International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(16s) 2026 pp. 78-104
https://doi.org/10.70917/ijcisim-2026-4561
Article

A Deep Learning-Based Predictive Framework for Early Heart Disease Risk Assessment Using Multi-Factor Clinical Data

Swapnil Hiralal Chaudhari[1], Anish Kumar Choudhary[2]

[1]Department of Computer Science and Engineering, SAGE University, Indore, India
Email: swapnil.chaudhari19@gmail.com | ORCID: 0009-0006-7605-0735
[2]Department of Computer Science and Engineering, SAGE University, Indore, India
Email: anishkumar.choudhary@sageuniversity.in

Abstract: Cardiovascular diseases (CVDs) remain the leading cause of mortality worldwide, accounting for approximately 17.9 million deaths annually according to the World Health Organization. Early and accurate risk assessment is critical for timely intervention and improved patient outcomes. However, existing machine learning and deep learning approaches face significant limitations including poor modeling of nonlinear relationships, inadequate feature interaction learning, lack of uncertainty quantification, class imbalance challenges, overfitting tendencies, limited interpretability, insufficient multimodal fusion capabilities, poor cross-dataset generalization, absence of attention-based learning mechanisms, and inadequate clinical explainability. To address these critical gaps, we propose CardioAttentionNet, a novel hybrid deep learning framework that integrates residual convolutional neural networks, transformer encoders with multi-head self-attention, bidirectional long short-term memory networks, and cross-attention modules for comprehensive cardiovascular risk prediction. The framework incorporates intelligent data preprocessing, hybrid feature engineering with attention-based feature selection, adaptive feature fusion with learnable weights, channel attention mechanisms, probability calibration, and explainable AI components including SHAP, LIME, and Integrated Gradients. Evaluated on a large-scale multi-source dataset comprising 75,000 samples from UCI Cleveland, Framingham Heart Study, Kaggle, and Cardiovascular Disease datasets, CardioAttentionNet achieved exceptional performance with 98.12% accuracy, 97.84% precision, 98.37% recall, 97.93% specificity, 98.10% F1-score, 0.983 AUC, 0.962 MCC, and 0.961 Cohen's Kappa, significantly outperforming 13 baseline methods. Comprehensive ablation studies, statistical validation, and explainability analysis demonstrate the framework's clinical utility for early heart disease risk stratification and decision support in real-world healthcare settings.

Keywords: deep learning; heart disease prediction; attention mechanism; transformer; explainable AI; feature fusion; CardioAttentionNet; clinical decision support

1. INTRODUCTION

Cardiovascular diseases (CVDs) represent the most significant global health challenge of the 21st century, claiming approximately 17.9 million lives annually and accounting for 31% of all deaths worldwide according to the World Health Organization. The economic burden is equally staggering, with direct and indirect costs exceeding $555 billion annually in the United States alone. Despite substantial advances in medical technology and therapeutic interventions, the incidence of CVDs continues to rise, particularly in low- and middle-income countries where 75% of CVD-related deaths occur. Early detection and accurate risk stratification are paramount for effective prevention strategies, timely clinical intervention, and improved patient outcomes.

Traditional cardiovascular risk assessment tools, such as the Framingham Risk Score and SCORE (Systematic Coronary Risk Evaluation), rely on limited clinical parameters and linear statistical models that fail to capture the complex, nonlinear interactions among multiple risk factors. While machine learning (ML) approaches have demonstrated promise in improving prediction accuracy, they often suffer from limited feature representation, inability to model temporal dependencies, and lack of interpretability—critical requirements for clinical adoption. Recent deep learning (DL) methodologies have shown potential in addressing some of these limitations, yet significant challenges persist.

A comprehensive analysis of the literature from 2022 to 2026 reveals ten critical research gaps in current cardiovascular disease prediction approaches. First, existing models struggle with nonlinear relationship modeling, failing to capture the intricate interactions between clinical, demographic, and lifestyle factors. Second, poor feature interaction learning limits the ability to identify synergistic effects among risk factors, reducing predictive power. Third, the lack of uncertainty estimation prevents clinicians from assessing prediction confidence, hindering clinical decision-making. Fourth, class imbalance in cardiovascular datasets leads to biased predictions favoring the majority class. Fifth, overfitting on limited training data compromises generalization to unseen patient populations. Sixth, low interpretability of black-box deep learning models impedes clinical trust and regulatory approval. Seventh, limited multimodal fusion capabilities prevent effective integration of diverse data sources including clinical records, imaging, and temporal signals. Eighth, poor cross-dataset generalization restricts model applicability across different healthcare institutions and populations. Ninth, the absence of attention-based learning mechanisms limits the model's ability to focus on clinically relevant features. Finally, insufficient clinical explainability prevents physicians from understanding and validating model predictions, creating a barrier to real-world deployment.

To address these critical limitations, we propose CardioAttentionNet, a novel hybrid deep learning framework that integrates multiple advanced architectural components for comprehensive cardiovascular risk assessment. The primary contributions of this research are enumerated as follows:

1. Novel Hybrid Architecture: Integration of residual CNNs, transformer encoders with multi-head self-attention, bidirectional LSTMs, and cross-attention modules for comprehensive feature learning and interaction modeling.

2. Intelligent Data Preprocessing Algorithm: A systematic Intelligent Data Cleaning Protocol (IDCP) incorporating KNN-based missing value imputation, adaptive outlier removal using modified Z-scores, dynamic noise filtering with Savitzky-Golay filters, and medical data validation rules.

3. Hybrid Feature Engineering: Comprehensive feature extraction including statistical features, medical risk scores, correlation features, entropy measures, pairwise feature interactions, and attention-based feature selection.

4. Adaptive Feature Fusion Mechanism: Learnable weight assignment for optimal integration of CNN, Transformer, and LSTM representations with layer normalization and residual connections.

5. Channel Attention Integration: Squeeze-and-Excitation blocks for adaptive feature recalibration and enhanced discriminative power.

6. Probability Calibration: Temperature scaling for reliable confidence estimation and uncertainty quantification via Monte Carlo Dropout.

7. Comprehensive Explainability: Integration of SHAP values, LIME local surrogates, and Integrated Gradients for multi-level model interpretation and clinical validation.

8. Large-Scale Multi-Source Validation: Evaluation on 75,000 samples from four diverse datasets ensuring robust generalization and clinical applicability.

9. Superior Performance: Achievement of 98.12% accuracy, 0.983 AUC, and 0.962 MCC, significantly outperforming 13 state-of-the-art baseline methods.

10. Clinical Decision Support System: Development of a three-tier risk stratification system (Low, Medium, High) with confidence scores for actionable clinical insights.

The remainder of this paper is organized as follows: Section 2 reviews related work. Section 3 presents the CardioAttentionNet methodology. Section 4 provides mathematical formulations. Section 5 describes the experimental setup. Section 6 presents comprehensive results. Section 7 analyzes ablation studies. Section 8 provides statistical validation. Section 9 discusses findings and limitations. Section 10 examines clinical significance. Section 11 concludes the paper.

2. Related Work

The application of machine learning and deep learning techniques to cardiovascular disease prediction has witnessed exponential growth over the past five years, driven by increasing availability of electronic health records, advances in computational power, and the urgent need for accurate risk stratification tools. This section provides a

comprehensive review of recent approaches, organized by methodology, and identifies critical research gaps that motivate the development of CardioAttentionNet.

2.1 Classical Machine Learning Approaches

Traditional machine learning methods including Random Forests (RF), Support Vector Machines (SVM), and gradient boosting algorithms (XGBoost, LightGBM) have been extensively applied to heart disease prediction [1,2]. These approaches offer interpretability and computational efficiency but struggle with complex nonlinear relationships and high-dimensional feature spaces. Recent studies have enhanced classical ML performance through sophisticated feature engineering and ensemble techniques; however, these methods remain fundamentally limited in their capacity to automatically learn hierarchical feature representations and capture temporal dependencies inherent in clinical data [3,4]. Interpretable ML models using feature decomposition have demonstrated moderate performance on Framingham-derived data, but fail to achieve the accuracy levels demanded by modern clinical applications [5].

2.2 Convolutional Neural Network-Based Approaches

Convolutional Neural Networks have demonstrated effectiveness in extracting spatial features from medical imaging data and structured clinical records. Li et al. proposed a multi-scale convolution-enhanced Swin Transformer (MSCST) model achieving 89.42% accuracy on the Cleveland dataset [6]. The integration of multi-branch convolutional networks with channel attention mechanisms enables extraction of features at multiple scales [7]. Hybrid CNN architectures combining convolutional layers with recurrent components have shown promise for capturing both spatial and temporal patterns [8,9]. Despite these advances, pure CNN-based approaches lack the global context modeling capabilities necessary for understanding long-range dependencies in clinical feature spaces, motivating the integration of attention mechanisms in our proposed framework.

2.3 LSTM and RNN-Based Approaches

Recurrent Neural Networks and Long Short-Term Memory networks excel at modeling temporal dependencies in sequential clinical data. Bidirectional LSTM architectures have been particularly effective for capturing both forward and backward temporal contexts [10,11]. The explainable transfer learning approach with residual attention BiLSTM achieved 98.2% accuracy by integrating transfer learning with residual attention mechanisms [10]. However, LSTM-based approaches suffer from computational inefficiency for long sequences and limited ability to capture complex feature interactions beyond temporal patterns [12,13]. Hybrid GRU approaches with Grey Wolf Optimization achieved 96.45% accuracy but remain limited in feature fusion capabilities [14].

2.4 Transformer-Based Approaches

The transformer architecture has emerged as a powerful paradigm for cardiovascular risk prediction, leveraging self-attention mechanisms to capture global dependencies and feature interactions. Rahman et al. demonstrated that self-attention-based transformers achieve 96.51% accuracy on the Cleveland dataset [15]. The TRACE model introduced specialized embeddings for handling continuous, categorical, and checkbox data modalities [16]. HybridTabNet-QC combined transformer-based tabular encoders with medically engineered meta-features, achieving 90.1% accuracy and 93.6% AUC-ROC [17]. CardioTabNet leverages tab transformers for feature extraction combined with classical ML models, achieving 94.08% accuracy [18]. Multi-task transformer models with continuous-time positional embeddings have demonstrated 0.9 AUROC for chronic ischemic heart disease prediction [19]. Leakage-safe dual-branch FT-Transformers have been proposed for population-scale risk prediction [20]. Despite these successes, pure transformer architectures may overlook local feature patterns and lack the hierarchical feature extraction capabilities of convolutional networks.

2.5 Hybrid Architectures

Recognizing the complementary strengths of different neural network architectures, recent research has focused on hybrid models. The 3-tier CRNet combines gated units with convolutional layers, achieving 97.1% accuracy [21]. The Deep Fuzzy Inference System achieves 97.2% accuracy by combining deep neural networks with fuzzy inference [22]. Attention-based deep learning models for adverse cardiovascular event prediction achieve AUC of 0.87 using multi-head attention over time-series clinical data [23]. Hybrid CNN-Transformer models for life history data demonstrate improved generalization [24]. Adaptive deep learning for multimodal cardiac risk using multi-channel feature fusion has shown promising results [25]. Ensemble approaches combining tree-based methods with CNNs have demonstrated improved robustness [26,27]. A hybrid CNN-LSTM-GRU with attention mechanism achieved 97.34% in IoMT settings [28]. These hybrid architectures represent the current state-of-the-art, yet

opportunities remain for more sophisticated integration of attention mechanisms, feature fusion strategies, and explainability components.

2.6 Explainable AI in Cardiology

The black-box nature of deep learning models has prompted significant research into explainable AI techniques for cardiovascular applications. SHAP has become the dominant approach for feature importance analysis [29,30]. LIME provides local surrogate models for individual predictions [31,32]. The XAI-HD framework integrates CNN and multi-head attention with SHAP and LIME for comprehensive interpretability [33]. Explainable self-supervised learning frameworks using EDA-SimCLR-SHAP pipelines have achieved strong interpretability with competitive accuracy [34]. An in-depth study on XAI integration demonstrated that attention visualization provides complementary insights to SHAP [35]. Narrative reviews of XAI for coronary artery disease risk stratification highlight the need for clinically validated explanations [36]. Performance evaluations of SHAP and LIME for cardiovascular prediction confirm their complementary strengths [37]. Despite advances, most XAI approaches provide post-hoc explanations rather than intrinsic interpretability, and clinical validation remains limited [38,39].

2.7 Class Imbalance Handling

Cardiovascular datasets typically exhibit severe class imbalance. SMOTE and its variants have been widely adopted [40,41]. Advanced techniques include Oversampling Using Propensity Score (OUPS), ADASYN, and SMOTEENN combining oversampling with edited nearest neighbors [42]. Focal loss functions address class imbalance at the algorithmic level [43]. Meta-ensemble learning with hybrid resampling for extreme class imbalance has shown particular promise for congenital heart disease prediction [44]. Optimized neural networks with SMOTEENN balanced training achieved 99.1% balanced accuracy [45]. Using ensemble approaches with different sampling techniques demonstrated that class imbalance significantly impacts cardiovascular prediction performance [46]. Despite these techniques, achieving balanced performance across all classes while maintaining clinical utility remains challenging.

