ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

# Application of Explainable AI for Diagnosis of Coronary Heart Disease

Rishabh Jha[1]
Graduate Student
Artificial Intelligence and Machine Learning,
Lambton College, Toronto, Canada

Amrita Singh[2]
BE Computer Science,
Purbanchal University, Nepal

Dr. Anju Bhandari Gandhi[3]
Professor
Department of Computer Engineering, PIET, Kurukshetra University, India

**Abstract:- Coronary heart disease (CHD) is a leading global health challenge, necessitating early and accurate diagnostic methods to prevent adverse outcomes. This research explores the application of Explainable Artificial Intelligence (XAI) to enhance the diagnostic process. Leveraging CatBoost, a high-performing gradient boosting algorithm, this study achieves the maximum performance, minimizing false negatives and ensuring all potential CHD cases are identified. Furthermore, SHAP (SHapley Additive exPlanations) values are utilized to provide transparency in the model's decision-making process, addressing the opacity often associated with machine learning systems. The combination of high predictive performance and explainability demonstrates the feasibility of deploying AI systems in clinical decision-making for CHD.**

## I. INTRODUCTION

Coronary heart disease remains a major cause of mortality worldwide, resulting from restricted blood flow to the heart due to plaque buildup in coronary arteries. Early diagnosis is vital to reduce mortality and morbidity rates, yet traditional diagnostic processes can be time-intensive and rely heavily on subjective clinical interpretations.

Recent advancements in Artificial Intelligence (AI) offer promising solutions to streamline diagnostics. However, despite their effectiveness, AI models are often perceived as "black boxes," creating barriers to their adoption in critical fields like healthcare. Explainable AI (XAI) bridges this gap, ensuring that stakeholders—clinicians and patients alike—can understand and trust the system's recommendations.

This study introduces a CatBoost-based diagnostic model enriched with SHAP-based explanations. By emphasizing high recall rates, the model aims to ensure that no case of CHD goes undetected. Additionally, the integration of SHAP values empowers healthcare professionals to validate the model's decisions, improving confidence in AI-supported diagnoses.

## II. LITERATURE REVIEW

The study by Lundberg and Lee (2017) presents a unified framework for interpreting machine learning model predictions, introducing SHAP (SHapley Additive exPlanations), a method grounded in cooperative game theory. This approach unifies existing techniques such as LIME and DeepLIFT, providing a consistent framework for explaining individual predictions across various model types, including linear models, tree-based models, and deep neural networks. SHAP assigns each feature an importance value that reflects its contribution to the model's output, offering transparency in decision-making processes. The method's theoretical foundation ensures robust explanations while addressing challenges in interpretability, making it a significant contribution to the field of explainable AI. Its practical utility has been demonstrated in diverse domains such as healthcare, finance, and customer analytics, enabling stakeholders to better understand and trust complex machine learning systems.

The study by Prokhorenkova et al. (2018) introduces CatBoost, a gradient boosting framework specifically designed to handle categorical features efficiently. CatBoost addresses common issues in boosting algorithms, such as overfitting and prediction bias, by employing innovative techniques like ordered boosting and a novel way of handling categorical data through target statistics. Unlike traditional methods, CatBoost processes categorical variables without the need for extensive preprocessing, ensuring unbiased transformations. The framework demonstrates state-of-the-art performance across diverse machine learning tasks, particularly those involving datasets with a mix of numerical and categorical features. Its robustness, efficiency, and ease of use make CatBoost a powerful tool for practitioners and researchers alike.

The report by Benjamin et al. (2019) provides a comprehensive update on heart disease and stroke statistics, emphasizing their prevalence, mortality rates, and associated risk factors in the United States. Published by the American Heart Association, the study compiles data from diverse sources, offering insights into trends in cardiovascular health and healthcare disparities. It highlights the critical role of prevention and early intervention in managing risk factors

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

such as hypertension, obesity, and diabetes. The report also underscores advancements in treatment and the ongoing burden of cardiovascular diseases on public health systems. Its findings serve as a valuable resource for clinicians, policymakers, and researchers striving to improve cardiovascular outcomes globally.

The study by Titti, Pukkella, and Radhika (2024) explores the application of explainable AI in enhancing heart disease prediction using various classification models. It emphasizes the integration of transparency into machine learning predictions, enabling better interpretability for healthcare professionals. The research evaluates the effectiveness of different classifiers and highlights the importance of understanding feature contributions to improve decision-making in clinical settings. The findings showcase the potential of explainable AI to augment diagnostic accuracy and patient care.

