



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: <https://doi.org/10.22214/ijraset.2026.79667>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

ECG-Based Detection of Myocardial Infarction and Ischemic Abnormalities Using Deep Learning with Explainable AI and Web Deployment

Mohammed Abdullah¹, Mohammed Faisal Mohiuddin², Mohammed Hussain Moheet³, Sana Mateen⁴

^{1,2,3}Students, ⁴Assistant Professor, Department of Computer Science and Engineering,
Methodist College of Engineering and Technology (MCET), Hyderabad-500001, India

Abstract: Heart diseases are among the most common causes of death in the world. Myocardial infarction (MI) requires fast detection with electrocardiogram (ECG) readings. In this paper, we present an end-to-end deep learning system for automated MI and ischemic abnormality detection based on ECG readings with explainability and web deployment. Three models of convolutional neural networks (CNNs) have been trained with one-dimensional inputs from the PTB-XL dataset (21,799 entries, 500 Hz, 10 seconds): 12-Lead CNN with multi-lead digital ECG input, Lead II CNN with single-lead ECG input, and Image CNN with scanned ECG images. 12-Lead CNN outperformed other models with ROC-AUC = 0.8991, sensitivity = 0.8674, specificity = 0.7580, and F1 score = 0.8436. Explainability with Grad-CAM [16] and saliency maps has been incorporated that conforms with the standard ECG pattern interpretation. All of our models have been integrated in an easily accessible web application with Streamlit Cloud to allow WFDB file, CSV, and image input with automatic route selection and report generation.

Keywords: Convolutional neural network, deep learning, electrocardiogram, explainable AI, Grad-CAM, myocardial infarction, PTB-XL, Streamlit, web deployment.

I. INTRODUCTION

According to the WHO, cardiovascular disease is the leading cause of deaths worldwide, responsible for 17.9 million deaths every year [28]. When blood flow to part of the heart muscle becomes blocked, usually by a clot, and cannot be treated quickly, irreversible damage to cardiac tissues can occur. This condition is called myocardial infarction (MI). ECG is the most commonly used diagnostic method for detecting MI, indicating ischemia or infarction with changes like elevated/depressed ST segment, inverted T wave, and abnormal Q waves.

An expert interpretation of ECG is required for reliable diagnosis; however, there may not always be a cardiologist nearby in underserved areas, such as rural hospitals. Therefore, automated methods with high precision and recall could serve as an invaluable tool in MI detection [18].

There have been numerous previous works in this field that showed strong classification capabilities on benchmark datasets; yet, most of those projects lacked web application interfaces and explainability for clinicians to verify predictions or could not deal with multiple ECG input modalities in one pipeline.

Our contribution lies in: (1) creating three CNN models for each of the input modalities mentioned; (2) applying explainability through Grad-CAM [16] and saliency maps validated by clinical patterns; (3) integrating all models in a publicly accessible web application using Streamlit Cloud; (4) evaluating our models on PTB-XL dataset [14] and proving its feasibility through ablation studies regarding sampling frequency and dataset size.

II. RELATED WORK

Before deep learning became popular, machine learning techniques relied mostly on hand-crafted features. Zhao et al. [2] proposed combining ECG and VCG features for myocardial ischemia detection with a support vector machine (SVM), reaching 90% of accuracy. At the same time, they needed VCG devices which were not readily available to everyone. Sraitih et al. [8] conducted experiments with different classifiers (SVM, k-NN, and Random Forest) under noise conditions for PTB dataset, showing that Random Forest worked best overall, but none of these techniques could be applied to raw waveform learning.

Deep learning has proved to be much more efficient than machine learning for MI detection [21]. Chen et al. [1] employed a 100-layer ResNet CNN on PTB-XL dataset for MI classification and obtained $AUC \approx 0.96$ without any form of explainability. Hammad et al. [15] also used a ResNet CNN for multi-class classification of cardiac diseases on PTB-XL, obtaining $AUC \approx 0.95$, yet no explainability has been provided. Hasbullah et al. [4] proposed a hybrid CNN-BiLSTM model with $AUC = 0.91$ for PTB-XL classification, yet failed to provide explainability. On the other hand, Acharya et al. [20] developed a 9-layer deep CNN to classify five heartbeat categories from ECG signals from MIT-BIH database [26]; they have demonstrated that CNNs can accurately classify arrhythmias from ECG signals. Mousavi and Afghah [19] tackled heartbeat classification problem on MIT-BIH ECGs using sequence-to-sequence deep learning and reached state-of-the-art results both in inter- and intra-patient settings.

