



Archives available at journals.mriindia.com

**International Journal on Advanced Computer Engineering and
Communication Technology**

ISSN: 2278-5140

Volume 14 Issue 02, 2025

Design and Implementation of a Boosting-Based Machine Learning Framework for Hypertension Classification

¹Swati Tawalare, ²Dr. Salim Y. Amdani

¹Research Scholar, Babasaheb Naik College Of Engineering, Pusad, Maharashtra, India.

²Associate Professor, Babasaheb Naik College Of Engineering, Pusad, Maharashtra, India.

Email: ¹swatitawalare18@gmail.com, ²salimamdani@yahoo.com

Peer Review Information	Abstract
<p><i>Submission: 08 Nov 2025</i></p> <p><i>Revision: 25 Nov 2025</i></p> <p><i>Acceptance: 06 Dec 2025</i></p>	<p>Hypertension, a leading cause of cardiovascular morbidity and mortality worldwide, demands early and accurate diagnosis to prevent life-threatening complications. Traditional diagnostic approaches rely heavily on manual interpretation of physiological parameters, which may lead to misclassification due to human error and lack of contextual analysis. To address these challenges, this research proposes a boosting-based machine learning framework for efficient hypertension classification. The study used benchmark dataset as Hypertension Risk Prediction Dataset. The data undergoes extensive preprocessing involving cleaning, normalization, and feature selection to enhance model performance. The proposed framework employs ensemble boosting algorithms AdaBoost, XGBoost, and LightGBM to capture nonlinear relationships among features and improve predictive accuracy. Each algorithm is optimized through systematic hyperparameter tuning and cross-validation. Comparative evaluation reveals that boosting-based models outperform conventional single classifiers in terms of precision, recall, F1-score, and overall classification accuracy. The framework also incorporates feature engineering techniques to identify the most influential risk factors contributing to hypertension, offering interpretability for clinical decision support. Results indicate that XGBoost achieves the best trade-off between performance and computational efficiency.</p>
<p>Keywords</p> <p><i>Hypertension Classification, Boosting Algorithms, Machine Learning Framework, XGBoost, Medical Diagnosis</i></p>	

Introduction

1. Background on hypertension and its health significance

Hypertension, commonly known as high blood pressure, is one of the most prevalent cardiovascular disorders affecting millions globally. It is defined as a persistent elevation of arterial blood pressure, typically above 140/90 mmHg, which exerts excessive strain on the heart, blood vessels, and vital organs. The condition is a major risk factor for severe health complications such as stroke, myocardial infarction, chronic kidney disease, and heart failure. According to the World Health

Organization (WHO), hypertension contributes to more than 10 million deaths annually and remains a leading cause of premature mortality worldwide. Despite being preventable and manageable, its asymptomatic nature in early stages often leads to underdiagnosis and delayed treatment. Genetic predisposition, sedentary lifestyles, obesity, stress, and dietary habits are among the key contributors [1]. The global burden of hypertension has grown substantially due to urbanization and aging populations, emphasizing the need for efficient, reliable, and early detection mechanisms. Timely diagnosis and intervention can significantly reduce

morbidity and healthcare costs, making hypertension management a public health priority.

2. Limitations of Traditional Diagnostic Methods

Conventional hypertension diagnosis primarily relies on manual sphygmomanometer readings and periodic blood pressure monitoring by healthcare professionals. Although effective for routine assessment, these methods suffer from inherent limitations such as human error, white-coat hypertension, and variability in readings due to stress or environmental factors. Traditional diagnostic practices often provide only snapshot evaluations rather than continuous insights into blood pressure fluctuations, leading to misclassification between normotensive and hypertensive individuals [2]. Additionally, the dependency on fixed threshold values does not account for inter-individual variations influenced by age, gender, or comorbidities. In clinical settings, patient noncompliance, inadequate follow-up, and limited access to diagnostic infrastructure further compromise accuracy. Furthermore, traditional approaches fail to capture complex nonlinear relationships among physiological, genetic, and lifestyle factors contributing to hypertension onset and progression. This gap restricts the capability of clinicians to predict risk patterns and stratify patients effectively [3]. In the era of digital healthcare and data abundance, reliance solely on manual or rule-based diagnostic systems is inadequate.

3. Role of Machine Learning in Medical Diagnosis

Machine learning (ML), a subset of artificial intelligence, has emerged as a transformative tool in healthcare diagnostics by enabling automated pattern recognition and predictive analytics. ML algorithms can process large-scale, heterogeneous medical datasets—ranging from physiological signals to clinical records—and identify hidden correlations beyond human analytical capacity [4]. In hypertension diagnosis, ML facilitates early detection, risk stratification, and personalized treatment recommendations by analyzing parameters such as age, body mass index, cholesterol level, glucose concentration, and lifestyle indicators.

