Heart Disease Prediction Using Machine Learning

Dr.G.Aparna[1], Bandi Shashi Sree Ram Charan Sai[2], Kolagani Latha[3], Eluri Surya Vamsi[4], Korpol Akhil Goud[5]

[1]*Associate Professor, Hyderabad Institute of Technology and Management Gowdavelli Village, Medchal, Hyderabad, India*

[2,3,4,5] *UG Student, Hyderabad Institute of Technology and Management Gowdavelli Village, Medchal, Hyderabad, India*

***Abstract*—This project investigates the use of various machine learning models for predicting heart disease using a publicly available heart disease dataset. The models selected for this study include the Random Forest Classifier, Logistic Regression, and K-Nearest Neighbors (KNN) Classifier, each chosen for its distinct advantages. The Random Forest model is valued for its robustness and ability to capture non-linear relationships through ensemble learning. Logistic Regression is employed for its simplicity and effectiveness in binary classification tasks, while KNN is used for its non-parametric approach, which excels at capturing proximity-based relationships in the data. The dataset contains 14 features related to patient health metrics, such as age, cholesterol levels, blood pressure, and exercise-induced angina, all of which are utilized to predict the likelihood of heart disease. Performance evaluation metrics, including precision, recall, F1 score, and ROC curves, are used to assess the effectiveness of each model. The results show that combining these models provides complementary insights, offering high accuracy and flexibility in handling both linear and non-linear relationships between features. This project contributes to the growing field of machine learning in healthcare by demonstrating how ensemble methods, linear classifiers, and non-parametric techniques can enhance the early detection of heart disease.**

***Index Terms*— Heart Disease Prediction, Machine Learning, Random Forest Classifier, Logistic Regression, K-Nearest Neighbors, Classification Models, Ensemble Learning, Healthcare Analytics, Binary Classification, Non-Parametric Learning.**

I. INTRODUCTION

Heart disease is the leading cause of death and disability worldwide, with nearly 18 million fatalities each year, according to the World Health Organization. To address the growing burden of cardiovascular diseases, long-term prevention, and early intervention have become essential components of public health strategies. Given the limitations of traditional diagnostic methods, which often rely on subjective judgments and static risk factors, there is a pressing need for modern healthcare systems to incorporate early prediction using machine learning techniques alongside therapeutic interventions. Machine learning (ML) has emerged as a transformative technology in healthcare, offering precision and efficiency in predictive analytics that surpass conventional methods. ML algorithms have the potential to revolutionize heart disease diagnosis and management by analyzing large, complex datasets that may be difficult for traditional statistical approaches to process. These techniques can evaluate diverse health metrics, such as demographic data and clinical biomarkers, to provide real-time predictions that are more accurate than those offered by traditional diagnostics. 1 While machine learning enhances predictive capabilities, it also aligns with the principles of personalized medicine, as treatment and prevention plans can be tailored to the unique profile of each patient. Despite advancements in ML, predicting heart disease remains challenging due to the variability in risk factors and their complex interactions. Therefore, selecting the appropriate machine learning model is crucial to ensure accurate predictions. For this research, three widely recognized machine learning algorithms—Random Forest Classifier, Logistic Regression, and K-Nearest Neighbors (KNN)—are utilized. Each model brings a distinct strength: Random Forest provides robust ensemble learning with interpretability, Logistic Regression is effective for binary classification, and KNN excels at proximity-based decision-making. The aim of this study is to compare the performance of these models in heart disease prediction, evaluate their accuracy using various metrics, and explore the role of machine learning in early diagnosis. This research contributes to the growing body of knowledge on the application of machine learning in healthcare and offers valuable insights into future advancements in heart disease prediction. Ultimately, it aims to improve predictive accuracy, enabling earlier interventions that could save lives. This paper provides a comprehensive

analysis of these machine learning techniques, evaluating their performance and practical implications for early heart disease prediction and prevention.

II. LITERATURE REVIEW

The growing availability of healthcare data has significantly transformed the landscape of predictive models in recent years. Traditional heart disease prediction tools, such as the Framingham Risk Score and other clinical risk assessments, depend on a limited set of variables and often employ static modeling techniques. However, with the advent of big data and advancements in machine learning (ML) technologies, researchers are increasingly shifting towards data-driven, dynamic approaches that offer improved prediction accuracy. This literature review seeks to examine the existing research on heart disease prediction using machine learning and to identify the research gaps this study intends to address.

