 32 filters of size 2×2 with stride 1 and padding 1, while Conv05 consists of 64 filters of size 3×3 with stride 2 and padding 2. The outputs of these layers are concatenated to form a $2 \times 2 \times 96$ feature map, followed by a dropout layer to prevent overfitting due to feature correlation.

Feature Fusion and Classification

The outputs from the stack and full branches are concatenated, producing a combined feature map of size $2 \times 2 \times 320$. A dropout layer is then applied to further reduce overfitting. To refine the feature representation and reduce computational complexity, a 1×1 convolutional layer with 256 filters is added. This is followed by a fully connected layer with 512 neurons to enhance the classification capacity of the model.

Finally, a fully connected layer with 4 output neurons is used, corresponding to the four target classes: NP (Normal Person), AH (Arrhythmia), MI (Myocardial Infarction), and HMI (History of Myocardial Infarction). A softmax activation function is applied at the output to produce probability distributions for multi-class classification.

5) Model Training and Workflow

As depicted in Figure 4, the ECG images undergo a preprocessing pipeline involving **cropping, resizing, and augmentation**, after which they are stored in an image datastore. The proposed CNN is trained using these preprocessed images, adjusting its learnable parameters through backpropagation. Upon completion of training, the model is evaluated on test ECG images to classify cardiac conditions into the predefined four categories.

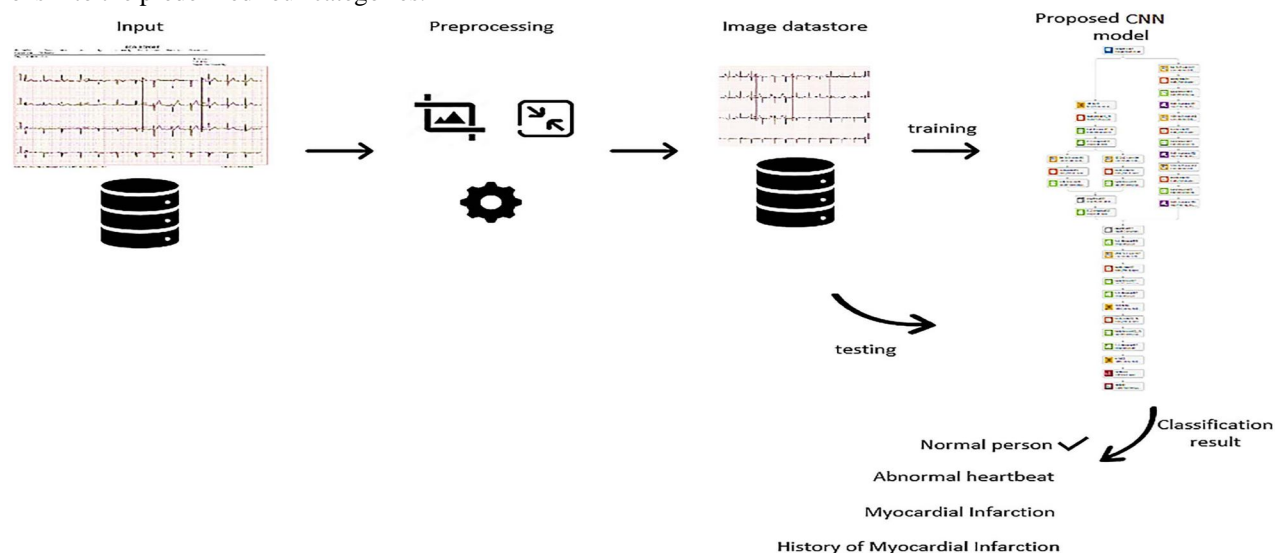


Fig. 4. Schematic of using the proposed CNN model for ECG images of cardiac patients' classification.

B. Pretrained Deep Learning Models

The pretrained deep NNs can be used for transfer learning, feature extraction, and classification. In this article, low-scaled SqueezeNet and AlexNet pretrained CNN networks that can be executed on a single CPU are used for transfer learning and feature extraction.

The transfer learning approach is commonly used with pre-trained deep NNs applied to a new dataset. Therefore, it could benefit from the pretrained network that has already learned a variety of features that can be transferred to other similar tasks. Most of the pretrained networks have been trained with more than a million images and can classify images into 1000 object classes. In applying the transfer learning approach, the final layers of the pretrained network are replaced with new layers to learn the specific features of the new dataset. Then, the model is fine-tuned by training it on a new training dataset with specific training parameters and testing its performance measure on a new test dataset.