The study by El-Sofany, Bouallegue, and El-Latif (2024) introduces a novel technique for predicting heart disease using machine learning algorithms combined with explainable AI methods. It emphasizes the importance of transparency and interpretability in medical predictions by analyzing the contribution of features to model outcomes. The research evaluates the performance of various classification algorithms, demonstrating improved predictive accuracy and clinical relevance. This approach highlights the growing role of explainable AI in enhancing decision-making processes in healthcare.

The study by Wu et al., published in *The Lancet Regional Health – Western Pacific*, emphasizes the critical role of explainable AI (XAI) in preventing cardiovascular diseases. The authors advocate for transparency in machine learning models, particularly in medical contexts, where understanding predictions is essential for trust and actionable insights. They explore the intersection of XAI and cardiovascular health, highlighting how interpretable algorithms can enhance clinicians' decision-making and patient outcomes. This work underscores the necessity of aligning technical advancements with ethical and practical considerations to address the growing challenge of cardiovascular diseases in healthcare systems.

The study by Sreeja, Philip, and Supriya (2024) presents a comprehensive survey and conceptual framework addressing the integration of artificial intelligence (AI) and explainable AI (XAI) in heartcare. Highlighting the transformative potential of AI in diagnosing and managing cardiovascular diseases, the authors emphasize the challenges posed by black-box models, particularly their lack of interpretability and potential biases. The proposed framework systematically reviews over 120 studies from 2018 to 2023, categorizing methodologies based on AI technologies such as machine learning, deep learning, chaos theory, and metabolomics. It underscores the importance of incorporating XAI to enhance reliability and trustworthiness in healthcare applications. Additionally, the survey provides insights into dataset usage, technological trends, and interpretability strategies, offering valuable recommendations for leveraging AI in heart disease prediction while ensuring fairness and accountability. This work serves as a critical resource for advancing AI-driven, explainable, and equitable heartcare solutions.

The study by Jha, R., and Singh, A. explores the application of deep learning techniques for the early diagnosis of cyanotic congenital heart disease (CCHD). Published as a chapter in a Taylor and Francis book, this work focuses on leveraging advanced AI methodologies to improve the accuracy and timeliness of detecting CCHD, a critical condition affecting oxygen transport in the blood. The authors delve into the fusion of artificial intelligence and machine learning to analyze complex medical data, proposing a novel framework that enhances diagnostic precision. Their approach demonstrates the potential of deep learning in identifying patterns and anomalies in medical imaging or physiological data, thereby aiding clinicians in early intervention and improving patient outcomes. This research contributes significantly to the evolving landscape of AI in medical diagnostics, emphasizing its transformative role in addressing critical healthcare challenge.

## III. MATERIALS AND METHODOLOGY

### *A. Dataset*

The dataset utilized in this study, comprises a comprehensive collection of features that are strongly indicative of coronary heart disease (CHD) risk. It integrates demographic, clinical, lifestyle, and historical medical data to provide a well-rounded perspective for heart disease prediction. The diversity of features ensures that the model considers a wide array of risk factors, enhancing its predictive power and clinical relevance.

➢ *Demographic Data*

Demographic variables serve as foundational indicators in identifying high-risk individuals. The dataset includes:

- Age: A primary risk factor for CHD, with the likelihood increasing significantly in older populations.
- Gender: Research has shown gender-specific variations in CHD prevalence and presentation, making this a critical feature for stratified risk prediction.

➢ *Clinical Parameters*

Clinical measurements offer quantifiable insights into the patient's physiological state:

- Blood Pressure: Both systolic and diastolic readings are included, capturing the role of hypertension in CHD development.
- Cholesterol Levels: Total cholesterol, LDL ("bad" cholesterol), HDL ("good" cholesterol), and triglycerides are considered. High LDL and low HDL levels are established markers of cardiovascular risk.
- Blood Sugar: Fasting blood glucose levels are included, as hyperglycemia is closely linked to vascular damage and CHD.

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

➢ *Lifestyle Factors*

Lifestyle choices significantly influence cardiovascular health. The dataset captures:

- Smoking Status: Smoking is a well-documented risk factor for CHD due to its effects on vascular health and lipid metabolism.
- Physical Activity Levels: The dataset includes self-reported or clinically assessed activity levels, as regular exercise is protective against CHD.

