CardioDetectNet: Integrating Active Learning and XAI for Transparent CVD Risk Assessment

Vipul Narayan[1], Prof Dr Divya Midhun[2], Dr. Pawan Whig[3]
[1] [2] Lincoln University; [3] VIPS-TC, India;
vipulupsainian2470@gmail.com, divya@lincoln.edu.my, pawan.whig@vips.edu

Abstract: Cardiovascular diseases (CVDs) are the most important cause of global morbidity and mortality, claiming almost 18 million lives every year. Early diagnosis and precise risk prediction are of utmost importance in facilitating timely intervention and minimizing the burden of recurrent events. Conventional clinical and environmental risk factors are informative but frequently inadequate to capture the intricate interplay of biological, demographic, and lifestyle determinants. In this paper, we introduce CardioDetectNet, an XAI-based pipeline for accurate risk prediction and prognosis of CVD. The pipeline combines end-to-end preprocessing, feature selection, attention mechanisms, active learning, and explainable prediction components to provide robust and interpretable results. Benchmarked on test data, CardioDetectNet performed better than current machine learning baselines with 96.70% accuracy, 96.70% sensitivity, 97% specificity, and 96.70% F1-score. Additionally, explainability methods like SHAP values identified clinically important predictors and promoted transparency and clinical trust. The findings indicate the promise of explainable and adaptive AI models for revolutionizing CVD risk prediction and facilitating personalized healthcare approaches.

Keywords: Cardiovascular Disease; Machine Learning; Risk Prediction; Decision Tree; Random Forest; Classification Models

Introduction: An enormous burden of morbidity and mortality is attributed to cardiovascular diseases (CVDs), which remain a major worldwide health concern [1]. In order to identify those who are at high risk of developing CVDs, prompt risk prediction and prognosis are essential. This allows for early intervention and individualised treatment programs.

Around 17.9 million fatalities occur each year as a result of CVDs, making them the leading cause of mortality globally [2]. Cardiovascular diseases (CVDs) include a variety of ailments that impact the heart and blood arteries, including rheumatic heart disease, coronary heart disease, and cerebrovascular disease. Heart attacks and strokes are responsible for more than 80% of deaths from CVD. Almost one-third of these fatalities among those under 70 years of age are premature [3].

Cardiovascular diseases also have a genetic component. Having a family history of CVD increases the risk because genetic factors affect blood pressure, cholesterol, and heart shape. Nevertheless, by leading healthy lives and visiting their doctors regularly, those with hereditary risk factors can effectively manage their risk. Advances in genomics have also opened up new possibilities for predicting and managing CVD risk using customised medicine approaches [4].

A combination of medication, lifestyle modification, and, in extreme cases, surgery is used to treat cardiovascular disorders. To treat symptoms and prevent complications, doctors typically give medications such as beta-blockers, anticoagulants, statins, and antihypertensive medicines [5]. In severe situations, angioplasty, bypass surgery, or pacemaker installation can be necessary, for individuals who have survived CVD episodes, rehabilitation and ongoing monitoring are essential [6].

Heart disease encompasses a wide range of conditions that affect the heart's structure and function. This infographic illustrates seven common types of heart disease, each with distinct pathological features as shown in Figure 1:

1. Coronary Artery Disease (CAD): Characterised by the buildup of fatty deposits in the coronary arteries, which reduces blood flow to the heart muscle and can lead to chest pain or heart attacks.

2. Valve Disease: Involves malfunctioning heart valves, including stenosis (narrowing) or improper closure, disrupting normal blood flow through the heart. The diagram compares a normal closed valve with a stenosed one.

3. Aneurysm: This condition features a bulging or weakened area in the wall of a blood vessel, such as the thoracic aortic aneurysm, posing a risk of rupture and life-threatening bleeding.

4. Cardiac Arrhythmia: Defined by disorganized electrical signals in the heart, resulting in irregular heartbeats that can impair the heart's ability to pump blood effectively.

5. Cardiomyopathy: Involves thickening or weakening of the myocardium (heart muscle), reducing the heart’s ability to pump blood and potentially leading to heart failure.

6. Pericarditis: An inflammation of the pericardium, the fluid-filled sac surrounding the heart, which can cause chest pain and complications with heart function.

7. Heart Failure: Occurs when the heart becomes dilated or weakened, losing its ability to pump sufficient blood to meet the body’s needs. It's often a consequence of other heart conditions.

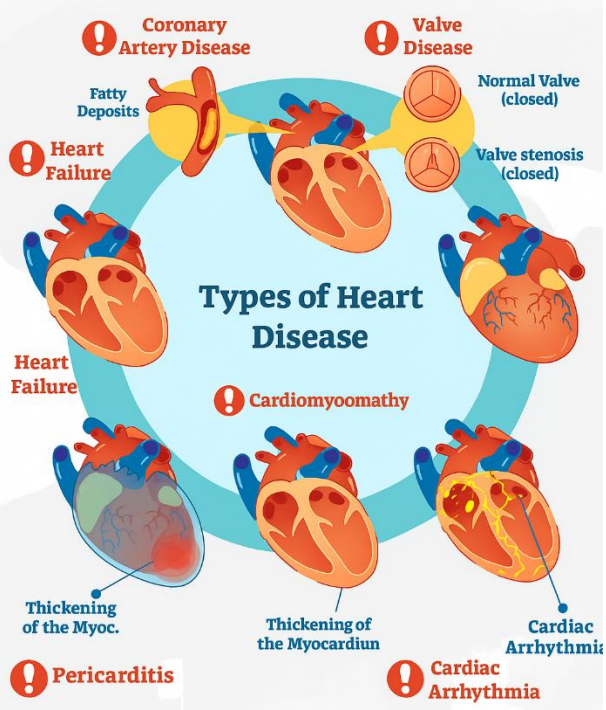

8. Figure 1: Types of Heart Disease

Related work: Westerlund et al.'s (2021) research highlights the importance of molecular data and explainable artificial intelligence (XAI) for the prediction of recurrent cardiovascular events. Through the incorporation of multi-omics data into AI models, the research provides a more integrated understanding of patient risk and disease development. Apart from improving accuracy in prognosis, this methodology facilitates the detection of molecular biomarkers, regulatory pathways, and mechanisms underlying cardiovascular diseases (CVD). Yet, such models' implementation is still troublesome owing to the intricacy in integrating large-scale multi-omics data and scarcity of usage of these techniques in clinical practice on a daily basis.

