Brain Stroke Prediction Using Gradient Boosting with Xai

S Mithilesh Reddy
Department of CSE, JNTUA College Of Engineering (Autonomous), Ananthapuramu, Andhra Pradesh, India

***Abstract*—Stroke is the largest cause of death and disability globally, making it a serious health concern. A disturbance in blood flow to the brain results in oxygen deprivation, damage to brain cells, and significant issues with speech, movement, and everyday tasks. Early detection plays a critical role in enabling timely medical intervention and improving patient outcomes. Conventional stroke prediction systems frequently use Explainable AI in conjunction with Machine Learning models to improve interpretability. However, many of these systems depend on conventional algorithms that may not fully exploit the capabilities of more advanced techniques. They also frequently face challenges related to class imbalance and may struggle to model the complex interdependencies among risk factors. This research presents an enhanced stroke prediction framework that integrates advanced gradient boosting methods with XAI to improve both accuracy and interpretability. A hybrid resampling method that blends Borderline, Tomek Links, and SMOTE. While optimized models involving CatBoost(Category Boosting), XGBoost (Extreme Gradient Boosting), LightGBM(Light Gradient Boosting Machine), and Gradient Boosting Machine (GBM) are utilized to ensure robust performance, SMOTE is used to correct data imbalance. This comprehensive approach supports more accurate stroke risk assessment and contributes to early diagnosis and informed clinical decision-making.**

***Index Terms*—Brain Stroke Prediction, Class Imbalance, Explainable AI (XAI), Gradient Boosting, Machine Learning.**

1. INTRODUCTION

As a prevalent death cause moreover lifelong disability globally, stroke has been a serious public health issue. Significant impairments in speech, mobility, and cognitive function are caused by oxygen starvation, brain cell death, and blood supply disruption to brain. As per estimates from the WHO (World Health Organization), strokes claim lives of about 13million individuals each year, moreover cause 5.5million deaths. Reducing mortality rates and limiting long-term consequences require early discovery and prompt medical intervention. Stroke risk is still hard to forecast with any degree of accuracy due to the complexity of the factors that contribute to it, which involve diabetes, heart disease, high blood pressure, smoking, as well as lifestyle choices. Traditional statistical models and ML approaches were widely utilized to assess stroke risk, but they often struggle to capture the intricate relationships between multiple risk factors and are limited in handling class imbalance issues within medical datasets.

1.1 GRADIENT BOOSTING IN STROKE PREDICTION

In recent years, gradient boosting techniques have gained popularity in medical research because of their capability to increase prediction accuracy and manage complex, nonlinear data patterns. Unlike conventional machine learning models, gradient boosting creates DT series (decision trees) sequentially, with each one aiming to correct the errors of the one before it. The process enhances the ability of the model to discover stroke risk patterns and generalise well across various datasets. This sequential refinement allows for the capture of subtle relationships within the data that might be missed by single models. Such models as XGBoost, CatBoost, LightGBM, and GBM demonstrated significant improvements in medical tasks related to prediction, particularly stroke risk assessment. The models perform better than conventional approaches since they have less bias and variance and can effectively process large datasets. Nevertheless, the fact that such models cannot be interpreted is a frequent criticism, even though they are highly accurate. This makes them hard to incorporate into clinical decision-making.

1.2 IMPORTANCE OF EXPLAINABLE AI IN HEALTHCARE

Interpretability has also become necessary because many advanced ML models are opaque, especially in sensitive areas such as healthcare. Clinicians would need not only proper prediction but also insight into the rationale behind the prediction to maintain transparency, trust, and responsible medical practice. Explainable Artificial Intelligence (XAI) is a response to this requirement, in which model behavior is rendered more interpretable. In this research, the methods of XAI have been combined with gradient boosting algorithms to bridge the interpretability-prediction accuracy gap. With SHAP (Shapley Additive Explanations) to quantify the value of each feature to a prediction, doctors can identify the factors that have the most impact on stroke risk. Furthermore, LIME (Local Interpretable Model-Agnostic Explanations) offers local surrogate models, which allow explaining individual predictions, offering case-specific knowledge. Its local interpretation, along with global interpretation, allows SHAP and LIME to allow physicians to make informed choices through providing personalized stroke risk assessments.

1.3 SYSTEM ARCHITECTURE FOR STROKE PREDICTION

The suggested stroke prediction system has a developed pipeline that is striving to achieve the highest accuracy, reliability, and interpretability. Once the health data of the patients is obtained using publicly available data sets, the process involves preprocessing data to standardize numerical variables, handle missing data, and encode categorical variables. Another major challenge is the natural class imbalance, which is that stroke cases are very underrepresented relative to non-stroke cases. A hybrid resampling approach is adopted, which is a combination of SMOTE, Tomek Links, and Borderline-SMOTE. Such a method not only enhances the representation of minority classes but also enhances the decision boundaries and minimizes noise and overfitting. Figure 1 represents the general organized procedure of the proposed stroke prediction system. This overall procedure helps to diagnose early and makes better medical intervention possible.