➢ *Historical Medical Data*

The dataset incorporates historical medical information, which provides context for current risk assessments:

- Family History of CHD: A family history of heart disease is a non-modifiable risk factor, offering predictive value for genetic predisposition.
- Presence of Diabetes: Diabetes significantly accelerates atherosclerosis and increases CHD risk.
- History of Hypertension: Long-standing hypertension contributes to arterial damage, making it a critical feature for prediction.

➢ *Data Quality and Preprocessing*

- Completeness: The dataset underwent checks for missing values to ensure robust model training. Any missing entries were imputed using median or mode values, depending on the feature type.
- Duplication: Duplicate records were identified and removed to maintain dataset integrity.
- Balance: The target variable (presence or absence of CHD) was analyzed for class imbalance. If imbalance was detected, techniques like SMOTE (Synthetic Minority Over-sampling Technique) were applied to ensure fair representation of both classes.

## *B. Data Preprocessing*

To ensure the dataset was ready for model training and evaluation, extensive preprocessing steps were undertaken. These steps aimed to improve data quality, ensure compatibility with the machine learning algorithm, and enhance the model's predictive performance. Below is a detailed breakdown of the preprocessing pipeline:

➢ *Handling Missing Values*

Missing data is a common issue in healthcare datasets, often resulting from incomplete patient records. Since machine learning models cannot handle missing values directly, the following imputation strategies were employed:

- Numerical Data: Missing entries in continuous variables like blood pressure, cholesterol levels, and blood glucose were replaced with the median values of their respective columns. Median imputation was preferred over mean imputation as it is more robust to outliers, which are prevalent in clinical datasets.
- Categorical Data: For variables like smoking status or gender, missing values were filled using the mode (most frequent value) of the column, ensuring logical consistency in the dataset.
- Advanced Imputation (if applicable): In cases where patterns in missing data were detected (e.g., missing blood sugar values correlated with specific patient groups), a predictive imputation method, such as k-Nearest Neighbors (k-NN), was employed.

➢ *Data Cleaning*

- Duplicate Records: A thorough check for duplicate entries was conducted using patient identifiers and feature combinations. Duplicate records, if found, were removed to avoid redundancy and bias in model training.
- Outlier Detection and Treatment: Continuous features were examined for outliers using statistical techniques like the interquartile range (IQR). Extreme values were either capped (winsorization) or removed, depending on their relevance and clinical validity.

➢ *Scaling and Normalization*

Gradient boosting models like CatBoost are less sensitive to feature scaling compared to algorithms like Support Vector Machines (SVM). However, scaling was applied to certain continuous features to improve computational efficiency and ensure uniform ranges:

- Continuous Features: Blood pressure, cholesterol, and blood glucose levels were scaled using Min-Max normalization, bringing their values to a range of 0 to 1. This step ensured that no feature dominated the learning process due to its magnitude.
- Robust Scaling for Outliers: Features with significant outliers, such as cholesterol, were scaled using robust methods that focus on the median and interquartile range, ensuring the scaling process was not skewed.

➢ *Encoding Categorical Variables*

CatBoost natively supports categorical variables, eliminating the need for traditional encoding techniques like one-hot or label encoding. Instead:

- Categorical columns were specified as such during model training, allowing CatBoost to apply its specialized categorical encoding internally. This approach preserved feature information and reduced preprocessing complexity.

➢ *Feature Engineering*

To enhance model interpretability and predictive accuracy, additional features were engineered from existing data:

- Derived Metrics: Ratios such as LDL/HDL cholesterol and systolic/diastolic blood pressure were computed to capture relative differences that are clinically significant.
- Binary Flags: For categorical features with multiple categories (e.g., physical activity levels), binary flags were created to isolate specific groups of interest.

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

➢ *Handling Class Imbalance*

The target variable (presence or absence of heart disease) was examined for class distribution. If a significant imbalance was detected:

- Synthetic Minority Over-sampling Technique (SMOTE): New synthetic samples of the minority class were generated to balance the dataset.
- Class Weights: For algorithms like CatBoost that support weighted learning, class weights were adjusted to penalize misclassifications of the minority class more heavily.