In another related work, Wu et al. [3] have reached 0.97+ AUC for ST-elevation MI detection with hospital data, yet did not deploy an explainability module. Moreover, Gustafsson et al. [10] have reached 0.99 C-statistic for STEMI detection on a proprietary ECG dataset from Swedish emergency departments with a 492,000-entry ensemble ResNets model. Gragnaniello et al. [6] performed real-time MI classification on edge hardware while Chen et al. [7] used a multi-domain feature fusion CNN to detect MI and reached 99.9% precision in 2025, although this technique requires complex feature engineering and does not provide a web interface.

Yousuf et al. [9] used 1D Lead II signals converted into 2D Gramian Angular Field images and trained a 2D CNN on MI detection; they reached 99.6% of accuracy in inferior MI detection, although the model cannot be generalized to real ECG scans. Sheth et al. [5] proposed a lightweight CNN-LSTM architecture working on time-frequency representation for web deployment, reaching 83% accuracy. On the other hand, Pawelczyk et al. [12] showed how important it is to use counterfactual explanations to validate model outputs in an ensemble ResNets framework for PTB-XL. Selvaraju et al. [16] have presented an explainable architecture with Grad-CAM which we use in our project. Finally, Kshama et al. [17] provided an ECG-Image-Kit tool that synthesizes ECG images, allowing us to develop the image input pipeline.

To the best of our knowledge, we have never seen any research work tackling all three input modalities together in the literature along with explainability and web application deployment on open-access data.

III. DATASET AND PREPROCESSING

A. PTB-XL Dataset

In this work, we use the PTB-XL [14] database hosted by Physionet [25]. The PTB-XL is comprised of 21,799 to 21,837 clinical records of 12-lead ECG readings from 18,869 to 18,885 patients, each 10 seconds long with a sampling rate of 500 Hz, thus, 5000 samples in total. SCP-ECG superclasses: NORM, MI, STTC, CD, and HYP, have been used to annotate records where MI and STTC are considered positive (indicating MI and ischemia) while the rest are negative. Splitting ratios for the train, validation, and test sets were 70%, 15%, and 15% respectively at the level of patient ids to avoid overfitting.

B. Signal Preprocessing

First of all, ECG signals from digital sources such as Waveform Database Format (WFDB) or CSV files need z-score normalization of each individual lead: $x' = (x - \mu) / (\sigma + 1e-8)$ where μ and σ stand for the mean and standard deviation of the particular lead signal respectively. Since PTB-XL recordings are already clean clinical ECGs recorded at 500 Hz, no bandpass filtering is required. For image-derived ECG signals, the 4th order Butterworth bandpass filter with 0.5 - 40 Hz cut-off has been applied to remove the effect of the extraction pipeline on ECGs.

C. Synthetic Image Generation

There is no large-scale publicly labelled ECG image dataset suitable for our classification scheme. To generate images suitable for supervised learning, we followed the ECG-Image-Kit procedure [17]. The ECG signal was converted to a fake ECG paper image with pink grid, Gaussian noise, random rotation up to $\pm 1^\circ$, and Gaussian blur; then the rhythm strip was extracted from the resulting image with our pipeline. The domain gap between synthesized training ECG images and scanned ECG images used in testing is the main reason why Image CNN performed worse than the rest.