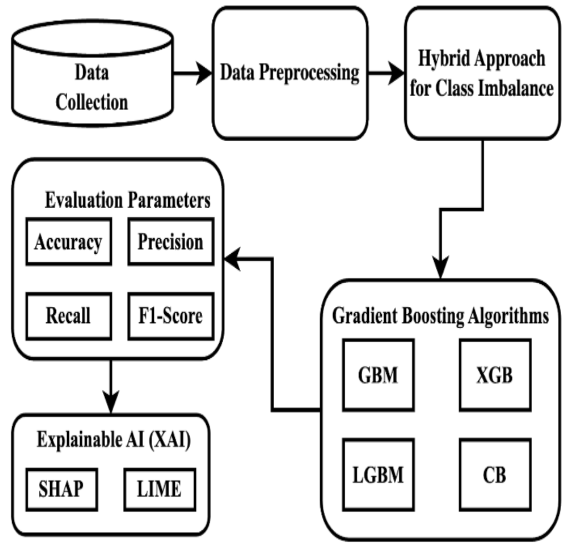

Figure 1: Proposed Framework

Gradient boosting models, GBM, XGBoost, LightGBM, and CatBoost, are trained on the balanced dataset to identify non-linear relationships between clinical risk factors and stroke occurrence. These models can discover complex non-linear relationships between features and have been especially well-suited for tabular medical data. Their performance is evaluated utilizing classification metrics, like precision, accuracy, recall, as well as F1-score, for ensuring comprehensive evaluation. To enhance the interpretability of model predictions, XAI approaches like SHAP as well as LIME have been applied after training. These techniques help medical professionals understand the importance of distinctive traits and the logic behind each prediction by offering both local as well as global explanations. Consequently, system not only becomes accurate but also transparent, reliable, and suitable for integration into clinical decision-making workflows.

This paper presents a novel approach to brain stroke prediction by integrating advanced gradient boosting models with explainable AI techniques. The proposed system addresses data imbalance and enhances model interpretability, aiming to deliver both accurate and clinically transparent predictions. Rest of paper is organised as follows: Section II examines previous work, Section III outlines technique and model creation, Section IV gives experimental data, Section V discusses findings and future scope, and Section VI concludes.

2. RELATED WORK

Stroke prediction has attracted considerable attention in both the medical and computational research communities due to its potential to reduce mortality and long-term disability. Traditional approaches primarily employed statistical models such as LR and rule-based expert systems, which relied on identifying isolated risk factors like age, hypertension, and diabetes. However, these methods often fell short in modeling the complex, nonlinear interactions among various stroke-related attributes. When ML and AI (artificial intelligence) became more popular, researchers started using more sophisticated methods that could learn complex patterns out of data. The models that applied SVM (support vector machines) and DTs were more effective than the classical statistical methods, but still failed to cope with generalizability to heterogeneous clinical data. Mridha et al. [1] have shown that simpler models like logistic regression and SVM are not enough in capturing complex features dependency, and therefore they perform poorly in real-life stroke prediction.

This drawback has resulted in a paradigm shift towards ensemble-based methods that harness the power of more than one classifier. Gradient Boosting and Random Forest (RF) ensemble models have been relatively successful in improving the predictive performance in healthcare, particularly on tasks involving stroke detection. Ushasree et al. [2] proposed ESPESM, a stacked ensemble technique that utilized several base classifiers and a meta-learner and showed a significant increase in the classification performance. In the same way, Bathla and Kumar [3] demonstrated that the performance of stroke identification can be enhanced by ensemble classifiers and effective feature selection, even in cases where the noisy or irrelevant features are available. Their research highlighted the significance of feature selection in lowering computational costs, enhancing model robustness, and lowering dimensionality—all crucial for medical data, which is frequently afflicted by redundant or weakly correlated information.

Even though the accuracy is a primary aim in the predictive modelling, it is not the only indicator of clinical application. Interpretability is also vital, particularly in the medical field where physicians must understand and trust the AI-based decisions. Nevertheless, most effective models, such as gradient boosting, tend to be opaque, and therefore, it is hard to derive valuable insights from their predictions. To solve this, researchers have resorted to XAI methods. Nazar et al. [9] reviewed the current research on HCI and XAI in healthcare and identified major strategies to make models transparent and clinicians trust them. Mridha et al. [1] used SHAP and LIME in their pipeline of stroke prediction to achieve both local and global transparency. These tools assist physicians in linking AI-informed recommendations to established medical information by visualizing how particular variables affect model outputs. Gawde et al. [4] further developed this idea by using XAI approaches that involve SHAP, LIME, PDP (Partial Dependence Plots), and ICE (Individual Conditional Expectation) and thus provided multidimensional interpretability to increase user confidence in AI systems.