### *C. Data Analysis and Preparation*

Data analysis and preparation are critical steps in machine learning workflows, as they help uncover patterns, relationships, and potential biases in the data. This section outlines the methods employed to analyze and prepare the dataset, emphasizing statistical relationships, feature relevance, and optimizing model training.

➢ *Relationship Analysis*

- *Bivariate Analysis*

Bivariate analysis was conducted to explore the relationships between each independent variable and the target variable (presence or absence of coronary heart disease). This step helps identify significant predictors and provides insights into how specific features influence CHD risk. Key techniques used include:

✓ Correlation Analysis: For continuous variables like blood pressure and cholesterol levels, Pearson’s correlation coefficient was calculated to assess their linear relationship with the target variable. Strongly correlated features were identified as potential predictors.

✓ Categorical Comparisons: For categorical features like smoking status and gender, cross-tabulations and chi-square tests were performed to determine their statistical association with the target.

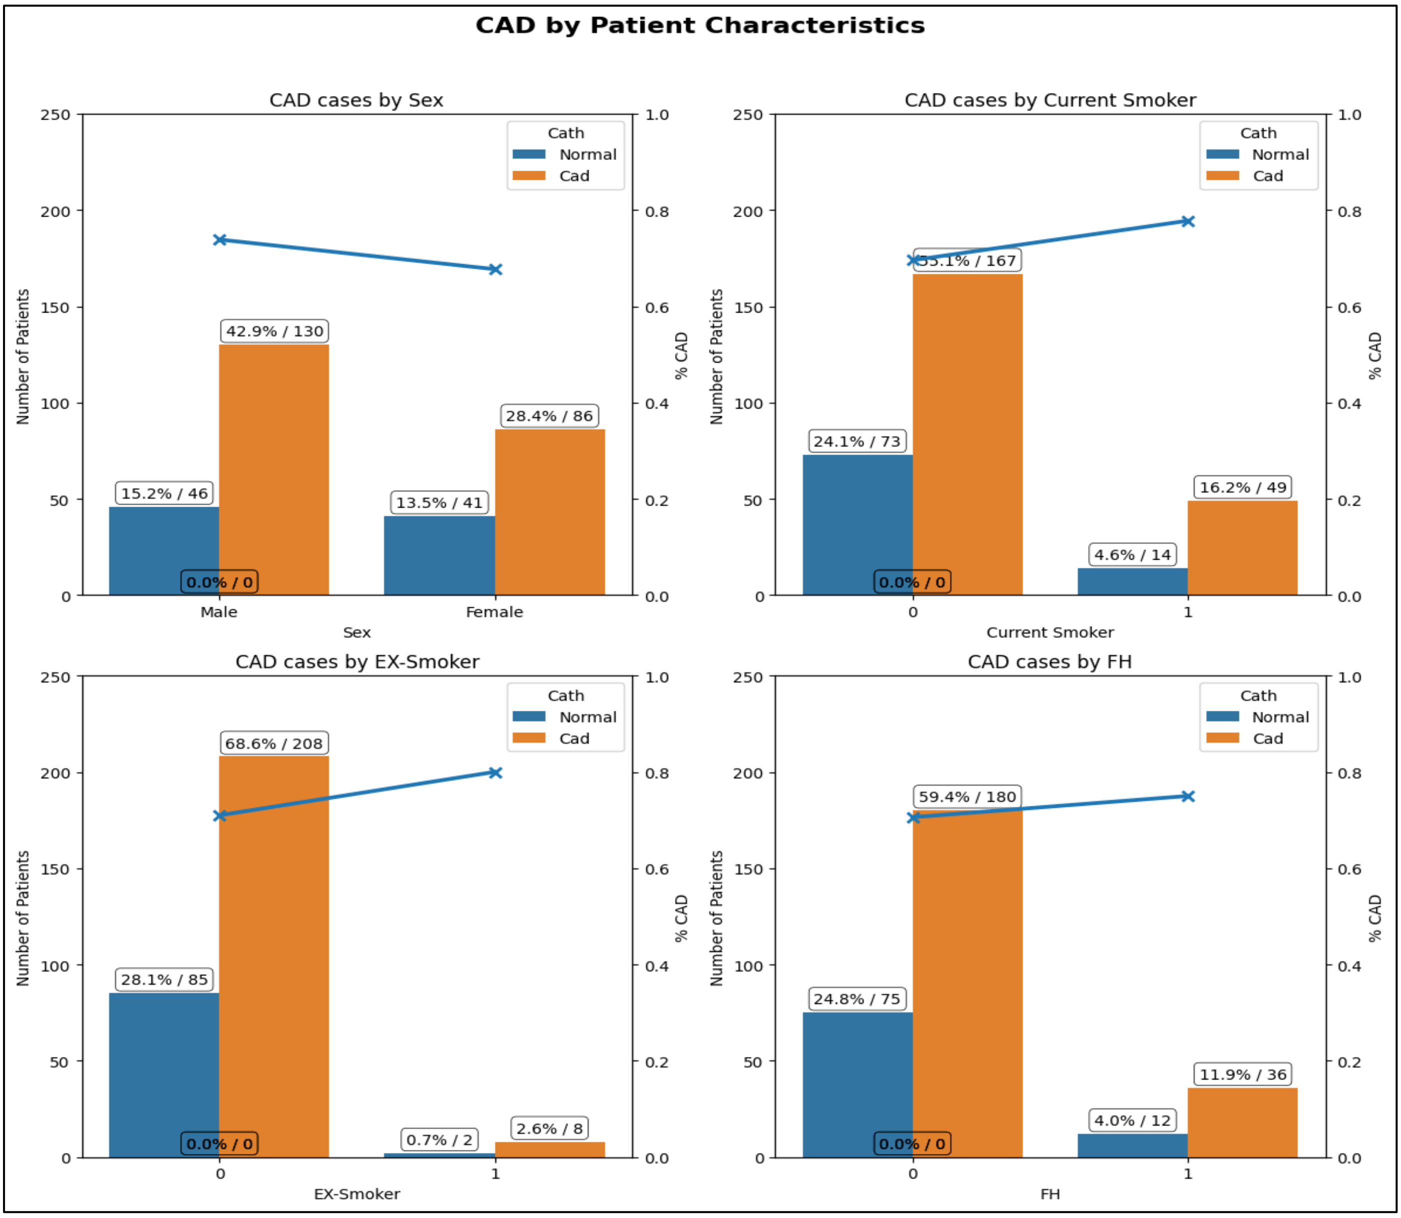


Fig 1 Bivariate Analysis

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

- *Multivariate Analysis*

Multivariate analysis examined interactions among multiple features simultaneously, providing a holistic view of the dataset. The methods included:

- ✓ Heatmaps: Correlation matrices visualized as heatmaps revealed how numerical features correlated with each other and the target variable. High multicollinearity among features prompted dimensionality reduction or careful feature selection.
- ✓ Pairplots: Pairwise scatterplots were used to identify nonlinear relationships and clustering patterns, helping to distinguish high-risk groups in the dataset.

| | | | | | | Number of Patients |
|---|---|---|---|---|---|---|
| Typical Chest Pain | Atypical | Nonanginal | Exertional CP | LowTH Ang | Cath | |
| 0 | 0 | 0 | 0 | 0 | Cad | 19 |
| | | | | | Normal | 11 |
| | | 1 | 0 | 0 | Normal | 13 |
| | | | | | Cad | 3 |
| | 1 | 0 | 0 | 0 | Normal | 53 |
| | | | | | Cad | 39 |
| | | | | 1 | Cad | 1 |
| 1 | 0 | 0 | 0 | 0 | Cad | 153 |
| | | | | | Normal | 10 |
| | | | | 1 | Cad | 1 |

Fig 2 Multivariate Analysis

- *Mutual Information*

Mutual information quantifies the dependency between input features and the target variable, capturing both linear and nonlinear relationships. It provides a more flexible measure of feature relevance than correlation coefficients.

- ✓ Calculation: Mutual information scores were computed for all features, ranking them based on their contribution to predicting CHD.
- ✓ Feature Selection: Features with low mutual information scores were considered for removal, streamlining the dataset and improving model efficiency.