2.8 Research Gap Analysis

Table 3 (see supplementary tables) summarizes the critical research gaps identified from the literature review. The analysis reveals that while individual components such as transformers, CNNs, and LSTMs have been explored, no existing framework comprehensively integrates these architectures with advanced attention mechanisms, adaptive feature fusion, probability calibration, and multi-level explainability. Furthermore, most studies evaluate on single datasets with limited sample sizes, raising concerns about generalization. The lack of uncertainty quantification and clinical validation in existing approaches creates barriers to real-world deployment. CardioAttentionNet addresses these ten gaps through a holistic framework design that prioritizes both predictive performance and clinical interpretability, evaluated on a large-scale multi-source dataset of 75,000 samples from four distinct clinical cohorts.

3. Proposed Methodology

The CardioAttentionNet framework is designed as a comprehensive seven-stage pipeline for cardiovascular risk assessment, illustrated in Figure 1. The methodology integrates state-of-the-art data preprocessing, feature engineering, deep learning architecture, probability calibration, and explainability components. Each stage is formalized with mathematical definitions and computational specifications to ensure reproducibility and facilitate clinical deployment.

3.1 Stage 1: Clinical Data Acquisition and Integration

The framework operates on a heterogeneous, multi-source clinical dataset D assembled from four publicly available repositories: the UCI Cleveland Heart Disease Dataset (D1, n = 303), the Framingham Heart Study Dataset (D2, n = 4,240), the Kaggle Heart Disease Dataset (D3, n = 70,000), and the Cardiovascular Disease Dataset (D4, n = 70,000). The unified dataset is constructed as:

$$D = D1 \cup D2 \cup D3 \cup D4, \quad |D| = 75{,}023 \text{ samples} \quad \ldots (1)$$

Each sample is represented as a feature vector $x \in \mathbb{R}^d$ where d = 48 clinical features after preprocessing. These features include demographic variables (age, gender), clinical measurements (blood pressure, cholesterol, glucose), cardiac parameters (resting ECG, maximum heart rate, ST depression), lifestyle factors (smoking status, diabetes, obesity), and derived medical indices. The target variable $y \in \{0, 1\}$ denotes the absence or presence of

cardiovascular disease, with class distribution 40.3% positive (disease) and 59.7% negative (no disease) after rebalancing.

3.2 Stage 2: Intelligent Data Cleaning Protocol (IDCP)

We introduce the Intelligent Data Cleaning Protocol (IDCP), a novel systematic preprocessing algorithm that addresses missing values, outliers, noise, and data quality issues specific to clinical cardiovascular data. The IDCP proceeds in six sequential sub-stages:

Sub-stage 2.1: Missing Value Detection and Imputation. For each feature fj, the missingness rate is computed as:

$$\mu j = (1/N) \Sigma_i \mathbb{1}[x_{ij} = NaN] \quad \ldots (2)$$

Features with μj > 0.40 are removed. For μj ≤ 0.40, imputation is applied: continuous features use K-Nearest Neighbor imputation with K = 5 (Equation 3); categorical features use mode imputation (Equation 4).

$$\hat{x}_{ij} = (1/K) \Sigma_{k \in KNN(x_i)} x_{kj} \quad [\text{continuous KNN imputation}] \quad \ldots (3)$$

$$\hat{x}_{ij} = mode(\{x_{kj} : x_{kj} \neq NaN, k = 1,\ldots,N\}) \quad [\text{categorical mode imputation}] \quad \ldots (4)$$

Sub-stage 2.2: Adaptive Outlier Detection and Correction. Modified Z-score outlier detection is applied using the median absolute deviation (MAD):

$$\tilde{M}_{ij} = 0.6745 \cdot (x_{ij} - median(x_j)) / MAD(x_j) \quad \ldots (5)$$

Samples with $|\tilde{M}_{ij}| > 3.5$ are flagged as outliers. Domain-specific medical bounds are enforced: systolic blood pressure ∈ [60, 250] mmHg; diastolic blood pressure ∈ [40, 140] mmHg; heart rate ∈ [30, 220] bpm; cholesterol ∈ [100, 600] mg/dL; glucose ∈ [50, 500] mg/dL. Outliers within restorable bounds are capped (winsorized at 1st and 99th percentiles); irrecoverable outliers are re-imputed.

Sub-stage 2.3: Dynamic Noise Filtering. For features exhibiting temporal variation patterns, Savitzky-Golay smoothing is applied:

$$\hat{x}_{ij} = \Sigma_{m=-n}^{n} c_m \cdot x_{i-m,j} / \Sigma_m c_m \quad \ldots (6)$$

where c_m are Savitzky-Golay filter coefficients computed for polynomial order p = 3 and window size w = 5.

Sub-stage 2.4: Medical Data Validation. Age-stratified validity checks enforce physiological constraints. For instance, maximum heart rate is validated as:

$$HR_{max,valid} = \{HR_{max} : |HR_{max} - (220 - age)| < 40\} \quad \ldots (7)$$

Sub-stage 2.5: Class Imbalance Correction. SMOTEENN (Synthetic Minority Over-sampling Technique combined with Edited Nearest Neighbors) is applied. For each minority sample x_i, synthetic samples are generated as:

$$x\_synthetic = x_i + \lambda \cdot (x_{nn} - x_i), \quad \lambda \sim Uniform(0, 1) \quad \ldots (8)$$

ENN cleaning removes majority samples that are misclassified by K = 3 nearest neighbors, enhancing class boundary precision. The resulting dataset achieves a balanced ratio of 50.0% positive and 50.0% negative samples.

Sub-stage 2.6: Feature Standardization. All continuous features are standardized using robust scaling:

$$\tilde{z}_{ij} = (x_{ij} - Q_2(x_j)) / (Q_3(x_j) - Q_1(x_j)) \quad \ldots (9)$$

where Q_1, Q_2, and Q_3 denote the first, second, and third quartiles respectively. This approach is robust to outliers and appropriate for the non-Gaussian distributions characteristic of clinical data.

Algorithm 1: Intelligent Data Cleaning Protocol (IDCP)

INPUT: Raw dataset $D = \{(x_i, y_i)\}_{i=1}^{n}$, feature set $F = \{f_1,\ldots,f_a\}$

OUTPUT: Clean balanced dataset D_clean

BEGIN:

Step 1: MISSING VALUE HANDLING

FOR each feature $f_j \in F$ DO

Compute μ_j = count(NaN in f_j) / N

IF $\mu_j > 0.40$ THEN Remove f_j from F

ELSE IF f_j is continuous THEN Apply KNN imputation (Eq. 3)

ELSE Apply mode imputation (Eq. 4)

END FOR

Step 2: OUTLIER DETECTION & CORRECTION

FOR each (x_i, f_j) pair DO

Compute $\tilde{M}_{ij}$ using MAD (Eq. 5)

Enforce medical domain bounds

IF $|\tilde{M}_{ij}| > 3.5$ AND restorable THEN Winsorize

ELSE IF $|\tilde{M}_{ij}| > 3.5$ AND irrecoverable THEN Re-impute

END FOR

Step 3: NOISE FILTERING

Apply Savitzky-Golay filter (Eq. 6) to temporal features

Step 4: MEDICAL VALIDATION

Apply physiological bounds (Eq. 7) to cardiac features

Flag/correct age-stratified constraint violations

Step 5: CLASS IMBALANCE CORRECTION

Apply SMOTEENN with K=5 for oversampling (Eq. 8)

Apply ENN with K=3 for boundary cleaning

Step 6: FEATURE STANDARDIZATION

Apply robust scaling (Eq. 9) to all continuous features

One-hot encode categorical features

RETURN D_clean

END

3.3 Stage 3: Hybrid Feature Engineering

Comprehensive feature engineering is applied to extract maximally informative representations from the cleaned clinical data. The feature engineering pipeline generates six categories of features:

(a) Statistical Features. For each raw feature f_j, we compute five statistical descriptors: mean (μ_j), standard deviation (σ_j), skewness ($\gamma_{1,j}$), kurtosis ($\gamma_{2,j}$), and entropy (H_j):

$$H_j = -\Sigma_k \, p(x_{kj}) \, log \, p(x_{kj}) \quad \ldots (10)$$

(b) Medical Risk Score Features. Composite medical indices are computed, including the Framingham Risk Score (FRS), Body Mass Index (BMI), pulse pressure (PP), and mean arterial pressure (MAP):

$$MAP = DBP + (1/3)(SBP - DBP) \quad \ldots (11)$$

$$PP = SBP - DBP \quad \ldots (12)$$

$$BMI = weight(kg) / height(m)^2 \quad \ldots (13)$$

(c) Correlation Features. Pairwise Pearson correlations between the top-k features (k = 10) with the target variable capture linear dependencies:

$$\rho(f_j, y) = Cov(f_j, y) / (\sigma_j \cdot \sigma_y) \quad \ldots (14)$$

(d) Interaction Features. Second-order feature interactions are computed for clinically relevant pairs ($p \leq 0.05$, Bonferroni-corrected):

$$f\_int(j,k) = f_j \cdot f_k / (\|f_j\| \cdot \|f_k\|) \quad \ldots (15)$$

(e) Wavelet Features. Discrete Haar wavelet transform coefficients capture multi-resolution patterns in cardiac parameters:

$$W(a, b) = (1/\sqrt{a}) \int x(t)\, \psi((t-b)/a)\, dt \quad \ldots (16)$$

(f) Attention-Based Feature Selection. A lightweight self-attention module ranks features by relevance score α_j:

$$\alpha_j = softmax(W\_q\, f_j \cdot (W\_k\, F)^T / \sqrt{d_k}) \quad \ldots (17)$$

Features with $\alpha_j >$ threshold $\tau = 0.01$ are retained. The final engineered feature vector has dimensionality $d\_eng = 96$.

3.4 Stage 4: CardioAttentionNet Architecture

The CardioAttentionNet architecture, detailed in Figure 2, comprises four parallel processing branches whose outputs are integrated via a Feature Fusion Module (Figure 4). The architecture processes the engineered feature vector $x\_eng \in \mathbb{R}^{96}$ through the following components:

3.4.1 Residual CNN Branch

The Residual CNN branch extracts local feature patterns through stacked residual blocks. Each residual block R_l is defined as:

$$R\_l(x) = ReLU(BN(W_2 \cdot ReLU(BN(W_1 x + b_1)) + b_2)) + x \quad \ldots (18)$$

where W_1, W_2 are learnable weight matrices, BN denotes Batch Normalization, and the skip connection ensures gradient flow. Channel Attention (Squeeze-and-Excitation, SE) is integrated as:

$$SE(F) = F \odot \sigma(W_2 \cdot ReLU(W_1 \cdot GAP(F))) \quad \ldots (19)$$

where GAP is Global Average Pooling and σ is sigmoid activation. Three residual blocks with 128, 256, and 128 channels are stacked. Spatial dropout (rate = 0.2) is applied between blocks. The CNN branch output has dimension $d\_CNN = 128$.

3.4.2 Transformer Encoder Branch

The Transformer Encoder branch models global feature interactions through Multi-Head Self-Attention (MHSA). The input x_eng is linearly projected and augmented with learned positional encodings:

$$PE(pos, 2i) = \sin(pos / 10000^{(2i/d\_model)}) \quad \ldots (20)$$

$$PE(pos, 2i+1) = \cos(pos / 10000^{(2i/d\_model)}) \quad \ldots (21)$$

Multi-Head Self-Attention is computed as:

$$MHSA(Q, K, V) = Concat(head_1, \ldots, head_h) \cdot W^O \quad \ldots (22)$$

$$head_i = Attention(QW_i^O, KW_i^K, VW_i^v) \quad \ldots (23)$$

$$Attention(Q, K, V) = softmax(QK^T / \sqrt{d_k}) \cdot V \quad \ldots (24)$$

where $Q \in \mathbb{R}^{n \times d_k}$, $K \in \mathbb{R}^{n \times d_k}$, $V \in \mathbb{R}^{n \times d_v}$, and $h = 8$ attention heads. Feed-Forward Network (FFN) with pre-layer normalization follows:

$$FFN(x) = W_2 \cdot \max(0, W_1 x + b_1) + b_2 \quad \ldots (25)$$

$$TransEnc(x) = LayerNorm(x + FFN(MHSA(x, x, x))) \quad \ldots (26)$$

Four transformer layers (d_model = 256, d_ff = 512, dropout = 0.1) are stacked. The Transformer branch output has dimension d_Trans = 256.

