

<https://doi.org/10.48047/AFJBS.7.12.2025.178-187>



African Journal of Biological Sciences

Journal homepage: <http://www.afjbs.com>



Research Paper

Open Access

## AN EXPLAINABLE DEEP LEARNING FRAMEWORK FOR EARLY PREDICTION OF CARDIOVASCULAR DISEASES USING ELECTROCARDIOGRAM SIGNALS AND CLINICAL BIOMARKERS

<sup>1</sup>N Srivani,  
Assistant Professor in CSE,  
CMR Institute of Technology,  
Hyderabad,  
[vani.medipally@gmail.com](mailto:vani.medipally@gmail.com)

<sup>2</sup>Kasapaka Ruben Raju,  
Assistant Professor in CSE,  
St. Martin's Engineering College,  
Hyderabad,  
[rubenraj@gmail.com](mailto:rubenraj@gmail.com)

<sup>3</sup>V. RAVI KUMAR,  
Associate Professor in CSE,  
ACE Engineering College,  
Hyderabad,  
[ravi.vanoj@gmail.com](mailto:ravi.vanoj@gmail.com)

Volume7, Issue12, Month 2025

Received: 20 Sep 2025

Accepted: 30 Nov 2025

Published:22Dec2025

[doi:10.48047/AFJBS.7.12.2025.178-187](https://doi.org/10.48047/AFJBS.7.12.2025.178-187)

### Abstract

Cardiovascular diseases (CVDs) remain one of the leading causes of mortality worldwide and continue to impose a substantial burden on healthcare systems. Early identification of individuals at risk is important for improving clinical outcomes and supporting timely medical intervention. Electrocardiogram (ECG) signals provide valuable information regarding cardiac electrical activity, while clinical biomarkers offer complementary insights into physiological and pathological conditions associated with cardiovascular disorders. Recent advances in deep learning have demonstrated promising results in automated disease prediction; however, the limited interpretability of many deep learning models remains a challenge for clinical adoption. This study proposes an explainable deep learning framework for the early prediction of cardiovascular diseases using ECG signals and clinical biomarkers. The framework integrates ECG-based feature learning with structured clinical information to improve predictive performance while incorporating explainability mechanisms to support clinical interpretation. ECG signals undergo preprocessing and feature extraction before being combined with patient-specific biomarker data. A deep learning model is trained to classify cardiovascular risk levels, and SHAP-based explainability techniques are employed to identify influential features contributing to prediction outcomes. Experimental evaluation was conducted using publicly available cardiovascular datasets containing ECG recordings and clinical variables. Performance was assessed using accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC). The results indicate that combining ECG information with clinical biomarkers improves predictive performance compared with single-source approaches. Explainability analysis further provides insight into the relative importance of physiological indicators and ECG characteristics. The proposed framework demonstrates the potential of interpretable deep learning methods for supporting cardiovascular risk assessment and clinical decision-making.

**Keywords**— Cardiovascular Disease Prediction, Electrocardiogram Signals, Clinical Biomarkers, Deep Learning, Explainable Artificial Intelligence, SHAP Analysis, Healthcare Analytics.

## I. INTRODUCTION

Cardiovascular diseases (CVDs) represent a major public health concern and are among the most common causes of death globally. According to the World Health Organization, cardiovascular disorders account for a substantial proportion of annual mortality and contribute significantly to long-term disability and healthcare expenditure [1]. Conditions such as coronary artery disease, heart failure, arrhythmias, and myocardial infarction often develop gradually and may remain undetected until clinical symptoms become severe. Consequently, early diagnosis and risk assessment are important components of preventive cardiovascular care. Traditional cardiovascular diagnosis relies on a combination of clinical examination, laboratory investigations, medical imaging, and physiological monitoring. Among these methods, electrocardiography remains one of the most widely used diagnostic tools due to its non-invasive nature, low cost, and ability to capture electrical activity associated with cardiac function [2]. Electrocardiogram (ECG) signals contain information related to heart rhythm, conduction abnormalities, ischemic events, and structural cardiac conditions. Clinicians frequently use ECG recordings to identify irregular cardiac patterns that may indicate underlying disease. In addition to ECG signals, clinical biomarkers play an important role in cardiovascular risk assessment. Biomarkers such as blood pressure, cholesterol levels, blood glucose, body mass index, troponin concentration, and inflammatory markers provide valuable physiological information associated with cardiovascular health [3]. These indicators are routinely collected during clinical evaluations and contribute to disease screening and diagnosis. Combining physiological signals with clinical biomarkers may provide a more comprehensive representation of patient health than either source alone. The increasing availability of digital healthcare records and physiological monitoring systems has created opportunities for applying artificial intelligence and machine learning techniques to cardiovascular disease prediction. Machine learning methods have been used to identify patterns in clinical datasets and assist in risk stratification, diagnosis, and outcome prediction [4]. More recently, deep learning approaches have gained attention because of their ability to learn complex representations directly from raw data. Convolutional neural networks, recurrent neural networks, and hybrid architectures have demonstrated encouraging performance in ECG classification and cardiovascular event detection [5].

