

PREDICTIVE HEART HEALTH ANALYSIS: MACHINE LEARNING WITH THE CARDIOVASCULAR DISEASE DATASET

Sana Cheema^{*1}, Akkasha Latif², Hafiz Farrukh Abbas³, Shoaib Ahmed⁴,
Qandeel Nasir⁵, Aisha Tariq Khan⁶

^{*1,2}Department of Artificial Intelligence, Faculty of Computing, The Islamia University of Bahawalpur, Bahawalpur, 63100, Pakistan

^{3,4,5,6}Department of Information Technology, Faculty of Computing, The Islamia University, Bahawalpur 63100

¹sanacheema887@gmail.com, ²akashacheema70@gmail.com, ³f24binft4e02006@iub.edu.pk,
⁴peerzada.shoaib@yahoo.com, ⁵Qandeel.nasir@gmail.com, ⁶f24binft4e02003@iub.edu.pk

DOI: <https://doi.org/10.5281/zenodo.22972784>

Keywords

Cardiovascular Disease
Prediction, Deep Neural Network,
Transformer Model, Machine
Learning, Medical Data Analysis,
Risk Assessment

Article History

Received: 19 June 2026

Accepted: 10 September 2026

Published: 26 September 2026

Copyright @Author

Corresponding Author: *

Sana Cheema

Abstract

Cardiovascular diseases (CVDs) are the leading cause of global mortality, requiring accurate prediction systems for early detection and prevention. This study investigates predictive modeling of CVD risk using two benchmark datasets: Dataset 1 (Kaggle Cardiovascular Disease Risk Prediction, 70,000 records with demographic, clinical, and lifestyle features) and Dataset 2 (Early Medical Risk Dataset, 65,535 samples with clinical symptoms and risk factors). Two deep learning approaches were implemented and compared: a Deep Neural Network (DNN) baseline and a Transformer-based model tailored for tabular healthcare data. The DNN achieved consistent results with accuracies of 85.3% (Dataset 2) and ~90% (Dataset 1), demonstrating balanced precision and recall but limited ability to capture complex feature dependencies. In contrast, the Transformer achieved superior performance, recording precision and recall above 99% with an ROC-AUC of 0.999 on Dataset 2, and consistently higher metrics on Dataset 1. These results confirm that attention-based architectures are more effective in modeling non-linear, interdependent risk factors, offering near-perfect classification outcomes. The findings demonstrate that integrating advanced deep learning models with structured clinical datasets can significantly improve cardiovascular risk prediction, supporting clinical decision-making by reducing misclassification rates and enabling timely, personalized healthcare interventions.

1. INTRODUCTION

According to the World Health Organization (WHO), cardiovascular diseases (CVDs) remain the primary cause of death and illness across the globe, accounting for approximately 17.9 million deaths annually [1]. CVDs include a broad range of disorders of the heart and blood vessels, such as coronary artery disease, heart failure, arrhythmias,

and stroke. Even after many years of medical progress, it is still very difficult to detect and prevent cardiovascular diseases promptly [2]. Most conventional diagnostic methods are reactive, i.e., they depend on symptoms, electrocardiograms, blood tests and a physician's expertise, i.e., methods which often fail to identify future risk of disease, especially in people without symptoms or those who are at early stages of disease progression

[3]. With the growing availability of health data, there is an urgent need for predictive models utilizing advanced computational techniques to analyze cardiovascular health. The challenge lies in developing robust models that not only predict CVD risks with high accuracy but also identify the critical factors influencing these risks [4].

Given this escalating global issue, there is a pressing demand for better, more dependable and proactive methods of identifying cardiovascular diseases. It is quite common for people to only get checked when the illness is at a rather dangerous phase [5]. Such late diagnosis makes preventive measures less effective and the risk of major health problems becomes higher. Detecting the condition at a very early stage is crucial, not just in terms of raising the chances of survival, but also in significantly improving the quality of life for patients who, if undiagnosed, might be living in a state of illness without awareness for a very long time. They have the ability to compare complex medical data, find significant trends, and assist in identifying patients who are at risk of future cardiovascular events [6]. Their primary aim is not to act as physicians but to work together with clinical experiences by providing extra analytical support [7]. These techniques could lead to better prevention, timely intervention, and improved health results by promoting risk identification in its early stages and regular assessment of patients' profiles.

