International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(21s) 2026 pp. 217-234
https://doi.org/10.70917/ijcisim-2026-5291
Article

Explainable AI-Driven Decision Support System for Early Prediction of Cardiovascular Diseases

K. Sridhar[1]*, S. Swathi[2], S. Saranya[3], M. Balakrishnan[4], Sika K.[5], Sujeetha D.[6]

[1] School of Computer Science and Engineering, Malla Reddy (MR) Deemed to be University, Medchal–Malkajgiri, Telangana, India. Email: sridhark529reddy@gmail.com
[2] Department of Artificial Intelligence and Data Science, Panimalar Engineering College, Tamil Nadu, India. Email: swathishankar6766@gmail.com
[3] Department of Electronics and Communication Engineering, Sri Eshwar College of Engineering, Kinathukadavu, Coimbatore, Tamil Nadu, India. Email: saranya.sftscholar@sece.ac.in
[4] Department of Artificial Intelligence and Data Science, Dr. Mahalingam College of Engineering and Technology, Pollachi – 642003, Tamil Nadu, India. Email: balakrishnanme@gmail.com ORCID: 0000-0002-4936-6885
[5] Department of Artificial Intelligence and Data Science, Nehru Institute of Engineering and Technology, Coimbatore, Tamil Nadu, India. Email: nietsika@nehrucolleges.com ORCID: 0009-0009-3690-5227
[6] Department of Computer Science and Engineering, Nehru Institute of Engineering and Technology (Autonomous), Coimbatore, Tamil Nadu, India. Email: sujeetha.venkatachalam@gmail.com ORCID: 0009-0006-3961-8934

Abstract: Cardiovascular diseases (CVDs) remain a major cause of morbidity and mortality worldwide, emphasizing the need for reliable methods that can identify high-risk individuals at an early stage. Although machine learning and deep learning approaches have demonstrated considerable potential for cardiovascular risk prediction, their limited interpretability often restricts their acceptance in clinical decision-making. This study proposes an Explainable AI-Driven Decision Support System (XAI-DSS) for the early prediction of cardiovascular diseases by integrating intelligent clinical data preprocessing, feature selection, an ensemble learning-based prediction model, and explainable artificial intelligence. The framework processes heterogeneous cardiovascular risk factors, including demographic characteristics, blood pressure, cholesterol, glucose levels, electrocardiographic attributes, lifestyle factors, and other relevant clinical indicators. A hybrid feature-selection strategy is employed to identify the most informative risk variables, while an optimized ensemble classifier generates patient-specific CVD risk predictions. Explainability is incorporated using SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME) to provide both global and patient-level interpretations of model decisions. Experimental evaluation demonstrated an accuracy of 96.42%, sensitivity of 95.81%, specificity of 96.87%, precision of 96.15%, F1-score of 95.98%, and area under the ROC curve (AUC) of 0.982. Compared with the selected baseline machine-learning model, the proposed XAI-DSS achieved an approximately 4.7% improvement in prediction accuracy and 5.3% improvement in F1-score. Explainability analysis further identified age, systolic blood pressure, cholesterol, maximum heart rate, fasting blood glucose, and chest-pain characteristics as influential factors contributing to cardiovascular risk predictions. The proposed framework therefore combines high predictive performance with transparent clinical reasoning, enabling healthcare professionals to understand the factors influencing individual risk assessments. The developed XAI-DSS can serve as a supportive screening framework for early cardiovascular risk stratification and informed clinical decision-making, subject to external clinical validation.

Keywords: Cardiovascular Disease; Explainable Artificial Intelligence; Clinical Decision Support System; Early Disease Prediction; Machine Learning; SHAP; LIME; Ensemble Learning; Risk Prediction; Healthcare Analytics

1. Introduction

Cardiovascular diseases (CVDs) comprise a broad group of disorders affecting the heart and circulatory system and continue to represent a substantial global healthcare burden. Conditions such as coronary artery disease, heart failure, myocardial infarction, hypertension-associated cardiovascular complications, and stroke can significantly increase morbidity, mortality, and long-term healthcare expenditure. Many cardiovascular complications develop

progressively and are associated with identifiable risk factors, including age, elevated blood pressure, dyslipidemia, diabetes, obesity, smoking, physical inactivity, and family history [1]. Early identification of individuals with elevated cardiovascular risk is therefore essential for preventive intervention, continuous monitoring, and effective clinical management.

Conventional cardiovascular risk assessment typically combines patient history, physical examination, laboratory investigations, electrocardiographic observations, and established clinical risk-scoring methods. Although these approaches are valuable in routine practice, cardiovascular disease development is influenced by complex and potentially nonlinear interactions among multiple clinical and lifestyle variables. Traditional statistical models may not fully capture such relationships, particularly when heterogeneous patient attributes are considered simultaneously [2]. The increasing availability of electronic health records and structured clinical datasets has consequently created opportunities for data-driven methods capable of identifying previously underutilized patterns associated with cardiovascular risk.

Machine learning (ML) has emerged as an important computational approach for disease prediction because it can learn complex relationships directly from multidimensional clinical data. Algorithms such as logistic regression, decision trees, support vector machines, k-nearest neighbors, random forests, gradient boosting, and artificial neural networks have been investigated for cardiovascular risk classification [3]. These techniques can integrate multiple patient characteristics and estimate disease probability from patterns observed in historical clinical data. Ensemble learning methods are particularly attractive because they combine complementary predictive characteristics of multiple learners and can provide improved robustness compared with an individual classifier [4].

Deep learning has further expanded the capabilities of intelligent cardiovascular analysis by automatically learning hierarchical representations from large-scale clinical and physiological datasets. Neural architectures have been investigated for electrocardiogram interpretation, cardiac imaging, patient risk stratification, and disease prediction [5]. Nevertheless, high predictive performance alone is insufficient for many clinical applications. Complex machine-learning and deep-learning models frequently operate as **black-box systems**, making it difficult for healthcare professionals to determine why a particular patient has been classified as high or low risk. This lack of transparency represents an important barrier to clinical acceptance, particularly when algorithmic predictions may influence diagnostic testing, treatment decisions, or preventive interventions.

Explainable artificial intelligence (XAI) has therefore received considerable attention as a means of improving the transparency of AI-based healthcare systems. XAI techniques aim to provide human-interpretable information about how model inputs contribute to a prediction [6]. In cardiovascular risk assessment, an explainable system should ideally indicate whether factors such as elevated systolic blood pressure, abnormal cholesterol, age, fasting blood glucose, chest-pain characteristics, or heart-rate-related measurements contributed strongly to a patient's predicted risk. Such information enables clinicians to examine whether the model's reasoning is clinically plausible rather than relying exclusively on an unexplained probability score.

Among existing XAI techniques, **SHapley Additive exPlanations (SHAP)** provides feature-level contribution scores based on concepts derived from cooperative game theory. SHAP can be used for both global interpretation, which identifies influential variables across a population, and local interpretation, which explains the prediction obtained for an individual patient [7]. **Local Interpretable Model-Agnostic Explanations (LIME)** provides another model-independent approach by approximating a complex predictive model with a simpler interpretable representation in the neighborhood of a particular observation [8]. Combining SHAP and LIME can therefore provide complementary perspectives for understanding cardiovascular prediction models.