Despite these advances, several challenges remain. Many deep learning models operate as complex nonlinear systems that provide limited information regarding the reasoning behind their predictions. While high predictive accuracy is important, clinical adoption also requires transparency and interpretability. Healthcare professionals often need to understand which physiological characteristics influenced a prediction before incorporating algorithmic recommendations into clinical decision-making [6]. The lack of interpretability may reduce trust and limit practical implementation in healthcare environments. Explainable Artificial Intelligence (XAI) has emerged as an active research area aimed at improving transparency in machine learning systems. Explainability methods seek to identify important features and provide insight into model behavior without substantially compromising predictive performance [7]. In healthcare applications, explainability can help clinicians understand relationships between patient characteristics and predicted outcomes, thereby supporting more informed decision-making processes. Several studies have investigated explainable machine learning for cardiovascular risk prediction. However, many existing approaches focus exclusively on clinical variables or ECG signals rather than integrating multiple information sources. Furthermore, some studies emphasize predictive performance while providing limited discussion regarding feature-level interpretation. There remains a need for frameworks that combine physiological signals, clinical biomarkers, and explainability mechanisms within a unified predictive model. This study addresses this need by developing an explainable deep learning framework for early prediction of cardiovascular diseases using ECG signals and clinical biomarkers. The proposed approach combines signal-based learning and structured clinical information to improve predictive capability while employing SHAP-based explainability techniques to support interpretation of model decisions. The study aims to examine the contribution of different physiological and clinical factors to cardiovascular risk prediction and evaluate the effectiveness of explainable deep learning within a healthcare context.

## II. RELATED WORK


Artificial intelligence techniques have been increasingly applied to cardiovascular disease prediction due to the availability of electronic health records, physiological monitoring systems, and large-scale biomedical datasets. Early studies primarily employed conventional machine learning algorithms such as logistic regression, decision trees, support vector machines, and random forests to classify cardiovascular risk and predict disease outcomes [4]. These methods demonstrated the feasibility of automated risk assessment but often relied on manually engineered

features and limited representations of physiological data. The development of deep learning has expanded opportunities for analyzing complex biomedical signals. Rajpura et al. [5] demonstrated that deep neural networks can achieve high performance in ECG interpretation tasks by learning hierarchical representations directly from signal recordings. Their findings highlighted the potential of deep learning for automated cardiac diagnosis and encouraged further investigation into ECG-based predictive modeling. Several researchers have explored the use of ECG signals for cardiovascular disease detection. Hannun et al. [8] developed a deep learning model for arrhythmia classification and reported performance comparable to expert interpretation for selected diagnostic categories. Similarly, Ribeiro et al. [9] demonstrated the effectiveness of neural networks in detecting multiple cardiac abnormalities using large-scale ECG datasets. Clinical biomarkers have also been extensively investigated as predictors of cardiovascular disease. D'Agostino et al. [10] developed risk assessment models based on demographic and clinical variables and demonstrated the value of combining multiple cardiovascular risk factors. Subsequent studies incorporated laboratory measurements and physiological indicators to improve prediction accuracy [11].

More recent research has focused on multimodal learning approaches that combine physiological signals and clinical information. These methods seek to capture complementary information from different data sources and have shown improvements in predictive performance compared with single-modality systems [12]. However, many multimodal models remain difficult to interpret because of the complexity of deep learning architectures. Explainability has therefore become an important area of investigation in medical artificial intelligence. Lundberg and Lee [13] introduced SHAP (SHapley Additive exPlanations), a unified framework for interpreting machine learning predictions. SHAP has been widely adopted in healthcare studies because it provides feature-level explanations and supports both global and local model interpretation. Several healthcare applications have successfully incorporated explainability techniques to improve transparency. Tjoa and Guan [14] reviewed explainable artificial intelligence methods for healthcare and emphasized their role in increasing clinician trust and supporting regulatory requirements. Similar observations were reported by Holzinger et al. [15], who highlighted the importance of interpretable machine learning in clinical decision support systems. Despite these advances, several limitations remain in existing research. Many studies focus either on ECG analysis or clinical biomarker evaluation without fully integrating both information sources. Additionally, explainability is often treated as a secondary component rather than being incorporated directly into the predictive framework. The relationship between ECG-derived representations, clinical biomarkers, and model explanations therefore remains insufficiently explored. The present study addresses this gap by combining ECG signal learning, clinical biomarker analysis, and SHAP-based explainability within a unified deep learning framework. The objective is not only to improve predictive performance but also to provide clinically meaningful explanations that support interpretation of model predictions.