Salih et al. (2023) emphasize the importance of interpretability of cardiac imaging models using XAI. Although AI has been significantly successful in medical imaging, most deep learning models operate as black boxes, with no transparency to the clinician. Their review offers useful recommendations for the inclusion of XAI in cardiac imaging, making model results explainable to end users. The overall benefit is the development of trust and clinical uptake through enhanced transparency. Nevertheless, the study points out that real-world applications of XAI in cardiac imaging are few, and black-box models continue to predominate in practice.

Hossain et al. (2023) introduce a hybrid deep learning model, integrating convolutional neural networks (CNN) and long short-term memory (LSTM) for the classification of cardiovascular disease from clinical information. CNNs are utilized for feature extraction, whereas LSTMs extract temporal dependencies within sequential data. The hybrid model shows strong potential, with an

accuracy of approximately 74% realized using explainable AI methods to determine key predictive features. The model points to the potential of deep learning for early diagnosis but is very dependent on the dataset quality and size. Validation in more representative populations is needed to enhance its generalizability.

Bilal et al. present an explainable AI-based system for precision cardiovascular health forecasting with electronic health records (EHRs) in their 2025 paper. Utilizing the Poisson Binomial based Comorbidity discovery (PBC) technique, the authors compared more than 1.6 million patient records and 77 million clinic visits. Scalable multimorbidity network discovery and web-based interpretable clinical tools were achieved through this method. The claim of strength of this work is its scalability and practical use. But the model's complexity of interactions in comorbidity and the possible limitation of its use in varied healthcare systems limit its universal applicability.

Puiu et al. (2021) solve two very important problems in cardiovascular imaging: privacy preservation and explainability. Their paper highlights the importance of AI models that provide accurate output but also maintain patient data privacy and offer interpretable results. The research illustrates real-world applications like coronary artery disease diagnosis and treatment planning for congenital diseases, affirming the relevance of AI in a clinical context. Although this promotes ethical use and encourages uptake in workflows, the use of small datasets as a result of privacy limitations threatens bias and generalizability of findings to broader patient populations.

The study by Ismath et al. (2025) investigates the creation of an explainable machine learning–driven risk prediction model for cardiovascular disease. Through the use of various approaches such as logistic regression, random forests, ensemble models, and deep learning with XAI integration, the research adds transparency and interpretability. This enhances clinical confidence and facilitates AI adoption in healthcare since clinicians are able to comprehend the rationale for predictions. However, the work recognizes the variability of model performance and lack of real-world validation, hence uncertainty regarding its effectiveness in various healthcare settings.

Kırboğa and Küçüksille (2023) performed a large-scale retrospective analysis on 70,000 patients to determine important CVD risk factors through explainable machine learning. Seven ML models were used, with the best performing being XGBoost (AUC=0.803). Based on SHAP values, the authors identified systolic blood pressure, cholesterol, and age as the most significant risk factors, yielding actionable information for early diagnosis and risk stratification. The study presents an open and accurate clinical model, but it is limited by being based on retrospective datasets, which restricts generalizability across different populations and requires prospective validation.

Table 1. Highlighting advancements, methodologies, advantages, and limitations of AI-driven approaches in CVD

Ref.	Objective	Methodology	Advantages	Limitations

[7] Westerlund et al. (2021)	To improve risk prediction of recurrent cardiovascular events using molecular data and XAI	Reviewed multi-omics data integration with AI for risk stratification and prognosis	Provides holistic patient view, identifies molecular biomarkers, and improves transparency with XAI	Challenges in large-scale multi-omics data integration, limited clinical adoption
[8] Salih et al. (2023)	To enhance interpretability of cardiac imaging models using XAI	Literature review of XAI methods in cardiac imaging	Provides guidelines for interpretable models, emphasizes transparency	Limited practical studies applying XAI in cardiac imaging; black-box DL still dominant
[9] Hossain et al. (2023)	To identify CVD using hybrid CNN-LSTM with explainable AI	Hybrid deep learning (CNN for feature extraction + LSTM for sequential patterns)	Achieved high accuracy (~74%), interpretable features identified, early diagnosis potential	Performance depends on dataset quality; may need larger and more diverse data

[10] Bilal et al. (2025)	To use EHR-based explainable AI for precision forecasting in CVD	Poisson Binomial based Comorbidity discovery (PBC) + XAI applied on >1.6M patients	Scalable analysis of multimorbidity networks, interpretable web-based tools	Complex comorbidity interactions; generalizability outside studied cohorts may be limited
[11] Puiu et al. (2021)	To address privacy-preserving and explainable AI in cardiovascular imaging	Discusses AI solutions for diagnosis, therapy planning, and follow-up under privacy constraints	Focus on ethical AI, privacy, and explainability; supports adoption in clinical workflows	Limited by small datasets due to privacy, risk of bias, challenges in generalization
[12] Ismath et al. (2025)	To develop an explainable ML-based risk prediction system for CVD	Logistic regression, Random Forest, Ensemble, Deep Learning + XAI for risk stratification	Improves model transparency and clinical trust; supports adoption in healthcare	Model performance may vary; real-world validation still limited