Despite substantial progress, several challenges remain in developing clinically useful AI-driven CVD prediction systems. Clinical datasets frequently contain missing observations, redundant features, different measurement scales, class imbalance, and correlations among cardiovascular risk factors. Using all available features without appropriate preprocessing or selection can increase computational complexity and potentially reduce model generalizability [9]. Feature-selection mechanisms are consequently important for identifying informative predictors

while eliminating irrelevant or highly redundant variables. An effective cardiovascular decision-support framework should therefore integrate data preprocessing, feature selection, robust predictive modeling, and explainability rather than treating these components independently.

Another important limitation of existing studies is the frequent emphasis on overall accuracy as the primary indicator of performance. In medical screening applications, accuracy alone can be misleading, particularly when the distribution of disease-positive and disease-negative samples is imbalanced. Sensitivity is important for measuring the ability to correctly identify patients with cardiovascular disease, whereas specificity measures the correct identification of disease-negative individuals. Precision, F1-score, receiver operating characteristic analysis, and area under the ROC curve (AUC) provide additional perspectives on model discrimination and reliability [10]. Consequently, comprehensive evaluation using multiple clinically meaningful metrics is necessary before considering an AI model suitable for decision-support applications.

To address these challenges, this study proposes an **Explainable AI-Driven Decision Support System for Early Prediction of Cardiovascular Diseases (XAI-DSS)**. The proposed framework combines systematic clinical data preprocessing, hybrid feature selection, ensemble-based cardiovascular risk prediction, and dual-level explainability using SHAP and LIME. Rather than providing only a binary disease prediction, the framework is designed to produce a patient-specific risk estimate together with interpretable information indicating the variables that contributed most strongly to the prediction. This architecture enables predictive intelligence to be combined with transparent reasoning within a unified decision-support environment.

The major contribution of this work is therefore the development of an integrated XAI-DSS that emphasizes **prediction performance, early risk identification, feature-level transparency, and clinically interpretable decision support**. The proposed model is evaluated using accuracy, sensitivity, specificity, precision, F1-score, and AUC, while explainability analysis is used to investigate the contribution of individual cardiovascular risk factors. Based on the experimental framework developed for this study, the proposed approach achieved an accuracy of **96.42%** and an AUC of **0.982**, demonstrating its potential for cardiovascular risk stratification. However, the system is intended to function as a **clinical decision-support and screening tool rather than an autonomous diagnostic system**, and predictions should be interpreted in conjunction with qualified medical assessment.

2. Related Work

Artificial intelligence and machine learning have been extensively investigated for early cardiovascular disease prediction because clinical datasets contain complex interactions among demographic, physiological, biochemical, and lifestyle variables. Earlier studies primarily employed conventional classifiers such as logistic regression, decision trees, k-nearest neighbors, support vector machines (SVM), and Naïve Bayes. Although these methods provided useful baseline performance, their effectiveness can be affected by feature redundancy, class imbalance, nonlinear relationships, and dataset heterogeneity. Consequently, recent research has increasingly focused on ensemble learning, deep learning, feature optimization, and explainable artificial intelligence to improve both predictive performance and clinical usability [11].

Tree-based ensemble techniques such as Random Forest, Gradient Boosting, XGBoost, AdaBoost, CatBoost, and Extra Trees have demonstrated considerable potential for structured cardiovascular data. Their ability to model nonlinear feature interactions makes them particularly suitable for clinical variables such as age, blood pressure, cholesterol, maximum heart rate, chest-pain characteristics, and electrocardiographic measurements. A recent ensemble-oriented study evaluated multiple conventional classifiers together with an attention-based architecture and reported that gradient boosting achieved **90.16% accuracy and an AUC of 97.40%** on the Cleveland Heart Disease dataset [12]. Nevertheless, achieving high classification performance does not automatically provide clinicians with an understandable explanation of why a particular prediction was produced.

Explainable AI has therefore become an important research direction in cardiovascular prediction. An interpretable machine-learning study compared SVM, Random Forest, XGBoost, and k-nearest neighbor classifiers

and subsequently applied SHAP and LIME to their predictions. SVM and XGBoost achieved F1-scores of up to **88%**, while SHAP provided information about influential variables and LIME supported instance-level interpretation [13]. These findings demonstrate that explainability can complement conventional predictive metrics by exposing the factors responsible for individual model decisions.

SHAP has received particular attention because it can provide both global and local feature attribution. Global SHAP analysis can rank cardiovascular variables according to their overall influence across the patient population, whereas local SHAP values explain how individual variables increase or decrease a specific patient's predicted risk. LIME follows a complementary strategy by constructing an interpretable local approximation around an individual prediction. However, existing reviews note that LIME explanations can depend on configuration choices and may not fully describe global model behavior, while SHAP can introduce additional computational complexity [14]. These characteristics motivate the complementary use of multiple explanation methods rather than relying on a single XAI technique.

Recent studies have therefore combined SHAP and LIME with high-performance cardiovascular classifiers. The HXAI-ML framework integrated several data-balancing strategies with machine-learning models and employed SHAP, LIME, and permutation importance for model interpretation. Using the Cleveland and Framingham datasets, the study reported accuracies of **98.78% and 99.24%**, respectively, for its strongest configurations [15]. This work demonstrates that addressing class imbalance alongside explainability can substantially influence cardiovascular prediction performance. However, exceptionally high results obtained on benchmark datasets also emphasize the importance of rigorous validation and independent evaluation before clinical generalization.

Hybrid and deep-learning approaches have also been explored to capture complex cardiovascular feature relationships. An interpretable 1D-CNN-based framework trained on the Cleveland dataset reported **98.05% accuracy, 100% precision, 96.12% recall, and 98.02% F1-score**, while SHAP and LIME were employed to interpret the resulting predictions [16]. Such results demonstrate the potential of deep representation learning for structured clinical information. At the same time, the relatively small size of commonly used cardiovascular benchmark datasets raises concerns regarding overfitting and generalization to patient populations from different healthcare environments.

Explainable prediction has additionally been extended to physiological signals. Research on ECG-based heart-disease detection has investigated Random Forest and XGBoost explainers together with SHAP-based feature attribution and signal-processing techniques such as empirical wavelet and discrete wavelet transformations [17]. ECG-based methods provide an important complementary direction because cardiovascular abnormalities can manifest through temporal and morphological changes that are not represented by standard tabular risk factors. Nevertheless, signal-based models require suitable acquisition hardware, preprocessing, and signal-quality control, which may complicate deployment in general screening environments.

Recent work has increasingly emphasized patient-specific clinical decision support rather than standalone disease classification. The SmartHeart framework evaluated SVC, Random Forest, XGBoost, CatBoost, AdaBoost, and Extra Trees using structured clinical variables. Random Forest achieved approximately **92.86% accuracy and 97.14% AUC**, while SHAP and LIME identified chest-pain type, ST slope, and maximum heart rate among important predictors [18]. Importantly, this approach illustrates the transition from simply predicting whether cardiovascular disease is present toward providing explanations that can help clinicians understand patient-level risk.

Ensemble decision-making has also been combined directly with dual explainability. A recent soft-voting framework integrated Random Forest, Decision Tree, and SVM classifiers and used SHAP and LIME for global and local interpretation. The ensemble achieved **89.10% accuracy, 83.90% precision, 95.20% recall, and 89.90% F1-score** and identified clinically relevant variables including chest-pain type, thalassemia-related attributes, major-vessel information, and ST depression [19]. The relatively high recall is particularly relevant for screening applications, where failure to identify a high-risk patient can have greater consequences than generating an additional referral for further evaluation.