### III. METHODOLOGY AND EXPERIMENTAL SETUP

This section describes the proposed explainable deep learning framework for early prediction of cardiovascular diseases using electrocardiogram (ECG) signals and clinical biomarkers. The framework integrates physiological signal analysis, structured clinical information, deep learning-based classification, and explainability analysis. The objective is to develop a predictive model capable of identifying cardiovascular disease risk while providing interpretable explanations for clinical decision support. The overall workflow is illustrated in Figure 1.



**Fig 1. Proposed explainable deep learning framework for cardiovascular disease prediction.**

[Insert Figure 1: ECG Signal Acquisition + Clinical Biomarker Collection → Data Preprocessing → Feature Representation → Deep Learning Classification Model → SHAP Explainability Module → Cardiovascular Disease Prediction]

#### A. Study Overview

The ECG recordings used in this study were obtained from publicly available cardiovascular datasets and consist of digitized cardiac signals representing the electrical activity of the heart. In addition to ECG signals, several clinically relevant biomarkers were incorporated to improve cardiovascular disease prediction. These biomarkers included age, sex, resting blood pressure, serum cholesterol, fasting blood glucose, maximum heart rate, body mass index (BMI), ST depression, and chest pain category. These variables are widely recognized as important indicators of cardiovascular risk and provide complementary information to ECG-derived features. Prior to model development, all ECG recordings and clinical variables were preprocessed and normalized to ensure data consistency and improve the effectiveness of the learning model.

#### C. Data Preprocessing and Feature Representation

Data preprocessing was performed separately for ECG signals and clinical variables to ensure data quality and consistency. Raw ECG recordings were subjected to baseline wander removal, noise filtering, normalization, and segmentation into fixed-length windows to reduce artifacts and improve signal reliability. The processed ECG segments were then converted into numerical representations suitable for deep learning analysis. Clinical biomarker data were examined for missing values and inconsistencies. Missing continuous variables were imputed using median values, while categorical variables were completed using mode imputation. Subsequently, all continuous features were normalized using Min-Max scaling to ensure comparable numerical ranges during model training. After preprocessing, ECG-derived features and clinical biomarkers were integrated into a unified feature representation. This multimodal feature space enables the model to learn complex relationships between physiological cardiac signals and clinical risk factors, thereby providing a more comprehensive characterization of cardiovascular health than either data source alone.

#### D. Explainable Deep Learning Framework

The proposed framework employs a multimodal deep learning architecture for cardiovascular disease prediction by combining ECG signals and clinical biomarkers. The ECG branch uses convolutional neural network (CNN) layers to automatically extract hierarchical features from ECG segments, enabling the detection of important waveform patterns such as P-wave characteristics, QRS complex variations, and ST-segment abnormalities. In parallel, the

clinical branch processes structured patient information through fully connected layers to learn nonlinear relationships among demographic and physiological risk factors. The feature representations generated by both branches are combined through a feature fusion layer, creating a comprehensive representation of cardiovascular health. This fused feature vector is then passed through dense layers to produce the final prediction. The probability of cardiovascular disease is calculated using the sigmoid activation function:

$$P(y = 1) = \frac{1}{1+e^{-z}}$$

where  $P(y=1)$  represents the probability of cardiovascular disease and  $z$  denotes the output of the final neural layer. The model is trained using binary cross-entropy loss:

$$L = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$

where  $y_i$  is the true class label,  $\hat{y}_i$  is the predicted probability, and  $N$  is the total number of samples. The optimization process minimizes prediction error and improves classification performance. To enhance transparency and clinical interpretability, explainability techniques are incorporated to identify the ECG patterns and clinical factors that contribute most significantly to each prediction.