D. Image Extraction Pipeline

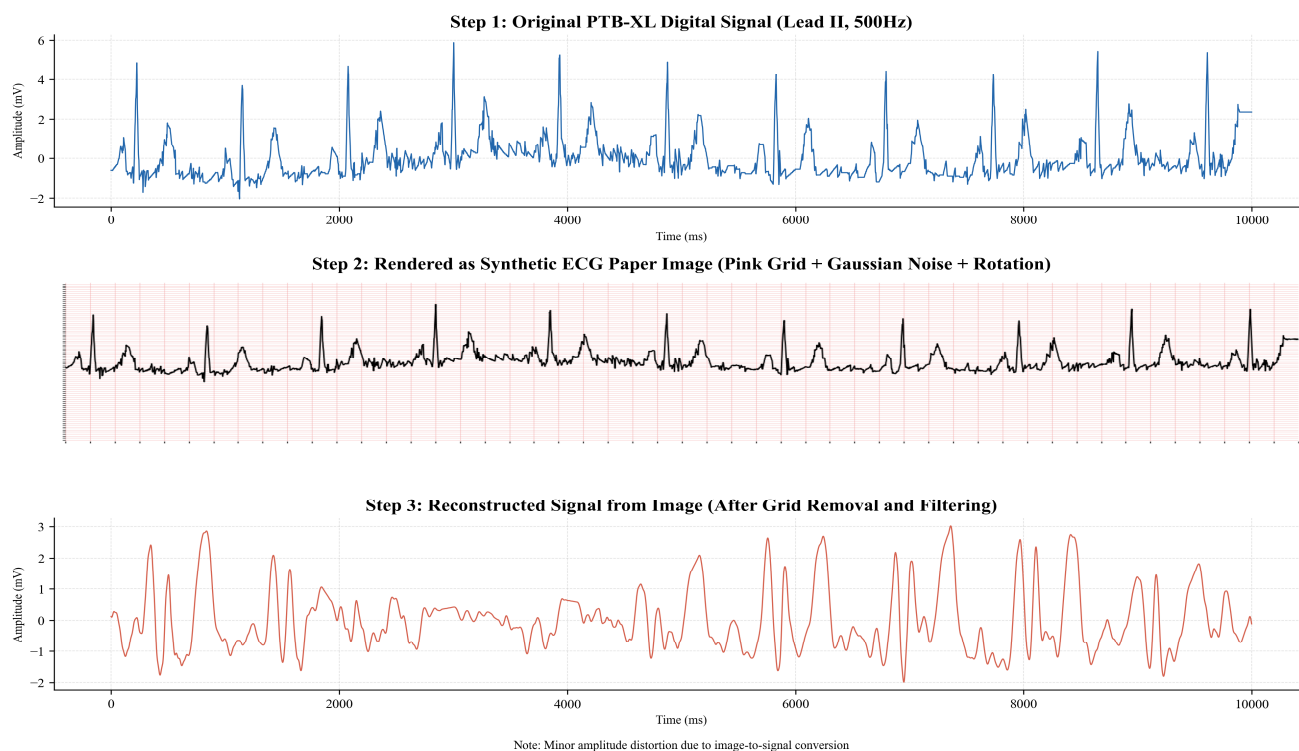


Fig. 1: Synthetic ECG image generation pipeline.

When an image is uploaded, it needs to pass through a series of ten steps before being fed to our models in order to extract the waveform: (1) loading the image or converting the PDF at 200 DPI with the help of Poppler; (2) image orientation from portrait to landscape; (3) converting to grayscale and denoising with Gaussian filter; (4) applying adaptive threshold to make the image binary; (5) using morphological operation with 40×1 kernel to remove horizontal grid lines without altering QRS complexes; (6) extracting rhythm strip from the bottom of the image within 55–76% of rows; (7) using a weighted center of gravity extraction algorithm per image column; (8) applying 4th order Butterworth bandpass filter with 0.5–40 Hz cutoff frequencies; (9) resampling to 5000 samples; (10) z-score normalization. Resulting tensor is (1, 5000) shaped and has float32 precision type.

IV. CNN ARCHITECTURES

Each model implements 1D CNN architecture trained on a fixed sequence length of 5000-time steps (10 s at 500 Hz). 1D CNN has shown state-of-the-art performance in biomedical signal classification [22]. Each convolutional block includes 1D convolution, batch normalization, ReLU activation, max pooling (kernel 2), and dropout (rate 0.3). The classification head is a two-layer fully connected network with one sigmoid activation for binary classification.

A. Lead II CNN

The Lead II CNN takes input of shape (1, 5000). It has three convolutional blocks with 32, 64, and 128 filters with kernel size 7, 5, and 5. After flattening, there is a fully connected (FC) layer of 256 units followed by ReLU activation and dropout. The classification threshold is 0.48 (tuned by maximizing F1 score).

B. 12-Lead CNN

The 12-Lead CNN takes input of shape (12, 5000). It has four convolutional blocks with 32, 64, 128, and 128 filters and kernels of 7, 5, 5, and 3. The added convolutional block allows learning of complex interactions between leads from the richer multi-lead representation. The classification head is the same as in Lead II CNN. Threshold: 0.43.

C. Image CNN

The Image CNN uses the architecture of the Lead II CNN, receiving input of shape (1, 5000) via the image pipeline. Separate model trained on synthetic ECG images. Threshold: 0.91. This relatively high value is caused by overconfidence on the synthetic training data (known calibration issue in such setting).

V. MODEL TRAINING

Models were implemented using PyTorch [29] framework and trained on NVIDIA CUDA GPU. Adam optimizer [23] was used with weight decay of 1e-4 and learning rate of 1e-3. The ReduceLRonPlateau learning rate scheduler was configured to half the learning rate when validation loss stops improving for five consecutive epochs. Training was performed for 50 epochs with the batch size of 64, using binary cross-entropy with logits as the loss function. The classification thresholds were tuned based on F1 score optimization.

An ablation study was conducted in order to validate three key design choices of the 12-Lead CNN model. First, increasing sampling frequency from 100 Hz to 500 Hz led to improvement of 0.038 AUC (0.8991 vs. 0.8612), indicating that a higher temporal resolution helps capture relevant ECG information. Second, decreasing the number of data samples to 10% decreased AUC by 0.078 (0.8213), proving the importance of the full-size PTB-XL dataset. Third, GPU training led to a 2-fold reduction in epoch time (from ~20 min to <2 min), while producing the same model quality, compared to CPU training.