The pretrained deep NNs can be used as a feature extraction tool without wasting time and effort on training. In this article, the extracted features from the pretrained networks are used to train traditional machine learning classifiers, namely SVM [41], K-NN [42], DT [43], RF [44], and NB [45]. The details of using the pretrained networks are explained in the next Sections.

C. KNN Algorithm

The K-Nearest Neighbors (KNN) algorithm is a popular supervised machine learning method used for both classification and regression problems. It predicts the output for a new data point by finding the 'k' nearest points from the training data based on a distance metric, usually Euclidean distance. The class most common among these neighbors is assigned to the new point (for classification), or their average is taken (for regression). KNN is easy to understand and effective for many applications because it makes no assumptions about the underlying data. However, it can be slow for large datasets since it needs to calculate distances for every prediction. The choice of 'k' greatly affects the model's accuracy, with too small or too large values leading to errors.

Steps

Here are the steps involved in the K-Nearest Neighbors (KNN) algorithm:

- Prepare the dataset: Collect and organize the labeled training data that will be used for comparison.
- Select the value of 'k': Decide how many neighboring points should be considered for making a prediction.
- Measure the distance: For a given test point, compute the distance between it and each point in the training set using a method like Euclidean distance.
- Identify the closest neighbors: Sort the distances and pick the top 'k' points that are nearest to the test point

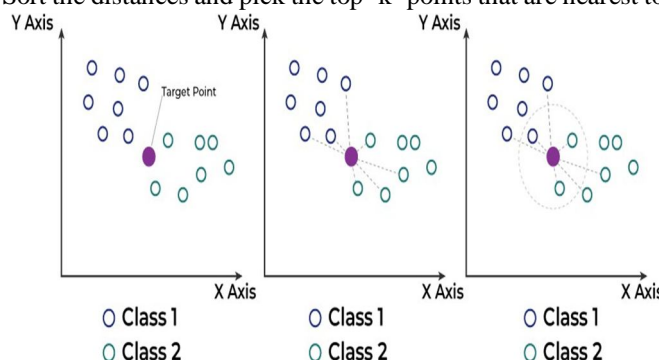


Figure 5: KNN Algorithm

- Make a decision: For classification tasks: find the most frequent class among the selected neighbors. For regression tasks: calculate the average of the values from the neighbors.
- Assign the output: Give the predicted label or value to the test data point based on the result from the neighbors.
- Evaluate the predictions: Optionally, assess the model's performance using test data or validation techniques.

III. EXPERIMENTS

The experimental evaluation of the proposed and referenced methods was conducted using the ECG Images Dataset of Cardiac Patients [3]. This dataset comprises 928 ECG records from distinct patients and is categorized into four clinically relevant classes, as outlined in Table and illustrated with sample images in Figure 6. The four diagnostic classes are:

- NP (Normal Person): Represents healthy individuals with no identifiable cardiac abnormalities.
- AH (Arrhythmia): Refers to conditions where the heart beats irregularly due to disrupted electrical impulses—either too fast, too slow, or erratically.
- MI (Myocardial Infarction): Commonly known as a heart attack, it occurs when the coronary arteries are blocked or narrowed, resulting in reduced or ceased blood flow and subsequent damage to the heart muscle.
- H. MI (History of Myocardial Infarction): Denotes patients who have recently experienced and recovered from a myocardial infarction.

The diversity and real-world complexity of the dataset make it a suitable benchmark for training and evaluating machine learning and deep learning models aimed at detecting and classifying cardiac abnormalities from ECG images.

The ECG image dataset utilized in this study includes **928 samples**, distributed across **four diagnostic categories**. The breakdown is as follows:

No.	Class	Number of Images
1.	Normal Person (NP)	284
2.	Abnormal Heartbeat (AH)	233
3.	Myocardial Infarction (MI)	239
4.	History of Myocardial Infarction (H. MI)	172
	Total	928

This class distribution highlights a moderately balanced dataset, suitable for multi-class classification tasks. The presence of both healthy and pathological categories allows for the development and evaluation of robust diagnostic models.

All experiments were conducted using MATLAB 2021b on a computing system running Windows 10 Pro (64-bit). The hardware configuration included an Intel® Core™ i7-4510U CPU @ 2.00 GHz, 8 GB of RAM, and a 4 GB NVIDIA GeForce 820M GPU, providing a suitable environment for training deep learning models on image data.