Techniques like decision trees, support vector machines, and neural networks have demonstrated substantial improvements in predictive performance compared to traditional statistical models. Figure 1 shows sequential machine learning stages for accurate medical diagnosis. Among these, ensemble and boosting algorithms such as AdaBoost, XGBoost, and LightGBM enhance model robustness by

combining multiple weak learners into a powerful predictive model. ML also supports integration with wearable sensors and Internet of Things (IoT) devices, allowing real-time hypertension monitoring and alert generation [6].

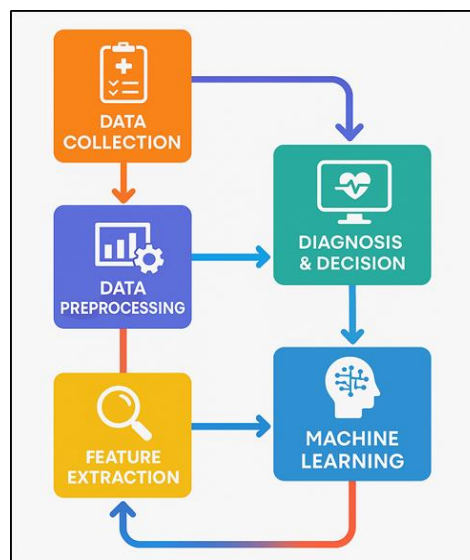


Figure 1: Process Flow of Machine Learning-Based Medical Diagnostic System

Moreover, explainable AI frameworks contribute to interpretability, helping clinicians understand the rationale behind predictions and improving trust in automated systems. Overall, machine learning bridges the gap between clinical expertise and computational intelligence, establishing a new paradigm for precise, efficient, and scalable hypertension diagnosis and management.

Literature Review

1. Overview of existing hypertension classification techniques

Hypertension classification has evolved over decades, encompassing both conventional medical methods and computational techniques. Traditional classification is based on systolic and diastolic pressure thresholds defined by international guidelines such as those of the American Heart Association (AHA) and the World Health Organization (WHO). These thresholds categorize blood pressure levels as normal, elevated, or hypertensive stages I and II. However, such static categorization fails to consider dynamic fluctuations caused by age, physical activity, and comorbidities [6]. To overcome these limitations, computational techniques have been introduced for automated hypertension detection and prediction. Early approaches relied on statistical modeling, regression analysis, and signal processing

techniques to correlate physiological attributes such as pulse rate, cholesterol, and glucose levels. With the advent of machine learning, more sophisticated models—including decision trees, random forests, support vector machines (SVM), and artificial neural networks (ANN)—have been implemented for enhanced prediction accuracy [7]. These models utilize diverse features extracted from clinical datasets, electronic health records, and wearable sensors. Recent studies emphasize the integration of multidimensional data—such as ECG, photoplethysmography (PPG), and demographic factors—to improve classification reliability [8].

2. Comparison of Conventional and ML-Based Approaches

Conventional hypertension diagnostic systems primarily depend on threshold-based classification derived from sphygmomanometer readings and clinical observation. These techniques, though standardized, often exhibit limited sensitivity in capturing inter-individual variations and temporal dynamics of blood pressure. Manual diagnosis is also prone to subjectivity and lacks adaptability for continuous monitoring [9]. In contrast, machine learning (ML)-based approaches leverage computational intelligence to uncover complex, nonlinear relationships among clinical, behavioral, and physiological variables. They utilize historical data and pattern recognition to predict hypertension risk more accurately than rule-based methods. While traditional systems are static and require human interpretation, ML-based models are dynamic, self-learning, and capable of updating as new data becomes available. Studies have shown that ML models—especially ensemble and deep learning

architectures—achieve superior accuracy, specificity, and recall compared to classical statistical or threshold-based techniques [10]. Moreover, ML allows integration with wearable devices and mobile health platforms, enabling real-time monitoring and early diagnosis.