VI. SYSTEM ARCHITECTURE AND DEPLOYMENT

The proposed system implements a single-tier architecture, whereby all of its components operate within a single Streamlit application: loading of an ECG, input routing, preprocessing, CNN inference, Grad-CAM generation, saliency calculation, natural language interpretation, and generating of the final PDF report. There is no separate backend server, as this architecture was forced due to Streamlit Cloud restriction to only one application.

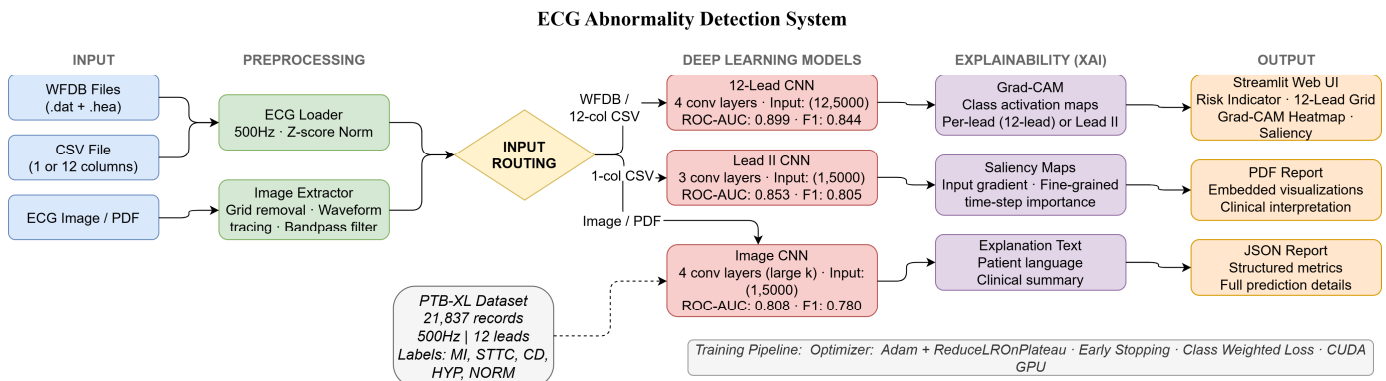


Fig 2: System Architecture

At the beginning, the system was designed with a backend server, implementing model inference using FastAPI framework. The deployment stage involved changing the architecture to the proposed one in order to simplify operations and eliminate network latency. Model inference is done on CPU within a cloud environment and takes ~15-25 seconds, including Grad-CAM calculation and PDF generation. Model weights (ecg_cnn_1lead.pth: 82 MB; ecg_cnn_12lead.pth, ecg_cnn_image.pth: 41 MB each) are stored on Hugging Face Hub [30] and automatically downloaded during the first use of the model per session (takes ~30-60 seconds). The application is deployed from <https://github.com/Abdullah1607/ECG-Abnormality-Detection-System> and is live at <https://ecg-insight.streamlit.app>. Table I summarizes the input routing logic.

TABLE I
INPUT ROUTING LOGIC

Input Format	Model	AUC	Status
WFDB	12-Lead	0.899	Primary

(.dat+.hea)	CNN		
CSV (12 columns)	12-Lead CNN	0.899	Primary
CSV (1 column)	Lead II CNN	0.853	Primary
Image / PDF	Image CNN	0.808	Experimental