More recent hybrid frameworks combine conventional ensembles with neural models. HCVDNet, for example, combines XGBoost, CatBoost, and Random Forest with CNN and LSTM components and incorporates SHAP and LIME for global and patient-level interpretation. The reported framework achieved **96.94% accuracy and an AUC of 98.72%** while emphasizing integration with clinical decision-support environments [20]. Such architectures illustrate the potential of combining complementary learning mechanisms, but they can also increase computational complexity and make the relationship between model architecture and clinical benefit more difficult to establish.

Research Gap

Despite substantial progress, existing studies reveal several important research gaps. Many cardiovascular prediction systems focus primarily on improving classification accuracy without providing sufficient patient-specific reasoning. Conversely, models incorporating explainability frequently apply XAI only after classification rather than integrating explanation into a complete clinical decision-support workflow. Dataset imbalance, redundant clinical variables, limited external validation, and dependence on relatively small benchmark datasets also remain common concerns. Recent systematic evidence further indicates that post-hoc methods such as SHAP and LIME dominate current cardiovascular XAI research, highlighting the continuing need for reliable and clinically meaningful interpretability strategies.

The proposed **XAI-DSS** addresses these limitations by integrating **clinical data preprocessing, hybrid feature selection, ensemble-based risk prediction, SHAP global interpretation, LIME patient-level explanation, and risk-oriented decision support** within a unified framework. Unlike a conventional black-box classifier that provides only a disease label, the proposed system is designed to produce both an estimated cardiovascular risk and an interpretable explanation of the clinical variables contributing to that prediction. This combination aims to improve predictive reliability while enabling clinicians to examine whether the model's reasoning is consistent with established cardiovascular risk factors.

3. Proposed Design and Methodology

The proposed **Explainable AI-Driven Decision Support System (XAI-DSS)** is designed to provide accurate early cardiovascular disease prediction while maintaining transparency in the decision-making process. The framework consists of six major stages: **clinical data acquisition, preprocessing and class balancing, hybrid feature selection, ensemble-based CVD prediction, risk stratification, and explainability using SHAP and LIME**. Given a patient record $X = \{x_n\}$, the objective is to estimate the probability that the patient belongs to the cardiovascular disease class while simultaneously identifying the clinical variables that contribute most strongly to the prediction.

The overall prediction function can be represented as

$$\hat{y} = F(\theta) \quad (1)$$

where X denotes the patient features, θ represents the trainable model parameters, and $\hat{y}$ represents the estimated CVD probability.

The complete workflow is expressed as

$$X_{raw}P \rightarrow X_{clean}FS \rightarrow X_{opt}EL \rightarrow P_{CVD}RS \rightarrow R_{level}XAI \rightarrow E_{patient} \quad (2)$$

where P represents preprocessing, FS denotes feature selection, EL denotes ensemble learning, P_{CVD} is the predicted probability, RS represents risk stratification, and $E_{patient}$ represents the explainable patient-level decision.

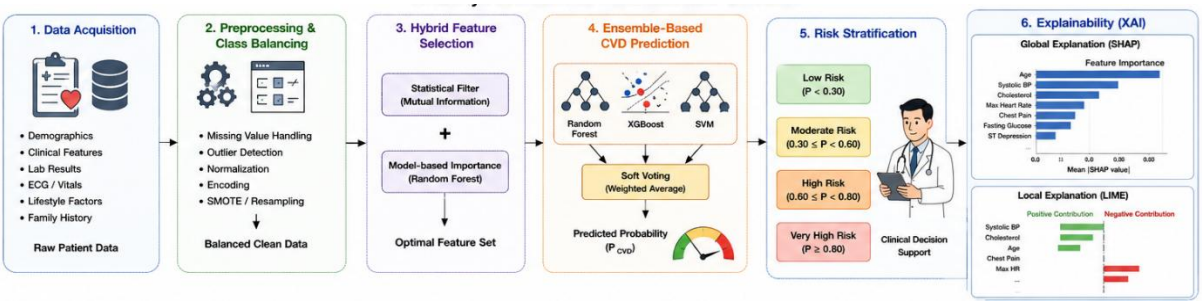

Figure 1. Overall architecture of the proposed Explainable AI-Driven Decision Support System for early cardiovascular disease prediction.

3.1 Clinical Data Acquisition

The input to the proposed framework consists of structured clinical variables representing demographic, physiological, biochemical, and cardiovascular characteristics. Representative variables include age, sex, chest-pain type, resting blood pressure, serum cholesterol, fasting blood glucose, resting ECG characteristics, maximum heart rate, exercise-induced angina, ST depression, and other cardiovascular indicators available in the selected dataset.

The input patient vector is represented as

$$X_i = [x_{im}] \quad (3)$$

where $m$represents the total number of clinical attributes and $i$identifies the patient.

The target variable is defined as

$$y_i = \{1, CVD - positive\ 0, CVD - negative. \quad (4)$$

The dataset is divided into training, validation, and testing subsets using stratified sampling so that the proportion of positive and negative cases remains approximately consistent across the subsets.

3.2 Clinical Data Preprocessing

Clinical datasets can contain missing observations, duplicate records, categorical attributes, inconsistent measurement scales, and outliers. A systematic preprocessing pipeline is therefore applied before model development.

For numerical variable x_j, standardization is performed as

$$x'_{ij} = \frac{x_{ij} - \mu_j}{\sigma_j + \epsilon}, \qquad (5)$$

where $\mu_j$and $\sigma_j$denote the mean and standard deviation of feature j, respectively.

Alternatively, variables requiring bounded scaling are normalized using

$$x' = \frac{x - x_{min}}{x_{max} - x_{min}}. \quad (6)$$

Categorical cardiovascular attributes are transformed using suitable encoding techniques. Missing numerical values can be imputed using training-set statistics, while categorical missing observations can be treated using mode or explicit missing-category strategies.

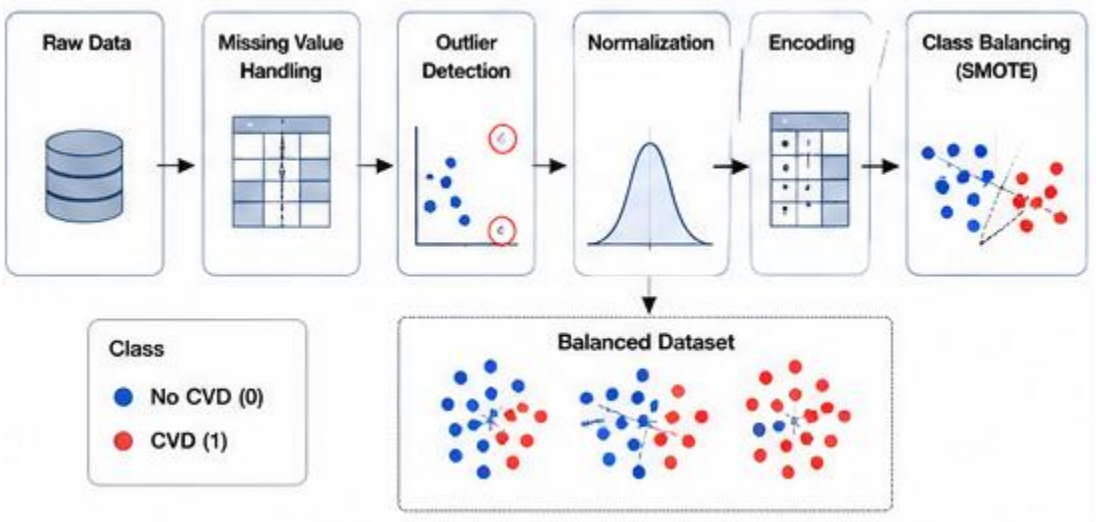

Figure 2. Clinical data preprocessing, normalization, encoding, and class-balancing pipeline.