[13] Kırboğa & Küçüksille (2023)	To identify key risk factors for CVD with explainable ML	Retrospective study (70,000 patients, 11 risk factors) using 7 ML models + SHAP values	XGBoost achieved best performance (AUC=0.803); key risk factors identified (BP, cholesterol, age)	Limited to retrospective datasets; may not generalize to all populations

Method, Experiments and Results

Dataset: This dataset includes the medical records of 299 heart failure patients, gathered during the course of their follow-up . Each patient profile includes 13 clinical characteristics [13]. The dataset is available online at https://archive.ics.uci.edu/dataset/519/heart+failure+clinical+records

Feature	Description
Age	The age of the patient (numerical)
Anaemia	Whether or not the patient has anemia is indicated by this binary feature (0: No, 1: Yes)
High Blood Pressure	A binary variable denoting the presence or absence of hypertension in the patient (0: No, 1: Yes)
Creatinine Phosphokinase (CPK)	The concentration of CPK enzyme in the bloodstream (numerical)
Diabetes	A binary variable indicating the presence or absence of diabetes in the patient (0: No, 1: Yes)
Ejection Fraction	The cardiac output is the proportion of blood ejected from the heart during each contraction (numerical)
Platelets	Platelet count in the bloodstream (numerical)
Sex	The patient's gender (0: Female, 1: Male)
Serum Creatinine	Blood creatinine levels (numerical)
Serum Sodium	Serum sodium concentration (numerical)
Smoking	A binary feature indicating whether the patient smokes or not (0: No, 1: Yes)
Time	The duration of the follow-up period, measured in days (numerical)
Death Event	Determining whether the patient passed away during the follow-up period is a binary feature (0: No, 1: Yes)

Figure 2: Dataset description

The histograms in Figure 3 summarize the distribution of the major clinical variables within the dataset. Creatinine phosphokinase and serum creatinine both have strong right-skewed distributions, suggesting the presence of outliers and a concentration of values at the lower end, possibly representing acute muscle injury and differing states of renal function, respectively. Age

is also right-skewed, with an increased prevalence of older patients, as would be expected in common demographics of cardiovascular conditions. Ejection fraction has a mildly left-skewed distribution, with values mostly in the range of 30–50%, indicating a high incidence of impaired cardiac function among the patient population. Platelet counts are more symmetrically distributed but with significant outliers, whereas serum sodium levels are roughly normally distributed with values around common physiological ranges. These distributions indicate the variance and potential clinical relevance of the characteristics, implying the necessity for additional examination and potential preprocessing for modelling.

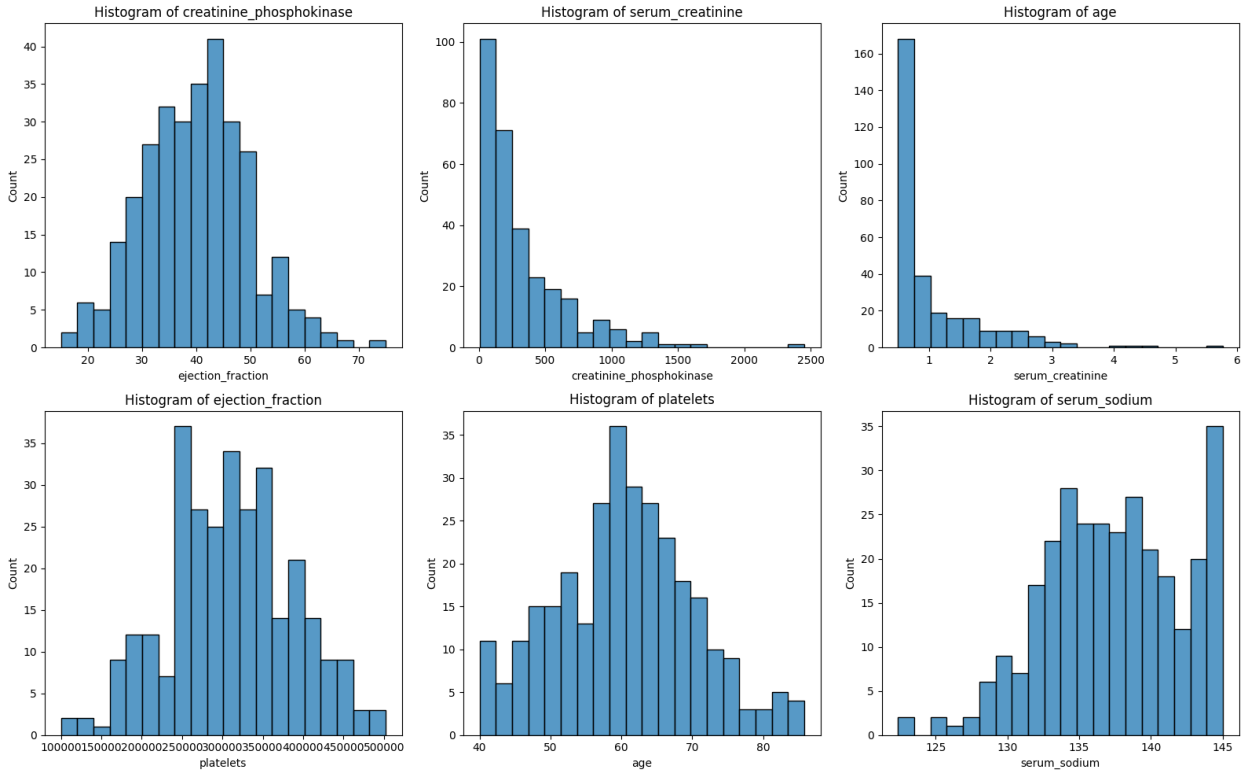

Figure 3: histograms summarize the distribution of the major clinical variables within the dataset