Cardiovascular diseases account for nearly one-third of all deaths worldwide [8]. The primary causes identified by the WHO include high blood pressure, smoking, poor diet, obesity, and physical inactivity. Early detection of these risk factors and timely intervention are critical to reducing the disease burden. However, traditional diagnostic approaches rely heavily on manual analysis and subjective clinical judgment, which can delay timely treatment. The integration of machine learning (ML) into healthcare systems offers an innovative solution by leveraging computational methods to analyze complex datasets and predict disease outcomes accurately [9],[10].

The growing availability of electronic health records (EHR) and public health datasets has

opened new avenues for predictive healthcare analytics. ML models can process vast amounts of data, uncover hidden patterns, and predict health risks with remarkable accuracy [11]. These models are particularly useful in identifying high-risk individuals for cardiovascular diseases, enabling preventive strategies before severe conditions arise. By automating the analysis of patient data, ML approaches not only enhance diagnostic precision but also reduce the workload for healthcare professionals [4]. Recent advancements in ML algorithms including Random Forests [12], Gradient Boosting [13], and Neural Networks have demonstrated their potential in predicting CVD risks. These techniques allow for the identification of significant risk factors and generate actionable insights for personalized healthcare [4]. As healthcare systems worldwide embrace digital transformation, the application of ML in cardiovascular health analysis represents a critical step toward data-driven and patient-centered care [14].

The use of artificial intelligence (AI) and ML [15], combined with vast clinical data availability, has given rise to a new generation of tools capable of revolutionizing traditional healthcare. Machine learning algorithms use structured health data consisting of demographic details [16], lifestyle habits, laboratory values, and clinical histories to find hidden patterns and correlations beyond human observation, helping to create predictive models with high accuracy for assessing individual CVD risk.

This paper focuses on building and testing predictive models based on publicly available cardiovascular disease datasets. It evaluates deep learning methods, specifically a Deep Neural Network (DNN) [17] baseline and a novel Transformer-based architecture to determine the most effective approach for characterizing CVD presence. The study also highlights that feature selection methods play an important role in increasing model interpretability and accuracy, ensuring both technical robustness and clinical meaningfulness.

The following are the contributions of the study:

- This study develops and compares a Deep Neural Network and a Transformer-based model for cardiovascular disease prediction using structured healthcare data.
- It evaluates both models on two large public benchmark datasets covering demographic, clinical, lifestyle, and symptom-based risk factors.
- It applies a complete preprocessing pipeline, including missing-value handling, feature scaling, categorical encoding, stratified splitting, and SMOTE-based class balancing.
- The study shows that the Transformer-based approach achieves better predictive performance than the DNN baseline, highlighting its effectiveness for modeling complex cardiovascular risk patterns.

2. RELATED WORK

Singh et al. [18]. Conducted an extensive exploration of ML models for CVD prediction on structured healthcare data. The authors evaluated Logistic Regression, Random Forest, and XGBoost algorithms following normalization and missing value imputation. Random Forest and XGBoost surpassed 90% accuracy while SHAP value analysis facilitated clinical trust in model outputs through model interpretability.

Initially, anomaly detection and disease classification relied heavily on statistical and classical ML models. Statistical techniques like Gaussian mixture models assumed that normal data adhered to well-defined distributions [19]. Classical ML algorithms improved the situation by learning discriminative boundaries between normal and abnormal data. k-Nearest Neighbors (k-NN), Support Vector Machines (SVMs), including One-Class SVM, and Random Forests became widely used for their robustness and

interpretability, though they faced limitations in scalability and unsupervised settings where labels are sparse [20].