The problem of class imbalance, in which many fewer stroke patients have been there than non-stroke cases, is a significant obstacle in stroke prediction tasks. This imbalance leads to biased models that often ignore minority class instances, which are the most clinically significant. Mridha et al. [1] tackled this by employing SMOTE (Synthetic Minority Oversampling Technique), Tomek Links, and Borderline-SMOTE to synthetically balance the dataset and refine decision boundaries. This hybrid resampling strategy not only improved sensitivity but also reduced the false-negative rate, making the model more reliable for identifying high-risk individuals. The same concern was echoed in comparative studies like those by Rahman et al. [5], who noted that class balancing, combined with feature engineering and ensemble methods like XGBoost and LightGBM, significantly improves prediction performance in imbalanced medical datasets.

From an epidemiological standpoint, understanding stroke prevalence and risk factors is essential for developing effective predictive systems. In their discussion of the rising incidence of stroke worldwide, Katan and Luft [6] emphasized the financial toll as well as the health consequences, particularly in low- and middle-income nations. The significance of scalable and easily accessible AI-powered solutions for early stroke detection is highlighted by their findings. In a large-scale study across Northern China, Xia et al. [7] identified key regional risk factors involving hypertension, diabetes, and obesity, noting significant variations in prevalence based on demographic and geographic factors. These findings support the idea that AI models should be adaptable to local population characteristics. On a global level, the WHO continues to report stroke as one of top 3 causes of death, reinforcing the urgency for early intervention tools powered by AI [8].

Beyond clinical and demographic considerations, the social context of stroke recovery is increasingly being recognized as critical. In their study of the association between social support and patient participation following a stroke, Elloker and Rhoda [10] discovered a substantial link between better recovery outcomes and emotional, instrumental, and informational assistance. Further, in a complementary study, the same authors proposed a predictive model emphasizing the role of family and community support in stroke rehabilitation, indicating that social factors are strong predictors of recovery success and should be integrated into future predictive systems Elloker and Rhoda [11]. These findings highlight the limitations of traditional models that solely focus on biological markers. Incorporating these nuanced social elements could lead to more comprehensive and personalized recovery plans for stroke patients. These are the features that are not taken into account in classical models but are necessary to provide a more complex and human-oriented measurement of the risk of stroke and the possibilities of recovery.

Regarding the availability of data, the Kaggle Stroke Prediction Dataset [12] becomes a new standard in stroke-related ML studies. It contains fundamental clinical characteristics that include BMI, age, glucose, blood pressure, type of work, heart disease, smoking, and marital status. The lack of genetic or chronic behavioral features makes the dataset a core training and validation resource for a stroke prediction algorithm since it is easily accessible and useful.

Machine learning has significantly improved with regard to stroke prediction, especially in ensemble modeling, class balancing, and explainability. Nonetheless, there are still issues of interpretability, integration into the real world, and individualization. This research suggests a universal framework that will improve interpretability with the help of gradient boosting models (XGBoost, LightGBM, CatBoost, GBM) and SHAP and LIME. Also, to address the imbalance, a hybrid resampling approach (SMOTE, Tomek Links, and Borderline-SMOTE) is used. This mixture strategy seeks to develop a more precise, open, and clinically beneficial stroke forecasting framework.

3. METHODOLOGY

3.1 DATA COLLECTION AND PREPROCESSING

Comprehensive patient records covering a broad range of clinical, demographic, and lifestyle factors determined to have an impact on the risk of stroke comprise the study's data. The main characteristics are age, gender, hypertension status, the presence of heart sickness, cigarette smoking, BMI and average glucose level, and type of employment. The variables play a very important role in the interpretation of personal health profiles as well as their relationship with stroke incidences. The data was taken from a well-known and publicly accessible healthcare repository, which makes it both genuine and trustworthy. The ability to predict stroke is a sensitive and high-impact problem, and thus, the dataset requires careful preprocessing to enhance the performance of the model and make it generalizable. It involves encoding of categorical variables, normalization of numeric data, and missing values. Table 1 summarizes in detail the properties of the dataset, such as the names of the attributes, as well as a short description of their comparative medical importance.

Table 1: Data Description

Feature	Description
Gender	Male, Female, Others
Age	Patient age (in years)
Hypertension	Presence of hypertension (0 = No, 1 = Yes)
Heart Disease	Presence of any heart-related condition (0 = No, 1 = Yes)
Ever Married	Marital status of the patient (Yes = married, No = single or not married)
Work Type	Employment category (Private, Self-employed, Govt, etc.)

Residence Type	Urban or Rural living area
Glucose Level	Average glucose concentration in the blood
Body Mass Index (BMI)	Indicator of body fat based on weight and height
Smoking Status	Classification into Never Smoked, Former Smoker, or Current Smoker
Stroke	Stroke occurrence is captured as the target variable (0 = No stroke, 1 = Stroke).

Partial or absent data may be a major drawback to the reliability and performance of the ML models, particularly when it comes to such applications as those connected with healthcare, where precision is a paramount issue. Consequently, the process of dealing with missing values is a crucial component of the data preparation process. When dealing with numerical variables in this particular investigation, such as BMI and average glucose level, mean imputation is employed. By substituting the mean of the known values for the missing values, this approach lowers the likelihood of bias while maintaining data consistency and the overall statistical distribution of every character. Categorical variables are encoded using the right encoding strategies depending on the character of the data to be prepared to train the model. Label encoding is employed to encode binary categorical data, including heart disease, hypertension, and gender. Each category is given a unique integer number, usually in the range of 0 and 1. This method will enable the model to make binary distinctions without any superfluous complexity. Conversely, the multi-class categorical data, which entails the smoking status and type of work, is modified through one-hot encoding. This approach does not introduce spurious ordinal relationships, which would otherwise distort the learning process, by explicitly defining separate binary columns associated with each category. A detailed overview of the encoding schemes applied to each categorical feature is provided in Table 2.