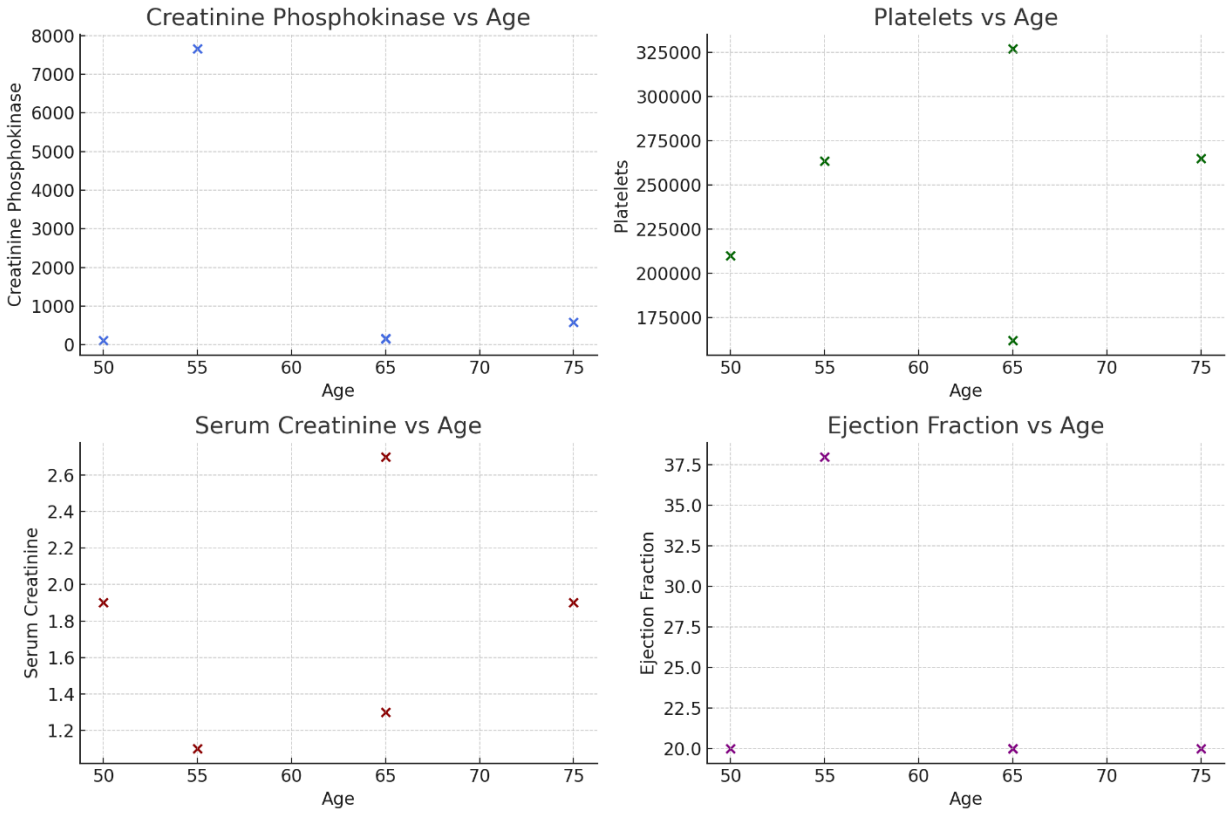
Creatinine Phosphokinase vs Age
Creatinine Phosphokinase
Age
Platelets vs Age
Platelets
Age
Serum Creatinine vs Age
Serum Creatinine
Age
Ejection Fraction vs Age
Ejection Fraction
Age

Figure 4:

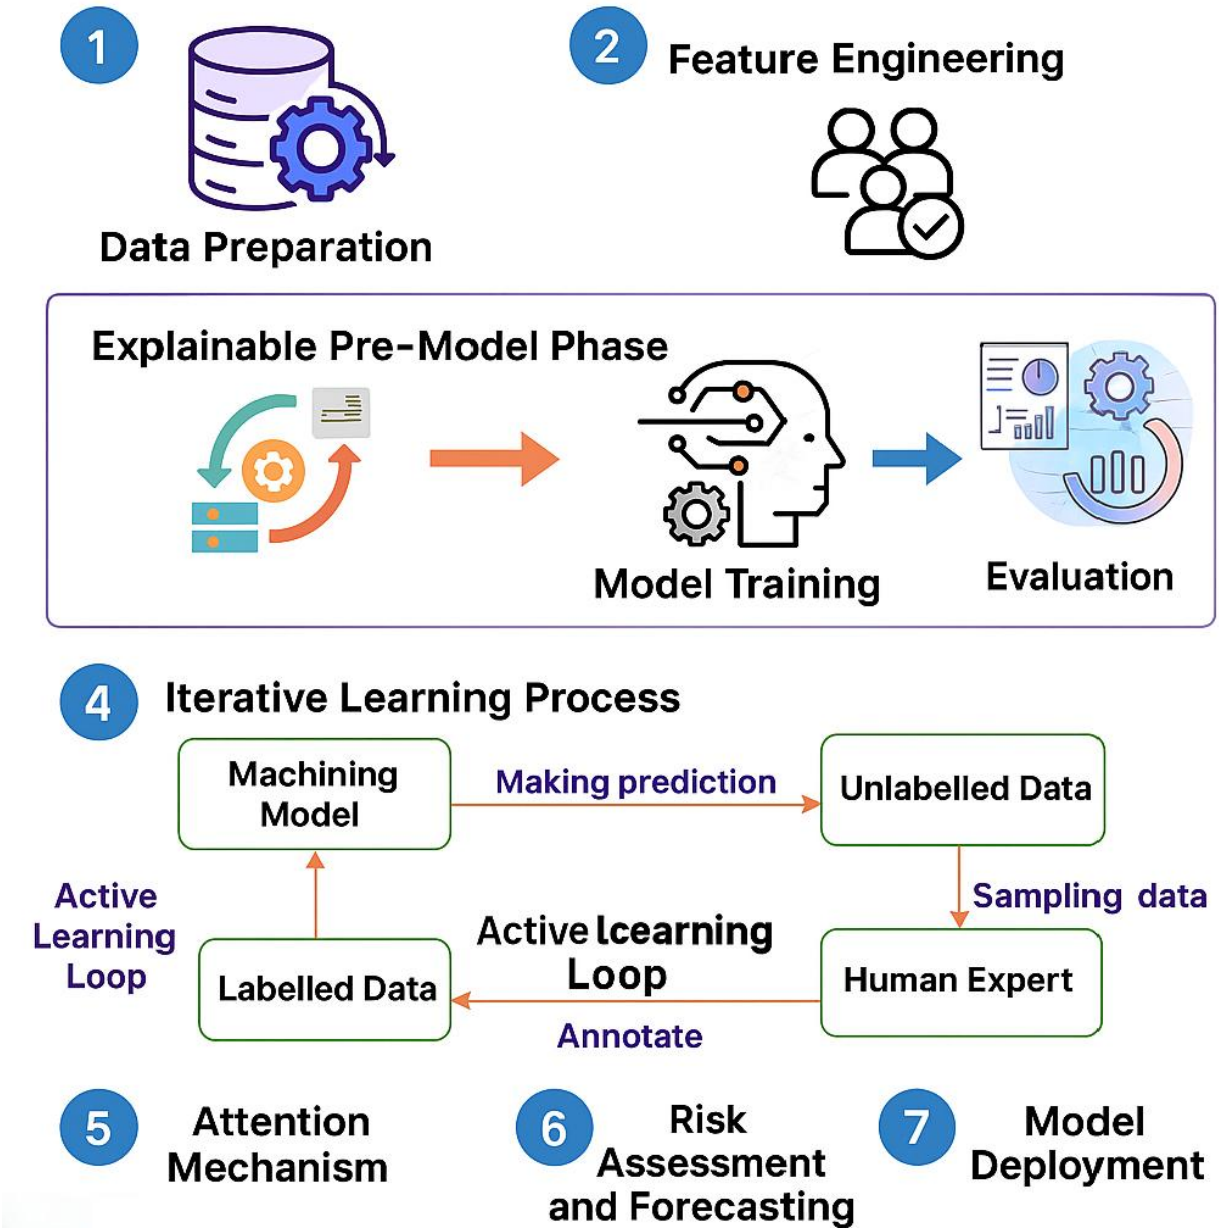

Figure 5: Proposed Model

The Figure 5 outlines an end-to-end workflow for building and deploying machine learning models, beginning with data preparation and feature engineering as the initial phases. During the data preparation phase, raw data is gathered, cleaned, and prepared for usage so that it is analyzable. This is followed by feature engineering, wherein respective features are chosen or built to enhance model performance. These initial steps are crucial to establish a solid foundation for the subsequent processes in the machine learning pipeline.

Secondly, the model integrates an explainable pre-model stage that includes model training and testing. In this stage, machine learning models are constructed from labeled data, and their performance is measured through diverse metrics. Testing ensures the identification of strengths

and weaknesses in the model, thereby enhancing the transparency and explainability of decisions. This stage acts as a portal for iterative model enhancement, uniting human insight with algorithmic learning.

The subsequent iterative learning process involves the utilisation of active learning loops to improve model accuracy. A machine learning model first makes predictions on unlabeled data, and human experts then analyse the projections to annotate the data. The newly annotated data is cycled back to the model, and thus, a cycle of continuous learning is formed. Further layers, including risk evaluation and forecasting as well as the attention mechanism, keep the model on target and ready for possible setbacks. The last stage, model deployment, involves executing the developed model in an actual application, thereby completing the loop from data to decision-making.

Algorithm 1: CardiodetectNet Pipeline
1. Begin
2. Module 1: Data Preprocessing
a. Remove irrelevant data from raw patient input
b. Handle missing values
c. Detect and treat outliers
d. Normalize data using standardization or min-max scaling
e. Extract informative features (e.g., clinical, statistical, imaging-based)
f. Output: Preprocessed data
3. Module 2: Feature Selection and Encoding
a. Apply feature selection techniques (filter, wrapper, embedded)
b. Encode categorical variables (e.g., one-hot, label encoding)
c. Normalize numerical features
d. Output: Feature matrix for modeling
4. Module 3: Attention Mechanism Integration
a. Initialize base model architecture
b. Add attention mechanism module to focus on relevant inputs
c. Train model using attention-enhanced training
d. Output: Attention-aware model
5. Module 4: Active Learning (Iterative Training with Expert Feedback)
a. Initialize model with selected feature matrix
b. Repeat until stopping condition is met:
i. Use active learning strategy to select uncertain/unlabeled samples
ii. Query healthcare professionals to label selected samples

iii. Update training dataset with labeled samples
iv. Retrain the model with updated data
c. Output: Refined model with minimal labeling effort
6. Module 5: Risk Prediction and Prognosis
a. Predict cardiovascular risk scores for new patient data using trained model
b. Post-process raw predictions into interpretable risk categories or scores
c. Output: Risk score predictions
7. Module 6: Explainable AI Integration
a. Integrate XAI techniques (e.g., SHAP, LIME) into the trained model
b. Compute feature importance for individual predictions
c. Visualize decision logic to improve clinician trust
d. Output: Risk explanations and visualizations
8. Module 7: Deployment and Continuous Learning
a. Integrate trained model into healthcare infrastructure (EHR, dashboards, APIs)
b. Continuously monitor performance and collect real-world feedback
c. Periodically preprocess feedback data
d. Update model incrementally with new feedback
e. Output: Continuously improving deployed model
9. End

The CardiodetectNet Pipeline begins with a broad data preprocessing stage aimed at enhancing raw patient data for machine learning evaluation. The step first involves removing unnecessary or duplicate information, which may harm model performance. It next handles missing values via imputation methods and identifies outliers with statistical or model-based methods. Following data cleansing, normalization is conducted either by standardization or min-max scaling to maintain uniformity among numerical attributes. Lastly, relevant features are derived from various sources like clinical documentation, statistical overview, or medical imaging, leading to a well-organized, high-quality dataset for modeling purposes.