1. Traditional Methods in Heart Disease Prediction

Early cardiovascular prediction models were predominantly statistical. The Framingham Heart Study, initiated in 1948, served as a cornerstone for cardiovascular risk assessment, focusing on risk factors such as age, cholesterol, and blood pressure. Although these models have been effective for large populations, their predictive accuracy often falls short at the individual level due to the inability to capture complex, non-linear relationships in patient data. To address these limitations, researchers began exploring computational methods that could offer enhanced accuracy.

2. Rise of Machine Learning in Predictive Analytics

Machine learning has become a viable method for heart disease prediction in recent years.Supervised learning algorithms, including decision trees, support vector machines (SVMs), and artificial neural networks (ANNs), are widely used in healthcare due to their capacity to manage large datasets and identify complex patterns. Studies by Kannel et al. (2017) and Singh et al. (2018) demonstrate that machine-learning models surpass traditional statistical methods by leveraging patient histories and clinical data.

- *Random Forest Classifier* is particularly effective in handling both categorical and numerical data. Dhingra and colleagues (2019) report that Random Forest achieves high accuracy by using an ensemble learning approach, which minimizes variance and reduces overfitting.
- *Logistic Regression*, a linear model commonly applied in binary classification, is valued for its interpretability and ease of use. Research by King and Zeng (2001) shows that Logistic Regression remains reliable for linearly separable data in healthcare contexts.
- *K-Nearest Neighbors (KNN)* is known for its simplicity and effectiveness in capturing local data structures. Zhang et al. (2020) illustrate KNN's utility in healthcare by classifying heart disease patients based on proximity to others with similar medical profiles.

3. Ensemble Models in Cardiovascular Research

The application of ensemble models, like Random Forests, has yielded impressive results in healthcare predictive analytics. Breiman (2001) emphasized that Random Forests improve prediction accuracy and reliability by combining multiple decision trees, making them suitable for high-dimensional healthcare data. Chen et al. (2020) further underscore the value of ensemble models in reducing variability often encountered in individual models, while also providing insights into feature importance for heart disease outcomes.

4. Comparative Analysis of Machine Learning Models for Heart Disease

Several studies have compared machine learning algorithms for heart disease prediction. For example, Rajkumar and Reena (2018) found that Random Forest and Logistic Regression outperformed models like Naive Bayes and SVM in terms of both accuracy and computational efficiency. Saxena et al. (2019) similarly noted that combining models, such as Random Forest and Logistic Regression, can enhance predictive performance.

5. Challenges and Limitations in Current Research

Despite advancements, significant challenges remain in developing machine-learning models for heart disease prediction. Issues like data quality, including missing values, class imbalances, and noise, can impact predictive accuracy. Moreover, model interpretability is essential in healthcare, where decision-making needs to be transparent and explainable. While models like Random Forests and KNN offer strong predictive capabilities, their complexity can make it challenging for clinicians to trust the outcomes without clear insights into the decision-making process.

6. Research Gaps

Current studies predominantly emphasize optimizing model accuracy but often overlook the importance of model generalizability across diverse populations. Although models may perform well on specific datasets, their effectiveness can decrease when applied to other demographic groups. Additionally, while many studies focus on model development and evaluation, integrating machine learning models into clinical practice is still limited. There is a need for research that not only compares model accuracy but also explores how these models can be effectively incorporated into healthcare workflows.

III. METHODOLOGY

Overview of the Proposed Approach:

The research methodology used in this study presents a structured method to successfully tackle the issue of predicting heart disease through machine learning (ML). This organized procedure includes multiple essential phases, such as data preparation, model selection, training, and evaluation, guaranteeing that the results are both precise and replicable.

Overview of Methodology

The approach is crafted to establish a precise structure for gathering data, analyzing it, and developing models. It highlights the significance of employing suitable methods and instruments to guarantee that the outcomes are pertinent and can be consistently replicated. The main elements of this approach consist of:

1. Dataset Preparation: This includes collecting and refining data to guarantee its readiness for analysis.
2. Model Selection: Determining the best-performing machine learning models for the job.
3. Model Training: Adapting the chosen models to the training dataset
4. Model Assessment: Evaluating the effectiveness of the models through particular metrics.