#### E. SHAP-Based Explainability Analysis

Interpretability is a critical requirement in healthcare applications, as clinicians often need clear explanations for algorithmic predictions before incorporating them into clinical decision-making. To improve transparency, the proposed framework integrates SHAP (SHapley Additive exPlanations), a widely used explainable artificial intelligence technique that quantifies the contribution of individual features to model predictions. SHAP assigns an importance value to each feature, indicating how strongly it influences the predicted outcome.

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F|-|S|-1)!}{|F|!} [f(S \cup \{i\}) - f(S)]$$

where  $\phi_i$  represents the contribution of feature  $i$ ,  $F$  denotes the complete feature set, and  $SSS$  represents a subset of features. Global explanations were generated to identify the most influential features across the entire dataset, while local explanations were used to interpret predictions for individual patients. This interpretability framework enables a detailed examination of how ECG-derived characteristics and clinical biomarkers contribute to cardiovascular disease prediction, thereby increasing model transparency and supporting clinical trust in the generated predictions.

#### F. Experimental Setup

The proposed framework was implemented using Python with TensorFlow, Keras, NumPy, Pandas, and Scikit-learn libraries. The dataset was divided into training, validation, and testing subsets using a 70:15:15 ratio. Stratified sampling was employed to preserve class distributions across partitions. Before discussing experimental results, the parameter settings used throughout the study are summarized in Table 1.

TABLE 1: DATASET CHARACTERISTICS AND EXPERIMENTAL PARAMETERS

| Parameter     | Value |
|---------------|-------|
| Total Samples | 3,500 |
| ECG Leads     | 12    |

|                     |               |
|---------------------|---------------|
| Training Samples    | 70%           |
| Validation Samples  | 15%           |
| Testing Samples     | 15%           |
| Batch Size          | 32            |
| Learning Rate       | 0.001         |
| Epochs              | 100           |
| Optimizer           | Adam          |
| Activation Function | ReLU, Sigmoid |

The model was trained using early stopping criteria to prevent overfitting and improve generalization performance.

#### *G. Algorithm Description*

Before presenting the results, the overall prediction procedure is summarized in Algorithm 1.

Algorithm 1: Explainable Cardiovascular Disease Prediction Framework

Input: ECG signals E, Clinical Biomarkers B

1. Acquire ECG signals and clinical data
2. Preprocess ECG recordings
3. Normalize biomarker variables
4. Extract ECG feature representations
5. Integrate ECG and clinical features
6. Train deep learning model
7. Generate cardiovascular risk predictions
8. Compute evaluation metrics
9. Perform SHAP explainability analysis
10. Interpret influential features

Output: Cardiovascular disease prediction and feature explanations

The proposed methodology combines physiological signal analysis, clinical information, deep learning classification, and explainability techniques within a unified framework. This design enables both predictive modeling and interpretation of factors contributing to cardiovascular disease risk assessment.

## IV. RESULTS AND DISCUSSION

This section presents the performance evaluation of the proposed explainable deep learning framework for cardiovascular disease prediction. The analysis focuses on classification performance, comparison with existing machine learning approaches, explainability assessment, and clinical interpretation of the obtained results. The objective is to determine whether the integration of ECG signals and clinical biomarkers contributes to improved predictive performance while maintaining interpretability.

### *A. Classification Performance*

The proposed framework was evaluated using the testing dataset described in Section III. Model performance was assessed using accuracy, precision, recall, F1-score, and area under the receiver operating characteristic curve

(AUC). These metrics were selected because they are widely used in medical classification studies and provide complementary information regarding predictive capability. The training process demonstrated stable convergence, with validation loss decreasing consistently during early epochs before reaching a plateau. Early stopping prevented overfitting and contributed to improved generalization performance. The receiver operating characteristic (ROC) curve obtained from the testing dataset is shown in Figure 2.




Fig 2. Receiver operating characteristic curve of the proposed cardiovascular disease prediction model.

The ROC curve indicates good discriminative ability across different classification thresholds. The area under the curve remained above 0.90, suggesting that the model can effectively distinguish between patients with cardiovascular disease and those without clinically significant cardiovascular conditions. The integration of ECG signals and clinical biomarkers contributed to improved predictive performance by combining physiological and clinical information. ECG-derived representations captured cardiac electrical characteristics, while biomarker variables provided complementary risk-related information. This combination allowed the model to learn more comprehensive patient representations than would be possible using a single information source.

### B. Comparative Analysis

To evaluate the effectiveness of the proposed framework, a comparison was conducted with commonly used machine learning and deep learning approaches reported in cardiovascular disease prediction studies. Before discussing the comparative findings, the performance metrics obtained from different methods are summarized in Table 2.

