The International Society of Data Scientists
Journal of Data Science and Artificial Intelligence
Vol. 5, No. 1. Online ISSN: 2831-4794

Modeling and Application of Explainable Artificial Intelligence for Stroke Prediction

Le Nhi Lam Thuy[a], Trinh Tan Dat[a], Pham The Bao[a,*]
[a]Faculty of Information Science, Saigon University, Ho Chi Minh City, Vietnam

ABSTRACT

Stroke remains a major global health concern, contributing signifi cantly to mortality and long-term disability. Early and accurate prediction can improve patient outcomes; however, traditional machine learning (ML) models often lack transparency. In this study, we develop a stroke prediction framework that combines machine learning algorithms with Explainable Artificial Intelligence (XAI) techniques to improve both predictive performance and model interpretability. By imple menting and comparing three ML algorithms—Random Forest, Support Vector Machine, and Logistic Regression—alongside two XAI methods, SHAP and LIME, this study offers a pathway toward interpretable and trustworthy AI in medical contexts. Experimental results on the held-out test set showed that Random Forest achieved the best performance, with 92.8% accuracy, 95.3% recall, 90.7% precision, 92.9% F1-score, and 92.8% ROC AUC. Furthermore, SHAP analysis identified age, average glucose level, and BMI as the most influential features, while LIME provided instance-level insights into individual predictions. The findings suggest that combining machine learning with explainability techniques can support more transparent stroke risk prediction and may assist clinical decision-making when further validated on larger and more diverse clinical datasets.

Keywords: Stroke Prediction, Explainable Artificial Intelligence, Machine Learning, Model Interpretability

I. INTRODUCTION

Stroke remains one of the leading causes of death and long-term disability worldwide, accounting for a substantial proportion of global mortality [1][2]. Stroke, also known as cerebrovascular accident (CVA), is one of the primary causes of mortality and long-term disability worldwide. This condition typically occurs due to the obstruc tion or rupture of arteries in the brain, resulting in significant impairment or loss of function in the affected brain regions. There are two main types of stroke: ischemic stroke, caused by blood vessel blockage, and hemorrhagic stroke, caused by bleeding within the brain [3]. Therefore, predicting and preventing stroke plays an important role in improving quality of life and reducing the burden on healthcare systems [4]. Artificial intelligence (AI), particularly machine learning (ML), has increasingly been applied to stroke risk prediction using clinical, demographic, and physiological data [5].

However, a major challenge in applying AI in healthcare is ensuring that clinicians can understand, interpret, and trust model predictions. Explainable AI (XAI) addresses this issue by providing clear and interpretable explanations of how predictions are generated by the models [6]. For instance, while models such as Random Forest can yield accurate predictions of stroke risk, explaining the key contributing factors behind these predictions can help physicians better under stand the results and make more informed treatment decisions [7].

Numerous approaches based on machine learning and deep learning have been de veloped and applied for stroke prediction, demonstrating promising results in enhanc ing early diagnosis and risk stratification [8][9][10]. Authors in [11] explored the growing role of artificial intelligence (AI), particularly machine learning (ML) and deep learning (DL), in enhancing ischemic stroke diagnosis and management through imaging anal ysis. They highlighted how AI-driven tools can support timely decision-making by im proving diagnostic accuracy and efficiency in acute stroke triage. They discussed vari ous clinical applications of ML and DL, such as detecting large vessel occlusion, esti mating time since stroke onset, automating ASPECTS scoring, lesion segmentation, and predicting treatment outcomes. The focus is on the integration of AI techniques with imaging modalities like CT and MRI to support stroke care. Dritsas and Trigka [12] presented utilized various machine learning algorithms—including Naive Bayes, Logistic Regression, K-nearest neighbors, Decision Trees, Random Forests, and Mul tilayer Perceptrons for stroke prediction. To address class imbalance, the authors em ployed the Synthetic Minority Over-sampling Technique (SMOTE) and evaluated the models using metrics

* ptbao@sgu.edu.vn

such as accuracy, precision, recall, F-measure, and AUC. The experimental results showed that a stacking ensemble model combining several base classifiers with logistic regression as the meta-classifier achieved the highest perfor mance.