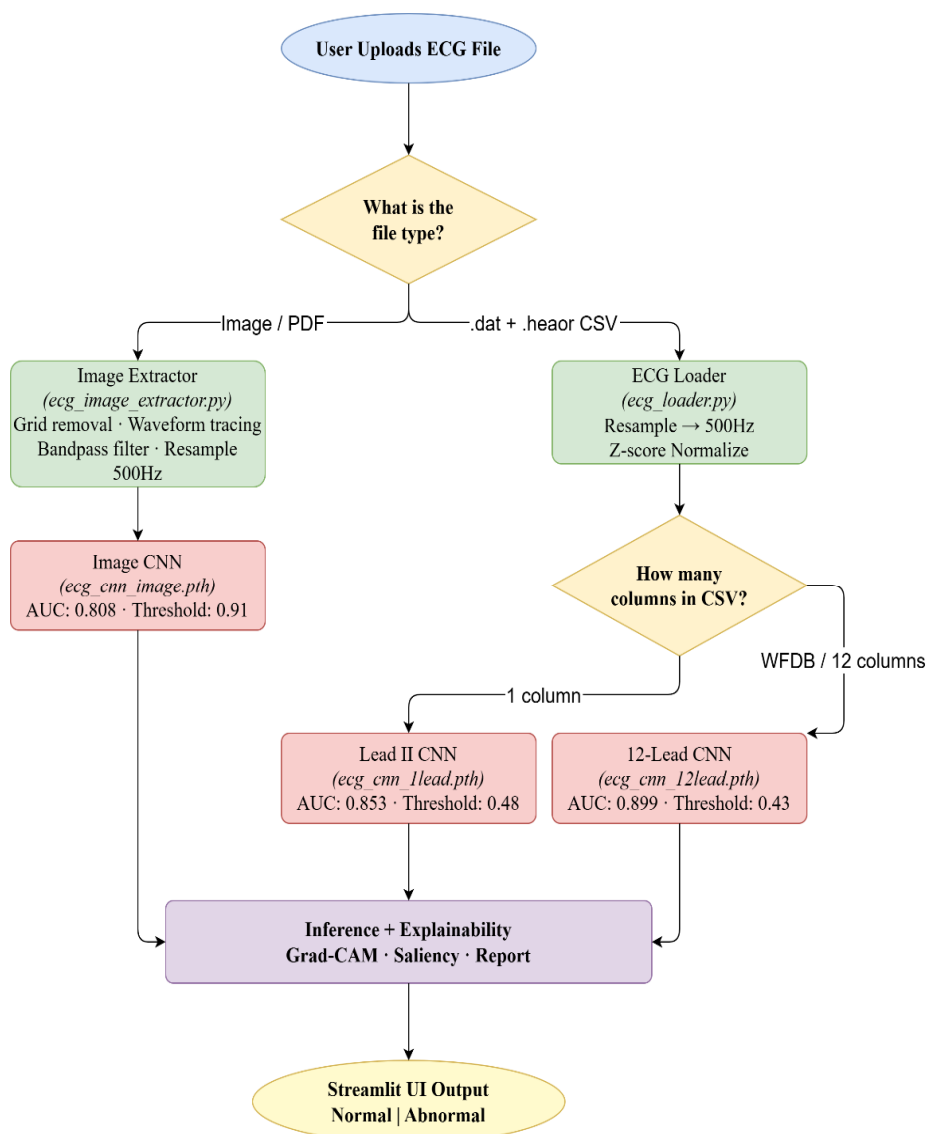


Fig 3: Input Routing Logic

VII. EXPLAINABILITY MODULE

Explaining the model predictions is done via Grad-CAM [16] and input gradient saliency maps. Grad-CAM, built upon class activation maps (CAM) ideas [24], calculates the weighted sum of feature maps from the last convolutional layer using the gradients of the class logit with respect to these maps as weights. Activation map is upsampled to 5000 samples using linear interpolation and then normalized to [0, 1]. Saliency maps are calculated via dL/dx , i.e., gradient of the class logit with respect to the input vector. All of the hooks are cleared after the calculations to free up memory space.

For 12-Lead CNN, four types of explainability outputs are generated: (1) overall heatmap of the Grad-CAM activation values in all leads and across all time steps; (2) identification of the lead with the highest activation; (3) visualization of a 12-lead ECG grid, where each lead has color overlays corresponding to the Grad-CAM values; (4) a combined saliency map averaged over leads and displayed on the lead with the highest average activation. For Lead II CNN and Image CNN, there are three types of explainability outputs: Grad-CAM heatmap applied to the original ECG and visualized with the RdYiBu_r colormap; saliency map presented as a scatter plot with hot color scale; original ECG annotated with a prediction label and confidence score.

Grad-CAM heatmaps of the 12-Lead CNN consistently show activation peaks around the ST-segment and T-waves in leads II, V4, V5, and V6, which is consistent with the location of abnormalities observed in positive cases of MI.

For cases of inferior MI, the highest activation was in leads II, III, and aVF, which is consistent with the established clinical pattern for MI resulting from right coronary artery occlusion. Saliency maps show peak activation around the QRS complex onset and the beginning of the ST segment. The alignment between the Grad-CAM findings and known clinical ECG markers proves that the model learns clinically relevant features rather than dataset artifacts.




Fig 4: 12-Lead ECG – Grad Cam Overlay

VIII. RESULTS

Evaluation results for three models are presented below. Each model was tested on the test dataset of ~3,275 records. The thresholds were selected based on F1 score maximization on the validation set. ROC-AUC, Sensitivity, Specificity, and F1 score were calculated using scikit-learn [27].