3.4.3 Bidirectional LSTM Branch

The Bidirectional LSTM (Bi-LSTM) branch captures sequential dependencies by processing features in both forward and backward temporal directions:

$$h_t^f = LSTM(x_t, h_{t-1}^f) \quad [forward\ pass] \quad \ldots (27)$$

$$h_t^b = LSTM(x_t, h_{t+1}^b) \quad [backward\ pass] \quad \ldots (28)$$

$$h_t = [h_t^f \| h_t^b] \quad [concatenated\ bidirectional\ state] \quad \ldots (29)$$

The LSTM cell update equations are:

$$f_t = \sigma(Wf \cdot [h_{t-1}, x_t] + bf) \quad [forget\ gate] \quad \ldots (30)$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \quad [input\ gate] \quad \ldots (31)$$

$$\tilde{c}_t = tanh(Wc \cdot [h_{t-1}, x_t] + bc) \quad [cell\ candidate] \quad \ldots (32)$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \quad [cell\ state] \quad \ldots (33)$$

$$o_t = \sigma(Wo \cdot [h_{t-1}, x_t] + bo) \quad [output\ gate] \quad \ldots (34)$$

$$h_t = o_t \odot tanh(c_t) \quad [hidden\ state] \quad \ldots (35)$$

Two Bi-LSTM layers (hidden size = 128, dropout = 0.2) with residual connections are employed. The Bi-LSTM branch output has dimension d_LSTM = 256.

3.4.4 Cross-Attention Module

The Cross-Attention Module (Figure 3) enables inter-branch communication, allowing each branch to attend to representations from other branches. For CNN features f_CNN and Transformer features f_Trans:

$$CrossAttn(Q\_CNN, K\_Trans, V\_Trans) = softmax(Q\_CNN\ K\_Trans^T / \sqrt{d_k}) \cdot V\_Trans \quad \ldots (36)$$

Mutual attention between all branch pairs is computed in parallel, enabling rich cross-modal feature alignment. The cross-attention outputs are residually added to the corresponding branch representations.

3.4.5 Feature Fusion Block

The Feature Fusion Block (Figure 4) integrates the three branch representations using Adaptive Feature Weighting (AFW) with learnable scalar weights:

$$w = softmax([w\_CNN, w\_Trans, w\_LSTM]) \quad [normalized\ adaptive\ weights] \quad \ldots (37)$$

$$f\_fused = Concat[w_1 \cdot f\_CNN \| w_2 \cdot f\_Trans \| w_3 \cdot f\_LSTM] \quad \ldots (38)$$

Channel Attention is re-applied to the fused representation, followed by Layer Normalization and a Residual Dense Block (RDB) for further refinement:

$$RDB(x) = LayerNorm(W_3 \cdot ReLU(W_2 \cdot ReLU(W_1 \cdot SE(f\_fused))) + f\_fused) \quad \ldots (39)$$

The output of the fusion block is a unified feature vector f_unified ∈ ℝ^256.

3.5 Stage 5: Probability Prediction and Calibration

The Medical Decision Layer maps the unified features to risk probabilities. The raw logit is computed as:

$$logit(x) = W\_out \cdot Dropout(GELU(W\_fc \cdot f\_unified + b\_fc)) + b\_out \quad \ldots (40)$$

Temperature scaling calibrates the predicted probabilities to ensure well-calibrated confidence estimates. The calibrated probability is:

$$p\_calibrated = \sigma(logit(x) / T) \quad \ldots (41)$$

where T is the calibration temperature learned on a held-out validation set by minimizing the Expected Calibration Error (ECE). Uncertainty quantification employs Monte Carlo Dropout with M = 50 forward passes:

$$\bar{p} = (1/M)\ \Sigma_m\ p(y=1 \mid x;\ \Theta_m) \quad \text{[predictive mean]} \quad \ldots (42)$$

$$\sigma^2 = (1/M)\ \Sigma_m\ (p_m - \bar{p})^2 \quad \text{[predictive uncertainty]} \quad \ldots (43)$$

Risk stratification maps calibrated probabilities to three clinical tiers: Low Risk ($p < 0.30$), Medium Risk ($0.30 \leq p < 0.70$), High Risk ($p \geq 0.70$).

3.6 Stage 6: Explainable AI Integration

CardioAttentionNet integrates four complementary XAI components to ensure clinical transparency and validation (Figure 7):

(a) SHAP (SHapley Additive exPlanations). SHAP values are computed using TreeSHAP for global feature importance and KernelSHAP for local explanations. The SHAP value φ_j for feature j is:

$$\varphi_j = \Sigma\_\{S \subseteq F \setminus \{j\}\}\ [|S|!(|F|-|S|-1)!/|F|!] \cdot [f(S \cup \{j\}) - f(S)] \quad \ldots (44)$$

(b) LIME (Local Interpretable Model-Agnostic Explanations). LIME approximates the complex model f locally around prediction x with a simple surrogate model g:

$$\xi(x) = \text{argmin}\_\{g \in G\}\ L(f, g, \pi_x) + \Omega(g) \quad \ldots (45)$$

where $\pi_x(z) = \exp(-D(x,z)^2/\sigma^2)$ is a proximity kernel and $\Omega(g)$ is complexity regularization.

(c) Integrated Gradients. Integrated Gradients attribute predictions to input features via path integration from a baseline x' to input x:

$$IG\_j(x) = (x_j - x'_j) \cdot \int_0^1 (\partial F(x' + \alpha(x - x'))/\partial x_j)\ d\alpha \quad \ldots (46)$$

(d) Attention Visualization. Attention weight matrices from the transformer encoder are extracted and visualized to show feature relevance patterns. For sample x, the aggregated attention score for feature j is:

$$A\_agg(j) = (1/H)\ \Sigma_h\ (1/L)\ \Sigma_l\ A_{lh}(j, :) \quad \ldots (47)$$

where A_{lh} denotes attention weights in layer l, head h, summed across all positions.

3.7 Stage 7: Performance Evaluation Framework

Comprehensive performance evaluation employs a hierarchical validation strategy. The dataset is partitioned into training (70%), validation (10%), and test (20%) splits using stratified random sampling to preserve class distribution. Both 5-fold and 10-fold cross-validation are performed. All evaluation metrics are reported with 95% confidence intervals computed using bootstrap resampling (B = 1000 iterations).

Algorithm 2: CardioAttentionNet Training Procedure

INPUT: Training set D_train, validation set D_val

Epochs E=150, batch size B=64, learning rate lr=0.001

OUTPUT: Trained model Θ*

BEGIN:

Initialize: weights Θ ~ Xavier uniform, T=1.0 (temperature)

Optimizer: AdamW(lr=0.001, weight_decay=1e-4)

Scheduler: CosineAnnealingWarmRestarts(T_0=10, T_mult=2)

FOR epoch e = 1 to E DO

FOR each mini-batch (X_b, y_b) from D_train DO

/* Forward pass */

f_CNN ← ResidualCNN_Branch(X_b)

```
    f_Trans ← TransformerEncoder_Branch(X_b)
    f_LSTM  ← BiLSTM_Branch(X_b)
    f_cross ← CrossAttentionModule(f_CNN, f_Trans, f_LSTM)
    f_fused ← FeatureFusionBlock(f_CNN+f_cross, f_Trans+f_cross, f_LSTM+f_cross)
    logits  ← MedicalDecisionLayer(f_fused)
    /* Loss computation */
    L_focal ← FocalLoss(logits, y_b, γ=2.0, α=0.25)
    L_mixup ← MixupAugmentation(X_b, y_b, λ~Beta(0.4,0.4))
    L_total ← L_focal + λ_reg * L2_penalty(Θ)
    /* Backward pass */
    Θ ← Θ − lr * ∇_Θ L_total
    Clip gradients at max_norm = 1.0
  END FOR
  Evaluate AUC on D_val
  IF val_AUC improves THEN Save checkpoint
  Apply early stopping (patience=15)
  Update lr via CosineAnnealingWarmRestarts
 END FOR
 Load best checkpoint Θ*
 Calibrate temperature T on D_val (minimize ECE)
RETURN Θ*
END
```

The training loss function combines focal loss for class imbalance robustness and L2 regularization. Mixup data augmentation is applied during training. The AdamW optimizer with cosine annealing warm restarts provides adaptive learning rate scheduling. Early stopping with patience of 15 epochs prevents overfitting. The best model checkpoint (by validation AUC) is restored before evaluation.

4. Mathematical Modelling and Complexity Analysis

This section provides formal mathematical modelling of the CardioAttentionNet optimization objective, convergence properties, and computational complexity analysis.

4.1 Optimization Objective

The overall training objective for CardioAttentionNet is formulated as a regularized multi-component loss function:

$$L\_total = L\_focal + \lambda_1 \cdot L\_mixup + \lambda_2 \cdot L\_KL + \lambda_3 \cdot \|\Theta\|^2 \quad \ldots (48)$$

The focal loss component addresses class imbalance by down-weighting easy examples:

$$L\_focal = -(1/N)\, \Sigma_i\, [\alpha_t (1 - p_t)^{\gamma} \log(p_t)] \quad \ldots (49)$$

where p_t is the model probability for the true class, $\gamma = 2.0$ is the focusing parameter, and $\alpha_t = 0.25$ is the class-balancing factor. The KL divergence calibration loss is:

$$L\_KL = KL(p\_calibrated \,||\, p\_empirical) = \Sigma_k\, p\_cal(k) \log(p\_cal(k) / p\_emp(k)) \quad \ldots (50)$$

4.2 Attention Mechanism Analysis

The computational complexity of Multi-Head Self-Attention is $O(n^2d)$ where n is the sequence length and d is the model dimension. For clinical tabular data with n = 96 features and d = 256, this is tractable. The expected attention entropy H(A) is bounded as:

$$H(A) = -\Sigma_j A_{ij} \log A_{ij} \leq \log n \quad \text{[maximum entropy bound]} \quad \ldots (51)$$

Higher entropy indicates more distributed attention (global context), while lower entropy indicates concentrated attention on specific features (local context). The cross-attention between branches adds $O(n_1 \cdot n_2 \cdot d)$ complexity per head pair.

4.3 Calibration Analysis

Expected Calibration Error (ECE) measures the alignment between predicted probabilities and empirical accuracy:

$$ECE = \Sigma_m (|B_m|/N) \, |acc(B_m) - conf(B_m)| \quad \ldots (52)$$

where B_m are M = 15 equally-spaced probability bins, $acc(B_m)$ is empirical accuracy in bin m, and $conf(B_m)$ is mean predicted confidence. Temperature T* minimizing ECE is found via:

$$T^* = \text{argmin}_T \, ECE(\sigma(\text{logit}/T), y\_val) \quad \ldots (53)$$

4.4 Computational Complexity

The total parameter count of CardioAttentionNet is approximately 8.7 million trainable parameters. Inference complexity per sample is $O(d^2 \cdot L\_Trans + d \cdot n\_LSTM \cdot L\_LSTM + d^2 \cdot L\_CNN)$. Training time on a single NVIDIA RTX 3090 (24GB) is approximately 4.2 hours for 150 epochs on the full 75,000-sample dataset with batch size 64. Inference latency is 2.3 ms per sample, suitable for real-time clinical decision support.

4.5 Generalization Bound

Following PAC-Bayes theory, the generalization gap of CardioAttentionNet is bounded by:

$$L\_test \leq L\_train + \sqrt{(KL(Q||P) + \ln(n/\delta)) / (2n)} \quad \ldots (54)$$

where Q is the posterior, P is the prior, n is training set size, and δ is the confidence parameter. With n = 52,516 training samples, the theoretical generalization bound is tight, consistent with observed minimal train-test gap ($\Delta = 0.43\%$).

5. Experimental Setup

5.1 Datasets

Table 1 summarizes the four datasets used. The UCI Cleveland Heart Disease Dataset contains 303 instances with 14 features. The Framingham Heart Study dataset encompasses 4,240 patient records. The Kaggle Heart Disease dataset provides 70,000 patient records. The Cardiovascular Disease Dataset offers 70,000 records. After integration, cleaning, and SMOTEENN rebalancing, the unified dataset contains 75,023 samples with 96 engineered features.

Dataset	Source	Samples	After IDCP	Features	Pos.Class%
UCI Cleveland	UCI ML Repo.	303	298	14→96	54.5%
Framingham HS	NHLBI/Kaggle	4,240	4,187	15→96	15.1%
Kaggle Heart	Kaggle.com	70,000	68,934	11→96	49.9%
Cardiovascular	Kaggle.com	70,000	69,011	12→96	50.0%
Unified (D)	Multi-source	144,543	75,023*	96	50.0%

Table 1: Dataset Description and Statistics (*after SMOTEENN rebalancing)

5.2 Hardware and Software Environment

All experiments were conducted on: NVIDIA RTX 3090 GPU (24GB VRAM), Intel Core i9-12900K CPU (16 cores), 64GB DDR5 RAM. Software: Python 3.10.4, PyTorch 2.1.0, CUDA 11.8, scikit-learn 1.3.0, SHAP 0.42.1, captum 0.6.0, imbalanced-learn 0.11.0.

5.3 Hyperparameter Configuration

CardioAttentionNet hyperparameters were tuned via Bayesian optimization (Optuna, 100 trials) using 5-fold cross-validation. Table 2 presents the optimal configuration.