Rahman, Hasan and Sarkar [13] presented a comprehensive investigation into stroke prediction using a combination of machine learning and deep neural network ap proaches. The authors applied a range of classification models including Random For est (RF), XGBoost, LightGBM, AdaBoost, Support Vector Machines (SVM), Logistic Regression, Naive Bayes, and Artificial Neural Networks (ANNs) with three and four layers. After data preprocessing—including handling imbalance via random over sampling, feature scaling, and principal component analysis (PCA)—the models were trained and evaluated using key metrics such as accuracy, precision, recall, F1-score, and AUC. The results demonstrate that the Random Forest classifier achieved the high est accuracy (99%) and AUC (1.0), outperforming all other machine learning methods. Authors in [14] applied machine learning (ML) techniques including Random Forest (RF), Support Vector Machines (SVM), K-Nearest Neighbor (KNN), Logistic Regres sion (LR), Naïve Bayes, Deep Neural Networks (DNN), and Artificial Neural Networks (ANN) to stroke prediction. Cheon, Kim and Lim [15] focused on using deep learning to predict mortality in stroke patients using non-imaging data such as medical service utilization and health behavior records. The authors proposed a deep neural network (DNN) combined with scaled principal component analysis (PCA) to extract and trans form 11 categorical and binary variables into continuous variables for improved model input. Their model achieved an area under the curve (AUC) of 83.48%, outperforming five other machine learning methods including random forest and AdaBoost.

In recent years, the integration of XAI with machine learning has gained significant traction in stroke prediction, aiming to improve both model accuracy and transparency. This trend reflects a growing need for interpretable AI systems in critical healthcare decisions [16][17]. In [18], the authors applied a machine learning-based framework for automated stroke prediction that emphasizes both accuracy and explainability. The au thors compared six ML classifiers and find that more complex models like Random Forest and XGBoost achieve the highest accuracies (up to 91%). To address the chal lenge of interpreting model decisions in healthcare, the study incorporates explainable AI (XAI) techniques to provide both global and local explanations of model outputs. Authors in [19] introduced an explainable machine learning pipeline for predicting stroke using an imbalanced clinical dataset of over 43,000 subjects. The study compares six classifiers and identifies the Multi-Layer Perceptron (MLP) as the best-performing model based on its low false-negative rate (18.6%) and high generalization perfor mance. To enhance transparency, the authors apply SHAP (Shapley Additive Explana tions) to interpret model predictions, revealing that age, BMI, and average glucose level are the most influential risk factors.

In this study, we propose to develop a stroke prediction system using machine learn ing models integrated with Explainable Artificial Intelligence (XAI) techniques. Spe cifically, we evaluate and compares three models—Random Forest, Support Vector Machine, and Logistic Regression—to identify the most suitable model for predicting stroke risk. The selected model is then combined with XAI methods, particularly SHAP and LIME, to provide transparent and interpretable explanations for each prediction. The main contributions of this study are summarized as follows:

- • We develop a machine learning-based stroke prediction framework using clinical, demographic, and physiological features.

  • We compare three commonly used classification algorithms, namely Random Forest, Support Vector Machine, and Logistic Regression, under the same preprocessing and evaluation setting.

  • We apply SHAP and LIME to explain the predictions of the best-performing model, providing both global feature importance and local instance-level interpretation.

II. MACHINE LEARNING METHODS FOR STROKE PREDICTION

2.1. Overview

In this study, we investigate the application of machine learning for stroke prediction. Furthermore, we apply explainable AI (XAI) techniques to perform a post hoc explain ability analysis of the best-performing ML model. Figure 1 shows the flowchart of our stroke prediction system.

First, we apply preprocessing techniques to handle missing values and normaliza tion. We apply the DecisionTreeRegressor from the scikit-learn framework to fill in missing values and use StandardScaler to normalize the data so that the mean is 0 and the standard deviation is 1. Next, the preprocessed data is fed into a machine learning

classification model, which is trained to distinguish between stroke and non-stroke cases. The model then produces an output representing the prediction result. Finally, the system includes an explainability component that helps interpret and understand the model's decision, providing transparency into how the prediction was made. In our study, we apply machine learning algorithms, including Logistic Regression (LR), Sup port Vector Machine (SVM), and Random Forest (RF), for stroke prediction systems.