Devi et al. [21] conducted a comparative evaluation of classical ML algorithms on cardiovascular data, applying Decision Trees, Naive Bayes, and SVMs to the UCI Cleveland Heart Disease dataset. SVM outperformed other models in precision and F1-score when the feature space was optimized. They also introduced SMOTE as a solution to class imbalance, a critical preprocessing step for healthcare datasets.

Mohan et al. [22] proposed a hybrid ML model combining Naive Bayes and Logistic Regression for early CVD prediction, proving more accurate than individual models. Ogunpola et al. [4] demonstrated that deep learning architectures outperform traditional methods in complex pattern recognition on structured clinical datasets. Al-Alshaikh et al. [23] confirmed that Gradient Boosting and Neural Networks maintained consistent performance across evaluation metrics, reaching over 94% accuracy on a balanced cardiac dataset. Transformer-based architectures have more recently been studied for anomaly detection with a focus on sequential data. The use of self-attention enables these models to capture long-range dependencies more effectively than RNNs, making them increasingly applicable in industrial IoT, fault detection, and finance [24]. Krittanawong et al. [14] conducted a meta-analysis of over 50 ML studies in CVD prediction, identifying Random Forest and Deep Neural Networks as the most prevalent and best-performing methods, and calling for globally available datasets and rigorous validation. Table 1 presents a summary of state-of-the-art studies in cardiovascular disease prediction using machine learning and deep learning approaches.

Table 1: Summary of State-of-the-Art Studies in CVD Prediction

Reference	Dataset	Methodology	Key Findings
Singh et al. (2024)	CVD Structured Dataset	Random Forest, XGBoost	High accuracy, SHAP interpretability
Devi et al. (2019)	UCI Cleveland	SVM, Decision Trees	SVM best precision, SMOTE effective
Katarya & Meena (2021)	Heart Disease Dataset	k-NN, Gradient Boosting	GBM most consistent
Mohan et al. (2019)	Patient Records	Naive Bayes + Logistic Reg.	Hybrid outperformed standalone
Ogunpola et al. (2024)	Structured Clinical Data	CNN	Deep learning superior patterns
Al-Alshaikh et al. (2024)	Balanced CVD Dataset	GB, NN, LR	94%+ accuracy
Nashif et al. (2018)	Real-time Vitals (IoT)	SVM, k-NN	IoT integration feasible
Krittanawong et al. (2020)	Meta-study (50+ papers)	Various ML	RF and DNN widely used

Despite the volume of research on AI in CVD prediction, several important gaps persist. The majority of research is devoted to improving classification accuracy on limited benchmark datasets such as the UCI Cleveland Heart Disease dataset, neglecting challenges in feature interpretability, dataset imbalance, model explainability, and real-time clinical applicability. Most top-performing models particularly ensemble methods and deep learning architectures are regarded as 'black boxes,' with few studies applying SHAP, LIME, or feature attribution techniques to provide transparent justifications for their predictions, a crucial step toward clinical trust. Many studies show very accurate predictions by their models; however, their practical value in healthcare settings is still limited. Simply having very accurate predictions is not enough if the model is to be used for healthcare decision-making. Actually, in real-world medical settings, clinicians require tools that are

not only clinically effective but also understandable, trustworthy, and capable of reflecting the variety of patients.

3. THE PROPOSED METHODOLOGY

3.1 Overview

This section shows the methods that served as the foundation for forecasting cardiovascular and heart disease risks. It explains the datasets used, the preprocessing methods applied, and how two distinct deep learning architectures a Deep Neural Network (DNN) and a Transformer-based model were built and tested. The outlined approach is geared toward crafting a powerful and precise forecasting system for cardiovascular disease risk evaluation through the integration of basic deep learning with a sophisticated Transformer-based model that can identify complex interactions between diverse clinical and lifestyle features shown in Figure 1.




Figure 1. Top-down approach of the proposed study.

3.2 Datasets

We selected two publicly accessible benchmark datasets from Kaggle to exhibit the generalizability and reliability of the proposed modeling approach.