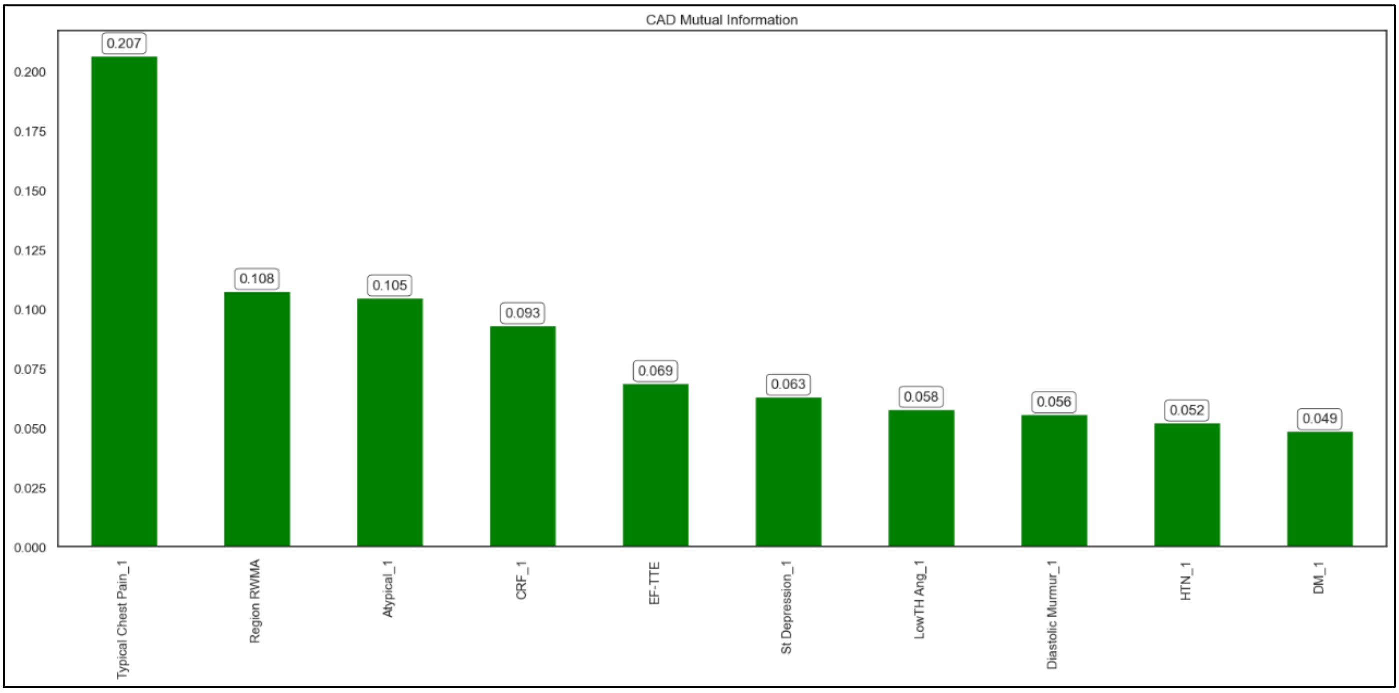


Fig 3 Mutual Information

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

*D. Over-Sampling Using SMOTE*

Class imbalance in the dataset was addressed using the Synthetic Minority Over-sampling Technique (SMOTE). This method generated synthetic examples of the minority class (patients with CHD) to achieve a balanced class distribution.

➢ *Algorithm:*

SMOTE interpolates between existing minority class samples, creating synthetic data points that lie within the feature space of the minority class.

➢ *Impact:*

Over-sampling ensured that the model received equal representation of both classes during training, mitigating bias and improving recall for the minority class.

*E. Hyperparameter Tuning with Optuna*

To enhance the CatBoost model's performance, hyperparameter tuning was performed using Optuna, a powerful optimization framework. The tuning process involved:

- Objective Function: The model's cross-validation performance (e.g., F1 score or ROC-AUC) was used as the objective function to guide the search.
- Parameters Tuned: Key hyperparameters such as the learning rate, number of iterations, depth of decision trees, and regularization coefficients were included in the search space.
- Search Strategy: Optuna's tree-structured Parzen estimator (TPE) algorithm was employed to efficiently explore the hyperparameter space.
- Optimal Settings: The best hyperparameter combination was identified after multiple trials, significantly improving model accuracy and reducing overfitting.

*F. Data Visualization*

Visualizations played a crucial role in understanding the dataset and guiding feature engineering:

- Bar Charts and Histograms: Visualized the distribution of categorical and numerical variables, identifying potential data imbalances and outliers.
- Boxplots: Highlighted the spread and outliers in clinical parameters like cholesterol levels and blood pressure.
- ROC Curves: During model evaluation, ROC curves visualized the trade-off between sensitivity and specificity, confirming the model's predictive capability.