A. Preprocessing of ECG Images

As illustrated in Figure 7, the original ECG images in the dataset contained non-informative header and footer regions, which do not contribute to the learning process. To improve the model's focus on meaningful content, all images were cropped to remove irrelevant sections. Subsequently, each image was resized to a uniform resolution of 227×227 pixels with 3 RGB channels, ensuring compatibility with the input layer of the CNN model.

B. Data Augmentation

To address dataset imbalance and improve model generalization, data augmentation techniques were applied, following the approach suggested in [47], [48]. Specifically, the ECG images were augmented through:

- Rotation
- Horizontal flipping
- Translation

These transformations increased the dataset size from 928 to approximately 4700 images, providing a richer and more diverse training set.

C. Training Configuration

The training hyperparameters used for the deep learning model are summarized in Table III. The model was trained using the Adam optimizer over 16 epochs, with a mini-batch size of 128. Given the importance of the learning rate (LR) in model convergence, various LR values were explored, as detailed in the subsequent results section.

- Iterations per epoch: 29
- Total training iterations: 464

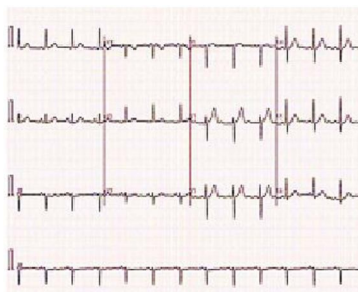


Fig 6.1 Sample from the ECG images dataset after performing cropping as a preprocessing.

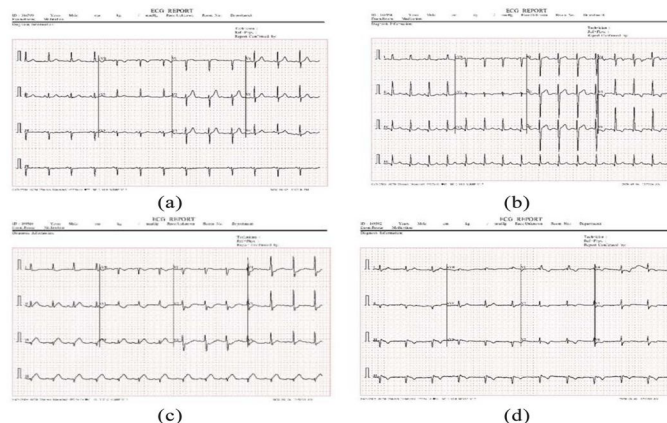
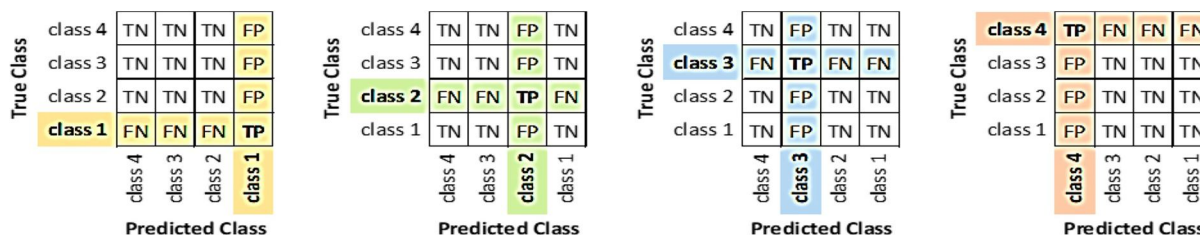


Fig. 6.2 Samples from the ECG images dataset. (a) NP. (b) AH. (c) MI. (d) H. MI.

IV. RESULTS

To assess the effectiveness of the proposed model, several standard evaluation metrics were employed: accuracy, precision, recall, F1-score, and the model's training and testing times. These metrics were computed based on the information derived from the confusion matrix, which summarizes the model's classification performance across the four target classes.

Figure 7 illustrates the structure of a four-class confusion matrix as applicable to this study, where each row represents the actual class and each column represents the predicted class.



True Class \ Predicted Class	class 4	class 3	class 2	class 1
class 4	TN	TN	TN	FP
class 3	TN	TN	TN	FP
class 2	TN	TN	TN	FP
class 1	FN	FN	FN	TP

True Class \ Predicted Class	class 4	class 3	class 2	class 1
class 4	TN	TN	FP	TN
class 3	TN	TN	FP	TN
class 2	FN	FN	TP	FN
class 1	TN	TN	FP	TN

True Class \ Predicted Class	class 4	class 3	class 2	class 1
class 4	TN	FP	TN	TN
class 3	FN	TP	FN	FN
class 2	TN	FP	TN	TN
class 1	TN	FP	TN	TN

True Class \ Predicted Class	class 4	class 3	class 2	class 1
class 4	TP	FN	FN	FN
class 3	FP	TN	TN	TN
class 2	FP	TN	TN	TN
class 1	FP	TN	TN	TN

Fig. 8. Semantic of the confusion matrices for four classes results.