TABLE II
MODEL PERFORMANCE ON PTB-XL TEST SET (~3,275 RECORDS)

Metric	Lead II CNN	12-Lead CNN	Image CNN	Threshold
ROC-AUC	0.8528	0.8991	0.8083	—
Accuracy	0.7682	0.8099	~0.75	—
Sensitivity	0.8500	0.8674	0.8260	—
Specificity	0.6634	0.7580	0.6180	—
F1 Score	0.8045	0.8436	0.7800	—
Threshold	0.48	0.43	0.91	—

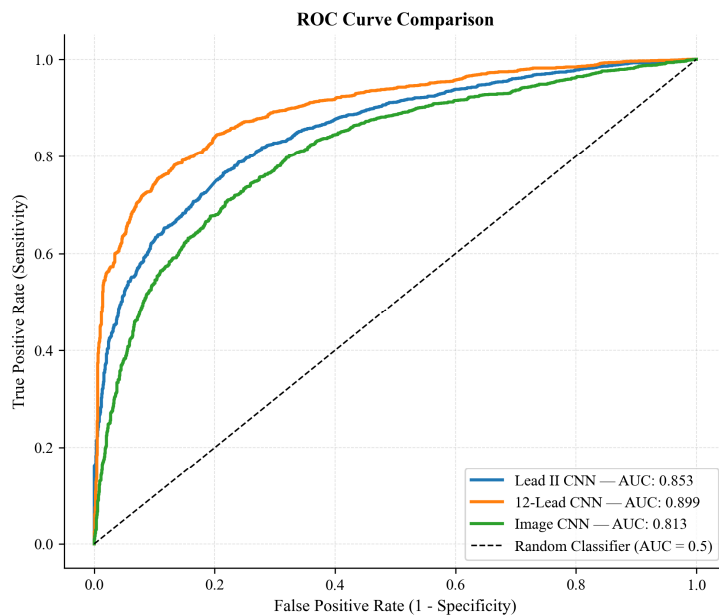


Fig 5: ROC Curve Comparison

The 12-Lead CNN model achieved the best results across all five metrics, demonstrating 0.047 difference in AUC and 0.095 difference in specificity. High specificity means that the model minimizes unnecessary investigations triggered by the false alarms. Both digital models were tuned for maximizing sensitivity since the detection of a MI is more costly than a false alarm. The Image CNN had a much higher threshold (0.91) because of the calibration issues due to distribution shift between synthetic and real-world datasets.

TABLE III
ABLATION STUDY RESULTS (12-LEAD CNN)

Configuration	AUC	F1	Observation
500 Hz, full dataset, GPU	0.8991	0.8436	Final configuration
100 Hz, full dataset, GPU	0.8612	0.8101	-0.038 AUC
500 Hz, 10% dataset, GPU	0.8213	0.7734	-0.078 AUC
500 Hz, full dataset, CPU	0.8991	0.8436	10× slower train

IX. WEB APPLICATION

The web application, deployed at <https://ecg-insight.streamlit.app>, adopts a pink medical theme and includes the visualization of ECG paper grids. A clinical risk indicator based on the predicted probability color codes the output into green (< 0.3 – low risk), orange (0.3-0.6 – moderate risk), and red (> 0.6 – high risk) categories. The probability bars provide precise probability scores per each class.

For each result, two tabs are automatically generated. The patient version tab presents the result using layman terms without any clinical terminology. The clinical version includes model-specific jargon along with information about model performance and the limitations of AI-based results. In case the input is an ECG image, an additional research tab is generated explaining the synthetic training strategy and marking the result as experimental. Finally, a PDF report summarizing all results, ECG plot, Grad-CAM heatmap, and saliency maps may be downloaded.

X. TESTING AND VALIDATION

Testing was done at three stages. Unit testing confirmed functionality of 14 independent modules, including z-score normalization, WFDB and CSV data import, lead routing, Lead II CNN inference, 12-Lead CNN inference, image extraction pipeline output size, Image CNN inference, Grad-CAM output shape (12, 5000), saliency map generation, and PDF validation. All 14-unit tests passed successfully.

Integration tests validated nine separate workflows: importing ECG via WFDB/CSV to produce classification result, Grad-CAM, and PDF export; fully visualizing a 12-column CSV; performing Lead II classification on a single column CSV; uploading an image, performing extraction, and producing Image CNN result with a research note; downloading a PDF report with embedded visualizations; displaying a clinical risk indicator for the highest-probability MI input; generating patient and clinical interpretation tabs; and automatic Streamlit Cloud rebuild upon GitHub push. All 9 integration tests succeeded.