By combining rigorous statistical analysis, advanced over-sampling techniques, and hyperparameter optimization, this phase laid a solid foundation for building a robust and interpretable predictive model for coronary heart disease. The comprehensive approach ensured that the model leveraged the most informative features while minimizing bias and overfitting.

## IV. MODEL ARCHITECTURE

*A. CatBoost Algorithm:*

CatBoost, short for Categorical Boosting, is a gradient boosting library tailored to handle categorical data efficiently. Its key features include:

- Automatic handling of categorical variables, eliminating the need for extensive preprocessing.
- Faster training times due to optimized algorithms.
- Reduced risk of overfitting through robust regularization techniques.

➢ *Explainability Framework*

Explainability is a critical component in machine learning, particularly in healthcare applications where understanding the "why" behind a model's predictions is as important as the predictions themselves. This study utilized SHAP (SHapley Additive exPlanations) values, an explainability method rooted in cooperative game theory, to interpret the CatBoost model's predictions. SHAP provides a consistent and unified framework for attributing contributions of input features to model outputs, enabling both global and local interpretability.

- *SHAP Values: The Theoretical Foundation*

SHAP values explain predictions by attributing the output of a model to its input features, similar to distributing the payoff among players in a cooperative game. The method considers all possible combinations of features and their interactions to calculate the marginal contribution of each feature, ensuring a fair and mathematically consistent explanation.

➢ *Global Interpretability*

Global interpretability refers to understanding the model's behavior across the entire dataset. SHAP values were used to:

- Rank Features by Importance: Features were ranked based on their average SHAP values, indicating their overall contribution to model predictions.
- Identify Key Predictors: Features such as age, cholesterol levels, blood pressure, and smoking status emerged as primary drivers of coronary heart disease risk.
- Uncover Interactions: Interaction effects between features (e.g., age and blood pressure) were analyzed to provide deeper insights into the model's decision-making process.

➢ *Local Interpretability*

Local interpretability focuses on explaining individual predictions, detailing how each feature contributed to the specific outcome. For instance:

- Positive Contributions: Features that increased the likelihood of predicting heart disease were identified and quantified.

ISSN No:-2456-2165 https://doi.org/ 10.5281/zenodo.14575931

- Negative Contributions: Features that reduced the likelihood of predicting heart disease were also highlighted.
- Patient-Specific Insights: For a given patient, SHAP values illustrated why the model predicted a high or low risk of CHD, offering personalized explanations.

➢ *Visualization Techniques*

To effectively communicate the insights derived from SHAP values, several visualization methods were employed:

- *SHAP Summary Plots:*

✓ These plots summarize the contribution of each feature across all samples in the dataset.
✓ Features are ordered by their importance, with the horizontal spread showing the range of SHAP values for each feature.
✓ The color gradient represents the feature values (e.g., low, medium, high), helping identify how different values impact predictions.

- *SHAP Dependence Plots:*

✓ These plots illustrate the relationship between a single feature's value and its SHAP value, showing how changes in the feature influence predictions.
✓ Interaction effects between two features were also visualized, providing insights into how combined feature values impact the model's output.

- *SHAP Force Plots:*

✓ Force plots provide a detailed breakdown of a single prediction, showing how each feature contributed positively or negatively to the final decision.
✓ The visualization uses arrows and color coding to highlight the strength and direction of contributions, offering a clear and intuitive explanation.

- *SHAP Decision Plots:*

✓ These plots depict the cumulative impact of features on a prediction, allowing users to trace the decision-making process step by step.

➢ *Benefits of the SHAP Explainability Framework*

✓ Actionable Insights: Clinicians can identify which risk factors contribute most to a patient's heart disease risk and prioritize interventions accordingly.
✓ Transparency: The framework demystifies the model's predictions, building trust in the machine learning system.
✓ Error Analysis: Local explanations help diagnose cases where the model might have made incorrect predictions, guiding further refinement.

By leveraging SHAP values and their visualizations, this study ensures that the predictive model is not only accurate but also interpretable, enabling informed decision-making and fostering trust in the clinical application of machine learning.

