© MAY 2026 | IRE Journals | Volume 9 Issue 11 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I11-1717397

Web-based Cardiovascular Disease Risk Prediction using Machine Learning

ARJUN SINGH[1], ARPITA YADAV[2], ANJALI SINGH YADAV[3], AMIT KUMAR[4]
[1,2,3,4]Department of CSE-AI, IIMT College of Engineering Greater Noida, India

Abstract- This paper presents a web-based cardiovascular disease (CVD) risk prediction system that combines supervised machine learning with an accessible user interface. An XGBoost classifier is trained on a cleaned Kaggle CVD dataset using eleven low-cost clinical and lifestyle features. The model achieves competitive performance (accuracy ~0.73, ROC-AUC ~ 0.80) and is deployed via a Streamlit application that provides probability-based risk categories for preliminary self-screening, illustrating an end-to- end pipeline from data preprocessing to cloud deployment.

I. INTRODUCTION

Cardiovascular disease (CVD) is one of the major reasons of death worldwide, taking 17.9 million lives approximately every year. It is essential to identify individuals at high risk so that they can modify their lifestyle. Over the past decades, there are many risk-prediction equations and clinical scoring systems have been developed that help detecting cardiology, including Framingham Risk Score, SCORE, QRISK etc. These models suggest a small set of established risk factors such as age, gender, blood pressure, cholesterol, glucose and smoking into a single numerical probability that guides treatment thresholds. Despite their effects, various limitations remain regarding the diverse population, reliance on clinical visits and laboratory tests, clinical experts requirement which are not always intuitive for patients. During the meantime, machine learning (ML) techniques have been developed as powerful tool for handling large, heterogeneous clinical datasets and discovering complex patterns associated with disease risk. These tools aided disease diagnosis faster and accurately, just by using some fixed algorithms. Supervised learning algorithms, including tree-based ensembles, support vector machines (SVM) and gradient boosting, have been demonstrated improved scores and accuracy as compared with conventional regression-based scores in researches. However, mostly models are trained and evaluated offline and not integrated into real-time tools that patients or normal user can access easily. Additionally, high-performing approaches rely on large number of features, including electrocardiography (ECG imaging) and specialized laboratory tests, which may not be easily available to the patients because of costs, time. Therefore there is a growing need for CVD-prediction systems that combine the strengths of modern technology like machine learning with the practical requirements of preventive care: use of low-cost, easily collectible features, robust performance, and user-friendly interface. This CVD predictor addresses these needs by designing, training and deploying a web-based cardiovascular disease risk predictor built around a XGBoost classifier. The model is trained on well-cleaned and processed Kaggle Cardiovascular Disease dataset with appropriate feature selection, using only low-cost predictors that can be obtained from normal clinical tests or can be self-reported as well. XGBoost is chosen for its ability to model non-linear relationships, handle mixed-type features and provide stable performance on tabular clinical data. The trained model achieves strong performance indicating that it can reliably distinguish between individuals with or without cardiovascular disease on the persisted test set. These modelling results are integrated into a Streamlit-based web application which accepts user inputs for age, gender, height, weight, blood pressure, cholesterol, glucose, smoking status, alcohol intake and physical activity. Based on these inputs, the application returns both a probability of cardiovascular disease and an interpretable categorical message highlighting if the user falls into a high-risk or low-risk category. The risk probability is shown using a progress bar and alert message in colour code, making the tool suitable for patient self-screening. The deployment of application on a public cloud platform makes it easier to access for users. This complete work shows an end-to-end pipeline from dataset and model development to a globally accessible risk-prediction service, illustrating how ML-based CVD models can move beyond offline experiments toward real-world impact.