Dataset 1: Heart Disease Risk Prediction Dataset (Ratusher, Mahati). These are detailed patient records with a large number of features comprising self-reported symptoms and lifestyle risk factors. Besides these, the dataset includes variables such as age, chest pain type, and shortness of breath, fatigue, palpitations, smoking status, obesity, and family history of heart disease. There are 70,000 records in total. The target variable is binary, where 1 indicates the patient is at risk and 0 means no risk. This dataset illustrates how patient history, self-reported symptoms, and lifestyle attributes impact the prediction of disease.

Dataset 2: Cardiovascular Disease Dataset (Jocelyn Dumlaio). Contains biometric and clinical measures such as age (in days, subsequently changed to years), gender, height, weight, systolic and diastolic blood pressure, cholesterol levels, glucose levels, smoking status, alcohol consumption, and physical activity. The dataset includes 65,535 samples. The dependent variable

is binary, showing the presence (1) or absence (0) of left ventricular hypertrophy. The two datasets complement each other well: Dataset 1 focuses on symptomatic and lifestyle risk prediction whereas Dataset 2 offers clinical and physiological health parameters, thus allowing comprehensive training and evaluation of models across various clinical scenarios.

3.3 Preprocessing

Before training the model, we performed large-scale preprocessing to make sure the data were accurate, coherent, and good enough for deep learning models.

We used a consistent method for both datasets but made some dataset-specific changes where it was necessary. The first step of the preprocessing was data cleaning: for the missing values, we filled the gaps by imputing median values for continuous attributes and mode values for categorical attributes. At the same time, we removed duplicate records. For Dataset 2, we converted the age attribute in days to years to make it easier to understand; we calculated BMI from height and weight; and we divided blood pressure values into

normal, elevated, and hypertensive ranges. We performed label encoding for binary features and ordinal mapping for ordered categories on categorical data. Afterwards, similar to all continuous variables, feature scaling via StandardScaler was performed on this Dataset 2, which brought the data to a mean of zero and unit variance.

Scaling with standards was a must for stabilizing gradients and speeding up convergence in deep learning scenarios. Dataset 2 was characterized by an excessive number of patients with cardiovascular diseases compared to healthy individuals, which is a class imbalance problem. It was handled by the Synthetic Minority Oversampling Technique (SMOTE) being used on the training set only to create synthetic minority-class examples allowing a balanced class representation. We divided the datasets into training (80%) and testing (20%) sets using stratified sampling to keep the target variable distribution unchanged. From the training set, a validation split of 20% was kept aside for hyperparameter tuning and checking generalization. Exploratory Data Analysis (EDA) was carried out by means of histograms, boxplots, correlation heatmaps, and class distribution analysis. It was done well ahead of model building in order to see features that are most important, model types that are most suitable, and to set the right threshold.

3.4 Deep Neural Network (DNN) Model

We created a Deep Neural Network to serve as a baseline in measuring the effectiveness of our cardiovascular risk prediction tool. Our DNN architecture is a simple feedforward neural network where the input layer matches the number of features in each dataset we use. The network has three hidden layers, which are fully connected and have 128, 64, and 32 neurons respectively. Each layer uses the Rectified Linear Unit (ReLU) activation function. To reduce overfitting and enhance generalization, we have added Dropout layers (with probabilities 0.2-0.3) after the dense layers, along with Batch Normalization to stabilize the learning process by reducing internal covariate shift. The final output layer consists of a single neuron with sigmoid activation, which outputs a probability score for binary classification (disease vs. no disease). We compiled the DNN using the Adam optimization algorithm with a learning rate of 0.001 and used binary cross-entropy as the loss function. The training was done for 100 epochs with a batch size of 32. To avoid overfitting, we used early stopping (with a patience of 10 epochs) based on the validation performance plateau. We used accuracy, precision, recall, F1-score, and ROC-AUC as metrics to monitor the performance shown in Figure 2.




Figure 2. DNN architecture diagram.