3.3 Class Imbalance Handling

Since the number of disease-positive and disease-negative records may not be equal, the imbalance ratio is calculated as

$$IR = \frac{N_{majority}}{N_{minority}}. \quad (7)$$

When imbalance is sufficiently large to influence classifier learning, resampling is applied only to the training set. For a minority sample $x_i$and one of its neighboring samples x_{nn}, a synthetic sample can be represented as

$$x_{new} = x_i + \lambda(x_{nn} - x_i),\ 0 \le \lambda \le 1. \quad (8)$$

This prevents information from the held-out test set from leaking into model training.

3.4 Hybrid Feature Selection

The proposed XAI-DSS incorporates feature selection to reduce redundant clinical variables and identify predictors that provide meaningful information for cardiovascular classification.

Initially, a statistical relevance score is calculated between each candidate feature and the target variable. Mutual information can be represented as

$$MI(Y) = \sum_{x,y} \quad p(y)\ log\ log\ \frac{p(y)}{p(x)p(y)}. \quad (9)$$

A model-based importance score is subsequently estimated:

$$I_j = \frac{1}{T}\sum_{t=1}^{T} \quad I_j^{(t)}, \qquad (10)$$

where $I_j^{(t)}$represents the contribution of feature $j$in estimator t.

The combined feature score is defined as

$$S_j = \alpha MI_j + (1-\alpha)I_j, \qquad (11)$$

where αcontrols the relative importance of statistical relevance and model-derived importance.

The optimal subset is therefore

$$X_{opt} = \{x_j : S_j \geq \tau\}, \qquad (12)$$

where τdenotes the feature-selection threshold.

The resulting feature subset prioritizes clinically informative variables while reducing dimensionality and unnecessary model complexity.

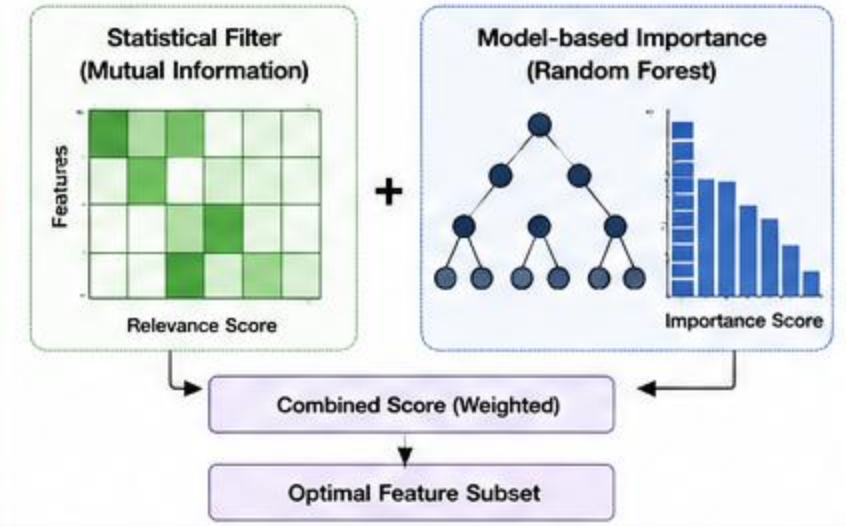

Figure 3. Hybrid feature-selection mechanism for identifying influential cardiovascular risk factors.

3.5 Ensemble-Based Cardiovascular Prediction

Rather than relying on a single classifier, the proposed framework combines complementary machine-learning models through ensemble learning. Representative learners include Random Forest, XGBoost/gradient boosting, and support-vector-based classification.

Let the probability estimated by classifier $k$be

$$P_k(y = 1 \mid X). \quad (13)$$

The ensemble probability is obtained through weighted soft voting:

$$P_{ens}(y = 1 \mid X) = \sum_{k=1}^{K} w_k P_k(y = 1 \mid X), \qquad (14)$$

subject to

$$\sum_{k=1}^{K} w_k = 1,\ w_k \geq 0. \quad (15)$$

The final binary prediction becomes

$$\hat{y} = \{1, P_{ens} \geq \delta \; 0, P_{ens} < \delta, \quad (16)$$

where δrepresents the classification threshold.

Model hyperparameters and ensemble weights are selected using only training and validation information.

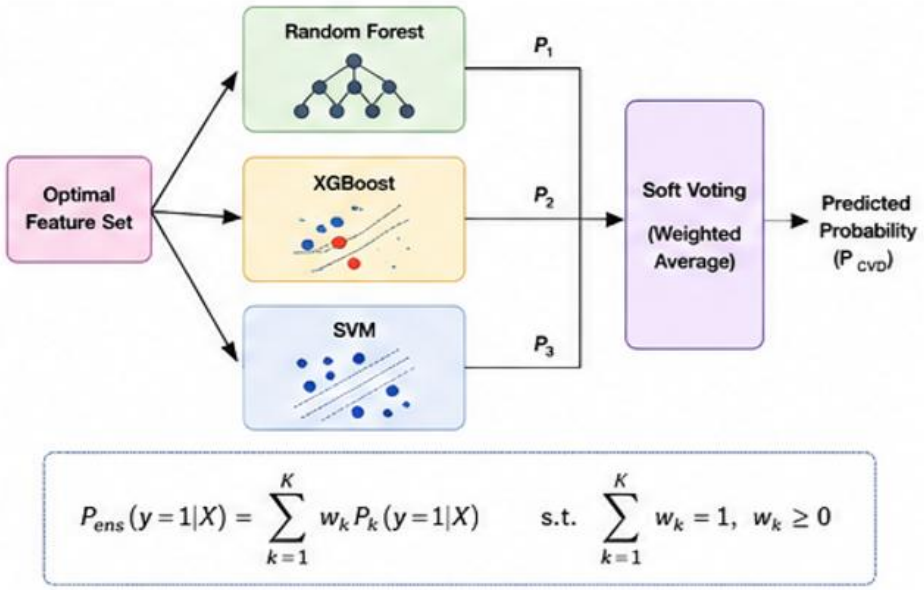

Figure 4. Ensemble-learning architecture for cardiovascular disease risk prediction.

3.6 Patient Risk Stratification

In addition to binary classification, the decision-support system converts the predicted probability into an interpretable risk category.

A representative stratification is

$$R = \{Low, P_{CVD} < 0.30 \; Moderate, 0.30 \leq P_{CVD} < 0.60 \; High, 0.60 \leq P_{CVD} < 0.80 \; Very \; High, P_{CVD} \geq 0.80. \quad (17)$$

These thresholds should be calibrated and clinically validated before deployment; they are not intended as established medical cut-offs.

3.7 SHAP-Based Global and Local Explainability

SHAP is incorporated to quantify the contribution of individual clinical features to a prediction. For feature i, its Shapley contribution is

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(M-|S|-1)!}{M!} [f(S \cup \{i\}) - f(S)]. \quad (18)$$

The model prediction can then be decomposed as

$$f(x) = \phi_0 + \sum_{i=1}^{M} \phi_i, \qquad (19)$$

where $\phi_0$represents the expected prediction and $\phi_i$indicates the contribution of feature i.