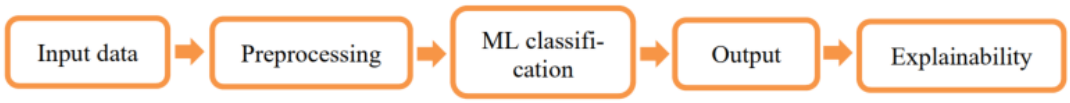


Figure 1. Our process for the stroke prediction system.

2.2. Logistic regresion model

Logistic regression is a statistical technique used for binary classification based on one or more independent variables [20]. In contrast to linear regression, which is applied to predict continuous outcomes in regression problems, logistic regression focuses on es timating the probability that an input falls into a specific class, usually represented as 0 or 1. It employs the logistic (sigmoid) function to convert the predicted values into a range between 0 and 1, which makes it well-suited for classification tasks.

In our system, it is employed to distinguish between patients who are likely to expe rience a stroke and those who are not, based on various input features such as age, hypertension, heart disease, and BMI. The model estimates the probability that a given input belongs to the positive class (stroke) using the sigmoid function.

One of the key advantages of Logistic Regression in this context is its interpretabil ity. The model coefficients provide insights into the relationship between each feature and the output, which can be valuable for medical decision-making. Additionally, Lo gistic Regression is computationally efficient and works well when the relationship be tween the features and the target is approximately linear. However, logistic regression may be limited in modeling complex, non-linear relationships.

2.3. Support vector machine model

Support Vector Machines (SVMs) [21] are a class of supervised learning algorithms widely used for tasks such as computer vision and pattern recognition. The fundamental idea behind SVMs is to find an optimal hyperplane that separates data into two distinct classes with the maximum possible margin.

In the linearly separable case, SVM constructs a linear decision boundary that max imizes the margin between the closest points of each class, known as support vectors. This leads to a convex quadratic optimization problem, which can be solved using the Lagrangian method. The final decision function determines the side of the margin on which new data points fall, classifying them accordingly.

When data is not linearly separable, SVM uses a "soft margin" approach by intro ducing slack variables that allow for some misclassification, controlled by a penalty parameter. This adjustment enables the model to handle noisy and overlapping data.

For non-linear classification problems, SVMs employ a kernel trick to implicitly map the input data into a higher-dimensional space where a linear separator can be applied. Common kernel functions include the linear kernel, polynomial kernel, and radial basis function (RBF) kernel. This method enables SVM to efficiently handle complex deci sion boundaries without explicitly computing the high-dimensional mappings.

2.4. Random Forest Model

A decision tree is a tree-based algorithm commonly used for classification tasks. However, decision trees are prone to overfitting, particularly when the dataset contains noise or outliers. To address these limitations, researchers have explored ensemble methods built on decision trees. Bagging, short for Bootstrap Aggregating, is a widely used en semble technique in machine learning. It enhances prediction stability and reduces var iance by training multiple models on various subsets of the training data, created through repeated sampling with replacement.

Random Forest is an ensemble learning method that used bagging with decision trees during training and outputs the mode of the classes of the individual trees [8, 21]. Each tree in a Random Forest is trained on a random subset of the data and a

random subset of features, which enhances the model's robustness and reduces overfitting compared to individual decision trees.

In stroke classification, Random Forest models have demonstrated strong predictive performance. They are especially useful for identifying complex, non-linear relation ships between input features. Furthermore, Random Forest models provide feature im portance rankings, which can help clinicians understand which factors are most predic tive of stroke occurrence or subtype.

2.5. Explainability

Explainable AI (XAI) [17] is a branch of artificial intelligence focused on developing AI models that can generate systems capable of explaining their decisions and behav iors in a transparent and understandable way to humans. XAI not only aims to improve transparency but also enhances trust and user acceptance of AI systems. It provides techniques, processes, and frameworks that allow people to trust and rely on machine made decisions. With traditional AI, users only receive the final decision after a ma chine learning process. However, with XAI, the reasoning behind the decision can be explained clearly, making the outcome more convincing. This is particularly important in healthcare, where clinicians need to understand not only whether a model predicts a high risk of stroke but also which clinical or demographic factors contribute to that prediction. In this study, we investigate Local Interpretable Model-Agnostic Explanations (LIME) [22] and Shapley Additive Explanations (SHAP) [23] to perform post hoc explainability analysis on the best-per forming machine learning model. These two methods are selected because they provide complementary perspectives: LIME focuses on individual-level explanations, whereas SHAP provides a broader view of feature importance across the dataset.