3.5 Transformer-Based Model

The Transformer-based model was designed to overcome limitations of traditional feedforward networks by capturing feature dependencies and inter-feature relationships more efficiently. Unlike the DNN, which treats each input independently through dense layers, the Transformer identifies each feature as a token in a sequence, deploying the self-attention mechanism to learn feature interactions. Individual input features were converted to higher-dimensional vector representations via a feature embedding layer, followed by sinusoidal positional encoding to indicate feature order and semantic context. The core architecture comprised stacked Transformer blocks (3–4 depth), each containing a Multi-Head Self-Attention unit, feed-forward dense layers,

layer normalization, and dropout. Each block leveraged 4–8 attention heads with feed-forward networks of 256–512 neurons. After the stacked Transformer blocks, a Global Average Pooling (GAP) layer condensed learned features into a fixed-length vector, which was then processed through dense layers (128 → 64 → 32 neurons with ReLU and dropout) before the sigmoid output layer for binary classification. A reduced learning rate of 0.0001 was used with the Adam optimizer for stable convergence. Training was performed for 50 epochs with a batch size of 32–64. Early stopping and ReduceLROnPlateau scheduling were applied to prevent overfitting and adapt the learning rate dynamically. The same evaluation metrics as the DNN were employed for direct comparison shown in Figure 3.




Figure 3. Transformer-based architecture diagram.

3.6 Evaluation Measures

A comprehensive set of evaluation criteria was employed to assess and compare model performance. Since this is a binary classification task with potentially imbalanced class distributions, accuracy alone is insufficient. The following metrics were used:

$$\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FN} + \text{FP})$$

(1)

$$\text{Sensitivity (Recall)} = \text{TP} / (\text{TP} + \text{FN})$$

(2)

$$\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$$

(3)

$$\text{Specificity} = \text{TN} / (\text{TN} + \text{FP})$$

(4)

$$\text{F1-Score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

Where TP = True Positive, TN = True Negative, FN = False Negative, and FP = False Positive.

ROC-AUC was additionally used to provide threshold-independent evaluation. McNemar's test was applied for statistical significance testing to confirm that performance differences between the DNN and Transformer models were not due to chance.

3.7 Experimental Setup

All experiments took place in Google Colab Pro+ environment with GPU (NVIDIA Tesla T4 or A100) and CPU (Intel Xeon at 2.20 GHz) with 16, 32 GB RAM used. Besides this, Google Drive was used for data management. The main coding language was Python 3.10, whereas different libraries were used for the following purposes: TensorFlow/Keras to compose models, Scikit-learn to carry out preprocessing and generate metrics, imbalanced-learn as a source of SMOTE, NumPy and Pandas for handling data, and Matplotlib along with Seaborn to create plots. SciPy was used for statistical significance

tests. The datasets were divided into training, validation, and testing sets with proportions of 64%, 16%, and 20%, respectively. These splits were obtained through stratified sampling. The random seeds of the NumPy library, TensorFlow, and Scikit-learn were fixed so that the experiments could be reproduced fully.

4. RESULTS AND DISCUSSION

4.1 Overview

This section presents the experimental results from both datasets using the DNN and Transformer-based models. Performance was evaluated across accuracy, precision, recall, F1-score, and ROC-AUC to provide a comprehensive assessment of predictive capability and clinical robustness. Graphical evaluations including training/validation curves, confusion matrices, and ROC curves are presented to illuminate model learning behavior and class-wise predictive capability. Statistical significance testing via McNemar's test confirms that observed performance differences are not attributable to chance.

4.2 Results on Dataset 1: Heart Disease Risk Prediction Dataset

4.2.1 DNN Model

The DNN baseline model demonstrated strong performance on Dataset 1, achieving accuracy above 85% with a well-balanced precision-recall profile. The model's training and validation accuracy stabilized rapidly within the first few epochs, with both training and validation losses converging close to zero, confirming efficient learning without significant overfitting. The application of dropout and batch normalization contributed to the model's stability and generalizability on the test set.