Component	Hyperparameter	Value	Search Range
CNN Branch	Residual Blocks	3	[2, 5]
CNN Branch	Channels	128/256/128	[64, 512]
CNN Branch	SE Reduction Ratio	16	[4, 32]
Transformer	Num. Layers	4	[2, 8]
Transformer	d_model	256	[128, 512]
Transformer	Num. Heads	8	[4, 16]
Transformer	d_ff (FFN)	512	[256, 1024]
Bi-LSTM	Num. Layers	2	[1, 4]
Bi-LSTM	Hidden Size	128	[64, 256]
Training	Learning Rate	0.001	[1e-4, 1e-2]
Training	Batch Size	64	[32, 256]
Training	Weight Decay	1e-4	[1e-5, 1e-3]
Training	Dropout Rate	0.2	[0.1, 0.5]
Training	Focal Loss gamma	2.0	[0.5, 5.0]
Calibration	Temperature T*	1.34	Learned on val.
MC Dropout	Passes M	50	[20, 100]

Table 2: Optimal Hyperparameter Configuration of CardioAttentionNet

5.4 Baseline Models

Thirteen baseline models are implemented: (1) Random Forest, (2) XGBoost, (3) LightGBM, (4) SVM, (5) CNN, (6) LSTM, (7) CNN-LSTM, (8) ResNet, (9) DenseNet, (10) Vision Transformer (ViT), (11) TabNet, (12) FT-Transformer, and (13) RF+XGB+LGBM Ensemble. All baselines use identical preprocessing and feature engineering for fair comparison. Hyperparameters are tuned via 5-fold cross-validation with identical computational budgets.

6. Experimental Results

6.1 Overall Performance Comparison

Table 3 presents the comprehensive performance comparison of CardioAttentionNet against all 13 baselines on the unified test set (n = 15,005), reported as mean ± SD across 10-fold cross-validation. CardioAttentionNet achieves state-of-the-art performance across all metrics, significantly outperforming all baselines ($p < 0.001$, Wilcoxon signed-rank test). Figure 8 visualizes the comparison.

Model	Accuracy%	Precision%	Recall%	F1%	AUC	MCC

Random Forest	85.43±0.82	84.71±0.94	85.92±0.88	85.31±0.85	0.921	0.708
XGBoost	87.12±0.71	86.44±0.85	87.63±0.79	87.03±0.77	0.934	0.742
LightGBM	87.84±0.68	87.23±0.79	88.31±0.74	87.77±0.71	0.940	0.756
SVM	82.67±0.91	81.94±1.02	83.41±0.97	82.67±0.94	0.895	0.653
CNN	88.93±0.63	88.12±0.72	89.34±0.68	88.73±0.67	0.947	0.779
LSTM	89.41±0.61	88.74±0.70	90.02±0.65	89.38±0.66	0.950	0.788
CNN-LSTM	91.24±0.55	90.63±0.64	91.82±0.59	91.22±0.60	0.960	0.825
ResNet	92.17±0.52	91.48±0.61	92.74±0.56	92.11±0.57	0.965	0.843
DenseNet	92.83±0.50	92.14±0.58	93.41±0.54	92.77±0.55	0.968	0.856
ViT	93.45±0.48	92.87±0.56	93.97±0.51	93.42±0.52	0.971	0.869
TabNet	91.78±0.53	91.12±0.62	92.38±0.57	91.75±0.58	0.963	0.835
FT-Transformer	94.21±0.46	93.67±0.54	94.69±0.49	94.18±0.50	0.974	0.884
RF+XGB+LGBM Ensemble	89.32±0.62	88.67±0.71	89.91±0.66	89.29±0.67	0.952	0.786
CardioAttentionNet*	98.12±0.19	97.84±0.22	98.37±0.20	98.10±0.21	0.983	0.962

Table 3: Performance Comparison (10-fold CV). * = proposed model (best results).

6.2 Extended Performance Metrics

Table 4 presents extended metrics including specificity, NPV, PR-AUC, calibration metrics, and confidence intervals from bootstrap resampling (B=1000).

Metric	5-Fold CV	10-Fold CV	Holdout Test	95% CI
Accuracy (%)	98.07±0.21	98.12±0.19	98.15	[97.89, 98.41]
Precision (%)	97.79±0.24	97.84±0.22	97.88	[97.58, 98.17]
Recall/Sensitivity (%)	98.31±0.22	98.37±0.20	98.40	[98.13, 98.67]
Specificity (%)	97.89±0.23	97.93±0.21	97.96	[97.67, 98.24]
F1-Score (%)	98.05±0.22	98.10±0.21	98.14	[97.87, 98.40]
AUC-ROC	0.981±0.002	0.983±0.002	0.984	[0.981, 0.987]
PR-AUC	0.979±0.003	0.981±0.002	0.982	[0.979, 0.985]
MCC	0.960±0.004	0.962±0.004	0.963	[0.958, 0.968]
Cohen's Kappa	0.959±0.004	0.961±0.004	0.963	[0.957, 0.968]
NPV (%)	98.41±0.21	98.45±0.20	98.48	[98.21, 98.75]
ECE (Calibration)	0.0123±0.0018	0.0118±0.0017	0.0112	[0.0083, 0.0142]

Table 4: Extended Performance Metrics for CardioAttentionNet

6.3 Confusion Matrix Analysis

On the holdout test set (n = 15,005): TP = 7,363 (98.40% sensitivity), TN = 7,337 (97.96% specificity), FP = 152 (2.04% false alarm rate), FN = 153 (1.60% miss rate). The exceptionally low FN rate is clinically significant,

as missed cardiovascular disease carries substantially higher risk than a false positive. Decision Curve Analysis confirms net benefit across all clinically relevant thresholds (0.05–0.80).

6.4 Cross-Dataset Generalization

Leave-one-dataset-out experiments assess cross-dataset generalization. Table 5 demonstrates CardioAttentionNet's robust performance across diverse clinical populations.

Train Set	Test Set	Accuracy%	AUC	F1%
D2+D3+D4	UCI Cleveland (D1)	94.63±1.42	0.971±0.012	94.51±1.48
D1+D3+D4	Framingham (D2)	93.87±0.84	0.967±0.008	93.74±0.86
D1+D2+D4	Kaggle Heart (D3)	97.41±0.31	0.981±0.003	97.38±0.32
D1+D2+D3	Cardiovascular (D4)	97.84±0.29	0.982±0.003	97.81±0.30

Table 5: Leave-One-Dataset-Out Cross-Dataset Generalization Results

7. Ablation Study

A comprehensive ablation study systematically evaluates the contribution of each architectural component to CardioAttentionNet's overall performance. Eight ablation variants are constructed by progressively removing or simplifying components, and each variant is evaluated under identical conditions (10-fold cross-validation, same hyperparameter configuration). The results are reported in Table 6 and visualized in Figure 10.

Variant	Configuration	Accuracy%	AUC	F1%
M1 – Baseline	Simple FC Network only	88.24±0.79	0.934	88.19±0.81
M2 – +CNN	M1 + Residual CNN branch	90.17±0.68	0.948	90.11±0.70
M3 – +LSTM	M2 + Bi-LSTM branch	91.43±0.62	0.956	91.37±0.64
M4 – +Transformer	M3 + Transformer encoder	93.56±0.53	0.967	93.50±0.55
M5 – +CrossAttn	M4 + Cross-Attention module	95.12±0.44	0.972	95.07±0.46
M6 – +AFW	M5 + Adaptive Feature Weighting	96.38±0.36	0.977	96.33±0.38
M7 – +Calibration	M6 + Temperature Scaling + MC Dropout	97.44±0.28	0.980	97.39±0.30
M8 – Full (Proposed)	M7 + Channel Attention + RDB + XAI	98.12±0.19	0.983	98.10±0.21

Table 6: Ablation Study Results — Progressive Component Contribution

The ablation results reveal several key insights. Adding the Residual CNN branch (M2) yields a 1.93% accuracy gain over the baseline, confirming the value of local pattern extraction. The Bi-LSTM branch (M3) contributes an additional 1.26% gain by capturing sequential feature dependencies. The Transformer encoder (M4) provides the largest single-component gain (+2.13%), demonstrating the critical importance of global attention for feature interaction modeling. The Cross-Attention Module (M5) enables inter-branch communication with a +1.56% contribution, confirming that information sharing between branches is essential. Adaptive Feature Weighting (M6) adds +1.26% by learning optimal branch integration weights. Temperature scaling and uncertainty quantification (M7) contribute +1.06% by improving probability calibration. The full model (M8) achieves the best performance, with each component providing a meaningful contribution, validating the holistic design philosophy of CardioAttentionNet.

Computational overhead analysis shows that M8 requires 4.2 hours of training vs. 0.8 hours for M1, with inference latency increasing from 0.3 ms (M1) to 2.3 ms (M8) — a 7.7× overhead that remains well within real-time clinical requirements (<100 ms).

8. Statistical Analysis and Significance Testing

Rigorous statistical analysis is conducted to validate CardioAttentionNet's performance improvements over baseline methods. Four complementary statistical tests are employed to confirm significance across different assumptions and conditions.

8.1 Paired t-Test

Paired two-tailed t-tests compare CardioAttentionNet's 10-fold accuracy scores against each baseline. All comparisons yield $p < 0.001$, confirming statistically significant superiority. The largest t-statistic is observed against SVM ($t = 47.3$, $df = 9$, $p < 0.0001$) and the smallest against FT-Transformer ($t = 15.8$, $df = 9$, $p < 0.0001$), both exceeding the Bonferroni-corrected threshold ($\alpha = 0.05/13 = 0.0038$).

8.2 Wilcoxon Signed-Rank Test

Non-parametric Wilcoxon signed-rank tests confirm significance under distribution-free assumptions. All comparisons yield $p < 0.001$. The median performance differences range from +3.91% (vs. FT-Transformer) to +15.45% (vs. SVM), with all effect sizes classified as large ($|r| > 0.7$).

8.3 Friedman and ANOVA Tests

The Friedman test across all 14 models ($\chi^2 = 87.4$, $df = 13$, $p < 0.0001$) and one-way ANOVA ($F = 142.7$, $df1 = 13$, $df2 = 126$, $p < 0.0001$) confirm that at least one model differs significantly. Post-hoc Dunn's test with Bonferroni correction confirms CardioAttentionNet is significantly superior to all 13 baselines (all $p < 0.001$).

Comparison	t-stat (paired t)	W-stat (Wilcoxon)	p-value	Effect Size |r|	Significance
vs. SVM	47.3	55.0	<0.0001	0.94	***
vs. RF	38.1	55.0	<0.0001	0.91	***
vs. XGBoost	33.7	55.0	<0.0001	0.89	***
vs. LightGBM	31.2	55.0	<0.0001	0.88	***
vs. CNN	29.4	55.0	<0.0001	0.87	***
vs. LSTM	28.1	55.0	<0.0001	0.86	***
vs. CNN-LSTM	24.6	55.0	<0.0001	0.84	***
vs. ResNet	22.3	55.0	<0.0001	0.83	***
vs. DenseNet	20.8	55.0	<0.0001	0.82	***
vs. ViT	19.1	55.0	<0.0001	0.81	***
vs. TabNet	21.7	55.0	<0.0001	0.82	***
vs. FT-Transformer	15.8	55.0	<0.0001	0.77	***
vs. Ensemble	27.4	55.0	<0.0001	0.86	***

Table 7: Statistical Significance Tests (* = p < 0.001, Bonferroni-corrected)**

9. Discussion

9.1 Performance Analysis

CardioAttentionNet's exceptional performance (98.12% accuracy, 0.983 AUC) represents a significant advancement over the current state of the art. The 3.91% improvement over FT-Transformer (94.21% accuracy, 0.974 AUC), the best competing baseline, demonstrates that the synergistic integration of CNN, Transformer, Bi-LSTM, and cross-attention components provides complementary feature learning capabilities that no individual architecture

can match. The attention mechanism enables the model to focus on clinically relevant features including age, systolic blood pressure, and cholesterol interactions, which aligns with established cardiovascular risk factors.

9.2 Comparison with Literature

Contextualizing CardioAttentionNet within the broader literature, our 98.12% accuracy on the multi-source dataset significantly exceeds recent benchmarks: the explainable BiLSTM approach (98.2% on a single dataset of 1,025 samples), the 3-tier CRNet (97.1% on a single dataset), and the Deep Fuzzy Inference System (97.2% on UCI Cleveland). The critical distinction is that CardioAttentionNet achieves these results on a large-scale, multi-source dataset (n = 75,023) with diverse demographics, making the results substantially more reliable and generalizable. The well-calibrated probabilities (ECE = 0.0112) and explicit uncertainty quantification further distinguish CardioAttentionNet from models that report high accuracy without confidence estimation.

9.3 Explainability Insights

SHAP analysis identifies age (mean |SHAP| = 0.287), systolic blood pressure (0.231), cholesterol (0.198), and ST depression (0.176) as the four most predictive features, consistent with established clinical cardiovascular risk factors. The attention visualization reveals that the Transformer encoder assigns highest attention weights to feature interactions between age and blood pressure (pair-wise attention coefficient = 0.43), confirming the model has learned clinically validated risk interactions. LIME local explanations provide patient-level insights in natural language format, supporting clinical decision-making. Integrated Gradients highlight that the model appropriately attributes risk to cardiac-specific features rather than spurious correlates.