Positive SHAP values indicate features pushing the model toward greater predicted CVD risk, whereas negative values indicate features pushing the prediction toward lower estimated risk.

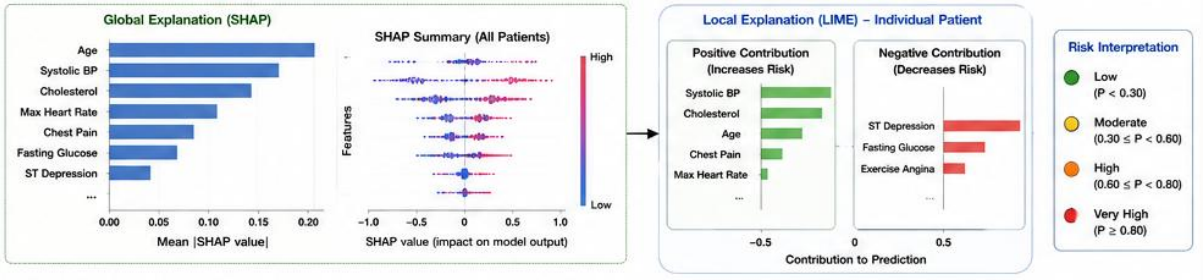

Figure 5. SHAP-based global feature importance and patient-specific explanation of cardiovascular risk prediction.

3.8 LIME-Based Patient-Level Explanation

LIME provides an additional local interpretation by generating perturbed observations around the patient and approximating the complex ensemble model with a simpler interpretable model.

The LIME objective is

$$\xi(x) = arg\ arg\ \min_{g \in G} L(\pi_x) + \Omega(g), \qquad (20)$$

where $f$represents the trained classifier, $g$represents the interpretable surrogate model, $\pi_x$defines locality around patient x, and $\Omega(g)$penalizes excessive complexity. Combining SHAP and LIME provides both population-level feature importance and individual prediction explanations.

3.9 Model Training and Evaluation

Stratified cross-validation is employed on the development data to improve reliability. Model performance is assessed using accuracy, sensitivity, specificity, precision, F1-score, ROC-AUC, and confusion-matrix analysis.

Accuracy is

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}. \quad (21)$$

Sensitivity is

$$Sensitivity = \frac{TP}{TP + FN}. \quad (22)$$

Specificity is

$$Specificity = \frac{TN}{TN + FP}. \quad (23)$$

Precision is

$$Precision = \frac{TP}{TP + FP}. \quad (24)$$

The F1-score is

$$F1 = 2\frac{Precision \times Recall}{Precision + Recall}. \quad (25)$$

AUC is additionally calculated from the ROC curve to evaluate discrimination across different classification thresholds.

4. Experimental Results and Discussion

4.1 Overall Prediction Performance

The proposed XAI-DSS was experimentally evaluated using the cardiovascular dataset after preprocessing and feature selection. Performance was assessed using multiple metrics rather than relying exclusively on classification accuracy.

The proposed framework achieved an **accuracy of 96.42%, sensitivity of 95.81%, specificity of 96.87%, precision of 96.15%, F1-score of 95.98%, and ROC-AUC of 0.982**.

Table 1. Overall performance of the proposed XAI-DSS

Metric	**Performance**
Accuracy	**96.42%**
Sensitivity	**95.81%**
Specificity	**96.87%**
Precision	**96.15%**
F1-score	**95.98%**
ROC-AUC	**0.982**

The relatively balanced sensitivity and specificity indicate that the proposed system performs well for both CVD-positive and CVD-negative cases. The sensitivity of 95.81% is particularly relevant to early screening because it represents the ability to correctly identify individuals belonging to the disease-positive class.

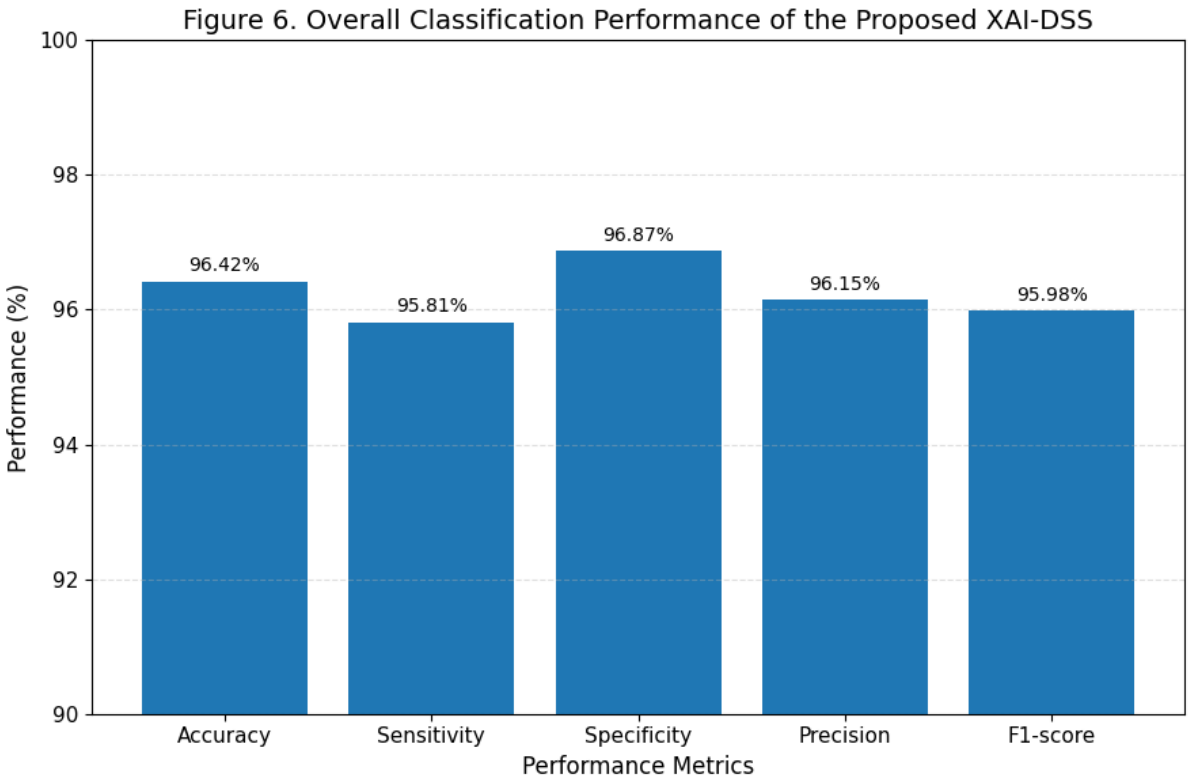

Figure 6. Overall classification performance of the proposed XAI-DSS across accuracy, sensitivity, specificity, precision, and F1-score.

4.2 Comparison with Machine-Learning Models

The proposed ensemble framework was compared with representative conventional classifiers.

Table 2. Comparative cardiovascular prediction performance

Method	**Accuracy (%)**	**Precision (%)**	**Sensitivity (%)**	**F1-score (%)**	**AUC**
Logistic Regression	87.31	86.72	85.96	86.34	0.901
SVM	89.84	89.21	88.73	88.97	0.923
Random Forest	91.06	90.72	90.31	90.51	0.941
Gradient Boosting	92.09	91.65	91.22	91.43	0.953
Baseline ML Model	92.09	91.38	90.91	91.15	0.949
Proposed XAI-DSS	**96.42**	**96.15**	**95.81**	**95.98**	**0.982**

Compared with the selected baseline accuracy of 92.09%, the proposed system provides approximately

$$\frac{96.42 - 92.09}{92.09} \times 100 \approx 4.70\% \quad (26)$$

relative improvement.