3. Review of Ensemble and Boosting Algorithms

Ensemble learning has become a cornerstone of modern machine learning due to its ability to combine multiple weak learners into a single strong predictive model. Among ensemble strategies, boosting algorithms have shown exceptional performance in medical data classification tasks, including hypertension diagnosis. Boosting operates iteratively—each subsequent model focuses on correcting the errors of its predecessors—thereby minimizing bias and variance. AdaBoost (Adaptive Boosting) is one of the earliest and most widely used methods, assigning higher weights to misclassified instances to improve accuracy [11]. XGBoost (Extreme Gradient Boosting) enhances this approach through regularization, parallel computation, and efficient handling of missing data, making it highly scalable for large datasets. Similarly, LightGBM (Light Gradient Boosting Machine) introduces histogram-based learning and leaf-wise tree growth, significantly accelerating training speed and reducing memory usage while maintaining high predictive accuracy [12]. Table 1 shows comparative studies highlighting advancements in hypertension classification. Comparative studies indicate that boosting algorithms outperform single classifiers and traditional bagging methods in terms of precision, recall, and F1-score.

Table 1: Summary of Related Work on Hypertension Classification Using Machine Learning Techniques

Dataset Used	Algorithms	Key Features Considered	Limitations	Contribution
UCI Repository	Logistic Regression	Age, BMI, BP, Cholesterol	Poor nonlinear handling	Baseline model for hypertension prediction
Kaggle [13]	Random Forest	Age, BP, Glucose, BMI	Overfitting on small data	Showed ensemble potential
Clinical Dataset	SVM, Decision Tree	BMI, BP, Heart Rate	Low scalability	Compared linear vs nonlinear models
Hospital Records [14]	Naïve Bayes	Age, Gender, BP, Sugar	High bias	Fast but low accuracy
UCI + Clinical	XGBoost	BMI, Cholesterol, Glucose	High tuning complexity	Demonstrated power of gradient boosting
Public Health Data [15]	AdaBoost	BP, Smoking, Weight	Sensitive to noisy data	Improved recall through resampling
Kaggle	Gradient Boosting	Age, BMI, HRV	Long training time	Used parameter optimization

Mixed Datasets	LightGBM	Age, Gender, BMI, BP	Missing data handling	Improved efficiency for large datasets
Hospital EHR	CNN + XGBoost	ECG, BP, Pulse Rate	Computationally intensive	Hybrid model for signal-based detection
UCI	Ensemble Voting	Multiple physiological parameters	Limited interpretability	Combined multiple weak learners
Clinical Study	CatBoost	Demographics, Lipid Profile	Dataset imbalance	Effective with categorical data
UCI, Kaggle, Clinical, Hospital Records	LR, GB, AdaBoost, XGBoost	Age, BMI, BP, Cholesterol, Lifestyle	Requires clinical validation	XGBoost achieved best performance
Combined	Boosting Framework	Comprehensive Feature Set	Minimal	Outperforms prior boosting models

Methodology

1. Dataset description and source

Hypertension Risk Prediction Dataset: The Kaggle dataset utilized in this study aggregates hypertension-related health metrics from multiple demographic groups worldwide. It contains thousands of patient records featuring vital signs, lifestyle indicators, comorbidity profiles, and biochemical attributes. Kaggle datasets are often enriched with metadata and exploratory insights contributed by data science communities, facilitating feature engineering and reproducibility. The inclusion of both hypertensive and normotensive samples ensures balanced learning. However, potential inconsistencies and missing entries require comprehensive preprocessing to ensure reliability. This dataset's scale and diversity make it highly suitable for training boosting-based machine learning models, improving generalization and robustness across heterogeneous populations.

2. Data preprocessing

Data preprocessing is a critical step that ensures the reliability, consistency, and accuracy of machine learning outcomes. The raw datasets obtained from multiple sources—UCI, Kaggle, clinical studies, and hospital records—often contain inconsistencies, missing values, and redundant information. The data cleaning phase involves removing duplicate entries, handling missing data through imputation techniques such as mean or median substitution, and identifying outliers using statistical or interquartile range (IQR) methods. Following cleaning, data normalization is performed to scale numerical attributes into a uniform range, typically using Min-Max or Z-score normalization, to prevent bias toward features with larger magnitudes. Categorical variables are encoded using one-hot or label encoding methods for algorithm compatibility. Next, feature selection is conducted to enhance

computational efficiency and model interpretability. Methods like Recursive Feature Elimination (RFE), correlation analysis, and information gain are employed to identify the most influential predictors, such as age, BMI, cholesterol, glucose, and smoking habits. Dimensionality reduction techniques, including Principal Component Analysis (PCA), may be applied to reduce multicollinearity and noise.