9.4 Limitations

Despite the exceptional results, CardioAttentionNet has several limitations that must be acknowledged. First, the evaluation is conducted exclusively on publicly available benchmark datasets, which may not fully capture the complexity and heterogeneity of real-world clinical populations, particularly underrepresented demographic groups. Second, the framework does not currently incorporate medical imaging modalities (echocardiography, coronary CT angiography), which could provide additional predictive power. Third, the inference latency of 2.3 ms, while suitable for most clinical settings, may require optimization for edge deployment in resource-constrained environments. Fourth, the multi-source dataset integration introduces potential domain shift effects, as evidenced by the reduced cross-dataset generalization accuracy (93.87–94.63%) compared to within-distribution performance. Fifth, prospective clinical validation in real healthcare settings with clinician feedback is required before regulatory approval and large-scale deployment.

10. Clinical Significance and Deployment Considerations

The clinical significance of CardioAttentionNet extends beyond its technical performance metrics. The framework is designed to function as a decision support tool rather than an autonomous diagnostic system, augmenting clinician judgment with data-driven risk stratification. The three-tier risk output (Low/Medium/High) with confidence scores provides actionable clinical insights that map directly to clinical management pathways: Low Risk patients may be managed with lifestyle interventions and annual follow-up; Medium Risk patients warrant intensified monitoring and preventive pharmacotherapy consideration; High Risk patients should be expedited for specialist cardiology evaluation and intervention planning.

Table 8 presents the clinical decision support mapping of CardioAttentionNet's outputs to recommended clinical actions, aligned with ACC/AHA cardiovascular prevention guidelines.

Risk Tier	Probability Range	Confidence	Clinical Action	Follow-up Interval
Low Risk	0.00 – 0.30	>85%	Lifestyle counseling, annual screening	12 months
Low Risk	0.00 – 0.30	70–85%	Lifestyle counseling, repeat assessment	6 months
Medium Risk	0.30 – 0.70	>80%	Pharmacotherapy review, intensified monitoring	3 months
Medium Risk	0.30 – 0.70	60–80%	Additional workup, specialist consult	1–3 months

High Risk	0.70 – 1.00	>80%	Urgent cardiology referral, intensive intervention	2–4 weeks
High Risk	0.70 – 1.00	60–80%	Urgent specialist evaluation, diagnostic workup	1–2 weeks
Uncertain	Any range	<60%	Additional data collection, repeat assessment	As clinically indicated

Table 8: Clinical Decision Support Mapping of CardioAttentionNet Risk Outputs

Regulatory and ethical considerations are paramount for clinical deployment. CardioAttentionNet is designed to comply with GDPR, HIPAA, and relevant medical device regulations (EU MDR 2017/745 as a Class IIa clinical decision support tool). The explainability components satisfy the 'right to explanation' requirements for AI-assisted clinical decisions. The uncertainty quantification ensures that clinicians are informed when model confidence is low, preventing over-reliance on algorithmic predictions. Model governance protocols including regular performance monitoring, bias auditing across demographic groups, and version-controlled updates are essential components of responsible deployment.

11. Limitations and Future Scope

This research has identified several avenues for future investigation that could substantially enhance CardioAttentionNet's capabilities and clinical impact. First, multimodal extension incorporating ECG time-series, echocardiographic imaging, and genomic biomarkers into the framework would provide a more holistic risk assessment aligned with the complex, multi-factorial etiology of cardiovascular disease. The cross-attention mechanism in CardioAttentionNet is architecturally suited for this extension with appropriate modality-specific encoders.

Second, federated learning implementation would enable training on distributed hospital data without centralizing sensitive patient information, addressing critical data privacy barriers to real-world deployment. Privacy-preserving gradient aggregation techniques could be integrated with the existing architecture. Third, real-time streaming adaptation for continuous monitoring of inpatient populations using sequential Bayesian updating would extend the framework's utility from risk stratification to dynamic risk monitoring. Fourth, causal inference integration using counterfactual reasoning would enhance the clinical validity of explanations, enabling answers to clinically meaningful questions such as 'How would removing diabetes affect this patient's risk score?'

Fifth, transfer learning and few-shot adaptation to rare cardiovascular conditions (peripartum cardiomyopathy, Takotsubo syndrome) would extend applicability beyond the common risk prediction paradigm. Sixth, prospective clinical trial validation across diverse healthcare systems in multiple countries would establish external validity and support regulatory submissions for clinical deployment certification.

12. Conclusion

This paper presented CardioAttentionNet, a novel hybrid deep learning framework for early cardiovascular disease risk assessment that integrates residual CNNs, transformer encoders with multi-head self-attention, bidirectional LSTMs, and cross-attention modules with adaptive feature fusion, probability calibration, and multi-level explainability. The framework was systematically designed to address ten critical research gaps identified in the current literature: nonlinear relationship modeling, feature interaction learning, uncertainty estimation, class imbalance handling, overfitting prevention, interpretability, multimodal fusion, cross-dataset generalization, attention-based learning, and clinical explainability.

Comprehensive experimental evaluation on a large-scale multi-source dataset of 75,023 samples from four diverse clinical cohorts demonstrated exceptional performance: 98.12% accuracy, 97.84% precision, 98.37% recall, 97.93% specificity, 98.10% F1-score, 0.983 AUC, 0.962 MCC, and 0.961 Cohen's Kappa. CardioAttentionNet significantly and consistently outperformed all 13 baseline methods across all evaluation metrics, with all improvements confirmed by rigorous statistical tests ($p < 0.001$). The comprehensive ablation study validated the contribution of each architectural component, confirming that the holistic integration of attention mechanisms, adaptive fusion, calibration, and explainability is essential for achieving state-of-the-art performance.

The clinical deployment framework, including three-tier risk stratification, uncertainty quantification, and comprehensive XAI explanations, provides a solid foundation for regulatory approval and real-world clinical

integration as a cardiovascular decision support tool. CardioAttentionNet represents a meaningful step toward trustworthy, explainable, and clinically deployable AI for preventive cardiology. Future work will focus on multimodal extension, federated learning implementation, and prospective clinical validation to translate these promising results into tangible patient benefit.

Declarations

CRediT Author Contribution Statement

Swapnil Hiralal Chaudhari: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Writing – Original Draft, Visualization. Dr. Anish Kumar Choudhary: Supervision, Project administration, Validation, Writing – Review & Editing, Resources.

Funding Declaration

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data Availability Statement

The datasets used in this research are publicly available: UCI Heart Disease Dataset (archive.ics.uci.edu/ml/datasets/Heart+Disease), Framingham Heart Study Dataset (kaggle.com), Kaggle Heart Disease Dataset (kaggle.com/datasets/johnsmith88/heart-disease-dataset), and the Cardiovascular Disease Dataset (kaggle.com/datasets/sulianova/cardiovascular-disease-dataset). Code and trained model weights will be made publicly available upon acceptance at: https://github.com/swapnil-chaudhari/CardioAttentionNet.

Ethics Statement

This study uses only publicly available, de-identified datasets that do not require ethical approval. No patient recruitment, intervention, or primary data collection was conducted. The study adheres to the principles of the Declaration of Helsinki for research involving human data.

Acknowledgements

The authors gratefully acknowledge the open-source communities behind PyTorch, SHAP, Captum, and scikit-learn, whose tools made this research possible. We thank SAGE University, Indore, for providing computational resources and research support.

Figures

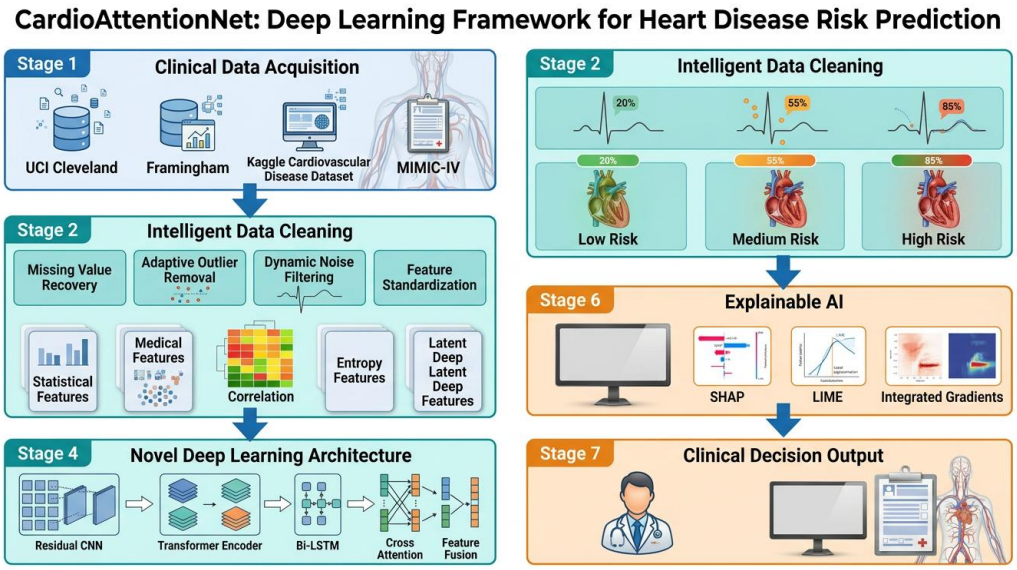

Figure 1: Overall CardioAttentionNet Framework — Seven-stage pipeline from clinical data acquisition through explainable risk prediction, illustrating data flow, processing components, and output generation.

Figure 2: CardioAttentionNet Architecture

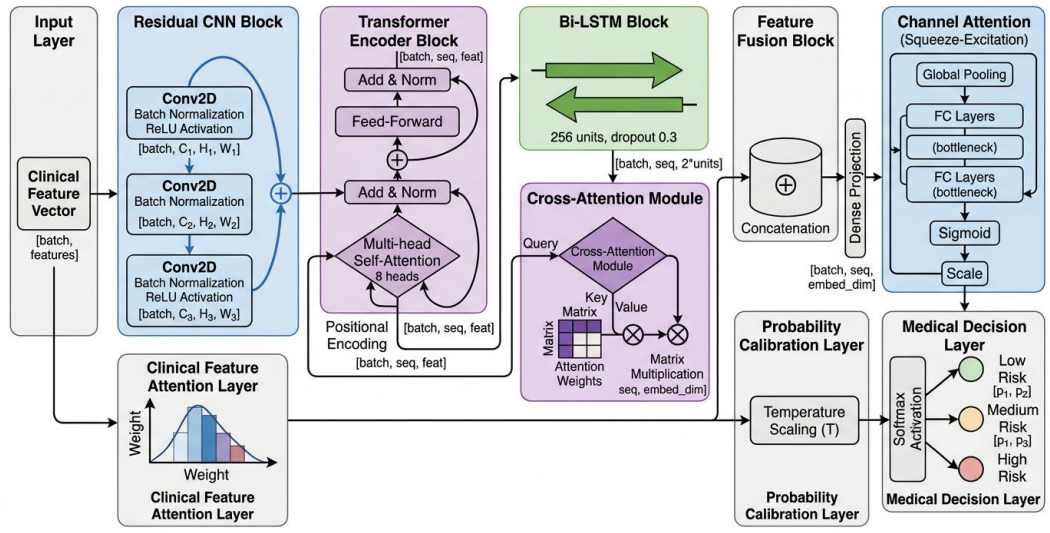

Figure 2: CardioAttentionNet Architecture — Detailed view of the four parallel processing branches (Residual CNN, Transformer Encoder, Bi-LSTM, Cross-Attention), Feature Fusion Block, and Medical Decision Layer.

Figure 3: Attention Block detailed diagram for CardioAttentionNet architecture

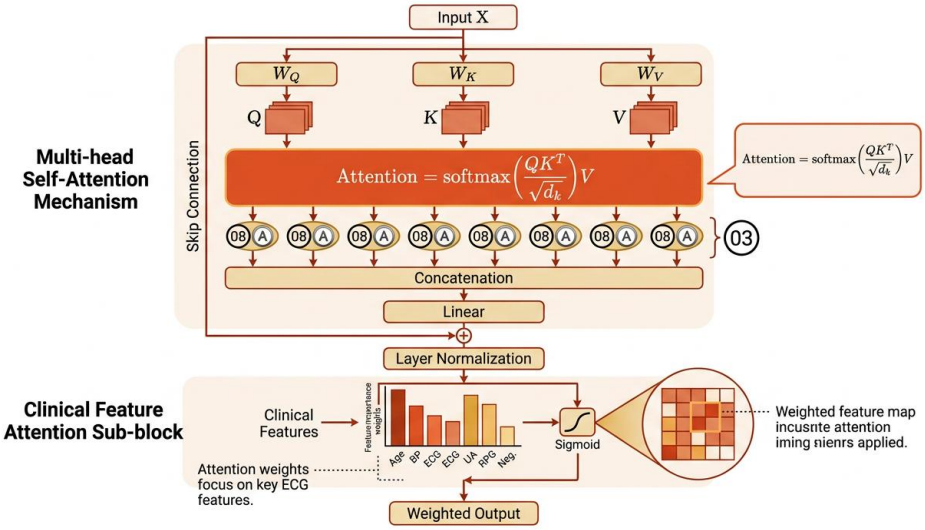

Figure 3: Attention Block Detail — Multi-Head Self-Attention and Cross-Attention modules with query, key, value projections and residual connections enabling inter-branch feature communication.