II. RELATED WORK

There have been many research works focused on estimating cardiovascular disease (CVD) risk using traditional statistical models. One of the earlier heart disease studies was the Framingham Heart Study. It gave multivariate functions like age, gender, blood pressure, lipids, smoking and diabetes to estimate 10-year risk of coronary diseases. The limitation of Framingham Risk score is that it may overestimate or underestimate risk in ethnic and geographic groups. Several other investigations using Australian, European and US datasets have compared Framingham, SCORE, QRISK, and ASCVD equations, generally reporting statistics in the range of 0.70-0.80 and highlighting systematic overestimation of risk. Systematic reviews of these tools demonstrate few models undergo rigorous external validation, and most tools needed clinical expertise as tools are implemented as clinician-facing calculators rather than as patient-oriented digital. Beside this, several works using national or hospital CVD databases found that tree-based ensembles achieved higher ROC-AUC than Framingham-type models. However, mostly ML models rely on a large number of features, including imaging and ECG parameters, and are evaluated in single-centre or research-only environments, limiting their applicability to resource-constrained or community settings. In parallel of model development, interest in mobile and web applications for CVD risk assessment has surged. The reviews of cardiovascular mHealth apps report that many tools implement traditional scores such as Framingham and ASCVD, but relatively very few modern ML algorithms are easy for patient understanding. Several existing applications provide little transparency regarding their model architectures, data sources, and validation methodologies, which raises concerns about their reliability and capacity to generalize across populations. However, previous research shows the impact of various machine learning techniques in cardiovascular disease (CVD) risk prediction, some unresolved defects are also uncovered, including the over dependence on extensive feature sets, deficient external validation, and the dearth of openly accessible, user-friendly deployment platforms. This study communicates these limitations by implementing a deployable and practical solution. A large CVD dataset, collected routinely and low-cos clinical tests requirements is publicly available on Kaggle. This dataset was used to train a XGBoost model which is further deployed as a publicly accessible Streamlit-based web application. This work prioritizes the minimal input requirements, transparent evaluation, and patient-friendly interface, aiming to bridge the gap between machine learning and real-world clinical applicability.

III. DATASET AND PRE-PROCESSING

This work uses the cardiovascular disease dataset which is publicly available through on Kaggle platform. The dataset contains anonymized clinical data which was collected from patients during routine medical check-ups. It consists of 70,000 records and each record is associated with a binary class label that indicates if the cardiovascular disease is present or absent (1/0). There are total of eleven predictive variables, including demographic information (age and gender), anthropometric measurements (height and weight), physiological measurements (systolic and diastolic blood pressure), biochemical indicators (cholestrol level and glucose level), and lifestyle-related factors (smoking status, alcohol consumption, and physical activity). Initially age is recorded in days, gender is recorded in categorical manner, and the remaining medical and daily lifestyle variables are showed using ordinal encoding or binary encoding as specified in the dataset documentation. The dataset available on Kaggle is stored in a semicolon separated format instead csv format. So, first we convert into comma separated value (csv) format during the pre-processing. Then we verify for the data types and attribute names. Further the data cleaning operations like missing values handling and inconsistent entries were performed to ensure the completeness and integrity of the input data. Age values converted into the years which improves interpretability and support downstream analysis. Gender values are replaced by binary classification of Male and Female categories. Cholestrol and Glucose levels are categorised as normal (“1”), above normal (“2”) and well above normal (“3”). These ordinal changes are kept to preserve their relative ranking for the learning algorithm. Daily lifestyle routine variables like smoking, alcohol consumption, and physical activity are recorded in binary classification representing the presence or absence of each behavior. After pre-processing, the dataset is partitioned into two parts as training and testing subsets in ratio of 80%-20%. Hierarchical sampling is applied for maintaining the original class distribution through both subsets. No additional resampling techniques are employed, as the dataset exhibits moderate class imbalance, which ensemble-based classifiers are generally capable of handling effectively. The processed dataset is saved as a separate file to ensure reproducibility and to provide a consistent input structure for both model training and web-based prediction.

IV. MODEL DESIGN AND METHODOLOGY

The primary aim of the proposed system is to develop a supervised machine-learning model capable of predicting the likelihood of cardiovascular disease (CVD) using commonly available clinical and lifestyle information and to deploy this model within an interactive web-based environment. The methodology is therefore organized into two closely related stages: first, the construction, training and evaluation of an Extreme Gradient Boosting (XGBoost) classifier using a pre- processed cardiovascular dataset; and second, the integration of the trained model into a Streamlit- based application that supports real-time inference from user-provided inputs. The prediction problem is formulated as a binary classification task in which each individual record is associated with a target variable indicating the presence (1) or absence (0) of CVD. The input feature set comprises eleven attributes, including age (in years), gender, height, weight, systolic and diastolic blood pressure, cholesterol level, glucose level, smoking status, alcohol consumption and physical activity. These variables are deliberately selected they can be obtained through routine clinical measurements or simple self-reported questionnaires, making the system suitable for preliminary risk screening and online self-assessment in general outpatient or community settings. All records are represented in a structured tabular format after data pre-processing, which includes basic cleaning, handling of implausible values and transformation of categorical variables into numerical encodings compatible with gradient-boosted decision trees. To train and assess the model, the cleaned dataset is randomly divided into training and testing subsets using an 80:20 split, with stratified sampling applied to maintain the original distribution of the target classes in both subsets so that minority and majority cases are proportionally represented. The training portion is used exclusively for model fitting and hyperparameter adjustment, while the held-out test set provides an unbiased estimate of predictive performance. Since the dataset is of moderate size and it does not show so much imbalance, we did not used SMOTE (Synthetic Minority Over-sampling Technique) or any random under-sampling. The robustness is pursued through ensemble boosting. Then before model training, we performed exploratory checks to verify feature ranges, anthropometric values, or to identify debatable blood pressure, and ensure if the dataset is consistent with original dataset documentation and medically trusted constraints.