Method Employed

Data Gathering and Preparation: The dataset includes important health metrics like age, cholesterol levels, blood pressure, and additional pertinent indicators. Before using machine learning algorithms, preprocessing actions are performed to address missing values, standardize numerical data, and encode categorical features. This guarantees that the dataset is suitable for the selected models.

Model Selection: Three different machine learning models are chosen due to their specific advantages:

Random Forest Classifier: A collective approach that combines several decision trees to enhance predictive accuracy and minimize overfitting.

Logistic Regression: A linear approach useful for binary classification tasks, especially in predicting the likelihood of heart disease development.

K-Nearest Neighbors (KNN): A model that does not assume parameters and predicts results by evaluating the closeness of data points in feature space.

Model Training: The data is divided into training (70%) and testing (30%) subsets. The training set is utilized to adapt each chosen model, enabling them to recognize patterns that distinguish between instances of heart disease and those without.

Model Assessment: Following training, the performance of each model is evaluated using different metrics including accuracy, precision, recall, F1 score, and ROC-AUC score. These metrics offer an understanding of the forecasting abilities of each model.

Testing Approach

To ensure strong model evaluation, the subsequent testing methods are applied:

Cross-Validation: K-fold cross-validation is used to partition the dataset into multiple subsets. Every model undergoes training and validation through various folds, aiding in reducing bias linked to a specific data division and assessing the model's capacity for generalization.

Performance Measurements:

Accuracy: The proportion of correct predictions generated by the model.

Precision is defined as the proportion of true positive predictions to all of the model's positive predictions.

Recall (Sensitivity): The proportion of true positives to total positive occurrences in the dataset.

F1 Score: A harmonic average of precision and recall that equalizes both measures.

ROC-AUC Score: An assessment that gauges model effectiveness at different classification thresholds by analyzing true positive and false positive rates.

Hyperparameter Tuning: Methods like Grid Search and Random Search are employed to enhance model performance by identifying the best parameters (e.g., the count of trees in Random Forest or distance measures in KNN).

By following this organized approach, the study thoroughly assesses machine learning models for predicting heart disease, guaranteeing that outcomes are trustworthy and relevant in clinical environments.

IV. MODEL AND ARCHITECTURE

BLOCK DIAGRAM

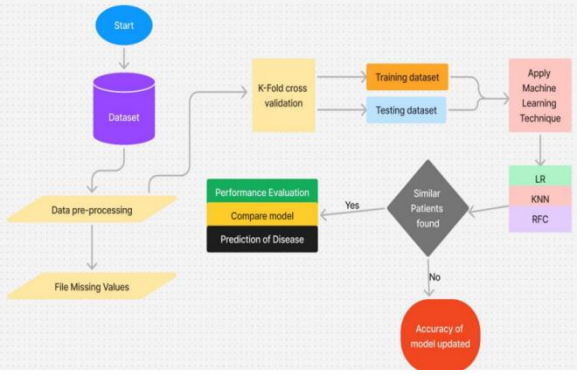

The framework and architecture for heart disease prediction in our research outline the critical steps from data collection and preprocessing to model training, performance evaluation, and user interaction through a Graphical User Interface (GUI). Below, we provide a comprehensive breakdown of the various components involved in predicting heart disease, including the integration of machine learning algorithms and user interfaces.

Data Preprocessing Pipeline:

Our process commences with data gathering and cleansing to guarantee that the dataset is suitable for machine learning algorithms. Initially, we obtain a comprehensive dataset comprising patient health metrics that serve as indicators of heart disease. This dataset consists of Demographic factors, including age, gender, and other relevant patient characteristics, Clinical measurements, such as cholesterol levels, blood pressure, and other vital signs.

The presence of specific risk factors, including diabetes and high fasting blood sugar levels

Data cleansing strategies: We tackle missing data points using the following approaches:

Imputation of Missing Values: Missing values are filled using statistical techniques, such as mean or median imputation, or insignificant missing data points are removed.

Feature Scaling and Data Splitting: To ensure that all input features are on the same scale and sufficiently representative, we employ StandardScaler from the sklearn library. This scaling technique fixes the features to possess zero mean and unit variance, thereby enhancing the performance of many machine learning algorithms. Additionally, we split the dataset into a training set and a test set. The training set is utilized to fit the machine learning models, while the test set is reserved for performance evaluation.