Table 2: Encoding Methods for Categorical Variables

Feature	Encoding Method	Encoded Values
Gender	Label Encoding	Female: 0, Male: 1, Other: 2
Ever Married	Label Encoding	No: 0, Yes: 1
Work Type	One-Hot Encoding	Govt_job: 0, Never_worked: 1, Private: 2, Self-employed: 3, Children: 4
Residence Type	Label Encoding	Rural: 0, Urban: 1
Smoking Status	One-Hot Encoding	Unknown: 0, Formerly Smoked: 1, Never Smoked: 2, Smokes: 3

Choosing the right features is essential to enhancing model performance by eliminating redundant and less significant features. In this research, a combination of statistical tests and correlation analysis is used to determine which characteristics are most pertinent to stroke prediction, as shown in Table 3. While Pearson correlation is used to evaluate numerical variables' relationship to the target variable, the Chi-Square test is employed to analyze the statistical significance of categorical features. This dual-pronged approach ensures a comprehensive assessment of feature relevance across different data types, leading to a more streamlined and effective feature set for the prediction model. The careful selection of these features directly contributes to the model's ability to accurately identify patterns associated with stroke risk.

Table 3: Feature Selection Results

Feature	ANOVA F-Score	Chi-Squared Score
Age	326.92	N/A
Glucose Level	90.50	N/A
BMI	7.76	N/A

Heart Disease	N/A	87.98
Hypertension	N/A	75.45
Ever Married	N/A	20.62
Smoking Status	N/A	3.37
Work Type	N/A	2.92
Residence Type	N/A	0.60
Gender	N/A	0.23

To ensure uniformity in feature magnitudes, numerical attributes such as BMI and glucose levels are normalized. Min-Max Scaling is applied to rescale values within a standard range of [0,1], which keeps learning of the model from being dominated by characteristics with greater numerical ranges. By using this normalizing strategy, gradient-based optimization becomes more stable, and all input variables contribute equally to model training.

3.2 HYBRID APPROACH FOR CLASS IMBALANCE

As seen in Figure 2, a combination resampling technique combining SMOTE, Tomek Links, and Borderline-SMOTE is employed to lessen the severe class imbalance frequently observed in stroke prediction datasets. SMOTE is the first step in the procedure, which improves the representation of the minority class by interpolating existing stroke cases to create synthetic examples. Following this, Tomek Links is applied to remove closely paired majority class instances that introduce ambiguity, thus improving class separability and reducing noise. Lastly, Borderline-SMOTE focuses on creating artificial samples close to the class boundary, strengthening the classifier's capacity to discriminate between examples that are stroke and those that are not. This combined strategy reduces overfitting, produces a more balanced dataset, and enhances the model's overall prediction ability.

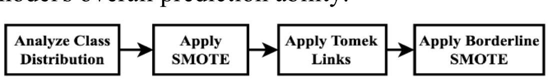

Figure 2: Hybrid Resampling Process

This three-step resampling strategy ensures a more balanced dataset, leading to improved model training and enhanced stroke detection accuracy. The impact of the resampling approach is illustrated in Figure 2. To provide a clearer view of the class distribution shift, Figure 3 compares the dataset before and after resampling. Initially, the dataset exhibits a stark imbalance, with only 4.3% stroke cases. After applying the hybrid resampling method, a 50-50 balance is achieved between stroke incidents and normal cases.

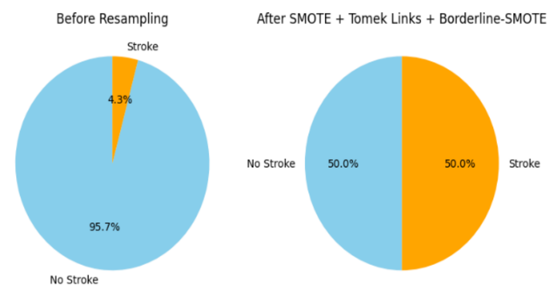

Figure 3: Class Distribution Before and After Resampling

3.3 MODEL DEVELOPMENT AND TRAINING

The development of a stroke prediction model includes data splitting and training multiple gradient boosting algorithms to evaluate and choose the model that classifies data most accurately. In the present research, four gradient boosting models Predictive performance was assessed for Gradient Boosting Machine (GBM), XGBoost, LightGBM, and CatBoost. 80% of the dataset was employed for training, while 20% was used for testing. A stratified distribution that preserves class proportions. A hybrid resampling technique (SMOTE, Tomek Links, and Borderline-SMOTE) was applied before splitting to mitigate class imbalance. Each model was trained using a systematic process, including data preprocessing, model initialization, training, and performance evaluation. GBM builds sequential decision trees, improving accuracy by minimizing residual errors, while XGBoost enhances efficiency through parallel processing and regularization techniques. LightGBM is appropriate for large datasets since it maximizes training time using a leaf-wise tree growth technique and histogram-based learning. CatBoost is tailored for categorical data, using ordered boosting and automatic feature encoding to reduce overfitting. Classification metrics including F1-score, recall, accuracy, and precision have been employed to assess the models. To determine which stroke prediction algorithm works best, this comprehensive comparative analysis was crucial for identifying the optimal model for clinical deployment.