LIME provides local explanations for machine learning model predictions by con structing a simple, interpretable model around a specific data point or individual case. LIME is model-agnostic, meaning it does not depend on the type of machine learning model being used. This allows it to be applied across a wide range of models. In the context of stroke prediction, LIME can help explain why a specific patient is classified as having a high or low stroke risk by identifying the features that most strongly influence that individual prediction. LIME has been successfully applied in various domains such as healthcare, finance, and other machine learning tasks, demonstrating its versatility and growing popularity. However, LIME has some limitations. It only indirectly addresses the problem of explainability by relying on surrogate models. Therefore, the quality of its explanations largely de pends on how well the surrogate approximates the original model. This may require dense sampling, potentially leading to high computational costs. Additionally, the sam pling process itself introduces uncertainty, which can result in inconsistent behavior and different explanations for the same input sample. Therefore, LIME explanations should be interpreted as useful local approximations rather than exact representations of the internal decision-making process of the original model.

SHAP is a global explainability method that applies game theory to evaluate the contribution of each feature across the entire dataset. It measures the impact of individ ual features on model predictions by assessing how the output changes when each fea ture is removed from the model. This helps users understand the importance of each feature in driving the model's predictions. SHAP provides a way to analyze model pre dictions through an additive approach, meaning that the overall explanation is con structed by summing the contributions of individual features. Compared with LIME, SHAP has a stronger theoretical foundation because it is based on Shapley values, which fairly distribute the prediction outcome among input features according to their marginal contributions. In this study, SHAP is used to identify the most influential risk factors in the Random Forest model, such as age, average glucose level, and BMI. Such global explanations are useful for examining whether the model's behavior is consistent with known clinical risk factors and for improving the interpretability of the prediction system. Both LIME and SHAP are model-agnostic and does not require knowledge of the underlying model structure or objective function, allowing it to be used with a wide range of machine learning models.

III. MACHINE LEARNING METHODS FOR STROKE PREDICTION

3.1. Dataset and Experimental Setup

The dataset used in this study is the Stroke Prediction Dataset provided by Kaggle, which contains 5,110 samples of patients along with 11 features [24]. These include demographic data (age, gender), behavioral information (smoking status), medical his tory (hypertension, heart disease), and physiological attributes (average glucose level, BMI). The target variable indicates whether a patient has had a stroke (stroke = 1) or not (stroke = 0). Table 1 presents a summary of the dataset description. In our study, stroke prediction is approached as a binary classification problem, wherein participants

are assigned to one of two classes: (i) the stroke label, including individuals who have experienced a stroke, and (ii) the non-stroke label, comprising individuals without a confirmed diagnosis of stroke.

Table 1. A summary of the stroke prediction dataset description

Attribute	Description
id	unique identifier
gender	"Male", "Female" or "Other"
age	current age
hypertension	No (0)/Yes (1)
heart_disease	No (0)/Yes (1)
ever_married	No/Yes
work_type	Children/ Government job/ Never_worked/ Private/ Self-employed
Residence_type	Rural/Urban
avg_glucose_level	average glucose level in blood
bmi	body mass index

Table 2 presents a summary of the descriptive statistics for six key numerical attrib utes in the stroke prediction dataset: age, hypertension, heart_disease, avg_glu cose_level, bmi, and stroke. The age variable shows a broad distribution, with values ranging from 0.08 to 82 years and a mean of 43.23, indicating a wide age spread among participants. The prevalence of hypertension and heart disease is relatively low, with only 10% and 5% of patients affected, respectively. This suggests that most of the da taset comprises individuals without these conditions, contributing to class imbalance. The average glucose level exhibits considerable variability, with a mean of 106.15 and a standard deviation of 45.28, reflecting a wide range in blood glucose levels. BMI, recorded for 4,909 individuals, has a mean of 28.89, implying that many participants are in the overweight category; however, the presence of extreme values (ranging from 10.30 to 97.60) may indicate data quality issues or outliers. The target variable, stroke, shows that only 5% of the patients had a stroke, reinforcing the need to address the class imbalance during model development. Overall, the dataset presents diverse patient characteristics and highlights the importance of careful preprocessing and evaluation when building predictive models.