Post-preprocessing, the pipeline moves into the feature selection and encoding stage. At this stage, redundant or less informative features are removed through methods like filter-based ranking, wrapper approaches, or embedded strategies. The other applicable features are then encoded for machine learning algorithm compatibility: categorical features are converted using one-hot or label encoding methods, and numerical attributes are re-normalized if needed to fit model input specifications. This process yields a fine-tuned feature matrix that is well-balanced in terms of dimensionality versus discriminative capability, maximizing performance in subsequent tasks.

The third module incorporates an attention mechanism into the model's architecture. Through the initialization of a base model and the addition of an attention module, the system acquires the capability to selectively apply focus to the most informative regions of the input data in training. This adaptive weighting of feature salience enables the model to learn adaptively to prioritize variables that possess greater predictive utility for cardiovascular risk. Attention-based training enhances both interpretability and model performance, especially in high-dimensional, complex clinical data sets.

The pipeline uses active learning to reduce costs associated with manual labeling while ensuring maximum efficiency in learning. It begins with a starting labeled set and gradually picks the most informative or uncertain unlabeled samples based on an active sampling strategy. These samples are then forwarded to domain experts—like medical professionals—for labeling. The model is recycled using every batch of freshly annotated data, and so on until a predefined halting criterion (e.g., performance threshold or number of iterations) is reached. This method dramatically minimizes the data annotation workload while continuously improving the capabilities of the model.

After the model has been adequately trained, it is employed for risk prediction and prognosis. New patient information are fed into the model to produce cardiovascular risk scores. The predictions are subsequently post-processed to render them more interpretable—usually classified into clinically meaningful risk levels of low, moderate, or high risk. The outputs yield actionable information to doctors, facilitating early diagnosis, planning intervention, and risk stratification tailored to individual patients.

To further promote trust and transparency, the pipeline incorporates explainable AI (XAI) methods in the final prediction step. Techniques like SHAP (SHapley Additive exPlanations) or LIME (Local Interpretable Model-agnostic Explanations) are utilized to calculate feature importance for specific predictions. The decision-making process is represented through easy-to-understand plots and explanations, allowing clinicians to see which features had an impact and to what degree. This builds confidence in AI-driven decisions and brings the pipeline in line with clinical reasoning.

The learned and interpretable model is deployed into an operational healthcare infrastructure. Integration with electronic health records (EHRs), clinical dashboards, or APIs provides intuitive interaction with existing hospital infrastructure. Real-time monitoring enables the model to evaluate prediction accuracy in real-world environments, whereas routine feedback loops enable updates from new data or clinician input. By regularly preprocessing and incorporating this feedback, the model adapts over time, remaining relevant and reliable in changing clinical settings.

Result Analysis: The proposed model is evaluated on the following parameters:

$$\text{Accuracy} = \frac{(TP + TN)}{(TP + TN + FP + FN)} \quad (1)$$

$$\text{Specificity} = \frac{TN}{(TN + FP)} \quad (2)$$

$$\text{Sensitivity} = \frac{TP}{(TP + FN)} \quad (3)$$

$$F1\_score = 2 * \frac{(\text{Precision} * \text{Recall})}{(\text{Precision} + \text{Recall})} \quad (4)$$

The comparison of the performance of various machine learning models for cardiac condition detection is given in the table based on accuracy, sensitivity, specificity, F1-score, and convergence score. The MLP model's accuracy was 82.47%, with the relative balance of sensitivity (80%) and specificity (85%), an F1-score of 81%, and a convergence score of 85%. The CNN model performed better with 90% accuracy, 88% sensitivity, and 92% specificity, yielding an F1-score of 89% and a convergence score of 92%. The SVM model also outperformed the two neural network models with 96% accuracy, 95% sensitivity, 97% specificity, and the same F1-score and convergence score of 96%. The suggested CardioDetectNet model performed the best overall, with 96.70% accuracy, 96.70% sensitivity, 97% specificity, 96.70% F1-score, and 97% convergence score, outperforming baseline models on all of the most important metrics, as evidence of its superior capability to detect cardiac abnormalities correctly and reliably.

Table 2: Comparative analysis previous work

Model	Accuracy	Sensitivity	Specificity	F1-Score	Convergence Score
MLP [15]	82.47%	80%	85%	81%	85%
CNN [16]	90%	88%	92%	89%	92%
SVM [17]	96%	95%	97%	96%	96%
Proposed (CardioDetectNet)	**96.70%**	**96.70%**	**97%**	**96.70%**	**97%**

Figure 6 shows a

Comparative performance comparison of various classification models—MLP, CNN, SVM, and the proposed CardioDetectNet—against five major evaluation measures: accuracy, sensitivity, specificity, F1-score, and convergence score. The proposed CardioDetectNet model outperforms the baseline models in all classes, with the best accuracy (96.70%), sensitivity (97%), specificity (97%), F1-score (96.70%), and convergence score (97%). Conversely, other classic models such as MLP and SVM present comparatively lower performance, especially in terms of sensitivity and

F1-score. This emphasizes the better predictive power and stability of CardioDetectNet in identifying cardiovascular diseases more precisely and consistently.