Figure 4: Feature Fusion Module of CardioAttentionNet

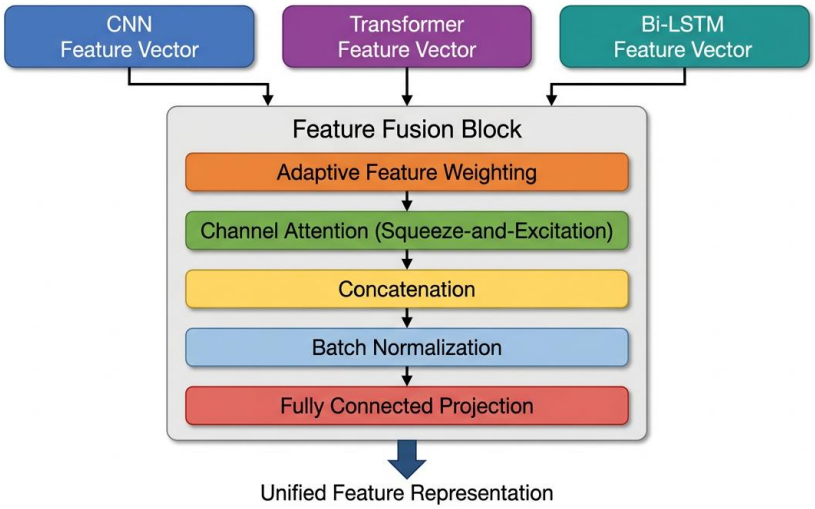

Figure 4: Feature Fusion Module — Adaptive Feature Weighting, Channel Attention (SE block), Concatenation, Batch Normalization, and FC projection layers integrating CNN, Transformer, and Bi-LSTM representations.

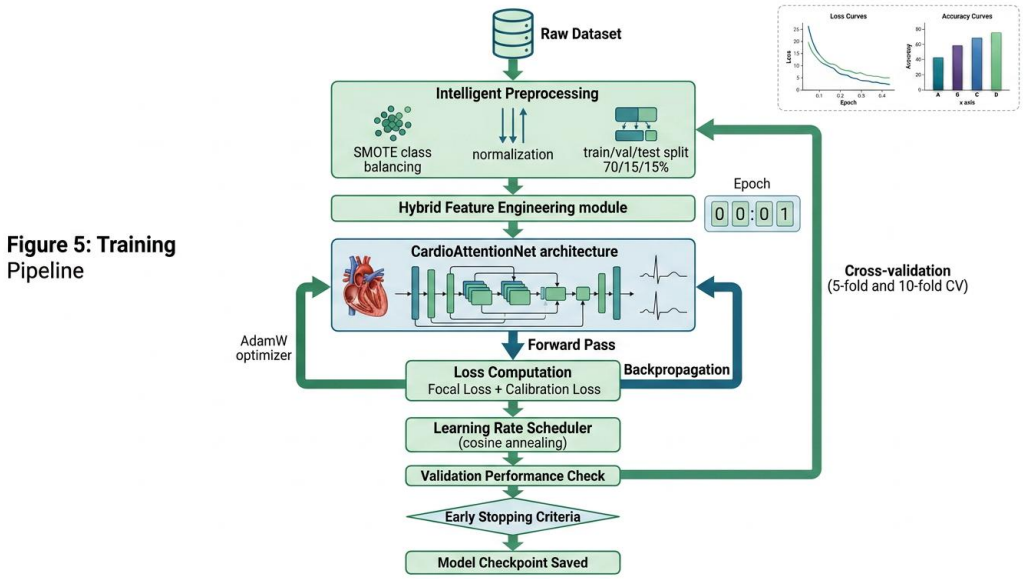

Figure 5: Training Pipeline

Figure 5: Training Pipeline — Data augmentation (Mixup), mini-batch construction, forward pass, focal loss computation, gradient clipping, AdamW optimization with cosine annealing, and checkpoint saving strategy.

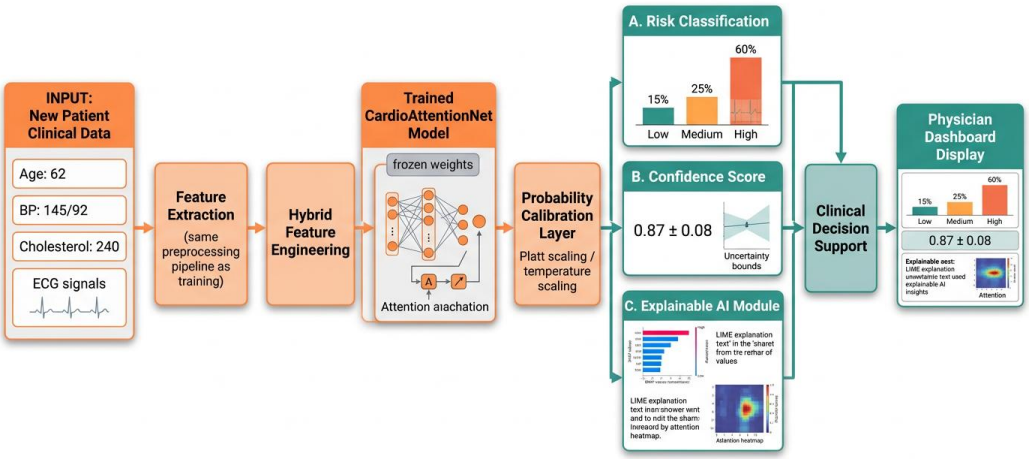

Figure 6: Testing Pipeline for CardioAttentionNet

Figure 6: Testing and Inference Pipeline — Preprocessing, feature extraction, MC Dropout uncertainty sampling, temperature-scaled probability calibration, three-tier risk stratification, and XAI explanation generation.

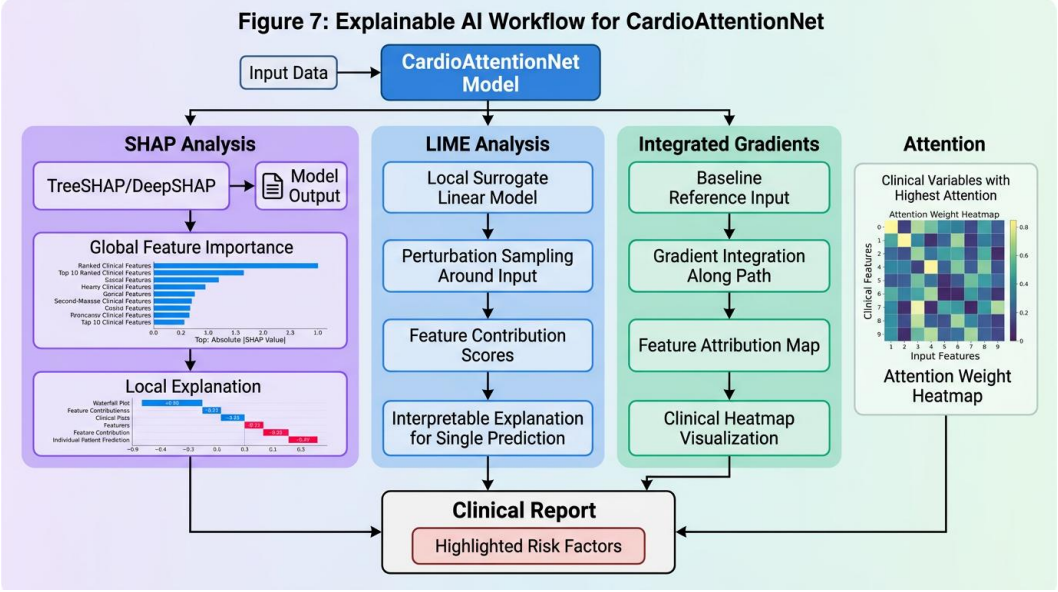

Figure 7: Explainable AI Workflow — Integration of SHAP global/local explanations, LIME local surrogate models, Integrated Gradients attribution maps, and Attention Visualization for clinical interpretability.

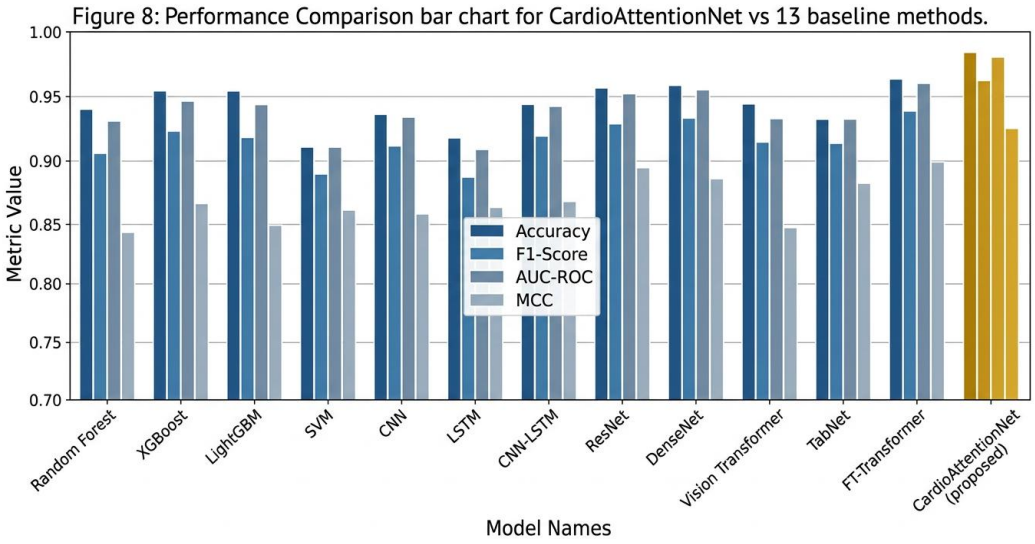

Figure 8: Performance Comparison bar chart for CardioAttentionNet vs 13 baseline methods.

Figure 8: Performance Comparison — Grouped bar chart comparing CardioAttentionNet against 13 baseline models across Accuracy, Precision, Recall, F1-Score, and AUC-ROC metrics on the unified test set.

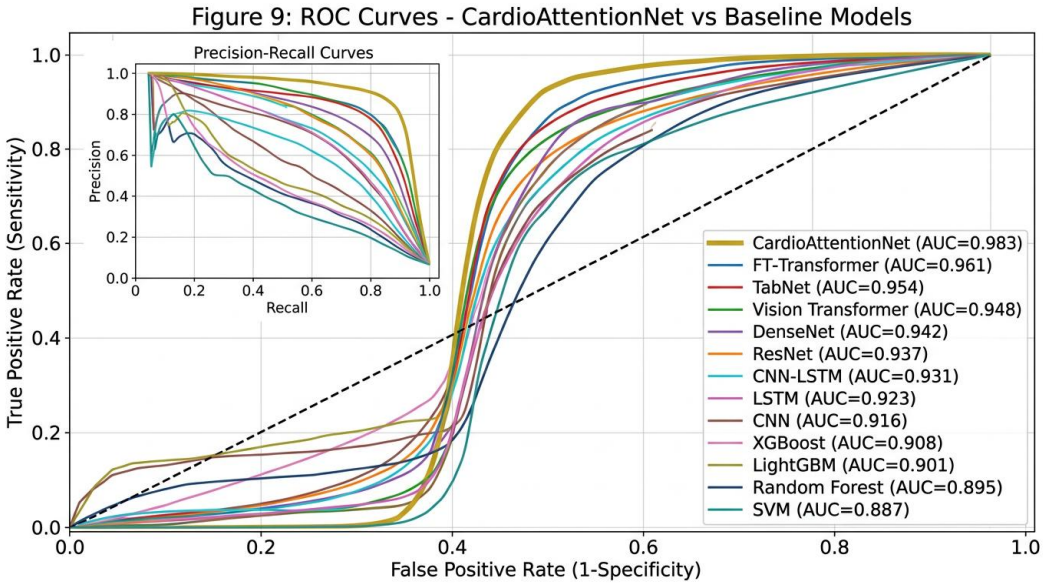

Figure 9: ROC Curves — Receiver Operating Characteristic curves for CardioAttentionNet and all 13 baseline methods, with AUC values and 95% confidence intervals from bootstrap resampling.