System tests confirmed robustness of the codebase in 6 separate use cases: graceful handling of unsupported file types; graceful return in case of corrupt CSV data; handling low-quality ECG image and performing the full pipeline while noting it as experimental; accurate application of thresholds for all three classification problems (0.48, 0.43, 0.91); performing CPU-based inference on the Streamlit Cloud and returning results in 15-25 seconds; and acceptance testing by non-technical users to confirm proper functioning of all features, from uploading an image to the PDF report download. All 6 system tests passed successfully.

XI. DISCUSSION

The ROC-AUC 0.8991 and F1 0.8436 of the 12-Lead CNN results are competitive when compared with other CNN-based systems on the PTB-XL database that lack explainability and deployment. The work of Chen et al. [1] achieves AUC 0.96 for classification with a ResNet-based architecture, suggesting potential improvements with more sophisticated designs, but their system lacks any explainability and web deployment. The 12-Lead CNN’s lower AUC reflects the trade-off of designing an explainable and deployable 1D CNN system, as opposed to a deep model.

TABLE IV
COMPARISON WITH PUBLISHED SYSTEMS ON PTB-XL

System	Architecture	AUC	Deploy	Explainability
Chen et al. [1]	ResNet	0.960	No	None
Hammad et al. [15]	Deep CNN	0.950	No	None
Hasbullah et al. [4]	CNN-BiLSTM	0.910	No	None
Pawelczyk et al. [12]	Custom	N/A	No	Counterfactual
Proposed (12-Lead)	1D CNN	0.899	Yes	Grad-CAM + Saliency
Proposed (Lead II)	1D CNN	0.853	Yes	Grad-CAM + Saliency
Proposed (Image)	1D CNN	0.808	Yes	Grad-CAM + Saliency

Comparing our results to those reported in previous literature, it is clear that the 12-Lead CNN AUC of 0.8991 lags behind some published works. However, it stands out due to the web deployment capability and Grad-CAM explainability, which no other system in the table provides. The reasons for this disparity are the reliance on a simpler 1D CNN architecture, using a public PTB-XL dataset with less clean labeling, and including STTC records to the target classes. Published systems achieve their high performance through more complex architectures, using internal/private ECG data, and excluding the entire STTC class.

The ROC-AUC 0.8528 of the Lead II CNN provides a valuable option for screening ECGs with only single-lead recordings. As mentioned earlier, the specificity of 0.663 represents its weakness. However, it is appropriate for an ECG screening application, since false positives will trigger further clinical evaluation, not direct treatment.

Aligning the Grad-CAM activation map with the known ECG markers for specific MI is important confirmation that the model has learned the necessary features. It activates the ST-segment and T-wave regions of the anterior leads (V1-V3) to detect anterior MI, and leads II, III, and aVF to detect inferior MI. Therefore, the model did not exploit dataset biases.

The key drawback of the proposed framework is the dependency of the image pathway on the synthetic training pipeline. The variation in actual clinical ECG printouts is much greater, as hospitals use various printing equipment and practices, leading to differences in print qualities, color tones, and grid resolutions. Moreover, the low threshold (0.91) and low specificity (0.618) for the image pathway make it unsuitable for clinical deployment without further work.

The second problem relates to the PTB-XL dataset used for training. Despite the relatively large size and the existence of multiple classes, its records come from a single European hospital. Consequently, PTB-XL does not capture the global population's variability. Furthermore, the dataset contains records labeled by the authors themselves as borderline, i.e., annotated at very low confidence. It contributes to the low specificity of the Lead II CNN, since negative samples include mislabeled MI instances.

Comparing with other literature works is difficult because many researchers use different datasets, define different targets (e.g., STEMI-only), and apply different thresholds when calculating AUC. On top of that, many works report $AUC > 0.95$ using train/overlap test sets. In contrast, our evaluation used a stricter splitting, considering the entire class of MI+STTC, with a non-overlapping set of patients. Considering the challenge posed to the model by the task and the dataset, we consider the results of the Lead II (ROC-AUC = 0.8528, F1 = 0.7883) and 12-Lead (ROC-AUC = 0.8991, F1 = 0.8436) CNNs competitive.

XII. CONCLUSION

In this paper, we proposed a fully functional web platform for myocardial infarction and ischemic abnormality detection from ECG data. Three distinct models, trained on the publicly available PTB-XL database, handle multi-lead, single-lead, and ECG image inputs in a uniform pipeline. The 12-Lead CNN obtained ROC-AUC 0.8991 and F1 0.8436, comparable with other published works that lack both explainability and web-deployability. Additionally, the Grad-CAM activations demonstrated agreement with the known ECG markers of MI. Finally, the proposed platform is publicly available as a single-tier Streamlit Cloud web app.