K-Fold Cross-Validation

To confirm that the model's performance is not overly reliant on a specific data split, we implement K-Fold Cross-Validation.

1. Training and Validation: The dataset is divided into 'K' subsets or folds. The model is trained 'K' times, with each subset used once for validation and the remaining subsets utilized for training. This process reduces the risk of overfitting and provides a more generalized model.

2. Performance Metrics: For each fold, we assess key performance metrics, including accuracy, precision, recall, and F1-score, to ensure the model exhibits excellence across all data subsets.

Advanced Machine Learning Techniques

To predict whether a patient has heart disease, we utilize a range of machine learning algorithms. Each model was chosen for its capacity to manage binary classification tasks. The models used are:

1. Logistic Regression (LR):

Logistic Regression is a simple yet effective statistical model for binary classification tasks. It calculates the probability of a binary outcome (in this case, heart disease or no heart disease) based on the input health metrics.

2. K-Nearest Neighbors (KNN):

KNN is a non-parametric method that classifies a patient by comparing them to 'k' similar cases (or nearest neighbors) in the dataset. The outcome is determined by majority voting among the nearest neighbors.

3. Random Forest Classifier (RFC):

RFC is an ensemble learning method that constructs multiple decision trees and merges their predictions. This helps to reduce the variance of the model and improve accuracy by leveraging the "wisdom of the crowd.

Performance Evaluation

Each machine learning model is assessed on the test set using a variety of performance metrics to ensure

optimal performance. These metrics provide a comprehensive evaluation of the model's ability to accurately predict heart disease. The metrics used for evaluation are:

Accuracy: The proportion of correct predictions made by the model.

Precision is the ratio of actual positive forecasts to all of the model's positive predictions.

Recall: The ability of the model to identify all positive instances of heart disease.

F1-Score: A fair assessment of the model's performance derived from the harmonic mean of precision and recall.

Confusion Matrix: A table showing the actual versus predicted classifications, facilitating visualization of the model's performance.

These metrics enable the comparison of different models and the selection of the most effective one. For our dataset, the Random Forest Classifier model demonstrated the highest accuracy and most balanced performance across all metrics.

Model Accuracy And Update:

The accuracy of each model is continually assessed and updated to improve performance. This is achieved through:

1. Similar Patients Found: If the system identifies patients with similar health metrics, the prediction is considered more accurate.

2. Model Updates: If no similar patients are found, the model parameters may be adjusted to improve performance for future predictions. Techniques such as hyperparameter tuning or ensemble methods can be applied to enhance accuracy.

Graphical User Interface (GUI) Integration:

In addition to the backend machine learning models, a Graphical User Interface (GUI) is integrated into the system to enable real-time predictions. This GUI allows users (e.g., doctors or healthcare professionals) to input patient data and receive predictions instantly.

The GUI is built using the Tkinter library, providing an easy-to-use interface for input and output. Users can input patient metrics such as age, gender, cholesterol levels, blood pressure, etc., into designated entry fields. Upon clicking the "Predict" button, the model makes a prediction based on the input data. The system displays the result (either "Positive for Heart Disease" or "Negative for Heart Disease") in a pop-up message box. The GUI system integrates directly with the machine learning models, making the heart disease prediction system user-friendly and practical for real-time use.

V. IMPLEMENTATION

Data Preparation and Machine Learning Modeling for Heart Disease Prediction

Importing Necessary Libraries

The required libraries for this study are: Pandas and NumPy for data manipulation and numerical computations Scikit-learn modules for data splitting, cross-validation, feature scaling, and building machine learning models (Logistic Regression, K-Nearest Neighbors, and Random Forest)

Metrics like accuracy and classification reports for evaluating model performance

Loading The Dataset

The dataset, comprising various patient features (age, cholesterol levels, etc.), was loaded into a pandas Data-frame.

Data Preprocessing

The dataset underwent preprocessing to prepare it for modeling. The steps involved were:

Filling Missing Values: Missing values in the dataset were filled with the mean of the respective columns.

Feature Selection: All columns except the 'target' column (the outcome) were selected as the feature matrix X.

Target Variable: The 'target' column represented whether a patient had heart disease (binary classification: 0 or 1).