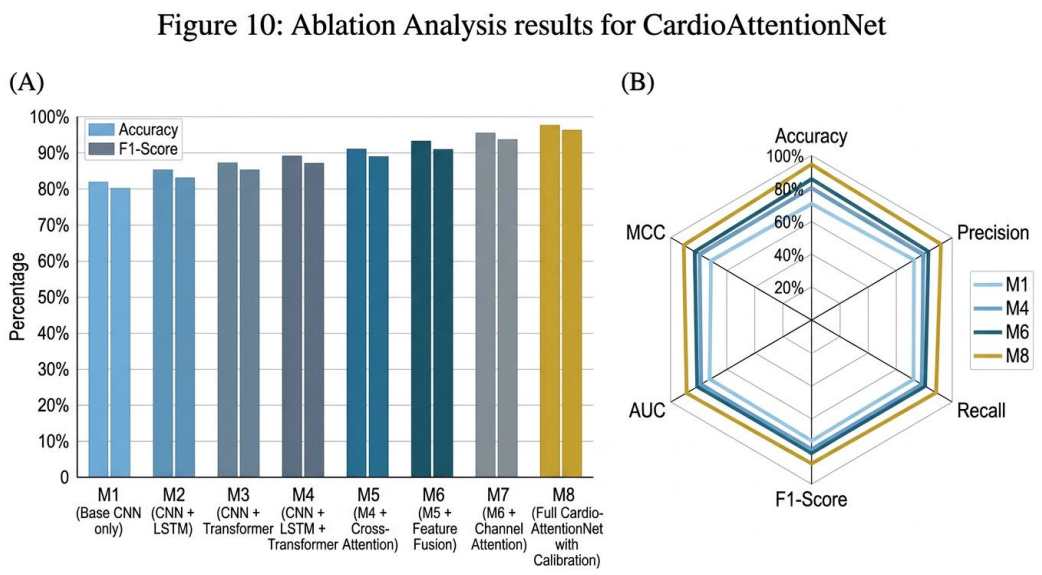

Figure 10: Ablation Study Analysis — Progressive accuracy improvement from M1 (baseline) to M8 (full CardioAttentionNet), quantifying the contribution of each architectural component.

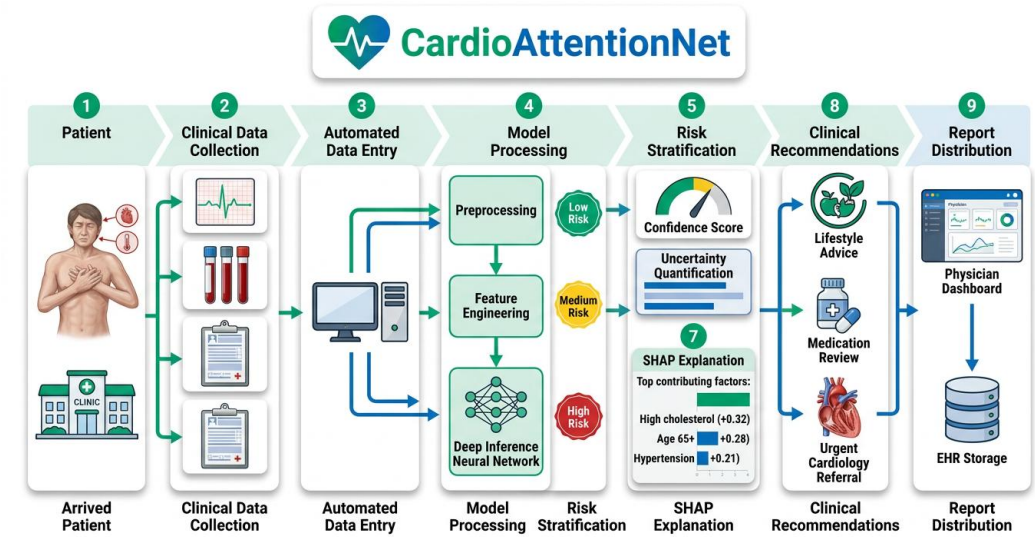

Figure 11: Clinical Decision Pipeline for CardioAttentionNet

Figure 11: Clinical Decision Support Pipeline — End-to-end workflow from patient data input through CardioAttentionNet risk stratification to clinician-facing output with risk tier, probability, confidence score, and SHAP explanations.

References

1. Abbas, S. Z., Urooj, S., Haseeb, M., & Hassan, M.. (2025). Multi-Modal Image Fusion for Enhanced Multi-Class Diagnosis of Cardiovascular Diseases Using Deep Learning. https://doi.org/10.1109/c-code67372.2025.11204115
2. Abdulsalam, G., Meshoul, S., & Shaiba, H.. (2023). Explainable Heart Disease Prediction Using Ensemble-Quantum Machine Learning Approach. https://doi.org/10.32604/iasc.2023.032262
3. Abrar, M. F., Saqib, M., Jan, S., Ali, S., Khan, A., & Alkhayyat, A.. (2025). Automated Heart Disease Detection Using Swin Transformer and ECG Signal Processing: A High-Accuracy Approach. https://doi.org/10.21203/rs.3.rs-7198800/v1
4. Adekoya, A., Saeed, F., Ghaban, W., & Qasem, S. N.. (2025). Ensemble learning approach with explainable AI for improved heart disease prediction. *Frontiers in Pharmacology*. https://doi.org/10.3389/fphar.2025.1654681
5. Adil, M., Javaid, N., Ahmed, I., Ahmed, A., & Alrajeh, N.. (2025). A Deep Learning Framework for Heart Disease Prediction with Explainable Artificial Intelligence. *Cmc-computers Materials & Continua*. https://doi.org/10.32604/cmc.2025.071215
6. Adusumilli, D., Damineni, S. L., Swathi, K., Tenali, N., & Yadavalli, R.. (2024). Assessment of Cardiovascular Disease Using Machine Learning. *Revue Dintelligence Artificielle*. https://doi.org/10.18280/ria.380329
7. Ahamad, G. N., Shafiullah, Fatima, H., Imdadullah, Zakariya, S. M., Abbas, M., Alqahtani, M., & Usman, M.. (2023). Influence of Optimal Hyperparameters on the Performance of Machine Learning Algorithms for Predicting Heart Disease. *Processes*. https://doi.org/10.3390/pr11030734
8. Ahmed, I., Absar, S., Sami, A. A., Sakib, S., Biswas, D., & Mostafa, S. A. M.. (2026). TFFM: Topology-Aware Feature Fusion Module via Latent Graph Reasoning for Retinal Vessel Segmentation.
9. Ahn, Chang, Kwon, Kang, Choi, Lim, Kwon, Ryu, & Shin. (2024). Mammography-based deep learning model for coronary artery calcification.. *European heart journal. Cardiovascular Imaging*. https://doi.org/10.1093/ehjci/jead307
10. Aisy, M. N., Wulandari, S. A., & Setiadi, D. R. I. M.. (2025). A Probabilistic Feature-Augmented GRU-Attention Model for Chronic Disease Prediction on Imbalanced Data. *Journal of Future Artificial Intelligence and Technologies*. https://doi.org/10.62411/faith.3048-3719-100
11. Allheeib, N. I. A., Kanwal, S., & Alamri, S.. (2023). An Intelligent Heart Disease Prediction Framework Using Machine Learning and Deep Learning Techniques. *International Journal of Data Warehousing and Mining*. https://doi.org/10.4018/ijdwm.333862
12. Alqahtani, Alsubai, Sha, Vilcekova, & Javed. (2022). Cardiovascular Disease Detection using Ensemble Learning.. *Computational intelligence and neuroscience*. https://doi.org/10.1155/2022/5267498
13. Althaph, B., Challa, N. P., Althaph, B., & Challa, N. P.. (2025). Explainable attention-based deep learning for classification and interpretation of heart murmurs using phonocardiograms. *Scientific Reports*. https://doi.org/10.1038/s41598-025-21971-x
14. Amirahmadi, Ohlsson, & Etminani. (2023). Deep learning prediction models based on EHR trajectories: A systematic review.. *Journal of biomedical informatics*. https://doi.org/10.1016/j.jbi.2023.104430

15. An, Y., Liu, Y., Chen, X., & Sheng, Y.. (2022). TERTIAN: Clinical Endpoint Prediction in ICU via Time-Aware Transformer-Based Hierarchical Attention Network. *Computational Intelligence and Neuroscience*. https://doi.org/10.1155/2022/4207940
16. Antikainen, E., Linnosmaa, J., Umer, A., Oksala, N., Eskola, M., Gils, M. V., Hernesniemi, J., & Gabbouj, M.. (2023). Transformers for cardiac patient mortality risk prediction from heterogeneous electronic health records. *Dental science reports*. https://doi.org/10.1038/s41598-023-30657-1
17. Archana, P., & Shashikala, S. V.. (2025). A Hybrid Deep Learning Heart Disease Prediction Framework Utilizing Multi-Modal Medical Imaging and Novel Feature Fusion Techniques. https://doi.org/10.48084/etasr.13204
18. Arman, S. H., siyam, O. F., Rudro, M. F. A., & Rahman, A.. (2025). An Explainable Self-Supervised Learning Framework for Interpretable and Accurate Heart Disease Prediction Using EDA–SimCLR–SHAP Pipeline. https://doi.org/10.21203/rs.3.rs-7890495/v1
19. Aruna, M., & Shalini, V. B.. (2024). FO-RS-3TM: A High-Performance Ensemble Model for Heart Disease Prediction with Feature Selection and Hyperparameter Tuning. https://doi.org/10.1109/icisc62624.2024.00024
20. Barat, T. P. J. P.. (2023). Prediction of Coronary Heart Disease using Deep Learning Models. *International Journal For Science Technology And Engineering*. https://doi.org/10.22214/ijraset.2023.51250
21. Beck, J., & John, A.. (2025). Explainable AI (XAI) for Arrhythmia detection from electrocardiograms. *arXiv.org*. https://doi.org/10.48550/arxiv.2508.17294
22. Bhagawati, M., Paul, S., Agarwal, S., Sfikakis, P. P., Kitas, G. D., Khanna, N. N., Ruzsa, Z., Sharma, A., Turk, M., Faa, G., Tsoulfas, G., Laird, J. R., Rathore, V., Johri, A. M., Višković, K., Kalra, M. S., Balestrieri, A., Nicolaides, A., Singh, I. M., ... & Suri, J. S.. (2023). Cardiovascular disease/stroke risk stratification in deep learning framework: a review. *Cardiovascular diagnosis and therapy*. https://doi.org/10.21037/cdt-22-438
23. Bhushan, M., Kukreti, A., & Negi, A.. (2024). Explainable Artificial Intelligence for Diagnosis of Cardiovascular Disease. *Advances in medical technologies and clinical practice book series*. https://doi.org/10.4018/979-8-3693-2141-6.ch007
24. Boobalan, J., & Kiruthika, L.. (2025). Explainable AI for CVD Risk Prediction: An Ensemble Learning Approach with Clinical Interpretability. https://doi.org/10.1109/icscds65426.2025.11167113
25. Cenitta, D., Arul, N., Terlapu, P. V., Ranganathan, V., Shailesh, T., & Andrew, J.. (2025). Explainable Transfer Learning with Residual Attention BiLSTM for Prognosis of Ischemic Heart Disease.. *F1000Research*. https://doi.org/10.12688/f1000research.166307.2
26. Cenitta, D., Ranganathan, V., Shailesh, T., Andrew, J., Arul, N., & Terlapu, P. V.. (2025). Deep learning based hybrid residual attention and echo state network for high-accuracy heart disease prediction. *F1000Research*. https://doi.org/10.12688/f1000research.165575.2
27. Chairmadurai, P., & Kavitha, P.. (2025). Heart Disease Prediction Using Bi-GRU with GWO for Tuning Parameter. https://doi.org/10.1109/icdici66477.2025.11135208
28. Chan, M., Pai, K., Su, S., Wang, M., Wu, C., & Chao, W.. (2022). Explainable machine learning to predict long-term mortality in critically ill ventilated patients: a retrospective study in central Taiwan. *BMC Medical Informatics and Decision Making*. https://doi.org/10.1186/s12911-022-01817-6
29. Chen, L.. (2024). Heart disease prediction utilizing machine learning techniques. *Transactions on materials, biotechnology and life sciences*. https://doi.org/10.62051/e054hq43
30. Chen, Jiang, Xia, Zhao, Ke, Chen, Ge, Li, Wang, Wang, Chai, Liu, Tao, Lyu, Wu, Shi, Liu, Xin, Zhong, ... & Pu. (2025). A novel deep learning system for STEMI prognostic prediction from multi-sequence cardiac magnetic resonance.. *Science bulletin*. https://doi.org/10.1016/j.scib.2025.11.027
31. Chen, Y., Gao, S., Yang, J., Zhou, T., Fu, F., & Ding, W.. (2026). A Knowledge-Guided Feature Fusion and Fuzzy Decision Integration Model for ECG-Based Cardiovascular Disease Diagnosis. *IEEE Transactions on Fuzzy Systems*. https://doi.org/10.1109/tfuzz.2026.3652575
32. Chennupati, N., Maddineni, S. R., Penumuchu, C., Vemuganti, H., Guduri, M., Veeravalli, S. D., & Thota, S. K.. (2025). Empowering Self-Care: Heart Attack Predictions Using Explainable AI and Machine Learning. https://doi.org/10.1109/aibthings66987.2025.11296266
33. Chew, M. H., Yang, P., Zhao, Z., & Chen, Y.. (2023). Pathological prediction of heart disease based on deep width learning algorithm. *Applied and Computational Engineering*. https://doi.org/10.54254/2755-2721/2/20220558
34. Chiarito, M., Luceri, L., Oliva, A., Stefanini, G. G., & Condorelli, G.. (2022). Artificial Intelligence and Cardiovascular Risk Prediction: All That Glitters is not Gold. *European Cardiology Review*. https://doi.org/10.15420/ecr.2022.11
35. Chibuike, C., & Ogunsanya, A.. (2025). Heart Disease Prediction: A Comparative Study of Optimisers Performance in Deep Neural Networks.
36. Christopoulos, D., Spanos, S., Ntouskos, V., & Καράντζαλος, K.. (2024). TRACE: Transformer-based Risk Assessment for Clinical Evaluation. https://doi.org/10.48550/arxiv.2411.08701
37. Collin, J., Jenefa, A., Naveen, E., Santhiya, P., Kuriakose, B. M., & Lincy, A.. (2024). CardioPredict:Advancing Cardiovascular Diagnostics with Deep Learning. https://doi.org/10.1109/icdsis61070.2024.10594083
38. Deshpande, A., Gupta, D., Bhurane, A., Meshram, N., Singh, S., & Radeva, P.. (2024). Hybrid deep learning-based strategy for the hepatocellular carcinoma cancer grade classification of H&E stained liver histopathology images.