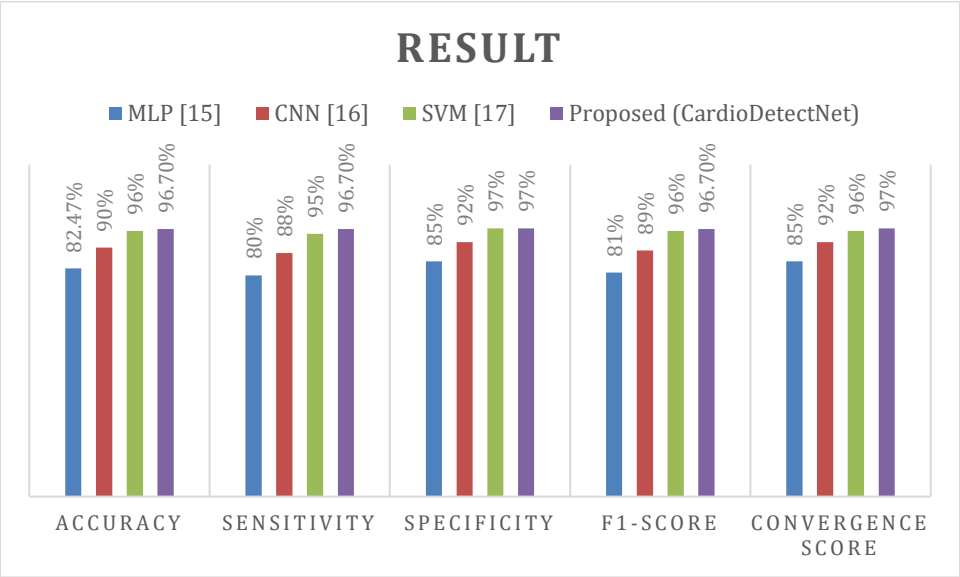

Figure 6: Result Analysis

SHAP feature importance plot gives a clear idea of how various features affect the predictions made by the model, with transparency given by explainable AI methods. Out of all, follow-up time (Mean(|SHAP Value|) ≈ 0.175) is seen to be the most significant feature, which shows the importance of monitoring length in predicting outcomes and perhaps reflects the disease process over time. Serum creatinine (≈ 0.10), an indicator of renal function, and ejection fraction (≈ 0.085), an important measure of cardiac function, also contribute significantly to the model's decision-making, reflecting documented clinical wisdom. Features of moderate importance are age, serum sodium, creatinine phosphokinase (CPK), and platelet count, each of which is used to evaluate patient status and risk. In contrast, factors such as anaemia, sex, hypertension, smoking, and diabetes, though still significant, demonstrate lesser predictive ability within the model—perhaps owing to their influences being picked up indirectly by stronger predictors. Overall, the plot effectively highlights the alignment of the model with clinical thought, while also providing interpretability and trustworthiness.