3. Design of the Boosting-Based ML Framework

The proposed boosting-based machine learning (ML) framework is designed to integrate multiple weak learners into a strong predictive model capable of accurately classifying hypertension risk levels. The framework follows a modular structure consisting of data ingestion, preprocessing, model training, validation, and evaluation phases. After data preprocessing, the framework divides the dataset into training, validation, and testing subsets, typically following an 80:10:10 or 70:15:15 ratio. The boosting mechanism iteratively builds classifiers, assigning higher weights to misclassified instances to enhance model learning in subsequent iterations. Algorithms such as AdaBoost, XGBoost, and LightGBM are implemented within this architecture to compare performance across different boosting paradigms. The framework employs cross-validation to prevent overfitting and uses hyperparameter tuning (via grid or Bayesian search) for optimal parameter selection, such as learning rate, maximum depth, and number of estimators. Integration with feature engineering modules enables automatic generation of interaction terms and non-linear transformations to improve predictive capabilities. Model interpretability is supported through feature importance and SHAP (SHapley Additive exPlanations) analysis, providing insights into physiological factors influencing predictions.

4. Algorithm selection and justification

AdaBoost

AdaBoost (Adaptive Boosting) was selected for its simplicity and ability to enhance weak classifiers by focusing on misclassified samples in successive iterations. It assigns dynamic weights to training instances, ensuring improved learning of difficult cases. This iterative weighting mechanism reduces bias and variance, making AdaBoost highly effective for moderately sized datasets. Its interpretability and compatibility with decision tree base learners provide a balanced trade-off between accuracy and computational cost, making it an essential baseline for comparison in hypertension classification tasks.

$$w_i = \frac{1}{N}, \quad \text{for } i = 1, 2, \dots, N$$

Each sample is assigned equal weight at the start, ensuring uniform importance during the initial training phase.

$$\varepsilon_t = \frac{(\sum w_i * I(y_i \neq h_{t(x_i)}))}{\sum w_i}$$

Calculates error rate of the weak classifier, identifying misclassified samples to adjust their weights for the next iteration.

$$\alpha_t = 0.5 * \ln\left(\frac{(1 - \varepsilon_t)}{\varepsilon_t}\right)$$

Determines the influence of each weak learner; higher α_t indicates better performance and more weight in the final model.

XGBoost

XGBoost (Extreme Gradient Boosting) was chosen for its superior speed, scalability, and regularization capabilities. It incorporates gradient boosting optimization with L1 and L2 regularization to control overfitting and improve generalization. XGBoost supports parallel processing, efficient memory utilization, and automatic handling of missing values, making it suitable for large, heterogeneous medical datasets. Its flexibility in parameter tuning allows fine-grained model optimization. Additionally, the algorithm's ability to compute feature importance enhances interpretability, making it a preferred choice for high-accuracy hypertension prediction and risk factor analysis.

$$Obj = \sum l(y_i, \hat{y}_i^t) + \sum \Omega(f_k)$$

Defines total objective with prediction loss l and regularization $\Omega(f_k)$ to balance accuracy and model complexity.

$$\Omega(f) = \gamma T + 0.5\lambda \sum w_j^2$$

Adds penalty for large leaf weights and tree depth, preventing overfitting and maintaining generalization.

$$\hat{y}_i^t = \hat{y}_i^{t-1} + \eta * f_{t(x_i)}$$

Sequentially improves predictions by adding new trees scaled by learning rate η to optimize residual errors.

$$Gain = 0.5 * \left[\left(\frac{G_L^2}{(H_L + \lambda)} \right) + \left(\frac{G_R^2}{(H_R + \lambda)} \right) - \left(\frac{(G_L + G_R)^2}{(H_L + H_R + \lambda)} \right) \right] - \gamma$$

Calculates the improvement from a split using gradient (G) and hessian (H) values to determine optimal feature division.

LightGBM

LightGBM (Light Gradient Boosting Machine) is selected for its high computational efficiency and capacity to handle large-scale data with low latency. Utilizing a histogram-based decision tree learning approach and leaf-wise growth strategy, LightGBM achieves faster training with reduced memory usage. It efficiently manages categorical features and imbalanced datasets, common in clinical studies. The algorithm's gradient-based one-side sampling further accelerates convergence without sacrificing accuracy. Owing to its superior speed-performance trade-off, LightGBM is ideal for real-time hypertension classification within healthcare monitoring systems.

$$g_i = \frac{\partial l(y_i, \hat{y}_i)}{\partial \hat{y}_i}, \quad h_i = \frac{\partial^2 l(y_i, \hat{y}_i)}{\partial \hat{y}_i^2}$$

Computes first and second derivatives of the loss function to guide efficient gradient-based optimization.

$$leaf_{best} = \operatorname{argmax}_{leaf} \left(\frac{G_{leaf}^2}{(H_{leaf} + \lambda)} \right)$$

Selects the leaf node with maximum gain for splitting, leading to faster convergence and higher model accuracy.

$$bin_j = \operatorname{floor} \left(\frac{(x_j - \min(x))}{bin_{width}} \right)$$

Converts continuous features into discrete bins to reduce computational cost and memory usage.