As the core predictive model, we chose an XGBoost classifier because of its proven effectiveness in structured medical data analysis and its robustness to sort of noisy features. It is also able to model complex non-linear relationships among risk factors. XGBoost builds an ensemble of decision trees in a sequence; each tree attempts to resolve the residual errors of the previous ones using gradient boosting. It also implements regularization terms to control overfitting while maintains high predictive power. Within the suggested model, the application uses the XGBClassifier interface with key hyperparameters including the number of boosting rounds (n_estimators) learning rate, maximum tree depth, subsampling ratios for rows and features, and L1/L2 regularization coefficients. Initially, we prepared a base model with a moderate number of trees and some default parameter values to obtain an acceptable balance between bias and variance. After that hyperparameters tuning done by adjusting n_estimators, learning rate, and max_depth meanwhile monitoring cross validated performance on the training set. This lightweight tuning approach is chosen so that model stability, training efficiency and deployability can be prioritize. The held-out test set is conducted for model evaluation using multiple performance metrics derived from confusion matrix. The confusion matrix records counts of true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN). Overall, the accuracy provides a general measure of correctness and it is computed as:

(1) Accuracy = (TP + TN) / (TP + TN + FP +FN)

To better capture performance on the positive class (of disease), the methodology, also reports as precision, is defined as:

(2) Precision = TP / (TP + FP)

It quantifies the proportion of predicted positives that are truly positive. Recall (sensitivity), is defined as:

(3) Recall = TP / (TP + FN)

Recall measures the fraction of actual CVD cases correctly identified by the model. The F1-score combines these two aspects (Recall and Precision) using the harmonic mean:

(4) F1 = 2 * (Precision * Recall) / (Precision+ Recall)

(5)

F1-score provides an only indicator that rebalances the extreme imbalances between precision and recall. It is particularly enlightening and beneficial if both missed detections and false alarms are clinically relevant. In addition to these threshold-dependent metrics, the area under the receiver operating characteristic curve (ROC-AUC) is computed by plotting true positive rate against false positive rate across a range of decision thresholds and numerically integrating the resulting curve, thereby summarizing the overall discriminative ability of the classifier independent of any fixed cut-off. Confusion matrices and these associated metrics are examined jointly, with special attention to false negatives because misclassifying high-risk individuals as low risk is more consequential than overestimating risk in a screening scenario. The final XGBoost model achieves strong overall performance, with balanced precision and recall and a ROC-AUC in the mid-0.8 range, indicating good separation between CVD and non-CVD cases on this dataset. Once training and evaluation are complete, the finalized XGBoost model is serialized to disk to enable efficient reuse during deployment without the need for retraining at inference time. In the Streamlit application, the same feature encoding, scaling assumptions and column ordering used during training are strictly enforced to avoid training–serving skew. User inputs such as age, cholestrol, blood pressure, glucose and lifestyle factors are collected through interactive form. These inputs are further validated for basic plausibility (numeric ranges and not null values). Then these are converted into appropriate numeric values and assembled into a single row pandas DataFrame. The feature names of DataFrame mus exactly match those used while training the model. Then this DataFrame is passed into the loaded XGBoost classifier to generate a prediction of probability of CVD for the given input features of a patient profile. In the next step, the output probability is mapped with a binary label considering a default decision threshold. It can also be categorized into perceptive risk levels like low, moderate or high risk for simpler interpretation by medically non-expert users. In the end, the predicted probability, classification result and accompanying text are presented using an intuitive and user-friendly interface. It emphasizes clarity and encourages users to understand the tool as a preliminary screening help rather than a diagnostic instrument. This deployment design supports transparency in risk communication, which enables fast experimentation with various different thresholds or vizualization elements. In this way, it creates future opportunities for extensions like multi-level risk classification, integration of additional clinical variables.

