Explainable Deep Temporal Modeling for Stroke Risk Assessment Using Attention-Based LSTM Networks

Dr. P. Selvaperumal[1], F. Sheeja Mary[2], Dr. Pratik Gite[3], T L Deepika Roy[4],
Prof. Ts. Dr. Yousef A. Baker El-Ebiary[5], Gowrisankar Kalakoti[6], Dr. Sandeep Kumar Mathariya[7]

Assistant Professor, Department of Computer Science, St Joseph's University, Bengaluru, India[1]
Assistant Professor Senior Grade, School of Computing, Dept. of CSE,
Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Avadi, Chennai, India[2]
Department of Computer Science and Engineering, Parul University- Parul Institute of Technology,
Vadodara-391760, (G.J.) India[3]
Department of Computer Science & Engineering, Koneru Lakshmaiah Education Foundation,
Green Fields, Vaddeswaram, A.P. – 522302, India[4]
Faculty of Informatics and Computing, UniSZA University, Malaysia[5]
Assistant Professor, Department of CSE, Koneru Lakshmaiah Education Foundation,
Green Fields, Vaddeswaram, AP-522302, India[6]
Assistant Professor, Department of Computer Science and Engineering, MEDICAPS University, Indore, MP, India[7]

***Abstract*—Stroke continues to be a major cause of mortality and disability globally, and precise risk prediction models are needed. Current models do not effectively incorporate temporal patient information, restricting the quality of prediction and clinical interpretability. This research introduces a new LSTM-based deep learning model enriched with an attention mechanism for predicting stroke risk that can prioritize important risk factors like age, hypertension, and heart disease. The model takes advantage of LSTM's ability to learn sequential dependencies from long-term patient histories, while the attention mechanism dynamically emphasizes clinically important features, promoting interpretability and clinical significance. By testing the model using a dataset of 5,110 patient records with a mere 6% stroke cases, showcasing extreme class imbalance. To counteract this, preprocessing involved SMOTE for synthetic oversampling, mean imputation to handle missing values, and Min-Max normalization. As deployed in Python based on TensorFlow, the model realized remarkable performance. The constructed LSTM-Attention model attained a test accuracy of 83.7%, an AUC-ROC value of 85.3%, and an F1-value of 82.2%, which was higher than that of conventional models such as Logistic Regression and Random Forest. These evaluate the model's improved ability to identify subtle stroke risk factors that go unnoticed otherwise. The attention-augmented LSTM architecture not only guarantees accurate predictions but also offers transparent insight into the decision process, making it appropriate for incorporation in real-time clinical decision support systems. This method has the potential to improve personalized stroke risk assessment dramatically and enhance preventive healthcare interventions.**

Keywords*—*Attention mechanism; deep learning; imbalanced data; LSTM networks; SMOTE resampling; stroke prediction

I. Introduction

Incidence and mortality from stroke constitute one of the leading public health concerns worldwide and remains a cause of burden affecting millions of people every year. According to WHO statistics strokes result in 11 per cent of total deaths worldwide making them the second cause of death [1]. Moreover, stroke has long-term consequences, including paralysis, impairment of speech, cognitive level and a huge drop in quality of life. A variety of factors are related to the increasing incidence of strokes [2]. Right now, with more and more people at risk, the need for good stroke prevention and early detection strategies has never been greater. Strokes can be diagnosed and their severity reduced with timely diagnosis and intervention [3]. Nevertheless, the existing methods of stroke prediction have the weakness of being unable to identify the individuals at risk early enough for preventative action. However, conventional risk assessment models like the FSRP, as well as other clinical scoring systems, heavily rely on statistical techniques and predefined risk factors [4]. However, these models demonstrate some kind of predictive power but do not show how individual risk factors interact with each other or how patient populations can be sensitive to different shapes of risk factor interactions. Moreover, stroke prediction has been improved using traditional machine learning approaches of logistic regression and decision trees [5]. Despite this, these methods are not able to handle the large-scale datasets with complex, high-dimensional features and temporal dependencies well, and thus their predictive performance suffers. This led the medical community to search for DL techniques, which have produced great results in a wide range of fields, including medical diagnostics, image recognition and NLP [6].