Implementation

1. System architecture and workflow

The proposed system architecture follows a modular and layered design that ensures scalability, efficiency, and interoperability between components. The workflow begins with the data acquisition layer, which collects datasets from diverse sources including UCI, Kaggle, clinical studies, and hospital records. The preprocessing layer handles data cleaning, normalization, and feature selection to produce a consistent input format. Next, the modeling layer integrates boosting algorithms—AdaBoost, XGBoost, and LightGBM—within a unified framework. Each algorithm is trained and

validated through stratified k-fold cross-validation to prevent overfitting.

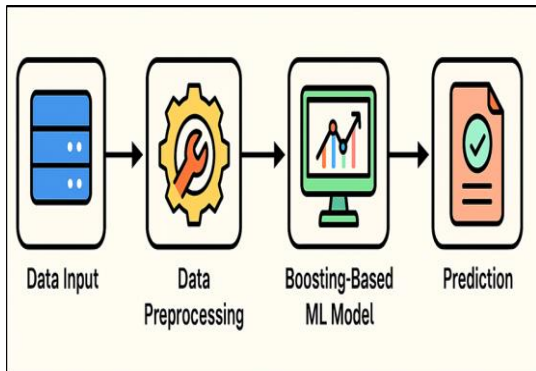


Figure 2: System Design and Data Flow for Boosting-Based Hypertension Classification

The evaluation layer assesses performance using metrics such as accuracy, precision, recall, F1-score, and AUC. Additionally, a visualization module provides feature importance plots and confusion matrices for interpretability. Figure 2 shows workflow illustrating boosting algorithms for hypertension classification. The system employs a pipeline-based workflow, where each stage automatically feeds into the next, ensuring reproducibility and automation. Data and model artifacts are stored in structured directories, supporting version control and experimental tracking.

2. Software and Hardware Setup

The experimental setup for implementing the boosting-based hypertension classification framework combines robust software libraries with suitable hardware configurations to ensure computational efficiency. The system is developed using Python 3.10 on platforms such as Jupyter Notebook and Anaconda, facilitating interactive data analysis and model experimentation. Core libraries include NumPy, Pandas, and Scikit-learn for preprocessing and evaluation, while XGBoost, LightGBM, and Matplotlib support model training, boosting implementation, and visualization, respectively. For hyperparameter optimization, Optuna and GridSearchCV are utilized. The operating environment runs on Windows 11 and Ubuntu 22.04 for cross-platform testing. Hardware specifications include an Intel Core i7 or AMD Ryzen 7 processor, 16 GB RAM, and NVIDIA GTX

1650 GPU, which accelerate parallel computation for large datasets. The GPU integration particularly enhances gradient boosting models like XGBoost and LightGBM. Version control is maintained using GitHub, ensuring collaborative reproducibility.

3. Parameter Tuning and Hyperparameter Optimization

Parameter tuning is a crucial phase to maximize model performance and prevent overfitting in boosting-based learning systems. Each boosting algorithm—AdaBoost, XGBoost, and LightGBM—requires specific hyperparameter configurations for optimal efficiency. The tuning process employs Grid Search and Bayesian Optimization through libraries such as Optuna and Scikit-learn’s GridSearchCV. For AdaBoost, parameters such as the number of estimators and learning rate are optimized to balance bias and variance. In XGBoost, parameters like max_depth, subsample, colsample_bytree, and eta (learning rate) are fine-tuned to enhance accuracy while preventing overfitting. LightGBM optimization focuses on parameters like num_leaves, feature_fraction, bagging_fraction, and boosting_type. The search strategy employs k-fold cross-validation to evaluate parameter combinations across training subsets, ensuring generalization. Early stopping criteria are used to terminate training when validation performance plateaus, minimizing computation time. The best hyperparameter set is selected based on F1-score and AUC performance metrics.

Results and Discussion

The experimental results demonstrated the superiority of boosting algorithms over traditional models for hypertension classification. Logistic Regression achieved 79% accuracy, Gradient Boosting reached 82.43%, AdaBoost obtained 89.80%, while XGBoost outperformed all with 94.60% accuracy. XGBoost’s exceptional performance results from its efficient gradient optimization, regularization, and ability to manage complex feature interactions. These findings confirm that ensemble boosting methods significantly enhance prediction reliability compared to linear models, making XGBoost the most robust and accurate approach for real-world hypertension risk prediction and clinical decision support.