Below outlines the mathematical definitions of the evaluation metrics, which are computed as follows:

To evaluate the classification performance of the proposed model, several well-established metrics were utilized, each derived from the values in the confusion matrix. The definitions of these metrics are provided below:

- Accuracy**
Represents the ratio of correctly classified instances (true positives and true negatives) to the total number of samples.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$
- Recall (Sensitivity or True Positive Rate)**
Measures the proportion of actual positives that are correctly identified by the model.

$$\text{Recall} = \frac{TP}{TP + FN}$$
- Precision (Positive Predictive Value)**
Indicates the proportion of positive predictions that are truly positive.

$$\text{Precision} = \frac{TP}{TP + FP}$$
- F1-Score**
A harmonic mean of precision and recall, balancing both metrics especially in cases of class imbalance.

$$\text{F1-Score} = \frac{2 \times \text{Recall} \times \text{Precision}}{\text{Recall} + \text{Precision}}$$

Where:

- TP = True Positives
- TN = True Negatives
- FP = False Positives
- FN = False Negatives

These metrics collectively provide a comprehensive assessment of the model's performance, especially in multi-class classification scenarios involving clinical datasets.

- Accuracy:** Measures the proportion of correctly classified instances among the total number of instances.
- Recall (Sensitivity):** Represents the proportion of actual positive samples that were correctly identified by the model. It evaluates how well the model captures the true class.
- Precision:** Indicates the proportion of predicted positive samples that are actually correct. It reflects the model's ability to avoid false positives.
- F1-Score:** A harmonic mean of precision and recall, providing a balanced measure that accounts for both false positives and false negatives.

These metrics collectively provide a comprehensive view of the model's predictive performance. They are especially important in healthcare applications, where the cost of misclassification can be high.

The training and inference times were also recorded to evaluate the computational efficiency of the model, which is critical for practical deployment, particularly in real-time or resource-limited clinical environments.

D. Transfer Learning Using Pretrained Networks

In this study, transfer learning was utilized by adapting two well-known pretrained deep learning architectures—AlexNet and SqueezeNet—originally trained on the ImageNet dataset, which contains over 1,000 object categories. These models were fine-tuned to classify ECG images into four diagnostic categories: Normal Person (NP), Abnormal Heartbeat (AH), Myocardial Infarction (MI), and History of Myocardial Infarction (H. MI).

For AlexNet, the final fully connected (FC) layer designed for 1000-class classification was replaced with a new FC layer consisting of four output neurons, corresponding to the four classes in our dataset. This new layer was followed by a softmax activation function to provide class probabilities.

In contrast, SqueezeNet—a lightweight architecture that avoids traditional fully connected layers—was modified by replacing its final convolutional layer, which originally had 1000 filters, with a new convolutional layer consisting of four 1×1 filters. This layer feeds into a new classification layer configured for our specific use case.

For both networks, the classification layers were replaced to align with the requirements of our task, while the core layers were retained to leverage pretrained knowledge.

Model Comparison and Performance Analysis

To evaluate the effectiveness of transfer learning versus the proposed model, all three architectures (AlexNet, SqueezeNet, and the proposed CNN) were trained on the same ECG dataset using different initial learning rate (LR) values: 0.01, 0.001, and 0.0001. The results of this comparative study are detailed in Table VI, where the following metrics were recorded:

- Accuracy (A.)
- Recall (R.)
- Precision (P.)
- F1-Score (F1)
- Training Time (T1)
- Testing Time (T2)

The best performance was achieved by the proposed CNN model, with an average accuracy of 98.23% when the initial learning rate was set to 0.0001, outperforming both AlexNet and SqueezeNet across most metrics.

E. Cross-Validation Results

To further validate the reliability of the proposed model, five-fold cross-validation was conducted. The performance measurements for each fold are summarized in Table VII. The bolded values in the table indicate the average performance across all five folds, demonstrating the model's stability and generalization capability across different data splits.

V. CONCLUSION

The titled "Cardiovascular Disease Detection in ECG Images: A Comprehensive Analysis of Machine Learning and Deep Learning Approaches" leverages state-of-the-art deep learning methodologies to automate the identification of cardiovascular abnormalities through analysis of ECG images. This system is designed with a dual-user architecture, supporting both super administrators and regular users, thus ensuring efficient data management and ease of interaction.