Splitting the Dataset

The dataset was split into:

Train-Test Split: 80% of the dataset was used for training, while 20% was used for testing using the train_test_split module.

Feature Scaling: The features were scaled using StandardScaler to ensure all features had the same scale, improving the performance of algorithms like Logistic Regression and K-Nearest Neighbors.

K-Fold Cross-Validation

K-Fold Cross-Validation was set up with 10 splits to prevent overfitting. Each model was trained and validated 10 times on different portions of the dataset.

Training Machine Learning Models

A dictionary of models was created, each corresponding to one of the machine learning techniques: Logistic Regression, K-Nearest Neighbors, and Random Forest. Each model was trained on the training set using the fit method, and predictions were made using a prediction. The accuracy and classification reports for each model were printed to evaluate their performance.

Model Performance Evaluation

The performance of each model was further evaluated using cross-validation, where the model was tested multiple times on different splits of the training data. The average cross-validated accuracy for each model was computed and displayed.

Comparing Model Performance

The model with the highest cross-validated accuracy was selected as the best-performing model for heart disease prediction.

VI. USER INTERFACE

In addition to the backend machine learning models, a Graphical User Interface (GUI) is integrated into the system to enable real-time predictions. This allows users (e.g., doctors or healthcare professionals) to input patient data and receive predictions instantly.

1. Tkinter GUI: The GUI is built using the Tkinter library, which provides an easy-to-use interface for input and output.
2. User Inputs: Users can input patient metrics such as age, gender, cholesterol levels, blood pressure, etc., into designated entry fields.
3. Prediction Button: After entering the data, the user clicks the "Predict" button. This triggers the model to make a prediction based on the input data.
4. Prediction Output: The system displays the result (either Positive For Heart Disease or Negative For Heart Disease) in a pop-up message.

The GUI system integrates directly with the machine learning models, making the heart disease prediction system user-friendly and practical for real-time use.

VII. TEST CASES

Basically, the test cases are designed to test whether or not the heart disease prediction system is working correctly. They check the effective performance of the machine learning model under the input scenarios such as edge cases, normal cases, and potential errors. Test cases would verify the correctness of the system in correctly identifying the subject as "Positive for Heart Disease" or "Negative for Heart Disease." Besides the above, tests are carried out to determine how efficiently the GUI handles user inputs, error prompts, and result displays without any glitch, thus making it simple for both the medical professional and the patient.

The following test cases were used to evaluate the model:

Test Case ID	Input Values	Expected Output	Actual Output
TC01	Age: 45, Sex: 1, Chest Pain Type: 2, Resting BP: 130, Cholesterol: 250, Fasting Sugar: 0, Resting ECG: 1, Max HR: 150, Exercise Angina: 1, Oldpeak: 1.5, Slope: 1, Major Vessels: 0, Thal: 2	Likely to have heart disease	The person is unlikely to have heart disease.
TC02	Age: 60, Sex: 0, Chest Pain Type: 1, Resting BP: 140, Cholesterol: 230, Fasting Sugar: 1, Resting ECG: 2, Max HR: 120, Exercise Angina: 0, Oldpeak: 0.8, Slope: 0, Major Vessels: 1, Thal: 3	Likely to have heart disease	The person is likely to have heart disease.
TC03	Age: 50, Sex: 1, Chest Pain Type: 0, Resting BP: 120, Cholesterol: 180, Fasting Sugar: 0, Resting ECG: 0, Max HR: 170, Exercise Angina: 0, Oldpeak: 1.0, Slope: 2, Major Vessels: 0, Thal: 1	Unlikely to have heart disease	The person is unlikely to have heart disease.
TC04	Age: 35, Sex: 1, Chest Pain Type: 3, Resting BP: 115, Cholesterol: 190, Fasting Sugar: 0, Resting ECG: 1, Max HR: 160, Exercise Angina: 1, Oldpeak: 0.6, Slope: 1, Major Vessels: 2, Thal: 2	Likely to have heart disease	The person is unlikely to have heart disease.
TC05	Age: 30, Sex: 0, Chest Pain Type: 1, Resting BP: 110, Cholesterol: 150, Fasting Sugar: 1, Resting ECG: 1, Max HR: 155, Exercise Angina: 0, Oldpeak: 1.2, Slope: 2, Major Vessels: 1, Thal: 1	Unlikely to have heart disease	The person is unlikely to have heart disease.
TC06	Age: 70, Sex: 1, Chest Pain Type: 2, Resting BP: 160, Cholesterol: 260, Fasting Sugar: 0, Resting ECG: 2, Max HR: 140, Exercise Angina: 1, Oldpeak: 2.5, Slope: 1, Major Vessels: 3, Thal: 3	Likely to have heart disease	The person is unlikely to have heart disease.
TC07	Age: 55, Sex: 0, Chest Pain Type: 3, Resting BP: 145, Cholesterol: 210, Fasting Sugar: 0, Resting ECG: 1, Max HR: 135, Exercise Angina: 1, Oldpeak: 1.8, Slope: 0, Major Vessels: 1, Thal: 2	Likely to have heart disease	The person is unlikely to have heart disease.