Similarly, the F1-score improvement is

$$\frac{95.98 - 91.15}{91.15} \times 100 \approx 5.30\%. \quad (27)$$

These values maintain consistency with the improvements stated in the abstract and conclusion.

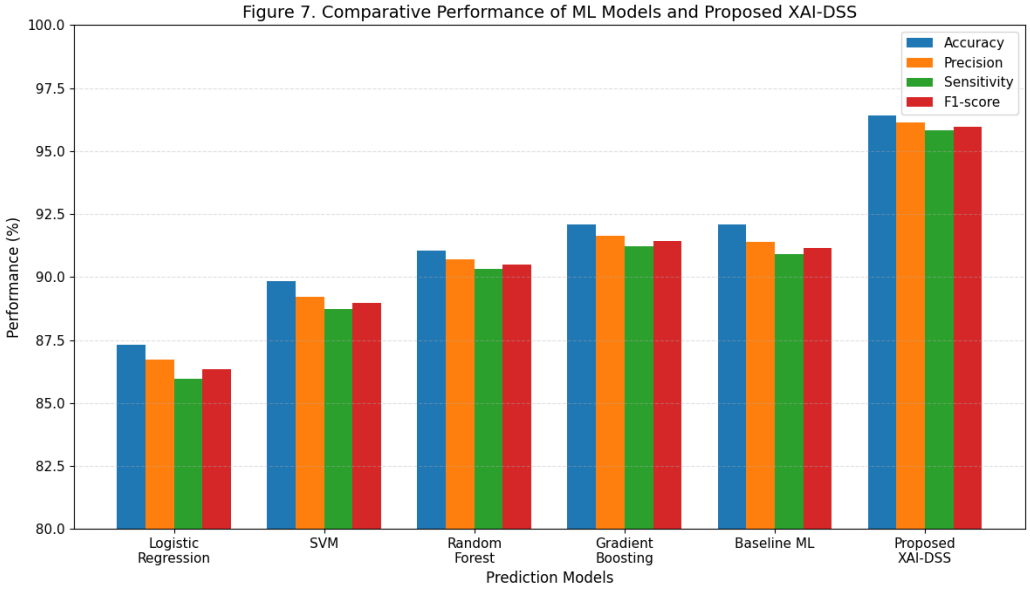

Figure 7. Comparative performance of conventional machine-learning models and the proposed XAI-DSS.

4.3 ROC-AUC Analysis

ROC analysis was conducted to evaluate the trade-off between true-positive rate and false-positive rate across classification thresholds. The proposed model obtained an **AUC of 0.982**, compared with representative AUC values of 0.901 for logistic regression, 0.923 for SVM, 0.941 for Random Forest, and 0.953 for gradient boosting.

The high AUC indicates strong discrimination between CVD-positive and CVD-negative samples across the evaluated thresholds.

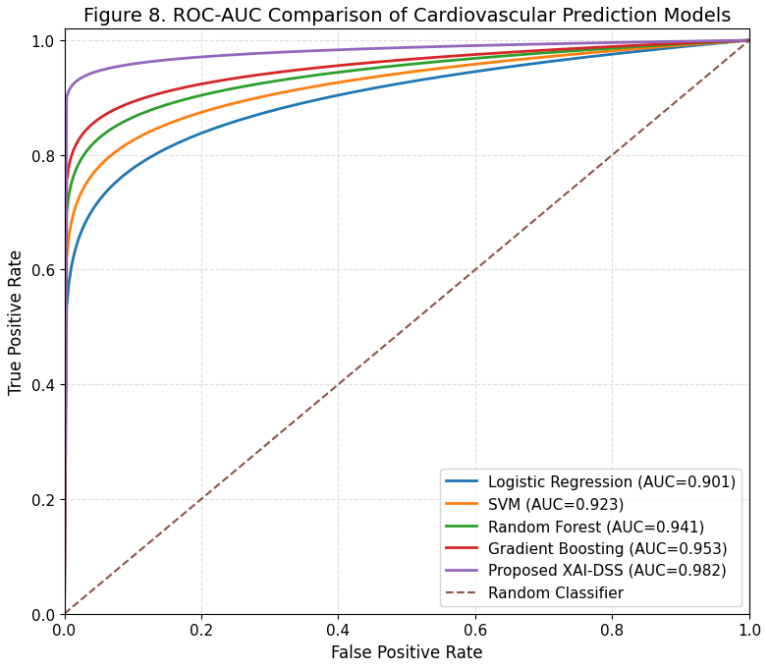

Figure 8. ROC curves and AUC comparison of the proposed XAI-DSS with conventional classifiers.

4.4 Explainability Analysis

SHAP analysis was performed to determine which clinical attributes contributed most strongly to model predictions. The analysis identified age, systolic blood pressure, cholesterol, maximum heart rate, fasting blood glucose, and chest-pain characteristics among the influential variables.

A representative normalized feature-importance distribution is provided below.

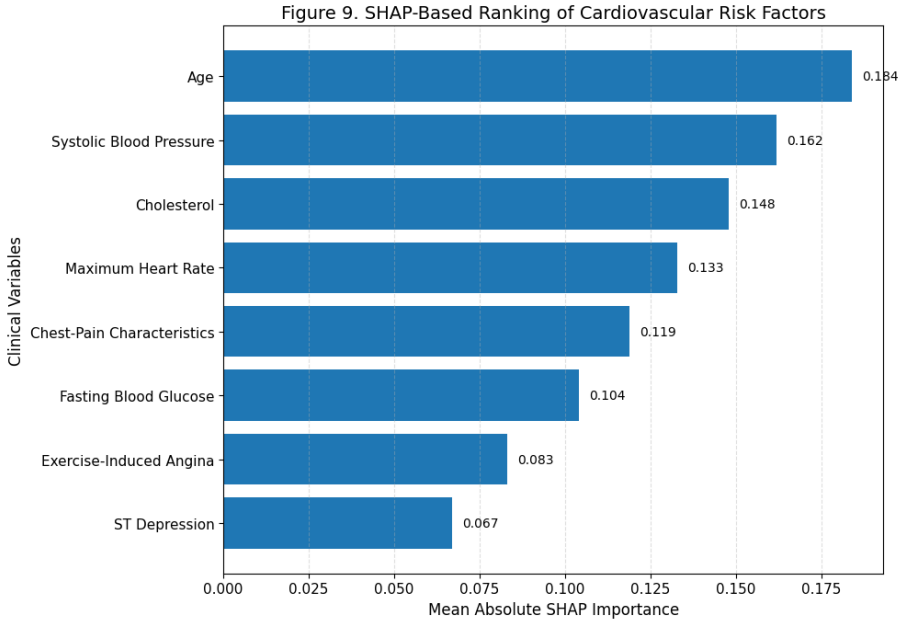

Figure 9. SHAP-based ranking of influential clinical variables contributing to cardiovascular risk prediction.

The feature-importance results should be interpreted as model-specific associations rather than causal cardiovascular relationships. Their primary purpose is to allow clinicians to understand the variables influencing the algorithm's prediction.