Table 2. Statistical description of stroke dataset in the study

Attribute	Count	Mean	Std	Min	25%	Median	75%	Max
age	5110	43.23	22.61	0.08	25.00	45.00	61.00	82.00
hypertension	5110	0.10	0.30	0.00	0.00	0.00	0.00	1.00
heart_disease	5110	0.05	0.23	0.00	0.00	0.00	0.00	1.00
avg_glucose_level	5110	106.15	45.28	55.12	77.25	91.89	114.09	271.74
bmi	4909	28.89	7.85	10.30	23.50	28.10	33.10	97.60
stroke	5110	0.05	0.22	0.00	0.00	0.00	0.00	1.00

std: standard deviation

In the experimental setup, several data preprocessing steps were implemented to en sure the dataset's quality and suitability for machine learning models. One of the pri mary challenges encountered was the presence of missing values in the BMI attribute. To address this, an imputation method was employed using the DecisionTreeRegressor from the scikit-learn library, which predicts missing values based on patterns learned from the available data. Following imputation, categorical variables were converted into numerical representations suitable for model training. Specifically, categorical fea tures were encoded using label encoding, which transforms textual categories into nu merical values while maintaining their distinct

identities. Lastly, all continuous numer ical features were standardized using the StandardScaler technique. This method re scales the features to have a mean of zero and a standard deviation of one, ensuring uniformity in feature scaling and enhancing model performance and convergence.

The dataset was divided into training and testing subsets using an 80:20 stratified split to preserve the original class distribution. The models were trained on the training set and evaluated on the held-out test set using accuracy, precision, recall, F1-score, and ROC AUC. Three machine learning models—Random Forest, Support Vector Machine, and Lo gistic Regression—were trained for stroke classification. Their performance was as sessed using key metrics: accurac, precision, recall, F1-score, and ROC AUC score. This comprehensive evaluation approach ensured reliable comparison and generaliza tion of the models across varying data splits.

3.2. Results and Discussion

Table 3 shows the performance of three machine learning models—Logistic Regres sion, Support Vector Machine (SVM), and Random Forest—for stroke prediction. The experimenal results shows that Random Forest yields the highest performance across all metrics. It achieves an accuracy of 92.8%, suggesting that the model correctly clas sified a large proportion of both stroke and non-stroke cases. More importantly, it ob tains a recall of 95.3%, which indicates its strong sensitivity to detecting actual stroke cases—a critical requirement in medical diagnostics where false negatives can have severe consequences. Its precision was also high at 90.7%, meaning it produced rela tively few false positives. The F1-score of 92.9%, which balances precision and recall, further confirms that Random Forest offers robust and reliable predictions. The ROC AUC score of 92.8% reflects excellent discriminatory ability between the stroke and non-stroke classes.

The SVM model performs moderately well, with an accuracy of 81.4% and a recall of 89.6%, indicating that it was relatively effective at identifying stroke cases. How ever, its precision was lower (76.9%), suggesting a higher rate of false positives com pared to Random Forest. The F1-score of 82.8% and ROC AUC score of 81.4% demon strate reasonable but not optimal balance and discriminative power.

Logistic Regression yields the lowest performance across all metrics. While it achieves a recall of 85.9%, its precision is only 76.6%, and the F1-score is 81.0%. Its accuracy and ROC AUC score, both at 79.8–79.9%, indicate more limited predictive capability. This model may struggle to capture complex, non-linear relationships in the data that more advanced models like Random Forest can exploit.

Table 3. Comparative performance of machine learning models on the stroke dataset (%)

Model	Accuracy	Recall	Precision	F1- score	ROC AUC score
Logistic Regression	79.8	85.9	76.6	81.0	79.9
SVM	81.4	89.6	76.9	82.8	81.4
Random Forest	92.8	95.3	90.7	92.9	92.8

The performance gap between Random Forest and the other models highlights the advantage of ensemble methods in handling heterogeneous feature spaces and captur ing complex interactions between variables. Its robustness and high recall make it par ticularly suitable for healthcare applications where the cost of misclassification can be significant. These findings support the selection of Random Forest for further interpret ability analysis using XAI methods such as SHAP and LIME.