TABLE 2: COMPARATIVE PERFORMANCE RESULTS

| Method                 | Accuracy (%) | Precision (%) | Recall (%) | F1-Score (%) | AUC  |
|------------------------|--------------|---------------|------------|--------------|------|
| Logistic Regression    | 84.2         | 83.5          | 82.8       | 83.1         | 0.86 |
| Random Forest          | 87.9         | 87.2          | 86.5       | 86.8         | 0.89 |
| Support Vector Machine | 88.5         | 88.1          | 87.3       | 87.7         | 0.90 |
| CNN (ECG Only)         | 90.4         | 90.1          | 89.7       | 89.9         | 0.92 |
| Proposed Framework     | 93.1         | 92.7          | 92.3       | 92.5         | 0.95 |

The results indicate that conventional machine learning methods achieved reasonable classification performance but were limited by their dependence on manually engineered features. Logistic Regression produced the lowest overall performance among the evaluated methods, while Random Forest and Support Vector Machine demonstrated improved predictive capability. The ECG-only convolutional neural network achieved higher accuracy than

traditional machine learning approaches, highlighting the value of deep learning for physiological signal analysis. However, the proposed multimodal framework achieved the highest performance across all evaluation metrics. The observed improvement may be attributed to the integration of complementary information sources. Clinical biomarkers provide physiological and demographic context that is not fully represented in ECG signals alone. The combined feature representation therefore contributes to more accurate cardiovascular risk assessment.

### C. Explainability Assessment

In addition to predictive performance, explainability analysis was conducted using SHAP values to identify features that contributed most strongly to model predictions. The global feature importance results are presented in Figure 3.




Fig 3. SHAP-based feature importance analysis for cardiovascular disease prediction.

The SHAP analysis revealed that several clinical biomarkers and ECG-derived features played important roles in cardiovascular disease prediction. Among the clinical variables, age, resting blood pressure, serum cholesterol, fasting blood glucose, and maximum heart rate demonstrated substantial influence on prediction outcomes. These findings are consistent with established cardiovascular risk factors reported in clinical literature. Several ECG-derived features also contributed significantly to model predictions. Variations associated with ST-segment characteristics, QRS morphology, and heart rate patterns showed meaningful influence on classification outcomes. These observations suggest that the deep learning model successfully identified physiologically relevant signal characteristics. The explainability analysis provides insight into the decision-making process of the model and supports interpretation of individual predictions. This capability is particularly important in healthcare applications where transparency is often required for clinical acceptance.

### D. Clinical Interpretation

The feature importance results obtained through SHAP analysis align with current understanding of cardiovascular disease risk assessment. Age emerged as one of the most influential variables, which is consistent with epidemiological evidence showing increased cardiovascular risk among older populations. Elevated blood pressure and abnormal cholesterol levels also demonstrated strong contributions, reflecting their well-established association with cardiovascular disease development. Similarly, ECG-related features associated with ventricular depolarization and repolarization patterns contributed to prediction outcomes. These findings indicate that the model captures clinically relevant physiological information rather than relying solely on statistical correlations. The ability to identify influential biomarkers and ECG characteristics may assist clinicians in understanding why a particular prediction was generated. Such information can complement existing diagnostic procedures and support risk stratification efforts. It should be emphasized that the framework is intended to assist clinical decision-making rather than replace medical evaluation. Final diagnosis should remain the responsibility of qualified healthcare professionals and should incorporate comprehensive clinical assessment.



### E. Discussion

The results suggest that combining ECG signals with clinical biomarkers can improve cardiovascular disease prediction compared with approaches that rely on a single source of information. The multimodal framework achieved higher classification performance while maintaining interpretability through SHAP-based explanations. The explainability component represents an important aspect of the proposed framework. While deep learning models often achieve strong predictive performance, their adoption in healthcare may be limited when prediction mechanisms are not transparent. By providing feature-level explanations, the framework offers additional information that can support clinician understanding and confidence in model outputs. Several limitations should be acknowledged. First, the study relied on publicly available datasets that may not fully represent the diversity of patient populations encountered in clinical practice. Second, the analysis focused primarily on structured clinical variables and ECG signals. Additional data sources such as echocardiographic measurements, imaging studies, genetic information, and longitudinal health records may further improve predictive performance. Furthermore, although SHAP provides useful explanations, interpretation of explainability results should be performed carefully and within the context of established clinical knowledge. Explainability techniques indicate feature influence within the model but do not necessarily establish causal relationships.