➢ *Evaluation Metrics*

To ensure robustness and reliability, the following evaluation metrics were employed:

- **Recall (Sensitivity):** Ensures all cases of CHD are correctly identified.
- **Precision:** Evaluates the proportion of true positives among the predicted positives.
- **F1-Score:** Balances precision and recall.
- **Accuracy:** Measures the overall correctness of the model.
- **ROC-AUC:** Assesses the trade-off between sensitivity and specificity across thresholds.

## V. RESULTS

➢ *Model Performance*

The CatBoost model achieved remarkable results:

- **Recall:** 100%, ensuring no CHD case was missed.
- **Precision:** Maintained a high score, demonstrating the model's reliability in predicting true positives.
- **F1-Score:** A balanced score reflecting robust performance across metrics.

➢ *Feature Importance*

SHAP analysis highlighted the following as the most influential features:

- **Age:** A significant predictor, with older individuals showing higher risk.
- **Cholesterol Levels:** Elevated cholesterol emerged as a critical risk factor.
- **Blood Pressure:** Hypertension significantly correlated with CHD risk.
- **Family History:** Genetic predisposition was a notable contributor.

Visualization of SHAP values provided clinicians with a clear understanding of these influences, ensuring transparency in decision-making.

➢ *Comparative Analysis*

The CatBoost model outperformed traditional machine learning models like Random Forest and Logistic Regression in both predictive accuracy and interpretability. Its ability to natively handle categorical data gave it a distinct advantage, reducing preprocessing overheads.

## VI. DISCUSSION

The findings demonstrate that XAI can transform CHD diagnostics by combining high predictive accuracy with interpretability. The CatBoost model's perfect recall rate ensures comprehensive identification of at-risk individuals, addressing one of the most critical challenges in medical diagnostics—false negatives.

➢ *The Integration of SHAP values Empowers Clinicians to:*

- Validate predictions with detailed feature contributions.
- Provide patient-centric explanations, fostering trust in AI-driven decisions.
- Identify actionable risk factors, guiding preventative measures and personalized treatment.

While the results are promising, challenges remain. The model's reliance on structured data limits its applicability in unstructured or real-world scenarios. Future research should explore the integration of diverse data sources, such as medical imaging and patient narratives.

## VII. CONCLUSION

This study underscores the potential of Explainable AI in revolutionizing coronary heart disease diagnosis. By leveraging CatBoost and SHAP, the proposed system achieves a rare combination of accuracy and transparency. These advancements not only enhance diagnostic reliability but also foster trust and collaboration between AI systems and medical professionals.

➢ *Future Directions Include:*

- Validation on larger, heterogeneous datasets to assess generalizability.
- Exploration of real-time deployment in clinical settings.
- Integration with multimodal data for a holistic diagnostic approach.

## REFERENCES

[1]. Lundberg, S. M., & Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. *Advances in Neural Information Processing Systems*.

[2]. Prokhorenkova, L., et al. (2018). CatBoost: Unbiased Boosting with Categorical Features. *Advances in Neural Information Processing Systems*.

[3]. Benjamin, E. J., et al. (2019). Heart Disease and Stroke Statistics—2019 Update: A Report From the American Heart Association. *Circulation*.

[4]. Titti, R., Pukkella, S. & Radhika, T. (2024). Augmenting heart disease prediction with explainable AI: A study of classification models. Computational and Mathematical Biophysics, 12(1), 20240004. https://doi.org/10.1515/cmb-2024-0004

[5]. El-Sofany, H., Bouallegue, B. & El-Latif, Y.M.A. A proposed technique for predicting heart disease using machine learning algorithms and an explainable AI method. Sci Rep 14, 23277 (2024). https://doi.org/10.1038/s41598-024-74656-2

[6]. Unveiling the black box: imperative for explainable AI in cardiovascular disease prevention Wu, Yanyi , The Lancet Regional Health – Western Pacific, Volume 48, 101145

[7]. M.U. Sreeja, Abin Oommen Philip, M.H. Supriya, Towards explainability in artificial intelligence frameworks for heartcare: A comprehensive survey, Journal of King Saud University - Computer and Information Sciences, Volume 36, Issue 6, 2024, 102096, ISSN 1319-1578, https://doi.org/10.1016/j.jksuci.2024.102096.

[8]. Jha, R., & Singh, A. (n.d.). Early diagnosis of cyanotic congenital heart disease using deep learning. In *Advances in artificial intelligence for healthcare* ( Taylor & Francis, USA) https://doi.org/10.1201/9781003493570-3.