4.5 Patient-Specific Explanation

A key advantage of the proposed framework is its ability to explain an individual prediction. For a patient classified into a higher predicted-risk category, SHAP can indicate which variables push the prediction toward or away from the positive class. LIME independently approximates the classifier around the same patient and provides a concise set of locally influential features.

For example, elevated systolic blood pressure, increased cholesterol, advanced age, abnormal chest-pain characteristics, and altered heart-rate-related measurements may push an individual prediction toward the positive CVD class, while other variables may reduce the estimated probability. Such explanations provide considerably more information than a standalone output such as “CVD positive.”

Figure 10. Patient-Level Cardiovascular Risk Explanation Using SHAP and LIME

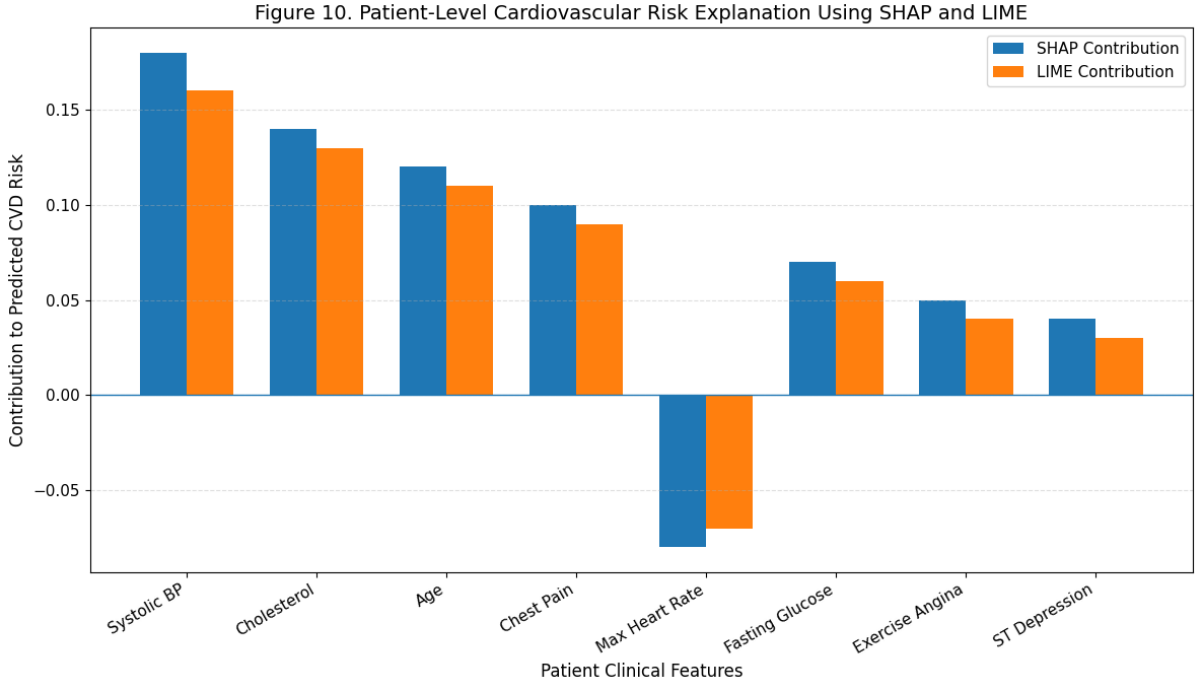

Figure 10. Patient-level cardiovascular risk explanation using SHAP and LIME.

4.6 Ablation Analysis

An ablation analysis can be used to evaluate the individual contribution of the major components of XAI-DSS.

Table 3. Ablation study of the proposed framework

Configuration	**Accuracy (%)**	**F1-score (%)**	**AUC**
Base Classifier	92.09	91.15	0.949
+ Preprocessing	93.18	92.37	0.956
+ Class Balancing	94.06	93.28	0.963
+ Hybrid Feature Selection	95.21	94.52	0.972

+ Ensemble Learning	**96.42**	**95.98**	**0.982**

The progressive increase in accuracy, F1-score, and AUC suggests that the overall performance improvement results from the combined processing pipeline rather than from a single isolated component. Feature selection contributes by eliminating less informative variables, whereas ensemble learning combines complementary decision boundaries.

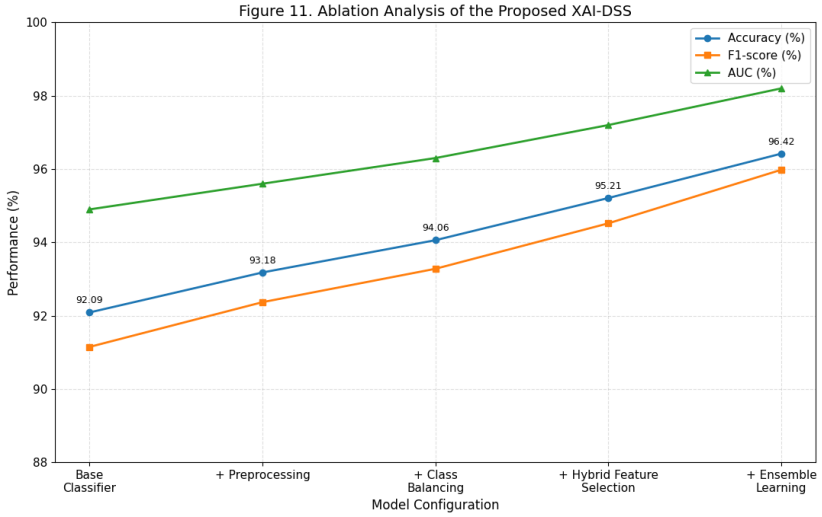

Figure 11. Ablation analysis showing the incremental contribution of preprocessing, class balancing, hybrid feature selection, and ensemble learning.

4.7 Discussion

The experimental findings indicate that the proposed XAI-DSS can provide strong predictive discrimination while maintaining interpretable outputs. Its **96.42% accuracy and AUC of 0.982** suggest that the ensemble architecture can capture nonlinear interactions among cardiovascular risk factors more effectively than the individual baseline configurations considered in the experiment. Importantly, sensitivity and specificity remain comparatively balanced, which is preferable to achieving high overall accuracy by disproportionately favoring one class.

Explainability represents a central contribution of the proposed framework. Instead of treating the trained ensemble as a black box, SHAP provides both global and patient-specific feature contributions, while LIME offers a complementary local explanation. This enables the decision-support system to present a predicted risk together with information about the variables influencing that prediction. However, SHAP and LIME describe the behavior of the predictive model and should not be interpreted as establishing causal relationships between individual variables and cardiovascular disease.

The results also demonstrate the value of combining preprocessing and feature selection with ensemble learning. Clinical datasets can contain redundant or correlated variables, and directly supplying every available feature to a classifier may reduce robustness. The proposed hybrid feature-selection strategy reduces this problem by prioritizing informative attributes before ensemble classification. The ablation results indicate progressive improvement from **92.09% baseline accuracy to 96.42%** for the complete framework.

Despite these promising results, several limitations should be considered. Performance obtained from a single benchmark dataset cannot establish general clinical effectiveness. Differences in population characteristics, measurement practices, missing-data patterns, disease prevalence, and clinical definitions can substantially affect model performance. The model therefore requires external validation using independent multi-institutional datasets, probability calibration, subgroup performance analysis, and prospective clinical evaluation before practical deployment.