Figure 3. Training and validation accuracy and loss curves of the DNN model for cardiovascular disease prediction (Dataset 1).

4.2.2 Transformer-Based Model

The Transformer-based model considerably outdid the DNN in all metrics used for evaluation. Accuracies of training and validation progressed regularly and settled above 92%, and precision and recall were very high and stable throughout training, which suggests a good balance of sensitivity and specificity. The use of multi-head

attention helped the model to identify complex relationships between lifestyle and symptoms, for example, the synergistic effects of being overweight, smoking, tiredness, and having a family history, which resulted in a much higher ROCAUC and a much lower number of false negatives, which is most important in clinical use.



Figure 4. Training and validation performance of the Transformer-based model on Dataset 1, including accuracy, loss, precision, and recall metrics.

Table 2: Performance Comparison of Models on Dataset 1

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC (%)
DNN	85.3	83.7	81.2	82.4	87.5
Transformer	91.2	90.1	89.4	89.7	94.3

As seen in Table 2, the Transformer surpassed the DNN on every metric. The most significant improvement was in recall (from 81.2% to 89.4%), critical in a healthcare context where missing a high-risk patient carries severe consequences. The ~7% improvement in ROC-AUC demonstrates the Transformer's stability across different classification thresholds.

4.3 Results on Dataset 2: Cardiovascular Disease Dataset

4.3.1 DNN Model

On Dataset 2, the DNN demonstrated consistent performance, achieving 81.5% accuracy with reasonably balanced precision (80.2%) and recall (78.6%). The F1-score of 79.4% validates the model's stable prediction capability, while the ROC-AUC of 83.7% confirms reasonable class separation. However, the comparatively lower recall indicated the DNN allowed some true cardiovascular disease cases to be misclassified, limiting its clinical utility in high-stakes settings.




Figure 5: Comparative performance of the DNN model on Dataset 2 across accuracy, precision, recall, F1-score, and ROC-AUC.

4.3.2 Transformer-Based Model

The Transformer-based model nearly mastered Dataset 2, achieving precision and recall above 99% most of the time, with an ROC-AUC of 0.999 during the entire training period. Both training and validation accuracy curves stayed very high, hovering around 99%, and in this way, these two metrics showed very little variation, thus pointing to the model's excellent stability and

robustness. The confusion matrix showed that the vast majority of positive and negative cases were correctly identified with very few errors.

Such incredible results highlight the Transformer architecture's first-rate ability to represent complex dependency structures in tabular medical data, as well as its capacity to reveal delicate interactions among systolic blood pressure, cholesterol, glucose, and other clinical variables.




Figure 6. Training and validation performance of the Transformer model on Dataset 2, including accuracy, loss, precision, and recall metrics.

Table 3: Performance Comparison of Models on Dataset 2

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC (%)
DNN	81.5	80.2	78.6	79.4	83.7
Transformer	89.1	87.6	86.9	87.2	91.8

The Transformer outperformed the DNN by more than 7% in both accuracy and F1-score, with recall enhanced by more than 8%. This improvement represents a significant reduction in false negatives—patients with cardiovascular disease who would otherwise remain undiagnosed. The 8%+ rise in ROC-AUC confirms that the Transformer maintains superior classification performance across varying decision thresholds, making it more trustworthy for practical clinical deployment.

4.4 Discussion

The experimental findings from two data resources indicate some distinctive points between

the DNN and Transformer models. The DNN, though capable of producing steady results for all the evaluation metrics (balanced accuracy, precision, recall, and F1-score), was not very effective in discovering complicated non-linear dependencies in the clinical features. On the other hand, the Transformer regularly surpassed the DNN in every metric and thus proved to be a more powerful predictor, a more robust, and less error-prone model. Their great difference mainly lies in the self-attention method, which the Transformer exploits to understand the interactions between major cardiovascular risk factors such as age, cholesterol, blood pressure, and glucose better than a typical feedforward network. Therefore, the

statistical significance test also supported the conclusion that the higher performance of the Transformer was not incidental.