3.4 EXPLAINABLE AI IN MODEL EXPLANATION

Although gradient boosting algorithms are quite accurate at predicting strokes, their interpretability is limited by their black-box nature, which makes it difficult for medical personnel to evaluate and trust AI-driven conclusions. In the healthcare domain, model transparency is crucial to ensure that predictions align with clinical reasoning. Explainable Artificial Intelligence (XAI) techniques help bridge this gap by offering insights into model behavior, enhancing accountability, and improving trust in machine learning predictions. Figure 4 illustrates the overall workflow for implementing XAI in this research, ensuring clarity in the integration of interpretability techniques with predictive modeling. In this research, two popular XAI techniques, SHAP and LIME, are used to understand the outcomes of the best-performing model. SHAP assigns a contribution value to each feature using the concept of cooperative game that demonstrates the impact of different factors on stroke prediction. Some visualizations, like SHAP summary and force plots, offer global and local interpretability. In contrast, LIME provides local surrogate models that approximate model predictions of specific instances, allowing healthcare professionals to understand the reasons behind labeling a patient with a low or a high risk. This integration of multiple approaches to research ensures that the models of stroke prediction are correct and explainable, which contributes to better reliability and feasibility of clinical decisions.

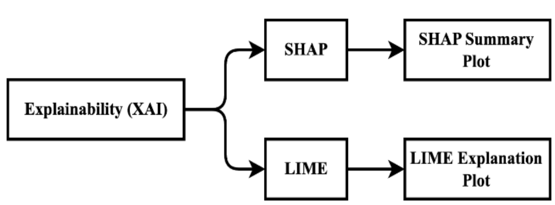

Figure 4: Workflow for XAI

In this part, the pipeline of stroke prediction is shown, including the essential steps of data preparation and class balancing, training, and explaining models. Preprocessing and hybrid resampling improve data quality and reduce class imbalance, while assessing gradient boosting models ensures strong predictive performance. Interpretable AI, such as SHAP and LIME, provides information on model decisions and may assist in making informed clinical decisions.

4. RESULTS

4.1 METRICS

Four important classification measures had been employed to determine the performance of trained models: F1-score, recall, accuracy, and precision. Accuracy measures a model's performance by calculating the percentage of correctly detected occurrences. The fraction of successfully predicted stroke cases among all expected cases is called precision or positive predictive value. Cases that guarantee dependability in risky projections. Recall assesses the model's capacity to recognise real stroke victims, reducing false negatives in the medical setting. The F1-score, accuracy, and recall harmonic mean is a balanced metric that takes into consideration both false positives and false negatives. This has been particularly beneficial in addressing the disparity in class. An egalitarian metric that takes into account both the quantity of false positives and false negatives is the F1-score. It is also very handy in solving class imbalance since this is been harmonic mean of both recall and precision.

4.2 EVALUATING AND COMPARING GRADIENT BOOSTING MODELS

In order to comprehensively determine the level of success of gradient boosting algorithms in stroke prediction, four models were tested: GBM, XGBoost, LightGBM, and CatBoost. Accuracy, precision, recall, and F1-score were the main measures employed to determine these models. These criteria were selected because it was necessary to take into consideration both each model's overall accuracy and how well it performed on unbalanced data, which is a common and significant issue in medical data. Life-threatening consequences occur in stroke prediction when a positive case is not properly identified. As shown in Figure 5, the accuracy of the models is compared visually in a bar graph, which allows them to have an overview of the overall performance of the models. Nonetheless, although accuracy is a pragmatic starting point, it is not a good measure of how well a model predicts minority classes like stroke cases. Therefore, it is necessary to investigate these models more closely with narrower metrics to comprehend their applicability to the real world of clinical practice.