V. WEB APPLICATION ARCHITECTURE

The deployment of the our developed cardiovascular risk prediction system is done using a lightweight web application architecture which integrates a web browser accessible user interface by Streamlit with a Python-based backend. The backend loads the trained ML model, executes prediction logic and manages application state. The frontend enables user interaction and result visualization. This application implements single page web app technology and uses a python framework that tightly couples frontend with backend (server-side callbacks). It helps in rapid development and does not need separate webpage codebases (HTML, CSS, JavaScript). The serialized XGBoost model load into memory when the application server starts and the model remains available for inference throughout the session lifecycle. User experience like changing input values or making a prediction request etc are transmitted to the backend, where first it validates the inputs and then applies the same feature encoding scheme that was used while training. Then it generates both probabilistic and categorized prediction outputs. This tightly integrated system design ensures low-latency inference and consistency in behavior across sessions.

The data flow starts when a user accesses the web application URL, it triggers initialization of session variables and rendering of a well-designed and structured input form. The user interaction form collects patient-related feature values including age, gender, height, weight, blood pressure, glucose and lifestyle attributes like doing workout daily or smoking status etc. After submission, the backend carries out three-layered pipeline: first, it applies basic validation and normalization checks to ensure physiologically plausible values. Second, the inputs are encoded as per the predefined schema (for gender, lifestyle factors and biological categories). And the third, encoded data are transmitted to the model for inference. These inputs are collected into a one-row pandas DataFrame with feature names and ordering to identical to those used while model fitting so that it aligns with the training phase. The UI shows prediction results using both numerical and visual elements. The CVD probability is presented alongside a corresponding categorical risk message, accompanied by a progress indicator and color cues to improve interpretability. It is clarified by informational notes that this system is intended for educational and screening purposes rather than clinical diagnosis.

For deployment, the application repository is connected to a managed cloud platform which builds and deploys the app automatically upon repository updates. A requirements text file tells for software dependencies and ensures reproductibility across environments. Basic operational checks like model load verification and schema consistency are done at startup to detect (if any) configuration issues at the earliest of process. Keeping security and privacy perspective in mind, the system does not take any personal identity data of users such as username or id verification. All the data taken is processed and display the results within the only active session and discarded after the session; no data from previous session is loaded in the latest one. The new session takes in the fresh data only.

VI. EXPERIMENTAL RESULTS AND DISCUSSION

After pre-processing and preaparing, the cardiovascular disease dataset was divided into two subsets as training and testing in 80:20 proportion. It was a stratified split in order to maintain consistency of positive and negative cases across both partitions. Then we trained an XGBoost classifier on the training subset using those 11 commonly available predictors for all people. These variables covers anthropometric measurements, demographic characteristics, blood pressure readings, biochemical indicators and daily life-style factors; with key hyperparameters such as number of estimators, maximum depth and learning rate tuned in such a manner to balance performance and stability. Evaluation is carried out on test set that shows that the final model attains the accuracy of ~0.735 (approximately). It means that about 3 quarters of test instances are classified correctly. Then the precision value for positive class (means positive CVD disease class) is ~0.754, which indicates that roughly 3 out of 4 individuals predicted as having the disease were actually CVD positive. The recall is around ~0.698, which shows that the model detects ~70% of all true CVD cases that are present in the test set. The F1-score is approximately ~0.724 shows a stable trade-off between precision and recall, which is also important in scenario of clinical screening where missed cases and unnecessary alarms, both must be controlled.