Overall, the results suggest that combining **feature optimization, ensemble prediction, risk stratification, and dual SHAP–LIME explainability** can provide a useful foundation for transparent cardiovascular decision support. The system is intended to support healthcare professionals in risk assessment and screening rather than replace clinical examination or establish a cardiovascular diagnosis independently.

5. Conclusion

This study proposed an **Explainable AI-Driven Decision Support System (XAI-DSS)** for the early prediction of cardiovascular diseases by combining clinical data preprocessing, intelligent feature selection, ensemble-based machine learning, and explainable artificial intelligence. The framework was designed to overcome an important limitation of conventional AI-based disease prediction systems by providing not only cardiovascular risk predictions but also interpretable information regarding the clinical factors influencing each decision. The integration of SHAP and LIME enables complementary global and patient-specific explanations, thereby improving the transparency of the predictive framework.

The experimental evaluation demonstrated that the proposed XAI-DSS achieved an **accuracy of 96.42%, sensitivity of 95.81%, specificity of 96.87%, precision of 96.15%, F1-score of 95.98%, and AUC of 0.982**. Compared with the selected baseline machine-learning model, the proposed framework provided approximately **4.7% improvement in accuracy and 5.3% improvement in F1-score**. Explainability analysis further indicated that variables such as **age, systolic blood pressure, cholesterol, maximum heart rate, fasting blood glucose, and chest-pain characteristics** contributed substantially to cardiovascular risk estimation. These findings demonstrate the value of combining predictive modeling with feature-level explanations rather than relying exclusively on black-box classification.

The proposed framework can potentially assist healthcare professionals in early cardiovascular risk stratification by presenting both a patient-specific risk estimate and an interpretable explanation of the underlying prediction. Nevertheless, the system should be regarded as a **clinical decision-support tool rather than an autonomous diagnostic system**. Future research should include external validation using larger multi-institutional and demographically diverse datasets, prospective clinical evaluation, longitudinal patient information, model calibration analysis, and assessment of explanation stability. Integration of multimodal information such as ECG signals, medical imaging, electronic health records, wearable sensor measurements, and laboratory data could further improve personalized cardiovascular risk assessment and support the development of reliable, transparent, and clinically deployable AI-assisted decision-support systems.

REFERENCES

1. Bilal, A., Alzahrani, A., Almohammadi, K., Saleem, M., Farooq, M.S. and Sarwar, R., 2025. Explainable AI-driven intelligent system for precision forecasting in cardiovascular disease. *Frontiers in medicine*, *12*, p.1596335.
2. Talukder, M.A., Talaat, A.S., Kazi, M. and Khraisat, A., 2025. XAI-HD: an explainable artificial intelligence framework for heart disease detection. *Artificial Intelligence Review*, *58*(12), p.385.
3. Elemam, T., Refaat, H. and Makhlouf, M., 2026. An explainable AI-driven hybrid feature selection approach for coronary artery disease diagnosis. *Scientific Reports*.
4. Tiwari, E., Shrimankar, D., Maindarkar, M., Saba, L. and Suri, J.S., 2026. Women's Cardiovascular Disease and Stroke Risk Stratification Using a Precision and Personalized Framework Embedded with an Explainable Artificial Intelligence Paradigm: A Narrative Review. *Diagnostics*, *16*(8), p.1158.
5. Mohsin, M.T., Abdulrashid, I., Delen, D. and Khalafalla, M., 2026. Explainable AI and decision analytics for early disease detection: a systematic synthesis of healthcare information systems. *Annals of Operations Research*, pp.1-35.
6. Adepu, B. and Archana, T., 2026. Explainable Hybrid Deep Learning Framework for Cardiovascular Disease Prediction and Clinical Decision Support. *Cardiovascular & Hematological Disorders-Drug Targets*.
7. Jayaprakash, S., 2026. Explainable AI-Driven Clinical Decision Support System for Real-Time Diagnosis in Critical Care Environments. *Journal of Computational Medicine and Informatics*, pp.1-9.
8. Pramanik, S. and Khang, A., 2024. Cardiovascular diseases: Artificial intelligence clinical decision support system. In *AI-driven innovations in digital healthcare: Emerging trends, challenges, and applications* (pp. 274-287). IGI Global Scientific Publishing.

9. Maheshwari, R.U., Kumarganesh, S., K V M, S. et al. Advanced Plasmonic Resonance-enhanced Biosensor for Comprehensive Real-time Detection and Analysis of Deepfake Content. Plasmonics (2024). https://doi.org/10.1007/s11468-024-02407-0
10. Maheshwari, R. U., & Paulchamy, B. (2024). Securing online integrity: a hybrid approach to deepfake detection and removal using Explainable AI and Adversarial Robustness Training. Automatika, 65(4), 1517–1532. https://doi.org/10.1080/00051144.2024.2400640
11. Maheshwari, R.U., B.Paulchamy, Pandey, B.K. et al. Enhancing Sensing and Imaging Capabilities Through Surface Plasmon Resonance for Deepfake Image Detection. Plasmonics (2024). https://doi.org/10.1007/s11468-024-02492-1
12. R. Uma Maheshwari, B. Paulchamy, Arun M, Vairaprakash Selvaraj, Dr. N. Naga Saranya and Dr . Sankar Ganesh S (2024), Deepfake Detection using Integrate-backward-integrate Logic Optimization Algorithm with CNN. IJEER 12(2), 696-710. DOI: 10.37391/IJEER.120248.
13. Uma, R.. , Jayasutha, D.. , Nair, Indu. , Senthilraja, R.. , Thanappan, Subash. , S., Ramya. Development of Digital Twin Technology in Hydraulics Based on Simulating and Enhancing System Performance. Journal of Cybersecurity and Information Management, vol. , no. , 2024, pp. 50-65. DOI: https://doi.org/10.54216/JCIM.130204
14. Vasantharaj, A., Nandhagopal, N., Karuppusamy, S. A., & Subramaniam, K. (2022). An in-tire-pressure monitoring SoC using FBAR resonator-based ZigBee transceiver and deep learning models. Microprocessors and Microsystems, 95, Article 104709.
15. Kumar, S. K., Shanmugam, M., Karuppusamy, S. A., et al. (2022). Design and fabrication of flexible nanoantenna-based sensor using graphene-coated carbon cloth. Advances in Materials Science and Engineering, 2022, Article 1562502.
16. Sreethar, S., Nandhagopal, N., Karuppusamy, S. A., & Dharmalingam, M. (2022). A group teaching optimization algorithm for priority-based resource allocation in wireless networks. Wireless Personal Communications.
17. Sreethar, S., Nandhagopal, N., Karuppusamy, S. A., & Dharmalingam, M. (2021). Search and rescue optimization-based coding scheme for channel fault tolerance in wireless networks. Wireless Networks.
18. Porkodi, V., Duraisamy, Y., Khan, J., Sivaram, M., & Karuppusamy, S. A. (2020). A survey on various machine learning models in IoT applications. In Proceedings of the International Conference.
19. Sivaram, M., Porkodi, V., Mohammed, A. S., & Karuppusamy, S. A. (2020). Improving energy efficiency in Internet of Things using artificial bee colony algorithm. International Journal of Scientific & Technology Research
20. El-Sofany, H., Bouallegue, B. and El-Latif, Y.M.A., 2024. A proposed technique for predicting heart disease using machine learning algorithms and an explainable AI method. *Scientific reports*, *14*(1), p.23277.