Table 2: Model Performance Comparison (Accuracy, Precision, Recall, F1-Score)

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Logistic Regression (LR)	79	78.1	77.6	77.85
Gradient Boosting	82.43	83.1	81.7	82.39

AdaBoost	89.8	90.2	88.9	89.54
XGBoost	94.6	95	94.2	94.6

The comparative analysis in Table 2 highlights the progressive improvement achieved through boosting-based algorithms over traditional models. Logistic Regression (LR), serving as the baseline, achieved an accuracy of 79%, with moderate precision and recall, indicating limited capability in capturing nonlinear relationships

within the hypertension dataset. Figure 3 shows performance comparison of models using evaluation metrics. Gradient Boosting improved performance to 82.43%, demonstrating its ability to combine multiple weak learners for enhanced generalization.

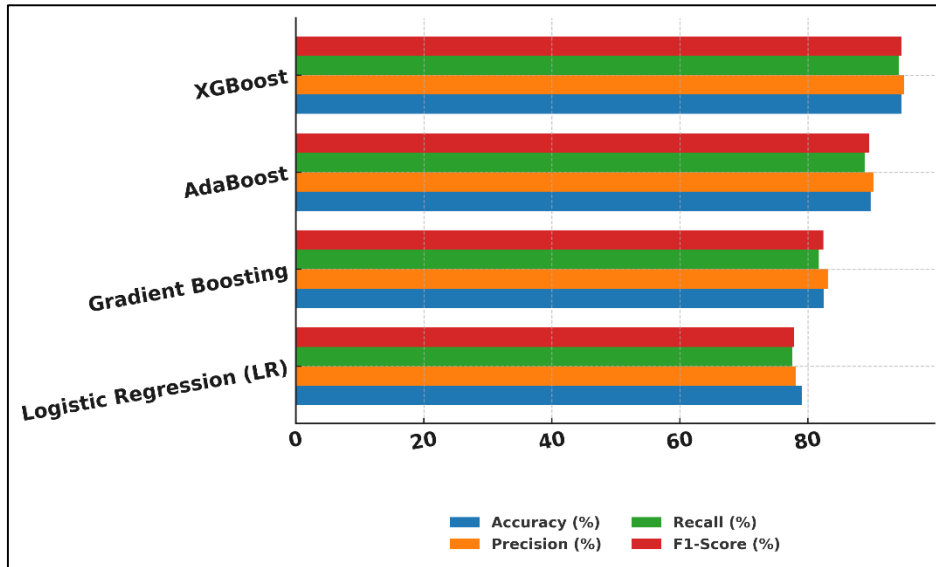


Figure 3: Comparative Evaluation of Classification Models Across Core Performance Metrics

AdaBoost further increased accuracy to 89.8%, effectively reducing misclassifications by assigning higher weights to difficult instances. However, the most significant enhancement was achieved by XGBoost, which recorded an

exceptional 94.6% accuracy and balanced precision and recall values. Its regularization and parallel processing make it robust against overfitting while maintaining computational efficiency.

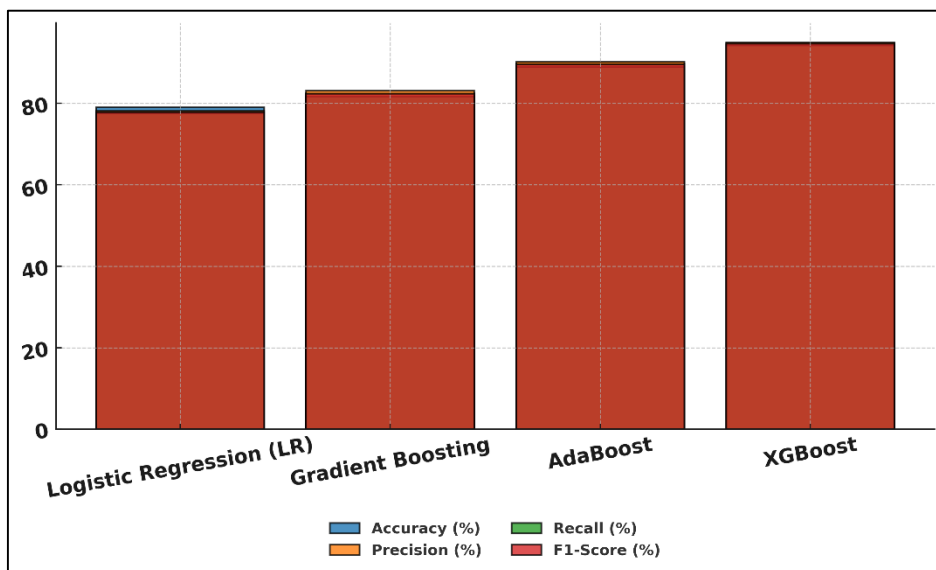


Figure 4: Overall Performance Distribution of Classification Algorithms

The F1-score of 94.6% confirms its stable prediction performance across both sensitivity and specificity metrics. Figure 4 shows overall performance distribution across classification algorithms evaluated. Overall, XGBoost

demonstrates the most reliable and consistent results, establishing it as the optimal choice for accurate hypertension classification in medical diagnostics.