Figure 2 shows SHAP values indicating the impact of individual features on stroke prediction based on the Random Forest model. Higher SHAP values indicate greater influence on the model's prediction of stroke. The vertical axis lists the input features such as age, BMI, average glucose level, gender, work type, hypertension, and heart disease. The horizontal axis represents the SHAP value, which quantifies the impact of each feature on the model's output. A SHAP value greater than zero indicates that the feature contributes positively to predicting a stroke, while a value less than zero sug gests it contributes to predicting no stroke. The color of each point reflects the actual value of the feature, with darker colors representing higher values and lighter colors indicating lower values. Each dot corresponds to one observation in the dataset.

Among all features, age has the strongest impact on the model's predictions. Higher age values (shown in dark blue) are consistently associated with an increased risk of stroke, while lower ages (lighter shades) tend to reduce the predicted risk. BMI also influences stroke prediction, although its contribution is more moderate. High BMI val ues slightly increase the likelihood of stroke, but many observations still cluster around a neutral impact. Average glucose level shows a widespread in SHAP values. High glucose levels tend to push the model toward predicting a stroke, indicating that ele vated blood sugar is an important risk factor. Conversely, low glucose levels contribute to a lower stroke risk. Gender has a relatively small effect, as its SHAP values are tightly concentrated around zero, suggesting that it does not substantially influence the model's predictions. Work type follows a similar pattern, with little variation in SHAP values and limited predictive power. Hypertension and heart disease show small posi tive contributions to stroke prediction when present. However, most of their SHAP val ues are close to zero, implying that these features are not consistently influential across the dataset. Overall, the plot shows that age and average glucose level are the most influential features in predicting stroke, followed by BMI. The remaining features, in cluding gender, work type, hypertension, and heart disease, play a less significant role.

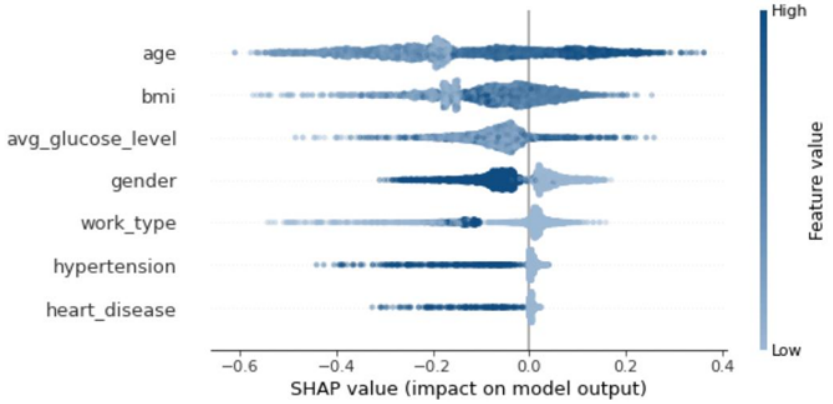


Figure 2. SHAP analysis of factors influencing stroke prediction based on the output of the Ran dom Forest model.

Figure 3 displays the explanation results from LIME for the Random Forest model. The LIME explanation visualizes how different features contributed to the Random Forest model's prediction that this individual is unlikely to have a stroke. The prediction probabilities show a 5% likelihood of stroke and 95% for "Other" (non-stroke), indi cating a strong model confidence in a negative stroke prediction. The feature contributions on the middle bar chart show which inputs pushed the pre diction toward or away from stroke. BMI (between 28.30 and 32.90) had the highest influence toward predicting stroke (0.08), followed by age between 45 and 61 (0.06), gender being 1 (0.05), and the absence of heart disease (0.05), though the latter two slightly reduced stroke likelihood. Low hypertension (0.03) and work type (0.03) also pushed the prediction toward non-stroke, along with relatively high glucose levels hav ing a small impact (0.02). The feature table on the right shows the actual values for the individual: BMI is 29.88, age is 61, heart disease and hypertension are both absent, gender is coded as 1 (likely male or female depending on the dataset), work type is 1, and average glucose level is 202.21. These values align with the feature contributions and reinforce why the model predicted a low stroke risk. In summary, despite risk in dicators like moderately high BMI and glucose level, protective factors such as age just at 61, no heart disease, no hypertension, and other demographic variables led the model to confidently classify this case as non-stroke.

