AI Cardiologist: Advancements In Supervised Learning For Heart Disease Prediction

Vinoth M[1], Dr.K.Annalakshmi[2]
Dept of Computer Applications
Dr. M.G.R. Educational and Research Institute
Deemed to be University, Chennai – 95, Tamil Nadu, India

Abstract- *Cardiovascular disease (CVD) continues to pose a significant global health challenge, demanding innovative approaches for early detection and prevention. This paper presents an AI Cardiologist system that leverages supervised machine learning techniques to predict heart disease with high accuracy. The proposed system integrates a Bagging Classifier ensemble method alongside a LeNet Convolutional Neural Network architecture to analyse multi-dimensional patient data—including demographics, clinical history, laboratory results, and ECG readings. The system is deployed as a full-stack web application using the Django framework, enabling clinicians to receive real-time, personalised risk assessments. Experiments conducted on the UCI Heart Disease (CARDIO) dataset demonstrate competitive accuracy. Future directions include integration of Explainable AI (XAI) and federated learning to enhance transparency and privacy preservation.*

Keywords: Heart disease prediction; supervised learning; Bagging classifier; convolutional neural network; LeNet; explainable AI; federated learning; cardiovascular disease.

I. INTRODUCTION

Cardiovascular disease (CVD) remains the leading cause of mortality worldwide, accounting for an estimated 17.9 million deaths annually [1]. Early and accurate detection is the most effective strategy to reduce this burden. However, traditional diagnostic methods often depend on invasive procedures, delayed symptom recognition, and the cognitive capacity of clinicians to synthesise large volumes of heterogeneous patient data.

Machine learning (ML) has emerged as a transformative tool in cardiology. Supervised learning algorithms—trained on labelled patient datasets—can uncover subtle, non-linear relationships between risk factors and disease outcomes that conventional statistical methods frequently miss. These models can integrate demographics, medical history, laboratory results, and electrocardiogram (ECG) signals into unified predictive frameworks.

This paper introduces the AI Cardiologist system: a full-stack application that combines a Bagging Classifier and a LeNet Convolutional Neural Network (CNN) to deliver personalised cardiovascular risk assessments. Beyond academic benchmarks, the system is deployed as a web application, making its predictive power accessible to healthcare professionals without requiring specialised hardware.

The remainder of this paper is organised as follows. Section II reviews related work. Section III describes the dataset and preprocessing pipeline. Section IV details the proposed system architecture. Section V presents experimental results. Section VI discusses ethical considerations and limitations. Section VII concludes with future directions.

II. RELATED WORK

Hossen et al. [2] applied Random Forest, Decision Tree, and Logistic Regression on the UCI Cleveland database, achieving up to 86% accuracy for heart disease classification. While their study established strong baseline performance, the models were not deployed into a functional application, limiting clinical utility.

The BIOWISH system [3] explored wearable inertial sensing—specifically seismocardiography (SCG) and gyrocardiography (GCG)—for biometric heart activity recognition. Although BIOWISH demonstrated the feasibility of mechanical cardiac signals as identifiers, it focused exclusively on identity verification rather than disease prediction and lacked a deployable diagnostic interface.

IoT-based frameworks pairing deep learning with remote physiological monitoring have been explored extensively [4]. However, such systems introduce severe cybersecurity vulnerabilities, as low-performance processing units are difficult to update against threats such as the Heartbleed exploit, compromising sensitive patient data.

The present work addresses these limitations by: (1) building a deployable full-stack application; (2) comparing multiple ML algorithms to select the most accurate predictor;

and (3) explicitly addressing explainability and bias as ethical priorities.

III. DATASET AND PREPROCESSING

A. Dataset

The system uses the CARDIO dataset, derived from the UCI Heart Disease Repository [5]. It contains 303 patient records with 13 input features and one binary target variable indicating the presence (1) or absence (0) of heart disease. Features include age, sex, chest pain type (cp), resting blood pressure (trestbps), serum cholesterol (chol), fasting blood sugar (fbs), resting ECG results (restecg), maximum heart rate achieved (thalach), exercise-induced angina (exang), ST depression (oldpeak), slope of the peak exercise ST segment (slope), number of major vessels (ca), and thalassemia type (thal).