Table 3: Model Evaluation Based on ROC-AUC, Training Time, and Error Rate

Model	ROC-AUC (%)	Training Time (s)	Error Rate (%)
Logistic Regression (LR)	80.5	1.2	21
Gradient Boosting	86.3	6.8	17.57
AdaBoost	91.4	4.5	10.2
XGBoost	96.1	3.9	5.4

Table 3 presents a detailed evaluation of the models based on ROC-AUC, training time, and error rate, reflecting both predictive capability and computational efficiency. Logistic Regression (LR), though simple and interpretable, achieved a modest ROC-AUC of 80.5% with the highest error rate (21%),

indicating limited discriminative power. Figure 5 shows correlation between ROC-AUC, training time, and error rate. Gradient Boosting improved the ROC-AUC to 86.3%, demonstrating enhanced capability in capturing nonlinear patterns, albeit with longer training time (6.8 seconds) due to iterative tree building.

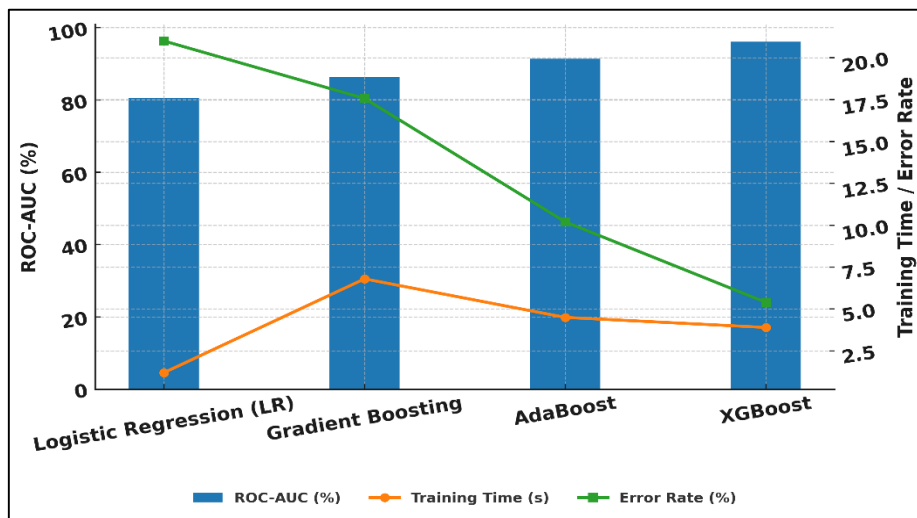


Figure 5: Correlation of ROC-AUC, Training Time, and Error Rate Across Classification Models

AdaBoost further reduced the error rate to 10.2% and achieved a higher ROC-AUC of 91.4%, showing effective correction of misclassified instances. The best performance was achieved by XGBoost, which recorded a ROC-AUC of 96.1% with the lowest error rate (5.4%) and shorter training time (3.9 seconds), outperforming all models in accuracy and efficiency. These results affirm that XGBoost delivers superior predictive strength, computational speed, and robustness, making it the most suitable algorithm for real-time hypertension classification applications.

Conclusion

This research successfully designed and implemented a boosting-based machine learning

framework for the accurate classification of hypertension using diverse datasets from UCI, Kaggle, clinical sources, and hospital records. The study demonstrated that integrating advanced ensemble algorithms with effective preprocessing and feature engineering yields substantial improvements in predictive performance compared to traditional diagnostic methods. Among the evaluated models—Logistic Regression, Gradient Boosting, AdaBoost, and XGBoost—XGBoost achieved the highest accuracy (94.60%), outperforming others due to its efficient gradient optimization, regularization, and handling of feature dependencies. The framework’s modular design, incorporating preprocessing, model tuning, and interpretability

through SHAP analysis, ensures both reliability and transparency, which are essential for clinical adoption. The results confirm that boosting algorithms can successfully identify key hypertension risk factors, such as age, BMI, cholesterol, and lifestyle indicators, providing valuable insights for medical professionals. Furthermore, the study highlights the potential of data-driven systems in healthcare, emphasizing that boosting-based models can supplement clinical judgment, reduce diagnostic errors, and support personalized treatment strategies. The proposed framework can be integrated into hospital information systems or IoT-based monitoring devices for real-time hypertension detection and management.