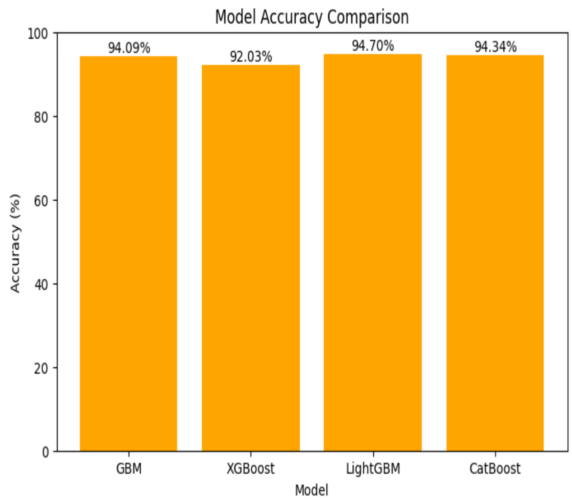

Figure 5: Model Accuracy Comparison

Table 4 shows a more detailed comparison of the models in terms of accuracy, precision, recall, and F1-score. This table can give more understanding of the weaknesses and strengths of every algorithm. To illustrate, LightGBM had the best recall in predicting non-stroke cases, which can be used to minimize false alarms and unnecessary interventions. On the other hand, CatBoost had a good precision-recall balance, which is specifically important when involves correctly identifying actual stroke cases and minimising false positives. Figures 6, 7, and 8 compare recall, precision, and F1-score between stroke and non-stroke classes, respectively, to provide a more comprehensive examination of the models' performance on a per-class basis. The efficiency of each model in identifying stroke cases through recall is displayed in Figure 6. The precision component, or the accuracy of the stroke predictions, is covered in Figure 7. The overall view in terms of F1-score, which is the sum of precision and recall, is provided in Figure 8. Such visual comparisons help make a more informed assessment and make the selection of models not only based on the statistical performance criteria but also based on clinical considerations when working with imbalanced healthcare data.

Table 4: Gradient Boosting Model Performance Comparison

Model	Precision (0)	Recall (0)	F1-Score (0)	Precision (1)	Recall (1)	F1-Score (1)
GBM	0.94	0.94	0.94	0.93	0.94	0.94
XGBoost	0.93	0.91	0.92	0.91	0.93	0.92
Light GBM	0.94	0.96	0.95	0.96	0.94	0.95
CatBoost	0.95	0.94	0.94	0.94	0.95	0.94