The platform enables super admins to monitor user activity, manage access, and oversee the system's overall performance. Simultaneously, end users can effortlessly upload ECG images and obtain detailed, AI-generated diagnostic reports. The integration of patient history further enriches the analysis, offering valuable clinical context that enhances the accuracy of diagnostic decisions.

The system employs advanced deep learning algorithms to automate the detection and classification process, significantly reducing the reliance on manual ECG interpretation. This not only accelerates diagnosis but also minimizes human error, thereby contributing to improved patient care and outcomes. As visualized in Figures 7.1 and 7.2, the model achieves high accuracy while maintaining efficient loss optimization, showcasing its robustness and reliability.

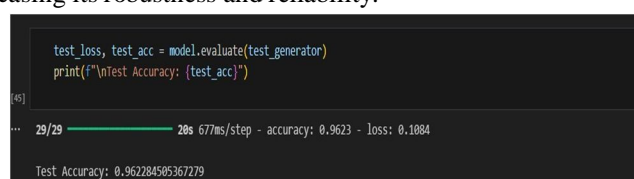


Figure 7.1: Model Accuracy

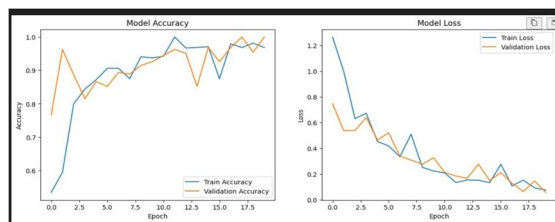


Figure 7.2: Model Accuracy vs Loss Accuracy

By streamlining cardiovascular assessments through an intelligent and accessible interface, this project democratizes the use of AI-powered diagnostics, extending its benefits to a wider spectrum of healthcare professionals, including those in resource-constrained settings.

In conclusion, the implementation of this system represents a significant advancement in the field of cardiac diagnostics. It delivers a powerful, scalable, and intelligent solution that merges cutting-edge technology with real-world clinical needs. Future developments will focus on further improving model performance, increasing dataset diversity, and integrating additional features to elevate diagnostic precision and system adaptability.

REFERENCES

- [1] M. Mohammad D Abdubaker, "Detection of Cardiovascular Disease from ECG Image Using Machine Learning and Deep Learning methods ",IEEE Transactions on Artificial Intelligence – 2 April 2023 vol 2
- [2] M. Swathy and K. Saruladha, "A comparative study of classification and prediction of Cardio-Vascular Diseases (CVD) using Machine Learning and Deep Learning tech- niques", ICT Express, vol. 8, no. 1, pp. 109-116, 2022
- [3] M. B. Abubaker and B. Babayigit, "Detection of Cardiovascular Diseases in ECG Images Using Machine Learning and Deep Learning Methods", IEEE Transactions on Artificial Intelligence, vol. 4, no. 2, pp. 373-382, April 2023
- [4] K. Fatema, S. Montaha, M.A.H. Rony, S. Azam, M. Z Hasan and M. Jonkman, "A ro- bust framework combining image processing and deep learning hybrid model to classify cardiovascular diseases using a limited number of paper-based complex ECG images", Biomedicines, vol. 10, no. 11, pp. 2835, 2022.
- [5] K. Prashant, P. Choudhary, T. Agrawal and E. Kaushik, "OWAE-Net: COVID-19 detec- tion from ECG images using deep learning and optimized weighted average ensemble technique", Intelligent Systems with Applications, vol. 16, pp.
- [6] Mohammed B. Abubaker and Bilal Babayigit "Detection of Cardiovascular Diseases in ECG Images Using Machine Learning and Deep Learning Methods"
- [7] World Health Organization (WHO), "Cardiovascular diseases," Jun. 11, 2021. Accessed: Dec. 27, 2021. [Online]. Available: <https://www.who.int/health-topics/cardiovascular-diseases>
- [8] Government of Westren Australia, Department of Health, "Com- mon medical tests to diagnose heart conditions," Accessed: Dec. 29, 2021. [Online]. Available: https://www.healthyswa.wa.gov.au/Articles/A_E/Common-medical-tests-to-diagnose-heart-conditions
- [9] M. Swathy and K. Saruladha, "A comparative study of classification and prediction of cardio-vascular diseases (CVD) using machine learning and deep learning techniques," ICT Exp., to be published, 2021. [Online]. Available: <https://doi.org/10.1016/j.ict.2021.08.021>



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)