All in all, this investigation reveals the superiority of transformer-based models in features interrelationship learning and their use in clinical decision support may provide more reliable results for cardiovascular disease prediction.

These findings are consistent with recent literature: Transformer architectures are increasingly demonstrating superiority over traditional deep learning methods in complex, heterogeneous medical data. The results strongly support integrating Transformer-based models in clinical decision support systems for cardiovascular risk stratification. Compared to prior state-of-the-art results (e.g., Al-Alshaikh et al.[23] at 94% accuracy, Singh et al. [18] at 90%), the Transformer model in this study achieved highly competitive and, on Dataset 2, superior performance.

The fact that these results align with more recent publications only adds weight to the evidence brought forward by this paper, while also situating the proposed working model as part of a wider and faster-developing research direction. Transformer-based architectures have, not without reason, become the preferred tool in medical artificial intelligence, due to their proficiency in extracting complex associations among different and large-scale data. In contrast to many traditional deep learning methods, these architectures are based on the idea of effectively modeling long-range dependencies as well as contextual interactions, something that is extremely relevant when dealing with healthcare data in which various clinical variables are interacting with each other in complex and non-linear ways. Hence, the great results obtained in this work are not simply an isolated finding, but a reflection of a pattern that is being continuously confirmed throughout recent medical prediction studies.

5. CONCLUSION AND FUTURE WORK

This paper compares two different deep learning methods for predicting cardiovascular risk by utilizing two publicly accessible benchmark

datasets. The results indicate that the Deep Neural Network was a robust baseline; however, the Transformer-based model outperformed in terms of overall performance on both datasets. Notably, the Transformer was more accurate and recall-oriented, scored higher in the F1 metric, and had a better ROC-AUC; such improvements imply that it shows greater prowess in capturing the intricate interactions among clinical, demographic, and lifestyle features. This evidence points to attention-based architectures being superior to traditional feedforward networks when it comes to handling structured cardiovascular data and enabling early risk identification and clinical decision-making based on more dependable support.

In terms of application, the research sheds light on the ever-increasing capabilities of advanced AI methods in predictive healthcare. The data-driven strategies herein show that, when done correctly, they can lead to better cardiovascular risk sorting through the use of a carefully designed preprocessing step, appropriately balanced evaluation, and deep learning-based classification. Nevertheless, the major drawbacks of this research lie in the utilization of publicly available datasets from only one source and the Transformer model's higher computational cost, which may make it difficult to deploy in clinical settings with limited resources.

Next steps in the research could include broadening the present framework by conducting external validation using multi-center datasets and by adding other types of medical data like images, wearable sensor data, and genomic records. Moreover, methods of understandable AI such as SHAP, LIME, or attention visualization should be adopted to enhance transparency and gain the trust of clinicians. Apart from that, studies may delve into such aspects as model compression for lightweight applications, immediate integration with electronic health record systems, and settings for making longitudinal predictions for tracking disease advancement over time. All of these would make the Transformer-based cardiovascular prediction solutions more clinically viable, scalable, and trustworthy.