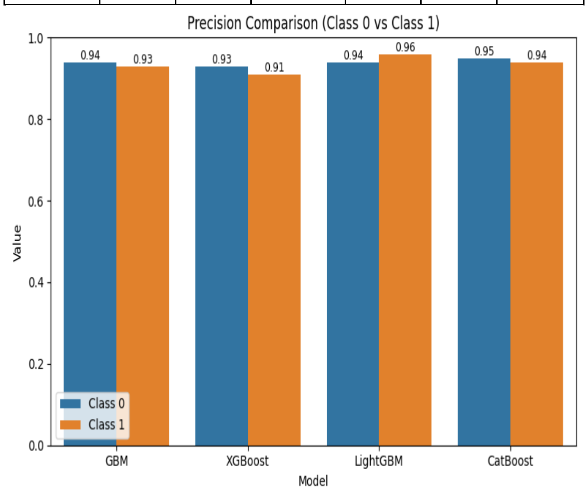

Figure 6 Precision Comparison for Stroke and Non-Stroke Cases

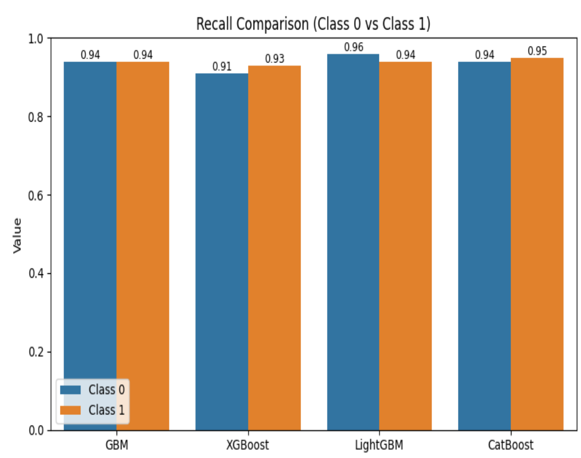

Figure 7: Precision Comparison for Stroke and Non-Stroke Cases

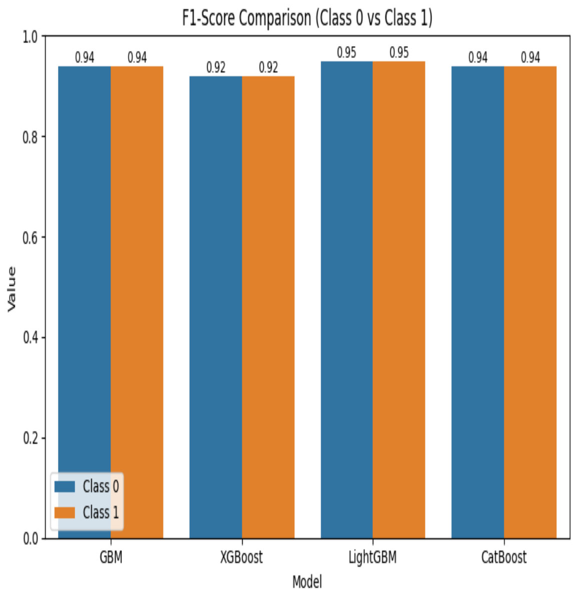

Figure 8: F1-Score Comparison for Stroke and Non-Stroke Cases

4.3 IMPLEMENTATION OF EXPLAINABLE AI

SHAP and LIME were the XAI methods that were applied to the improvement of the interpretability as well as transparency of the model. As it demonstrates how each feature adds to the justification of model predictions, the SHAP summary plot highlights the importance of each characteristic (Figure 9). This global interpretability aids in the confirmation that the predictions of the model are made on meaningful patterns and not on irrelevant data, which is vital in healthcare settings where explainable decisions should be drawn. SHAP provides a good explanation basis as the contributions scores are fairly distributed across the features in a manner that encourages reliability of the model. Instead, LIME is interested in local interpretability in that it produces explanations of individual predictions. It estimates the complicated model with a simplified one based on a particular case, indicating which characteristics affected the result of the particular case. Figures 10 and 11 show LIME interpretations of patient cases, showing how the model came to the conclusion on the stroke prediction. These explanations assist medical workers in evaluating the reasoning of the model and finding out whether it matches clinical expectations. The combination of SHAP and LIME ensures both high-level transparency and instance-level clarity, making the model suitable for real-world medical use.

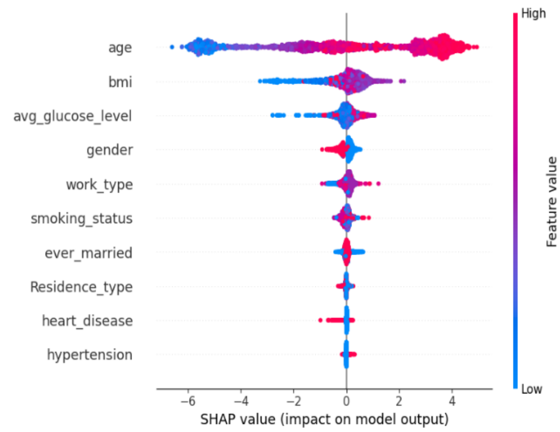

Figure 9: SHAP Summary Plot for Feature Importance

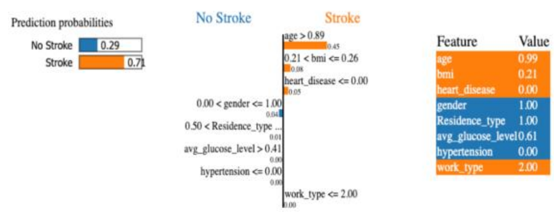

Figure 10: LIME Explanation

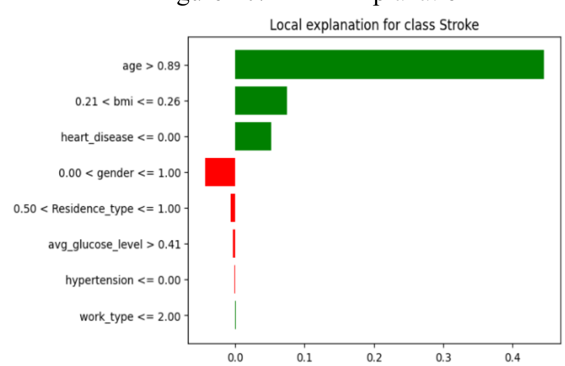

Figure 11: LIME Explanation Plot

4.4 DEPLOYING WITH STREAMLIT WEB APP

To improve accessibility and practical usability, the final stroke prediction model has been deployed as an interactive web application using the Streamlit framework. This user-friendly tool allows clinicians and researchers to input patient-specific details as well as receive real-time stroke risk predictions, making advanced machine learning capabilities more accessible to non-technical users. The program includes explainability features powered by SHAP and LIME in addition to prediction results. These features offer clear visual insights into the significance of each input factor. These justifications promote open decision-making and increase confidence in AIassisted diagnosis. The application's graphical user interface, illustrated in Figure 12, is designed for ease of use, offering a streamlined and informative experience for interpreting results and understanding the underlying rationale behind each prediction.

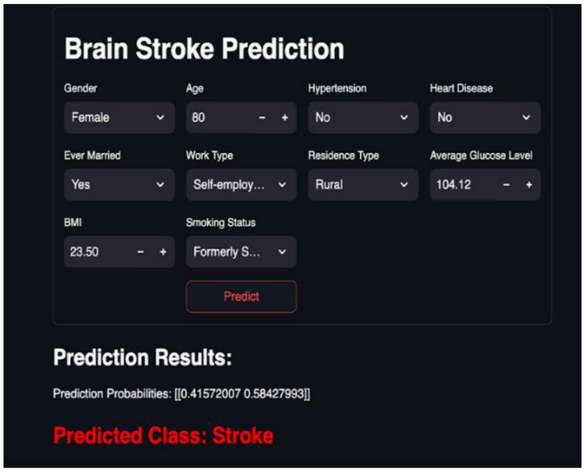

Figure 12: Streamlit Web Application UI for Stroke Prediction

5. DISCUSSION AND FUTURE SCOPE

Using gradient boosting models (GBM, XGBoost, LightGBM, and CatBoost), this work effectively created a brain stroke prediction system that was implemented through Streamlit to provide real-time predictions. Interpretability of the model was enhanced through XAI approaches, including LIME as well as SHAP. Enabling a more thorough comprehension of feature contributions. The models demonstrated high accuracy, with LightGBM and CatBoost performing the best. Age, BMI, and average blood sugar were important predictors that matched recognized medical conclusions. Additionally, methods for preparing data, such as feature encoding, addressing missing values, and class balancing with SMOTE and Tomek Links, improved model generalization and fairness.