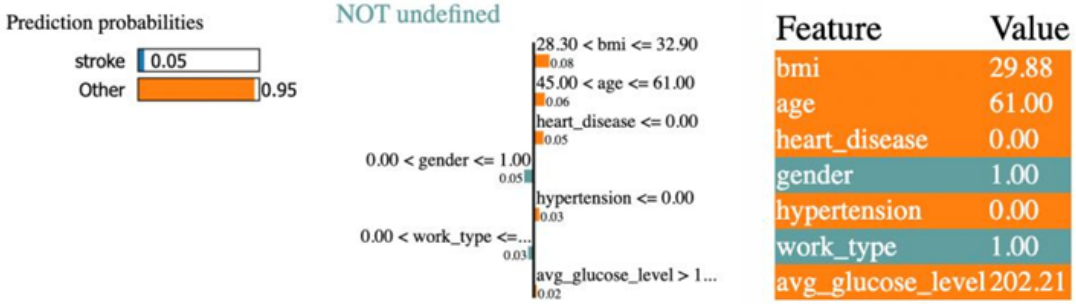


Figure 3. Explanation results from LIME for the Random Forest model.

IV. CONCLUSION

In this study, we investigated the effectiveness of combining machine learning algo rithms with Explainable Artificial Intelligence (XAI) techniques for stroke prediction in a transparent and interpretable manner. Experimental results showed that the Ran dom Forest classifier showed better performance across all metrics compared to Logistic Regression and Support Vector Machine. The integration of SHAP and LIME provided meaningful insights into the model's predictions, with SHAP identifying age, average glucose level, and BMI as the most influential features, and LIME offering instance-level explanations that enhance model transparency. However, the study has certain limitations. The presence of class imbalance, while addressed through basic pre processing, could be better managed with advanced techniques such as SMOTE. Fur thermore, the exclusion of deep learning models leaves unexplored the potential bene fits of more complex architectures in handling high-dimensional or large-scale data. Although the use of a stratified train–test split provides an initial assessment of model performance, future studies should employ repeated cross-validation or external validation to obtain more robust estimates of generalization performance. In addition, future research should consider validating the approach on more diverse datasets, in corporating advanced class-balancing strategies, and comparing results with deep learning-based models. Additionally, developing hybrid explainability systems that combine both global and local interpretations could further support clinical decision making.

Acknowledgments. This research was supported by Saigon University under Fund Grant No. CS.2025.B2.029.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

REFERENCES

[1] M. O. Owolabi, A. G. Thrift, S. Martins, W. Johnson, J. Pandian, F. Abd-Allah, C. Varghese, A. Mahal, J. Yaria, H. T. Phan, G. Roth, S. L. Gall, R. Beare, T. G. Phan, R. Mikulik, B. Norrving, and V. L. Feigin, "The state of stroke services across the globe: Report of World Stroke Organization–World Health Organization surveys," International Journal of Stroke, vol. 16, pp. 889–901, 2021.

[2] V. L. Feigin, M. Brainin, B. Norrving, S. Martins, R. L. Sacco, W. Hacke, M. Fisher, J. Pandian, and P. Lindsay, "World Stroke Organization (WSO): Global Stroke Fact Sheet 2022," International Journal of Stroke, vol. 17, pp. 18–29, 2022.

[3] E. Salvadori, G. Papi, G. Insalata, V. Rinnoci, I. Donnini, M. Martini, C. Falsini, B. Hakiki, A. Romoli, C. Barbato, P. Polcaro, F. Casamorata, C. Macchi, F. Cecchi, and A. Poggesi, "Comparison between ischemic and hemorrhagic strokes in functional outcome at discharge from an intensive rehabilitation hospital," Diagnostics, vol. 11, 2021.

[4] E. S. Donkor, "Stroke in the 21st century: A snapshot of the burden, epidemiology, and quality of life," Stroke Research and Treatment, vol. 2018, Art. no. 3238165, 2018.

[5] P. Sajda, "Machine learning for detection and diagnosis of disease," Annual Review of Biomedical Engineering, vol. 8, pp. 537–565, 2006.

[6] V. Hassija, V. Chamola, A. Mahapatra, A. Singal, D. Goel, K. Huang, S. Scardapane, I. Spinelli, M. Mahmud, and A. Hussain, "Interpreting black-box models: A review on explainable artificial intelligence," Cognitive Computation, vol. 16, pp. 45–74, 2024.