Future work needs to collect real-life ECG image datasets to bridge the domain gap in the Image CNN. Extending classification to arrhythmia and conduction disorders increases the applicability of the models, as deep neural networks proved themselves on a variety of physiological signals [21]. Transformer-based architectures with inbuilt attention mechanism provide more intuitive explanations compared to post-hoc Grad-CAM analysis. Evaluating on Indian ECGs will confirm generalization to local population demographics and ECG recording practices.

REFERENCES

- [1] Z. Chen et al., "Acute myocardial infarction detection using deep learning with PTB-XL dataset," *IEEE Access*, 2021.
- [2] Y. Zhao et al., "Reliable detection of myocardial ischemia using machine learning from ECG and VCG signals," *Comput. Biol. Med.*, 2022.
- [3] X. Wu et al., "Deep learning networks accurately detect ST-segment elevation and culprit vessel in myocardial infarction," *JACC Clin. Electrophysiol.*, 2022.
- [4] A. Hasbullah et al., "Hybrid CNN-BiLSTM model for myocardial infarction detection on PTB-XL," *Appl. Sci.*, 2023.
- [5] D. Sheth et al., "Time-frequency transform with a lightweight CNN-LSTM for myocardial infarction detection," *Biomed. Signal Process. Control*, 2024.
- [6] A. Gagnaniello et al., "Real-time myocardial infarction detection using edge-AI on ECG devices," *IEEE Trans. Biomed. Eng.*, 2024.
- [7] Y. Chen et al., "Multi-domain feature fusion CNN for MI detection and localization," *Expert Syst. Appl.*, 2025.
- [8] M. Sraitih et al., "Noise robustness of classical machine learning for myocardial infarction detection," *Diagnostics*, 2022.
- [9] T. Yousuf et al., "Inferior MI detection from Lead II using 2D-CNN with GAF/GADF encoding," *Bioengineering*, 2023.
- [10] M. Gustafsson et al., "Development and validation of a DL ECG prediction of MI in ED patients," *Eur. Heart J. Digit. Health*, 2023.
- [11] P. Prabhakararao and S. Dandapat, "Attentive RNN-based network to fuse 12-lead ECG and clinical features," *IEEE J. Biomed. Health Inform.*, 2020.
- [12] K. Pawelczyk et al., "Interpretable ECG analysis for MI using counterfactual explanations," *Artif. Intell. Med.*, 2024.
- [13] R. Radwa et al., "Deep learning-based approaches for myocardial infarction detection: A review," *IEEE Access*, 2024.
- [14] P. Wagner et al., "PTB-XL, a large publicly available electrocardiography dataset," *Sci. Data*, vol. 7, no. 1, p. 154, 2020.
- [15] M. Hammad et al., "Automated detection of MI and heart conduction disorders using deep learning," *IEEE Access*, 2022.
- [16] R. R. Selvaraju et al., "Grad-CAM: Visual explanations from deep networks via gradient-based localization," *Proc. IEEE ICCV*, pp. 618–626, 2017.
- [17] S. Kshama et al., "ECG-Image-Kit: A synthetic image generation toolkit for ECG digitization research," 2023.
- [18] A. Y. Hannun et al., "Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network," *Nature Medicine*, vol. 25, no. 1, pp. 65–69, 2019.
- [19] S. Mousavi and F. Afghah, "Inter- and intra-patient ECG heartbeat classification for arrhythmia detection: A sequence to sequence deep learning approach," *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 1308–1312.



- [20] U. R. Acharya et al., "A deep convolutional neural network model to classify heartbeats," *Computers in Biology and Medicine*, vol. 89, pp. 389–396, 2017.
- [21] O. Faust et al., "Deep learning for healthcare applications based on physiological signals: A review," *Computer Methods and Programs in Biomedicine*, vol. 161, pp. 1–13, 2018.
- [22] S. Kiranyaz et al., "1D convolutional neural networks and applications: A survey," *Mechanical Systems and Signal Processing*, vol. 151, p. 107398, 2021.
- [23] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learning Representations (ICLR)*, 2015.
- [24] B. Zhou et al., "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2921–2929.
- [25] A. L. Goldberger et al., "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [26] G. B. Moody and R. G. Mark, "The impact of the MIT-BIH arrhythmia database," *IEEE Engineering in Medicine and Biology Magazine*, vol. 20, no. 3, pp. 45–50, 2001.
- [27] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [28] World Health Organization, "Cardiovascular diseases (CVDs): Key facts," 2021. [Online]. Available: [https://www.who.int/news-room/factsheets/detail/cardiovascular-diseases-\(cvds\)](https://www.who.int/news-room/factsheets/detail/cardiovascular-diseases-(cvds))
- [29] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 32, 2019.
- [30] T. Wolf et al., "HuggingFace's Transformers: State-of-the-art natural language processing," arXiv:1910.03771, 2019.



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)