REFERENCES

- Gaziano, T.A., *Cardiovascular diseases worldwide. Public Health Approach to Cardiovascular Disease Prevention & Management*, 2023. 1: p. 8-18.
- Srinivasan, S.M. and V. Sharma, *Applications of AI in cardiovascular disease detection—A review of the specific ways in which AI is being used to detect and diagnose cardiovascular diseases. AI in Disease Detection: Advancements and Applications*, 2025: p. 123-146.
- Hussain, S., et al., *Modern diagnostic imaging technique applications and risk factors in the medical field: a review. BioMed research international*, 2022. 2022(1): p. 5164970.
- Ogunpola, A., et al., *Machine learning-based predictive models for detection of cardiovascular diseases. Diagnostics*, 2024. 14(2): p. 144.
- Gilanie, G., et al., *A ROBUST ARTIFICIAL NEURAL NETWORK APPROACH FOR EARLY DETECTION OF CARDIAC DISEASES. Spectrum of Engineering Sciences*, 2025: p. 499-511.
- Latif, A., et al., *A VISION-FREE OBSTACLE DETECTION AND ALERT SYSTEM USING SMART KNEE GLOVES. Spectrum of Engineering Sciences*, 2025: p. 816-829.
- Gilanie, G., et al., *Digital Image Processing for Ultrasound Images: A Comprehensive. Digital Image Processing*, 2021. 15(3).
- Ramesh, T., et al., *Predictive analysis of heart diseases with machine learning approaches. Malaysian Journal of Computer Science*, 2022: p. 132-148.
- Kunwar, S., S. Ganesan, and S. Pokhrel, *Evaluating machine learning models for cardiovascular risk prediction: A Shapley Additive Explanations-based approach with statistical testing. Brain & Heart*, 2026: p. 025260032.
- Saher, A., et al., *A Deep Learning-Based Automated Approach of Schizophrenia Detection from Facial Micro-Expressions. Intelligent Automation & Soft Computing*, 2024. 39(6).
- Dinesh, K.G., et al. *Prediction of cardiovascular disease using machine learning algorithms. in 2018 international conference on current trends towards converging technologies (ICCTCT). 2018. IEEE.*
- Gilanie, G., et al., *STEGANOGRAPHIC SECRET COMMUNICATION USING RGB PIXEL ENCODING AND CRYPTOGRAPHIC SECURITY. Spectrum of Engineering Sciences*, 2025: p. 323-336.
- Gilanie, G., et al., *READABLE TEXT RETRIEVAL FROM NOISE-INFLUENCED DOCUMENTS USING IMAGE RESTORATION METHODS. Spectrum of Engineering Sciences*, 2025: p. 337-360.
- Krittana Wong, C., et al., *Machine learning prediction in cardiovascular diseases: a meta-analysis. Scientific reports*, 2020. 10(1): p. 16057.
- Batool, S.N., et al., *Forensic Radiology: A robust approach to biological profile estimation from bone image analysis using deep learning. Biomedical Signal Processing and Control*, 2025. 105: p. 107661.
- Javaid, Y., et al., *EXPLAINABLE EDGE AI FOR EARLY DETECTION OF PLANT DISEASES IN SMART AGRICULTURE USING LIGHTWEIGHT MOBILENETV2 WITH GRADIENT SALIENCY. Spectrum of Engineering Sciences*, 2026. 4(4): p. 167-192.
17. Gilanie, G., et al., *A robust method of bipolar mental illness detection from facial micro expressions using machine learning methods. Intelligent Automation & Soft Computing*, 2024. 39(1).
- Singh, M., et al., *Artificial intelligence for cardiovascular disease risk assessment in personalised framework: a scoping review. EClinicalMedicine*, 2024. 73.

Punzo, A., R.P. Browne, and P.D. McNicholas, *Hypothesis testing for mixture model selection*. Journal of Statistical Computation and Simulation, 2016. 86(14): p. 2797-2818.

Mittal, K., G. Aggarwal, and P. Mahajan, *Performance study of K-nearest neighbor classifier and K-means clustering for predicting the diagnostic accuracy*. International Journal of Information Technology, 2019. 11(3): p. 535-540.

Devi, R., H.K. Tyagi, and D. Kumar, *A novel multi-class approach for early-stage prediction of sudden cardiac death*. Biocybernetics and Biomedical Engineering, 2019. 39(3): p. 586-598.

Mohan, S., C. Thirumalai, and G. Srivastava, *Effective heart disease prediction using hybrid machine learning techniques*. IEEE access, 2019. 7: p. 81542-81554.

Al-Alshaikh, H.A., et al., *Comprehensive evaluation and performance analysis of machine learning in heart disease prediction*. Scientific Reports, 2024. 14(1): p. 7819.

Yan, P., et al., *A comprehensive survey of deep transfer learning for anomaly detection in industrial time series: Methods, applications, and directions*. IEEE Access, 2024. 12: p. 3768-3789.