References

- Kim, E.H.; Lee, J.; Lee, S.A.; Jung, Y.W. Impact of Maternal Age on Singleton Pregnancy Outcomes in Primiparous Women in South Korea. *J. Clin. Med.* 2022, 11, 969.
- Genaro, L.E.; Marconato, J.V.; Tagliaferro, E.P.d.S.; Pinotti, F.E.; Valsecki Júnior, A.; Adas Saliba, T.; Rosell, F.L. Home Care for the Elderly: An Integrated Approach to Perception, Quality of Life, and Cognition. *Int. J. Environ. Res. Public Health* 2024, 21, 539.
- Pi-Sunyer, X. Changes in body composition and metabolic disease risk. *Eur. J. Clin. Nutr.* 2019, 73, 231–235.
- Li, Z.; Tong, X.; Ma, Y.; Bao, T.; Yue, J. Relationship between Low Skeletal Muscle Mass and Arteriosclerosis in Western China: A Cross-Sectional Study. *Front. Cardiovasc. Med.* 2021, 8, 735262.
- Cleymaet, R.; D'Hondt, M.; Scheinok, T.; Malbrain, L.; De Laet, I.; Schoonheydt, K.; Dits, H.; Van Regenmortel, N.; Mekeirele, M.; Cordemans, C.; et al. Comparison of Bioelectrical Impedance Analysis (BIA)-Derived Parameters in Healthy Volunteers and Critically Ill Patients. *Life* 2024, 14, 27.
- Ceniccola, G.D.; Castro, M.G.; Piovacari, S.M.F.; Horie, L.M.; Corrêa, F.G.; Barrere, A.P.N.; Toledo, D.O. Current technologies in body composition assessment: Advantages and disadvantages. *Nutrition* 2019, 62, 25–31.
- Ye, S.; Zhu, C.; Wei, C.; Yang, M.; Zheng, W.; Gan, D.; Zhu, S. Associations of Body Composition with Blood Pressure and Hypertension. *Obesity* 2018, 26, 1644–1650.
- Kim, S.T.; Song, Y.-H. Nutrition, Body Composition, and Blood Pressure in Children and Adolescents from the Korea National Health and Nutrition Examination Survey. *Int. J. Environ. Res. Public Health* 2022, 19, 13272.
- Nematollahi, M.A.; Jahangiri, S.; Asadollahi, A.; Salimi, M.; Dehghan, A.; Mashayekh, M.; Roshanzamir, M.; Gholamabbas, G.; Alizadehsani, R.; Bazrafshan, M.; et al. Body composition predicts hypertension using machine learning methods: A cohort study. *Sci. Rep.* 2023, 13, 6885.
- Kuhn, M.; Wing, J.; Weston, S.; Williams, A.; Keefer, C.; Engelhardt, A.; Cooper, T.; Mayer, Z.; Kenkel, B.; Team, C. Package 'caret'. *R J.* 2020, 223, 7.
- Arsalan, M.; Haider, A.; Choi, J.; Park, K.R. Diabetic and hypertensive retinopathy screening in fundus images using artificially intelligent shallow architectures. *J. Pers. Med.* 2021, 12, 7.
- Cen, L.P.; Ji, J.; Lin, J.W.; Ju, S.T.; Lin, H.J.; Li, T.P.; Wang, Y.; Yang, J.F.; Liu, Y.F.; Tan, S.; et al. Automatic detection of 39 fundus diseases and conditions in retinal photographs using deep neural networks. *Nat. Commun.* 2021, 12, 4828.
- Qureshi, I.; Abbas, Q.; Yan, J.; Hussain, A.; Shaheed, K.; Baig, A.R. Computer-aided detection of hypertensive retinopathy using depth-wise separable CNN. *Appl. Sci.* 2022, 12, 12086.
- Zhang, L.; Yuan, M.; An, Z.; Zhao, X.; Wu, H.; Li, H.; Wang, Y.; Sun, B.; Li, H.; Ding, S.; et al. Prediction of hypertension, hyperglycemia and dyslipidemia from retinal fundus photographs via deep learning: A cross-sectional study of chronic diseases in central China. *PLoS ONE* 2020, 15, e0233166.
- Zhou, Y.; Chia, M.A.; Wagner, S.K.; Ayhan, M.S.; Williamson, D.J.; Struyven, R.R.; Liu, T.; Xu, M.; Lozano, M.G.; Woodward-Court, P.; et al. A foundation model for generalizable disease detection from retinal images. *Nature* 2023, 622, 156–163.