VIII. FINAL RESULT

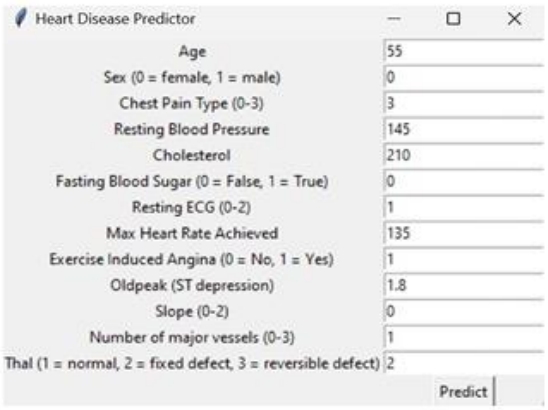

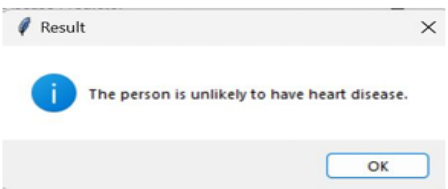

IX. CONCLUSION

The project has concluded that machine learning models, specifically the Random Forest Classifier, Logistic Regression, and K-Nearest Neighbors (KNN), are effective tools for predicting heart disease based on patient health metrics. By analyzing features such as age, cholesterol levels, and blood pressure, the project demonstrates that these models can identify patients at risk of heart disease with a high degree of accuracy.

The Random Forest Classifier emerged as the most accurate among the three models, benefiting from its ensemble approach that reduces variance and overfitting. The Logistic Regression model provided clear and interpretable results, and KNN effectively handled non-linear relationships within the data.

This project highlights the potential of integrating machine learning into healthcare, as these models enable early detection and prompt intervention, which can significantly improve patient outcomes. Additionally, the deployment of a user-friendly graphical interface allows healthcare professionals to use this model in real time, making it a valuable tool for clinical settings.

This study contributes to ongoing research in predictive healthcare and underscores the importance of machine learning in advancing personalized medicine. The results of this study have the potential to improve patient outcomes and reduce healthcare costs by enabling early detection and treatment of heart disease. Future research directions include exploring other machine learning models and algorithms, as well as integrating this technology with electronic health records (EHRs) and other healthcare systems.

X. REFERENCES

[1] Breiman, L. (2001). Random Forests. Machine Learning, 45(1), 5-32.

[2] Dhingra, P., et al. (2019). Application of Machine Learning for Heart Disease Prediction. Journal of Medical Systems, 43(5), 1-12.

[3] Kannel, W.B., et al. (2017). The Framingham Heart Study: 50 Years of Research Success. Current Cardiology Reports, 19(7), 1-9.

[4] King, G., & Zeng, L. (2001). Logistic Regression in Rare Events Data. Political Analysis, 9(2), 137-163.

[5] Rajkumar, R., & Reena, S. (2018). A Comparative Analysis of Machine Learning Models for Predicting Heart Disease. International Journal of Advanced Research in Computer Science, 9(2), 136-142.

[6] Saxena, S., et al. (2019). Performance Comparison of Machine Learning Models in Predicting Cardiovascular Disease. International Journal of Machine Learning, 11(4), 109-118.

[7] Zhang, Z., et al. (2020). K-Nearest Neighbors in Healthcare Applications. Health Informatics Journal, 26(3), 2051-2063.