[7] A. Esteva, A. Robicquet, B. Ramsundar, V. Kuleshov, M. DePristo, K. Chou, C. Cui, G. Corrado, S. Thrun, and J. Dean, "A guide to deep learning in healthcare," Nature Medicine, vol. 25, pp. 24–29, 2019.

[8] A. Hassan, S. G. Ahmad, E. U. Munir, I. A. Khan, and N. Ramzan, "Predictive modelling and identification of key risk factors for stroke using machine learning," Scientific Reports, vol. 14, Art. no. 11498, 2024.

[9] L. Schwartz, R. Anteby, E. Klang, and S. Soffer, "Stroke mortality prediction using machine learning: Systematic review," Journal of the Neurological Sciences, vol. 444, Art. no. 120529, 2023.

[10] R. Karakis, K. Gurkahraman, G. D. Mitsis, and M.-H. Boudrias, "Deep learning prediction of motor performance in stroke individuals using neuroimaging data," Journal of Biomedical Informatics, vol. 141, Art. no. 104357, 2023.

[11] A. Cho, L.-N. Do, S. K. Kim, W. Yoon, B. H. Baek, and I. Park, "Application of machine learning and deep learning in imaging of ischemic stroke," Investigative Magnetic Resonance Imaging, vol. 26, pp. 191–199, 2022.

[12] E. Dritsas and M. Trigka, "Stroke risk prediction with machine learning techniques," Sensors, vol. 22, 2022.

[13] S. Rahman, M. Hasan, and A. K. Sarkar, "Prediction of brain stroke using machine learning algorithms and deep neural network techniques," European Journal of Electrical Engineering and Computer Science, vol. 7, pp. 23–30, 2023.

[14] R. Mia, S. Khanam, A. Mahjabeen, N. H. Ovy, D. Ghimire, M.-J. Park, M. I. Begum, and A. S. M. S. Hosen, "Exploring machine learning for predicting cerebral stroke: A study in discovery," Electronics, vol. 13, 2024.

[15] S. Cheon, J. Kim, and J. Lim, "The use of deep learning to predict stroke patient mortality," International Journal of Environmental Research and Public Health, vol. 16, 2019.
[16] R. Alkhanbouli, H. M. A. Almadhaani, F. Alhosani, and M. C. E. Simsekler, "The role of explainable artificial intelligence in disease prediction: A systematic literature review and future research directions," BMC Medical Informatics and Decision Making, vol. 25, Art. no. 110, 2025.
[17] V. Shobana, S. Maheshwari, M. Savithri, S. S. Ramasamy, and N. Kumar, "Integrating explainable machine learning (XAI) in stroke medicine: Opportunities and challenges for early diagnosis and prevention," in 2025 International Conference on Machine Learning and Autonomous Systems (ICMLAS), pp. 461–469, 2025.
[18] 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. 52288–52308, 2023.
[19] C. Kokkotis, G. Giarmatzis, E. Giannakou, S. Moustakidis, T. Tsatalas, D. Tsiptsios, K. Vadikolias, and N. Aggelousis, "An explainable machine learning pipeline for stroke prediction on imbalanced data," Diagnostics, vol. 12, 2022.
[20] D. Dey, M. S. Haque, M. M. Islam, U. I. Aishi, S. S. Shammy, M. S. A. Mayen, S. T. A. Noor, and M. J. Uddin, "The proper application of logistic regression model in complex survey data: A systematic review," BMC Medical Research Methodology, vol. 25, Art. no. 15, 2025.
[21] F. Asadi, M. Rahimi, A. H. Daeechini, and A. Paghe, "The most efficient machine learning algorithms in stroke prediction: A systematic review," Health Science Reports, vol. 7, Art. no. e70062, 2024.
[22] G. Marvin, D. Jjingo, J. Nakatumba-Nabende, and M. G. R. Alam, "Local interpretable model agnostic explanations for online maternal healthcare," in 2023 2nd International Conference on Smart Technologies and Systems for Next Generation Computing (ICSTSN), 2023, pp. 1–6.
[23] Y. Nohara, K. Matsumoto, H. Soejima, and N. Nakashima, "Explanation of machine learning models using Shapley additive explanation and application for real data in hospital," Computer Methods and Programs in Biomedicine, vol. 214, Art. no. 106584, 2022.