However, this research has certain limitations, including a relatively small dataset, limited features (excluding genetic predisposition and lifestyle factors), and class imbalance, which, despite mitigation efforts, remained a challenge. Gradient boosting models are also computationally expensive and need to be deployed in an optimized way. In addition, the model has not been clinically validated, which implies that it cannot currently be applied as an independent diagnostic instrument. These limitations will be critical in improving the reliability and practicality of the model.

To improve in the future, there is a possibility to increase the data size, add more risk factors, and consider the hybrid methods of combining deep learning and gradient boosting. The medical credibility will have to be achieved through clinical validation, i.e., testing in the real world and cooperation with medical experts. Also, the usability can be improved by optimizing the deployment plans with the help of serverless cloud computing and enhancing the accessibility of users with the help of multilingual support and mobile integration. With such improvements, the system will help in timely stroke detection, enhanced patient care, and clinical decision making.

6. CONCLUSION

This research has been able to come up with a Brain Stroke Prediction System based on improved Gradient Boosting Models (GBM, XGBoost, LightGBM, and CatBoost) and make it available as a Streamlit Web App to access in real-time. Integrating the XAI approaches of SHAP and LIME, greater clarity was achieved, and predictions could be explained to clinicians and researchers. The system had a high predictive accuracy, and the LightGBM and CatBoost models were the most accurate. The major risk factors, that involve age, BMI, and average blood sugar levels, were found to significantly predict stroke. In line with the known medical science. The research also indicated the improvement of feature encoding, class imbalance, and data pretreatment in improving the performance of the model. However, limitations such as limitations on the dataset, lack of certain risk variables, and the need for clinical validation indicate that these aspects might be in need of further enhancement. Future plans are to increase the dataset, add new risk factors, make the model more efficient, and incorporate the system into clinical processes. This research advances AI-driven risk assessment, healthcare decision-making, as well as early stroke detection by utilizing explainability and AI.

REFERENCES

[1] K. Mridha, S. Ghimire, J. Shin, A. Aran, M. M. Uddin, and M. F. Mridha, "Automated Stroke Prediction Using Machine Learning: An Explainable and Exploratory Study with a Web Application for Early Intervention," IEEE Access, vol. 11, pp. 52288–52301, 2023.

[2] D. Ushasree, A. V. Praveen Krishna, and C. M. Rao, "Enhanced stroke prediction using stacking methodology (ESPESM) in intelligent sensors for aiding preemptive clinical diagnosis of brain stroke," Measurement: Sensors, vol. 33, p. 101108, 2024.

[3] P. Bathla and R. Kumar, "A hybrid system to predict brain stroke using a combined feature selection and classifier," Intelligent Medicine, vol. 4, pp. 75–82, 2024.

[4] S. Gawde, S. Patil, S. Kumar, P. Kamat, K. Kotecha, and S. Alfarhood, "Explainable Predictive Maintenance of Rotating Machines Using LIME, SHAP, PDP, ICE," IEEE Access, vol. 12, pp. 29345–29356, 2024.

[5] S. Rahman, M. Hasan, and A. K. Sarkar, "Brain Stroke Prediction Using Machine Learning and Deep Learning Techniques," European Journal of Electrical Engineering and Computer Science (EJECE), vol. 7, no. 1, pp. 23–29, Jan. 2023.

[6] M. Katan and A. Luft, "Global Burden of Stroke," Seminars in Neurology, vol. 38, no. 2, pp. 208–211, 2018.

[7] X Xia et al., "Prevalence and risk factors of stroke in the elderly in Northern China: data from the National Stroke Screening Survey," Journal of Neurology, vol. 266, pp. 1449–1458, 2019.

[8] World Health Organization (WHO), "The top 10 causes of death," Fact Sheet, Aug. 2024. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes-of-death.

[9] M. Nazar, M. M. Alam, E. Yafi, and M. M. Su'ud, "A Systematic Review of Human–Computer Interaction and Explainable Artificial Intelligence in Healthcare with Artificial Intelligence Techniques," IEEE Access, vol. 9, pp. 153316–153334, 2021.

[10] T. Elloker and A. J. Rhoda, "The relationship between social support and participation in stroke: A systematic review," African Journal of Disability, vol. 7, pp. 1–9, Oct. 2018.

[11] T. Elloker and A. Rhoda, "The Role of Social Support in Stroke Recovery: A Predictive Model," J. Rehabil. Stud., vol. 15, no. 4, pp. 45-60, Apr. 2018.

[12] Kaggle, "Stroke Prediction Dataset," [Online]. Available: https://www.kaggle.com/datasets/fedesoriano/stroke-prediction-dataset.