In addition to these threshold-dependent metrics, the area under the Receiver Operating Characteristic Curve (ROC-AUC) is approximately 0.801, demonstrating that the classifier has good discriminative ability across a wide range of decision thresholds and can generally rank individuals with CVD higher in risk than those without the condition. Taken together, these results suggest that the XGBoost model provides reliable predictive performance using only low-cost, routinely collected inputs and offers a solid basis for an accessible risk-assessment tool. Inspection of the confusion matrix reveals that most misclassifications are concentrated near the decision boundary, where patients with and without cardiovascular disease exhibit very similar risk profiles. In particular, a subset of individuals with CVD are assigned relatively low predicted probabilities and are therefore labelled as negative at the default threshold, constituting clinically important false negatives, whereas many false positives fall into an intermediate-risk zone where one or more risk factors are elevated but not strongly distinctive. From a practical standpoint, false negatives are more concerning because they may delay further diagnostic work-up or lifestyle counselling, while false positives primarily increase follow-up workload but still keep potential high-risk individuals under closer observation. To mitigate this issue, the deployed web application presents continuous probability scores alongside the binary prediction, allowing users and clinicians to interpret borderline values (for example, probabilities around 0.4–0.6) as signals of uncertainty and potential risk rather than definitive reassurance. When viewed against traditional clinical risk scores whose reported C-statistics often fall between 0.70 and 0.80 in external cohorts, an ROC-AUC of about 0.80 indicates that the proposed model is at least competitive with established tools while using a comparable set of simple predictors. This performance gain is rationally imputable to XGBoost's ability to model non-linear effects and interactions among variables like blood pressure, age and metabolic indicators. These variables are not fully captured by classical linear scoring systems. However, this study has significant limitations which questions the integrity of the application: the experiments are restricted to only one publicly available dataset, so it is required an external validation on independent patient cohorts.; and the current focus is on the predictive accuracy rather than rich interpretability or casual insight. Regardless these constraints, the experimental data shows that a boosted-tree ensemble, which is combined with careful preprocessing and delivers probability based outputs through a web interface, it forms a strong and practical foundation for real-world cardiovascular risk screening. A table showing the results of our model is shown below:

Table I. Results obtained from XGBoost Model testing

Metrics	Values
Accuracy	0.7348571428571429
Precision	0.7535515750463249
Recall	0.6975414522584333
F1-score	0.7244655581947743
ROC-AUC	0.801208904476377

VII. COMPARATIVE ANALYSIS AND NOVELTY

Traditional cardiovascular risk assessment tools, such as the Framingham Risk Score and the ASCVD pooled cohort equations, are well established in clinical practice. These models estimate long-term risk using a limited set of predictors, including age, sex, blood pressure, cholesterol levels, diabetes status, and smoking history, typically within regression-based frameworks. Although widely validated, their reported discriminatory performance generally falls within a C-statistic range of approximately 0.70 to 0.80. Moreover, multiple studies have shown that these scores may miscalibrate when applied to populations that differ from the cohorts on which they were developed. In addition, most conventional calculators are designed primarily for clinical use and present results as static numerical outputs, offering limited support for patient-centric interpretation or interaction. Recent machine-learning research has explored more flexible modelling approaches for cardiovascular risk prediction. Algorithms such as Random Forests, gradient boosting methods, and neural networks have demonstrated improved discrimination compared to linear models when trained on structured clinical data. Several studies report higher ROC-AUC values for Random Forest classifiers relative to logistic regression baselines. However, many of these investigations rely on institution- specific datasets, incorporate a large number of laboratory or imaging variables, and remain confined to offline evaluations. Consequently, their applicability to routine screening or public use is often restricted.

The present work differentiates itself in three key aspects. First, the model is trained exclusively on low-cost, routinely obtainable features, including basic demographics, anthropometric measures, blood pressure readings, categorical laboratory indicators, and simple lifestyle variables. This approach with least possible input design improves the feasibility in primary-care and self-assessment contexts while attaining a ROC-AUC of ~0.80, which comparatively equivalent or slightly exceeds the discrimination oftenly done with traditional risk scores. Second, this study highlights deployment instead of purely methodological approach by integrating the trained model into a web application which is accessible publicly. The proposed system uses fact-based approach using real data unlike some fixed equation based calculator. It can be updated or modified if new evidence is available. Third, person is not required to be medically expert to use this tool for self-assessment. One who is aware with general terminology of the condition can easily use our tool.

Finally, a outstanding feature is in the way how the risk is communicated to the person who is performing a test or patient himself. The application delivers ongoing probability assessments paired with easy-to-understand visual indicators. This approach helps users grasp various levels of risk more effectively than simply using a 'YES/NO' classification system. Overall, these features collectively establish the suggested system as a useful, user-focused addition that enhances current methods for evaluating cardiovascular risk. A comparative linear chart is shown below:

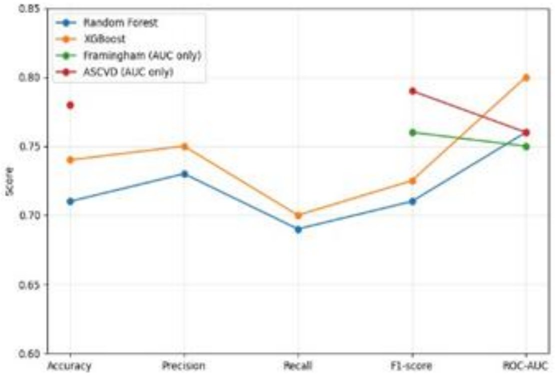

Fig. 1. Comparison of Various Models' Performance

VIII. CONCLUSION AND FUTURE WORK

This study developed an end-to-end cardiovascular disease risk prediction framework that integrates a XGBoost classifier with a publicly available web application. Trained on the Kaggle Cardiovascular Disease dataset using eleven routinely collected demographic, clinical, and lifestyle variables, the model demonstrated reliable predictive capability, achieving strong accuracy, F1-score, and ROC-AUC on unseen test data. A cloud-hosted Streamlit interface offers an accessible, user-friendly tool for a trained model's preliminary screening and cardiovascular risk assessment, complementing clinical judgement in resource-limited settings. Future work will focus on achieving higher accuracy of AI tools by integrating more features like ECG imaging and heartbeat rate, receiving data using IOT sensors and integrating system in a wrist band with display for showing the information. Enhancing interpretability through explainable AI methods, such as SHAP, and evaluating usability through clinician-guided pilot studies will be essential steps toward responsible real-world adoption.

IX. ACKNOWLEDGEMENT

We would like to thank the Department of Computer Science and Engineering (Artificial Intelligence), IIMT College of Engineering, Greater Noida, for providing laboratory facilities, computational resources and continuous support throughout this project. Special gratitude is extended to project guide, Mr. Amit Kumar, for his valuable technical guidance, encouragement and constructive feedback at every stage of model development, web application design and manuscript preparation. We also acknowledge the contributors of the open-source tools and libraries used in this work, including Python, scikit-learn, pandas, NumPy and Streamlit, as well as the Kaggle community for making the Cardiovascular Disease dataset publicly available for research and educational purposes.

REFERENCES

[1] Matheny, M., McPheeters, M. L., Glasser, A., Mercaldo, N., Weaver, R. B., Jerome, R. N., ... & Tsai, C. (2011). Systematic review of cardiovascular disease risk assessment tools.

[2] Matheson, M. B., Kato, Y., Baba, S., Cox, C., Lima, J. A., & Ambale-Venkatesh, B. (2022). Cardiovascular risk prediction using machine learning in a large Japanese cohort. Circulation Reports, 4(12), 595-603.

[3] Ward, A., Sarraju, A., Chung, S., Li, J., Harrington, R., Heidenreich, P., ... & Rodriguez, F. (2020). Machine learning and atherosclerotic cardiovascular disease risk prediction in a multi-ethnic population. NPJ digital medicine, 3(1), 125.

[4] Yang, L., Wu, H., Jin, X., Zheng, P., Hu, S., Xu, X., ... & Yan, J. (2020). Study of cardiovascular disease prediction model based on random

forest in eastern China. Scientific reports, 10(1), 5245.

[5] Dritsas, E., Alexiou, S., & Moustakas, K. (2022). Cardiovascular Disease Risk Prediction with Supervised Machine Learning Techniques. ICT4AWE, 1, 315-321.

[6] Dalal, S., Goel, P., Onyema, E. M., Alharbi, A., Mahmoud, A., Algarni, M. A., & Awal, H. (2023). Application of machine learning for cardiovascular disease risk prediction. Computational Intelligence and Neuroscience, 2023(1), 9418666.

[7] Ambale-Venkatesh, B., Yang, X., Wu, C. O., Liu, K., Hundley, W. G., McClelland, R., ... & Lima, J. A. (2017). Cardiovascular event prediction by machine learning: the multi-ethnic study of atherosclerosis. Circulation research, 121(9), 1092-1101.

[8] Rasheed, S., Kumar, G. K., Rani, D. M., Kantipudi, M. P., & Anila, M. (2024). Heart disease prediction using gridsearchcv and random forest. EAI Endorsed Transactions on Pervasive Health and Technology, 10.

[9] Singh, Y. K., Sinha, N., & Singh, S. K. (2016, November). Heart disease prediction system using random forest. In International Conference on Advances in Computing and Data Sciences (pp. 613- 623). Singapore: Springer Singapore.

[10] Chavez-Ecos, F. A., Chavez-Ecos, R., Vergara Sanchez, C., Chavez-Gutarra, M. A., Agarwala, A., & Camacho-Caballero, K. (2024). Mobile health apps for cardiovascular risk assessment: a systematic review. Frontiers in Cardiovascular Medicine, 11, 1420274.