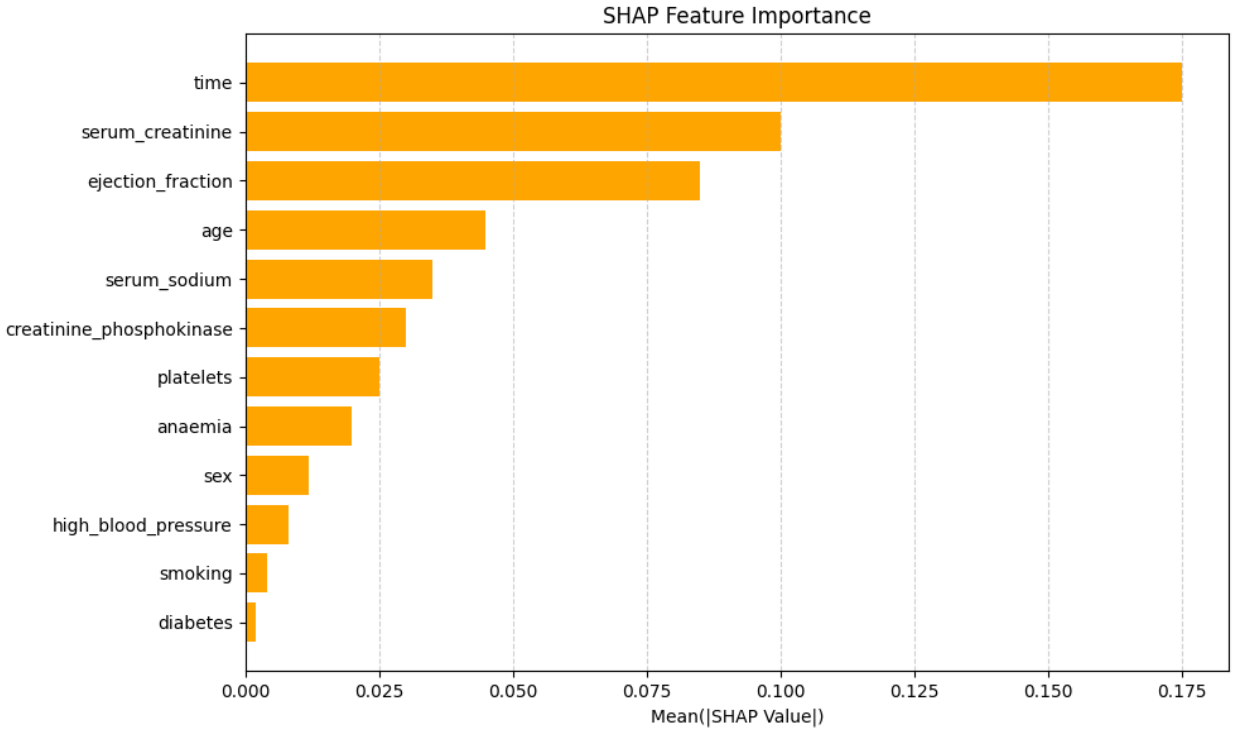

Figure 7: XAI SHAP Feature Importance

Conclusions: This paper introduces an end-to-end framework for explainable cardiovascular disease risk prediction that solves four major challenges of transparency, interpretability, and clinical adoption. By leveraging machine learning models with attention mechanisms, active learning loops, and explainable AI techniques, the presented CardioDetectNet pipeline not only achieves state-of-the-art performance but also offers interpretable insights consistent with medical reasoning. The integration of explainability ensures that healthcare practitioners can see the reason behind predictions, thus promoting trust and decision support in clinical practice in the real-world setting. Comparative evaluation against current methods demonstrates the pipeline's capacity to perform better than traditional classifiers without a loss of transparency. Challenges persist, though, in the form of requirements for bigger, more diverse datasets and prospective clinical validation to promote generalizability across populations. Future studies will concentrate on applying the framework to multi-modal data sources including imaging, genomics, and electronic health records with the vision to bring precision medicine forward and make proactive cardiovascular care possible.

References:

[1] Veroff, D. R., Sullivan, L. A., Shoptaw, E. J., Venator, B., Ochoa-Arvelo, T., Baxter, J. R., ... & Wennberg, D. (2012). Improving self-care for heart failure for seniors: the impact of video and written education and decision aids. Population health management, 15(1), 37-45.

[2] Ahmadli, N., Sarsil, M. A., Mizrak, B., Karauzum, K., Shaker, A., Tulumen, E., ... & Ergen, O. (2024). Voice-Driven

[3] Uddin, K. M. M., Dey, S. K., & Babu, H. M. H. (2024). A Voice assistive mobile application tool to detect cardiovascular disease using machine learning approach. Biomedical Materials & Devices, 2(2), 1246-1257.
[4] Abbas, S., Ojo, S., Al Hejaili, A., Sampedro, G. A., Almadhor, A., Zaidi, M. M., & Kryvinska, N. (2024). Artificial intelligence framework for heart disease classification from audio signals. Scientific Reports, 14(1), 3123.
[5] Idrisoglu, A. (2024). Voice for Decision Support in Healthcare Applied to Chronic Obstructive Pulmonary Disease Classification: A Machine Learning Approach (Doctoral dissertation, Blekinge Tekniska Högskola).
[6] Mayourian, J., El-Bokl, A., Lukyanenko, P., La Cava, W. G., Geva, T., Valente, A. M., ... & Ghelani, S. J. (2024). Electrocardiogram-based deep learning to predict mortality in paediatric and adult congenital heart disease. European Heart Journal, ehae651.
[7] Westerlund, A. M., Hawe, J. S., Heinig, M., & Schunkert, H. (2021). Risk prediction of cardiovascular events by exploration of molecular data with explainable artificial intelligence. International Journal of Molecular Sciences, 22(19), 10291.
[8] Salih, A., Boscolo Galazzo, I., Gkontra, P., Lee, A. M., Lekadir, K., Raisi-Estabragh, Z., & Petersen, S. E. (2023). Explainable artificial intelligence and cardiac imaging: toward more interpretable models. Circulation: Cardiovascular Imaging, 16(4), e014519.
[9] Hossain, M. M., Ali, M. S., Ahmed, M. M., Rakib, M. R. H., Kona, M. A., Afrin, S., ... & Rahman, M. H. (2023). Cardiovascular disease identification using a hybrid CNN-LSTM model with explainable AI. Informatics in Medicine Unlocked, 42, 101370.
[10] Bilal, D. A., Alzahrani, A., Almohammadi, K., Saleem, M., Farooq, M. S., & Sarwar, R. (2025). Explainable AI-driven intelligent system for precision forecasting in cardiovascular disease. Frontiers in Medicine, 12, 1596335.
[11] Puiu, A., Vizitiu, A., Nita, C., Itu, L., Sharma, P., & Comaniciu, D. (2021). Privacy-preserving and explainable AI for cardiovascular imaging. Studies in Informatics and Control, 30(2), 21-32.
[12] Ismath, F., Turcanu, C., & Sobnath, D. (2025, August). Predicting Cardiovascular Disease with Machine Learning: An Explainable AI Approach. In Proceedings of the AAAI Symposium Series (Vol. 6, No. 1, pp. 235-243).
[13] Kırboğa, K. K., & Küçüksille, E. U. (2023). Identifying cardiovascular disease risk factors in adults with explainable artificial intelligence. Anatolian journal of cardiology, 27(11), 657.
[14] Mall, P. K. (2025). Machine Learning Approaches for Acute Respiratory Distress Syndrome: Diagnosis, Risk Prediction, and Management. SGS-Engineering & Sciences, 1(1).
[15] Pal, M., Parija, S., Panda, G., Dhama, K., & Mohapatra, R. K. (2022). Risk prediction of cardiovascular disease using machine learning classifiers. Open Medicine, 17(1), 1100-1113.
[16] García-Ordás, M. T., Bayón-Gutiérrez, M., Benavides, C., Aveleira-Mata, J., & Benítez-Andrades, J. A. (2023). Heart disease risk prediction using deep learning techniques with feature augmentation. Multimedia Tools and Applications, 82(20), 31759-31773.
[17] Kecman, V. (2005). Support vector machines–an introduction. In Support vector machines: theory and applications (pp. 1-47). Berlin, Heidelberg: Springer Berlin Heidelberg.