B. Preprocessing Pipeline

Raw medical data frequently contains missing values, duplicate records, and inconsistent data types. The preprocessing pipeline, implemented with Python's Pandas and NumPy libraries, performs the following steps: (1) missing value removal via dropna(); (2) duplicate elimination; (3) data type verification (float vs. integer variables); and (4) class imbalance correction via Random Over-Sampling using imblearn.

After oversampling, the dataset is split 80/20 into training and testing sets using stratified sampling (random_state=42) to preserve class proportions. Features are verified to be free of null values prior to model training.

IV. PROPOSED SYSTEM ARCHITECTURE

A. Overall Architecture

The AI Cardiologist follows a three-tier architecture: (1) Presentation Layer—HTML/CSS/JavaScript frontend collecting patient parameters and displaying predictions; (2) Application Layer—Django backend hosting the trained models, routing API requests, and executing inference; and (3) Data Layer—SQLite database storing user credentials and historical prediction records.

The data pipeline flows from raw input through a mandatory preprocessing gate, into either the Bagging Classifier (for structured tabular inputs) or the LeNet CNN (for ECG image inputs), and finally to the output page presenting the risk assessment.

B. Machine Learning Module: Bagging Classifier

Bootstrap Aggregating (Bagging) is an ensemble meta-algorithm that reduces model variance by training multiple base classifiers on different bootstrapped subsets of the training data and aggregating their predictions through majority voting [6]. In this work, scikit-learn'sBaggingClassifier is instantiated with Decision Trees as base estimators.

For a dataset D with N samples, Bagging creates T bootstrap samples D_1, ..., D_T each of size N (with replacement). A classifier h_t is trained on each D_t. Final prediction: H(x) = argmax_c sum_{t=1}^{T} I(h_t(x) = c). This mechanism reduces variance without increasing bias, addressing the high-variance characteristic of individual decision trees and improving generalisation to unseen patient profiles.

C. Deep Learning Module: LeNet Architecture

For ECG image classification, the system employs a LeNet-5 inspired CNN [7]. The architecture comprises: (1) Convolutional layer (32 filters, 3x3 kernel, ReLU activation); (2) Max Pooling (2x2); (3) Second Convolutional layer (128 filters, 3x3, ReLU); (4) Max Pooling (2x2); (5) Flatten; (6) Dense layer (256 units, ReLU); (7) Output Dense (4 units, Softmax).

The model is trained with the RMSprop optimiser and categorical cross-entropy loss for 100 epochs. A ModelCheckpoint callback saves the best model weights (monitor='accuracy') to LeNet1.h5. Input images are rescaled to 224x224 pixels with data augmentation (shear, zoom, horizontal flip) to improve robustness.

Batch normalisation and Dropout (rate = 0.4 in the manual architecture variant) are employed to stabilise training and prevent overfitting. The ReLU activation f(x) = max(0, x) introduces the non-linearity required to learn complex physiological marker interactions.

D. Deployment

The trained Bagging model is serialised with joblib (CARDIO.pkl) and loaded into the Django application at startup. The LeNet model is loaded via TensorFlow/Keras (LeNet1.h5). Django routes patient form submissions to the appropriate model, executes inference, and returns risk classifications. SQLite records all predictions for audit and historical review.

V. EXPERIMENTAL RESULTS

A. Evaluation Metrics

Model performance is evaluated using Accuracy, Precision, Recall, F1-Score, and Confusion Matrix analysis. For a medical diagnostic tool, minimising False Negatives (FN)—patients with disease predicted as healthy—is paramount, as these represent life-threatening missed diagnoses.

TABLE I

Performance Comparison of Classification Algorithms

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Logistic Regression	83.6	84.2	83.1	83.6
Decision Tree	81.0	80.8	81.3	81.0
Random Forest	86.9	87.1	86.7	86.9
Bagging Classifier	89.3	89.5	89.1	89.3
LeNet CNN (DL)	91.4	91.8	91.0	91.4

B. Cross-Validation

K-Fold cross-validation (k=5) is applied to the Bagging Classifier to ensure generalisation. The mean cross-validated accuracy across folds is reported to prevent overfitting to a single train-test split. Stratified splitting ensures each fold preserves the original class distribution.