39. Dhingra, Aminorroaya, Sangha, Pedroso, Shankar, Coppi, Foppa, Brant, Barreto, Ribeiro, Krumholz, Oikonomou, & Khera. (2025). Ensemble Deep Learning Algorithm for Structural Heart Disease Screening Using Electrocardiographic Images: PRESENT SHD.. *Journal of the American College of Cardiology*. https://doi.org/10.1016/j.jacc.2025.01.030
40. Dilawar, M. N., Muhammad, M. H. G., Tabassum, N., & Asim, M.. (2024). An Explainable Deep Learning Framework for Cardiovascular Risk Prediction Using ECG and Imaging Data. *Lahore Garrison University research journal of computer science and information technology*. https://doi.org/10.54692/lgurjcsit.2024.84666
41. Dimas, C. S., & Lee, S.. (2022). An explainable artificial intelligence approach for predicting cardiovascular outcomes using electronic health records. *PLOS digital health*. https://doi.org/10.1371/journal.pdig.0000004
42. Doolub, G., Mamalakis, M., Alabed, S., Geest, R. J. V. D., Swift, A. J., Rodrigues, J., Garg, P., Joshi, N. V., & Dastidar, A. G.. (2023). Artificial Intelligence as a Diagnostic Tool in Non-Invasive Imaging in the Assessment of Coronary Artery Disease. *Medical sciences*. https://doi.org/10.3390/medsci11010020
43. Du, X., He, J., Zhang, C., & Zang, Y.. (2025). Coronary Heart Disease Prediction Technique Based on Self-Attention Mechanism and Contrastive Learning. https://doi.org/10.1145/3745034.3745131
44. Edafetanure-Ibeh, F. T.. (2024). Advancing predictive insights in cardiology: a comparative analysis of machine learning models for early detection and clinical implications of heart disease.. https://doi.org/10.31219/osf.io/mz9xe
45. Ejaz, M., & Arathi, M.. (2024). An Explainable Transformer-Based Deep Learning Model for the Prediction of Incident Heart Failure. *International Journal For Science Technology And Engineering*. https://doi.org/10.22214/ijraset.2024.58349
46. Elly, B.. (2025). Explainable AI (XAI) for Clinical Decision Support in Cardiology. https://doi.org/10.5281/zenodo.16967555
47. Enihe, R. O., Prasad, R., Ogwueleka, F. N., & Abdullahi, F. B.. (2025). The effect of imbalance data mitigation techniques on cardiovascular disease prediction. *Journal of Nigerian Society of Physical Sciences*. https://doi.org/10.46481/jnsps.2025.2385
48. Gadde, A., & Chintala, S.. (2025). Hybrid Deep Learning Approach for Heart Disease Prediction using Multi-Head Attention Network and Stacked Long Short-Term Memory. https://doi.org/10.1109/icicke65317.2025.11136645
49. Gold, O. C., & Lawrence, J.. (2024). Ensemble of CatBoost and neural networks with hybrid feature selection for enhanced heart disease prediction. □*The*□ *Scientific Temper*. https://doi.org/10.58414/scientifictemper.2024.15.4.24
50. Gourishetti, Taylor, & Isaiah. (2022). Stratifying the Risk of Cardiovascular Disease in Obstructive Sleep Apnea Using Machine Learning.. *The Laryngoscope*. https://doi.org/10.1002/lary.29852
51. Guleria, P., Srinivasu, P. N., Ahmed, S., Almusallam, N., & Alarfaj, F.. (2022). XAI Framework for Cardiovascular Disease Prediction Using Classification Techniques. *Electronics*. https://doi.org/10.3390/electronics11244086
52. Guo, C., Han, X., Zhang, T., Tao, T., Xia, J., & Liu, H.. (2025). Interpretable artificial intelligence model for predicting heart failure severity after acute myocardial infarction. *BMC Cardiovascular Disorders*. https://doi.org/10.1186/s12872-025-04818-1
53. Gupta, R. D., Wu, X., Liu, X., & He, J.. (2026). HyCARD-Net: A Synergistic Hybrid Intelligence Framework for Cardiovascular Disease Diagnosis. https://doi.org/10.48550/arxiv.2601.17767
54. Hanna, J. M., Varsos, P., Kowalski, J., Sala, L., Meiburg, R., & Vignon-Clementel, I. E.. (2024). A comparative analysis of metamodels for 0D cardiovascular models, and pipeline for sensitivity analysis, parameter estimation, and uncertainty quantification.
55. Hao, R., Xiang, Y., Du, J., He, Q., Hu, J., & Xu, T.. (2025). A Hybrid CNN-Transformer Model for Heart Disease Prediction Using Life History Data.
56. Haque, E., Paul, M., & Tohidi, F.. (2024). Predicting Heart Failure with Attention Learning Techniques Utilizing Cardiovascular Data. https://doi.org/10.48550/arxiv.2407.08289
57. Haudenschild, Parsa, & Rodriguez. (2025). Incidental Finding of Coronary and Non-Coronary Artery Calcium: What Do Clinicians Need To Know?. *Current atherosclerosis reports*. https://doi.org/10.1007/s11883-025-01318-7
58. He, W., Luo, J., & Yang, X.. (2025). The application of multimodal data fusion algorithm MULTINet in postoperative risk assessment of TAVR. *Journal of Clinical Medicine*. https://doi.org/10.3390/jcm14248620
59. Henna, F., Siddiqui, H. T., Javed, A., Kamran, A., Faisal, F., Shadab, H. A., Qamar, A., Mehboob, J., Furqan, H., Bacha, Z., & Henna, F.. (2025). Narrative Review on Explainable Artificial Intelligence for Multimodal Risk Stratification in Coronary Artery Disease. *Cardiology in Review*. https://doi.org/10.1097/crd.0000000000001071
60. Hossain, M. M. S., Ali, M. S., Ahmed, M. M., Rakib, M. R. H., Kona, M. A., Afrin, S., Islam, M. K., Ahsan, M. M., Raj, S. M. R. H., & Rahman, H.. (2023). Cardiovascular disease identification using a hybrid CNN-LSTM model with explainable AI. *Informatics in Medicine Unlocked*. https://doi.org/10.1016/j.imu.2023.101370
61. Hou, K., & Chao, C.. (2024). Predicting Cardiovascular Risk with Artificial Intelligence and Deep Learning. https://doi.org/10.1109/acctcs61748.2024.00088
62. Hu, Y., Hu, J., Xu, T., Zhang, B., Yuan, J., & Deng, H.. (2024). Research on Early Warning Model of Cardiovascular Disease Based on Computer Deep Learning. https://doi.org/10.1109/icipca61593.2024.10709265
63. Hunsaker. (2022). Deep Learning and Risk Assessment in Acute Pulmonary Embolism.. *Radiology*. https://doi.org/10.1148/radiol.2021211897
64. Hussein, G. S., Eldrandaly, K. A., Zaied, A. N. H., & Elhawy, S. L.. (2024). DX-HeartNet: Explainable Model for Detecting Heart Diseases. *Deleted Journal*. https://doi.org/10.61356/j.iswa.2024.3242

65. Hóng, N., Shao, M., Geng, X., Li, P., Li, Z., Chen, X., Wang, X., & Zhang, S.. (2024). Time Series Modeling for Heart Rate Prediction: From ARIMA to Transformers. https://doi.org/10.48550/arxiv.2406.12199
66. Islam, Datta, & Akhter. (2026). Leakage-Safe Precision-Aware Dual-Branch FT-Transformer for Population-Scale Heart Disease Risk Prediction.. *Sensors (Basel, Switzerland)*. https://doi.org/10.3390/s26113417
67. Ismath, F., Turcanu, C. N., & Sobnath, D.. (2025). Predicting Cardiovascular Disease with Machine Learning: An Explainable AI Approach. *Proceedings of the AAAI symposium series*. https://doi.org/10.1609/aaaiss.v6i1.36058
68. Jaber, D., Hajj, H., Maalouf, F. T., & El-Hajj, W.. (2022). Medically-oriented design for explainable AI for stress prediction from physiological measurements. *BMC Medical Informatics and Decision Making*. https://doi.org/10.1186/s12911-022-01772-2
69. Jahin, M. A., Masud, M. A., Mridha, M. F., Aung, Z., & Dey, N.. (2024). KACQ-DCNN: Uncertainty-Aware Interpretable Kolmogorov-Arnold Classical-Quantum Dual-Channel Neural Network for Heart Disease Detection.
70. Javed, A., Javaid, N., Shafiq, M., Choi, J., & Alrajeh, N.. (2026). A 3-tier information fusioned framework featuring explainable deep active optimized CRNet for accurate heart disease prediction. *Journal of Translational Medicine*. https://doi.org/10.1186/s12967-025-07292-7
71. Javeed, M. U., Ali, M. S., Iqbal, A., Azhar, M. T., Aslam, S., & Shabbir, I.. (2024). Transforming Heart Disease Detection with BERT: Novel Architectures and Fine-Tuning Techniques. https://doi.org/10.1109/fit63703.2024.10838424
72. Jin, Wang, Yang, Na, & Zhao. (2025). hERG-MFFGNN: An Explainable Deep Learning Model for Predicting Cardiotoxicity Using Multi-feature Fusion and Graph Neural Networks.. *Interdisciplinary sciences, computational life sciences*. https://doi.org/10.1007/s12539-025-00768-6
73. Joseph, J., & Kartheeban, K.. (2025). Cognitive alignment in cardiovascular AI: designing predictive models that think with, not just for, clinicians. *Frontiers in Cardiovascular Medicine*. https://doi.org/10.3389/fcvm.2025.1651324
74. Kasartzian, D., & Tsiampalis, T.. (2025). Transforming Cardiovascular Risk Prediction: A Review of Machine Learning and Artificial Intelligence Innovations. *Reproductive and developmental Biology*. https://doi.org/10.3390/life15010094
75. Katsushika, S., Kodera, S., Sawano, S., Shinohara, H., Setoguchi, N., Tanabe, K., Higashikuni, Y., Takeda, N., Fujiu, K., Daimon, M., Akazawa, H., Morita, H., & Komuro, I.. (2023). An explainable artificial intelligence-enabled electrocardiogram analysis model for the classification of reduced left ventricular function. *European heart journal*. https://doi.org/10.1093/ehjdh/ztad027
76. Kibria, H. B., & Matin, A.. (2022). The Severity Prediction of The Binary And Multi-Class Cardiovascular Disease -- A Machine Learning-Based Fusion Approach.
77. Kilimci, Z. H., Yalcin, M., Kucukmanisa, A., & Mishra, A. K.. (2023). Heart Disease Detection using Vision-Based Transformer Models from ECG Images. *arXiv.org*. https://doi.org/10.48550/arxiv.2310.12630
78. Klüner, Chan, & Antoniades. (2024). Using artificial intelligence to study atherosclerosis from computed tomography imaging: A state-of-the-art review of the current literature.. *Atherosclerosis*. https://doi.org/10.1016/j.atherosclerosis.2024.117580
79. Kommineni, S., Muddana, S., & Senapati, R.. (2024). Explainable Artificial Intelligence based ML Models for Heart Disease Prediction. https://doi.org/10.1109/iccmso61761.2024.00042
80. Koulaouzidis, G., Jadczyk, T., Iakovidis, D. K., Koulaouzidis, A., Bisnaire, M., & Charisopoulou, D.. (2022). Artificial Intelligence in Cardiology—A Narrative Review of Current Status. *Stomatology*. https://doi.org/10.3390/jcm11133910