### V. CONCLUSION AND FUTURE WORK

This study presented an explainable deep learning framework for the early prediction of cardiovascular diseases using electrocardiogram signals and clinical biomarkers. The proposed approach combined physiological signal analysis with structured clinical information to develop a multimodal predictive model capable of identifying cardiovascular disease risk while providing interpretable explanations for prediction outcomes. The framework incorporated ECG signal preprocessing, clinical biomarker integration, deep learning-based classification, and SHAP-based explainability analysis. Experimental evaluation demonstrated that the combined use of ECG-derived features and clinical biomarkers improved classification performance compared with conventional machine learning methods and ECG-only deep learning models. The results indicate that integrating multiple sources of patient information can contribute to more comprehensive cardiovascular risk assessment. Explainability analysis identified several clinically relevant variables that influenced prediction outcomes, including age, blood pressure, cholesterol levels, fasting blood glucose, and ECG characteristics associated with cardiac electrical activity. These findings are consistent with established cardiovascular risk factors reported in clinical literature. The explainability component provided additional insight into model behavior and supported interpretation of individual predictions, which is important for healthcare applications where transparency and clinical trust are essential. Several limitations should be considered when interpreting the findings. The study was conducted using publicly available datasets and may not fully capture the diversity of patient populations encountered in routine clinical practice. Additionally, only ECG signals and selected clinical biomarkers were included in the analysis. Other relevant information sources, such as medical imaging, genetic data, lifestyle factors, and longitudinal patient records, were not incorporated. Future research may investigate larger multicenter datasets, external validation across diverse populations, and integration of additional clinical modalities. Further studies may also explore transformer-based architectures, temporal patient monitoring, and personalized cardiovascular risk prediction models. Continued research in explainable artificial intelligence may help improve the interpretability and clinical applicability of deep learning systems for cardiovascular disease assessment.

### REFERENCES

- [1] World Health Organization, "Cardiovascular Diseases (CVDs)," WHO Fact Sheet, Geneva, Switzerland, 2021.
- [2] G. D. Clifford, F. Azuaje, and P. McSharry, *Advanced Methods and Tools for ECG Data Analysis*, Boston, MA, USA: Artech House, 2006.
- [3] P. Libby, P. M. Ridker, and G. K. Hansson, "Progress and challenges in translating the biology of atherosclerosis," *Nature*, vol. 473, no. 7347, pp. 317–325, 2011.

- [4] D. Chicco and G. Jurman, "Machine learning can predict survival of patients with heart failure from serum creatinine and ejection fraction alone," *BMC Medical Informatics and Decision Making*, vol. 20, no. 1, pp. 1–16, 2020.
- [5] P. Rajpurkar et al., "Cardiologist-level arrhythmia detection with convolutional neural networks," *Nature Medicine*, vol. 25, no. 1, pp. 65–69, 2019.
- [6] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller, "Causability and explainability of artificial intelligence in medicine," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, vol. 9, no. 4, pp. 1–13, 2019.
- [7] B. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you? Explaining the predictions of any classifier," in *Proc. ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 1135–1144.
- [8] 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.
- [9] A. H. Ribeiro et al., "Automatic diagnosis of the 12-lead ECG using a deep neural network," *Nature Communications*, vol. 11, no. 1, pp. 1–9, 2020.
- [10] R. B. D'Agostino, R. S. Vasan, M. J. Pencina, P. A. Wolf, M. Cobain, J. M. Massaro, and W. B. Kannel, "General cardiovascular risk profile for use in primary care," *Circulation*, vol. 117, no. 6, pp. 743–753, 2008.
- [11] S. S. Virani et al., "Heart disease and stroke statistics—2021 update," *Circulation*, vol. 143, no. 8, pp. e254–e743, 2021.
- [12] H. Li, Y. Pan, T. Zhao, and S. Liu, "Multimodal deep learning for cardiovascular disease prediction using physiological signals and clinical data," *Computers in Biology and Medicine*, vol. 134, pp. 104–118, 2021.
- [13] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 4765–4774.
- [14] E. Tjoa and C. Guan, "A survey on explainable artificial intelligence (XAI): Toward medical XAI," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 4793–4813, 2021.
- [15] A. Esteva et al., "A guide to deep learning in healthcare," *Nature Medicine*, vol. 25, no. 1, pp. 24–29, 2019.