C. Test Data

Table II presents a representative sample from the test dataset used for model evaluation.

TABLE II

Sample Test Data

Patient	Age	Gender	Chol	HR	ECG	Risk
P-001	52	M	212	168	1	Medium
P-003	45	F	200	155	0	Low
P-004	65	M	240	175	1	High
P-006	38	F	185	140	0	Low

VI. ETHICAL CONSIDERATIONS AND LIMITATIONS

The deployment of AI in clinical settings raises critical ethical considerations. First, algorithmic bias: if training data under-represents specific demographics (e.g., ethnicity, age group), the model may produce inaccurate risk scores for those populations, worsening health disparities. Mitigation requires diverse, representative training datasets.

Second, the 'black box' problem: deep learning models such as the LeNet CNN do not inherently explain their predictions. Clinicians are understandably reluctant to alter treatment plans based on opaque model outputs. Explainable AI frameworks—SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations)—can deconstruct complex model outputs, assigning explicit feature weights so a physician can see, for example, that a high-risk classification was driven by a specific combination of HbA1c level, resting heart rate, and family history.

Third, patient privacy: patient data must be protected under frameworks such as HIPAA and GDPR. Federated learning—training models across multiple hospitals without centralising raw data—offers a promising privacy-preserving solution.

Current limitations include: (1) reliance on a single dataset (UCI CARDIO), limiting demographic diversity; (2) absence of real-time wearable data integration; and (3) the current system does not yet incorporate XAI outputs in the clinical interface.

VII. CONCLUSION AND FUTURE WORK

This paper presented the AI Cardiologist, a full-stack system combining a Bagging Classifier ensemble method with a LeNet Convolutional Neural Network for cardiovascular disease prediction. The system achieved up to 91.4% accuracy on the UCI CARDIO dataset and is deployed as a Django web application, making it platform-independent and accessible to clinical users without specialised hardware.

The Bagging Classifier demonstrated superior performance over individual classifiers (Logistic Regression, Decision Tree, Random Forest) by reducing model variance through bootstrap aggregation. The LeNet CNN module further improved accuracy on ECG image classification tasks through automated feature extraction and spatial pattern recognition.

Future work will focus on three key directions: (1) Integration of Explainable AI (XAI) using SHAP values to provide clinicians with feature-level reasoning for each prediction; (2) IoT and wearable synergy to enable continuous 24/7 cardiovascular monitoring; and (3) Federated learning to train models across multiple hospitals without compromising patient data privacy. These enhancements will evolve the AI

Cardiologist from a research prototype into a comprehensive, clinically validated cardiovascular health management platform.

REFERENCES

[1] World Health Organization, Cardiovascular diseases (CVDs) Fact Sheet, WHO, Geneva, 2023.
[2] M. D. A. Hossen, T. Tazin, and S. Khan, "Supervised Machine Learning-Based Cardiovascular Disease Analysis and Prediction," Mathematical Problems in Engineering, vol. 2021, pp. 1–14, 2021.
[3] D. Jarchi et al., "BIOWISH: Biometric Recognition Using Wearable Inertial Sensors Detecting Heart Activity," IEEE Sensors Journal, vol. 22, no. 8, pp. 8207–8218, 2022.
[4] A. Amodei, "A Measurement Approach for Inline Intrusion Detection of Heartbleed-Like Attacks in IoT Frameworks," Journal of Cybersecurity Research, vol. 9, pp. 112–128, 2023.
[5] D. Dua and C. Graff, UCI Machine Learning Repository – Heart Disease Dataset, University of California, Irvine, 2026. [Online]. Available: http://archive.ics.uci.edu/ml
[6] L. Breiman, "Bagging Predictors," Machine Learning, vol. 24, no. 2, pp. 123–140, 1996.
[7] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," Proceedings of the IEEE, vol. 86, no. 11, pp. 2278–2324, 1998.
[8] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.