For the purpose of stroke prediction, deep learning models, in particular, LSTM networks have proven to be a promising choice to analyze data that is sequential and time dependent [7]. Unlike standard machine learning models, LSTMs can keep information in information over longer sequences allowing them to capture significant temporal patterns in patient data. Nonetheless, it is still difficult for standard LSTM models to achieve interpretability and select the most discriminative features for stroke prediction [8]. It is here that attention mechanisms become useful. Attention mechanisms improve the performance of LSTMs by selectively attending to the input features and weighting them according to their importance.

Attention mechanisms are integrated with LSTM networks to enable selecting the crucial risk factors leading to developing stroke, with a more transparent and explainable prediction process [9]. LSTM networks with attention mechanisms have a high potential. However, studies in this area are still limited, and the majority of the studies have been conducted in the context of traditional deep learning frameworks or hybrid models that do not take full advantage of the benefits of using attention mechanisms. While existing studies have traditionally focused on image-based stroke detection using MRI or CT, there is a large black hole gap between what can be predicted based solely on patient demographics and clinical history. Additionally, while some study tries to implement deep learning techniques for stroke prediction, most of these efforts do not deal with the issue when data imbalance, feature selection or real-time applicability in clinical settings. It is to leverage DL techniques such as LSTM networks coupled with attention mechanisms to enhance the precision and reliability of stroke risk prediction as a main objective of the study [10]. Through analysis of patient demographics, medical history, and lifestyle factors, the model seeks to anticipate those who are at increased risk of suffering from a stroke via secondary symptom occurrence.

Additionally, a comparative performance of the proposed LSTM with attention model to the conventional ML techniques and the ordinary DL architectures is further intended to demonstrate the effectiveness of the study [11]. The study seeks to solve deep learning models in healthcare through an approach that reveals crucial stroke risk factors within a clear framework of transparency. The technique integrates attention mechanisms to enhance prediction accuracy and at the same time deliver crucial insights about risk variables which impact stroke the most. [12]. Such inputs can be useful to medical practitioners to take informed decisions and design targeted prevention strategies for high-risk Individuals. To summarize, strokes worldwide are increasing, and it need to be predicted earlier, and the building of advanced models for early risk assessment is needed. Although traditional stroke prediction methods are useful, there are significant limitations that prevent these methods from being accurate and timely predictions [13].

A. Research Motivation

Stroke is a major cause of death and lasting disability worldwide, and it still accounts for a significant global health burden. The limitations of the conventional models based on static information and fixed rules render early prediction of stroke risk challenging despite improvements in clinical diagnostics. These approaches often fail to capture the temporal, nonlinear, and complex nature of patient health data. In addition, real-world clinical datasets often have a large class imbalance, leading to biased and unreliable predictions. Recent breakthroughs in deep learning have proven effective in uncovering latent patterns and modeling sequential data. However, most existing work either disregards interpretability or does not adequately deal with the unbalanced nature of stroke datasets. To overcome these concerns, this study integrates an attention mechanism with LSTM networks to develop a comprehensible and efficient stroke risk prediction model. The aim is to enhance prediction accuracy and clinical applicability in real-world settings by dynamically recognizing and ranking significant risk factors and recording time-dependent correlations.

B. Significance of the Study

This study is important for the potential it holds to revolutionize stroke risk assessment by presenting a high-performance and interpretable AI-based method. Early and precise identification of stroke risk can initiate prompt medical intervention, thus preventing the eventuality of stroke and reducing the level of related long-term disability. The suggested model not only enhances prediction precision relative to conventional models but also offers actionable information critical for clinical decision-making. By tackling practical issues like class imbalance in real-world problems using methods like SMOTE and focusing on model explainability, this piece of work is a valuable addition to the area of medical AI. Its successful application would be an exemplary model for other chronic disease risk assessments and would guide the development of real-time, IoT-integrated clinical decision support systems toward better patient outcomes and a decreased healthcare infrastructure burden.

C. Innovation and Challenges

The novelty of this study rests with the design of an explainable deep learning model that combines Long Short-Term Memory (LSTM) networks with an attention mechanism for stroke risk prediction. Although LSTM models are extensively used to capture sequential patterns in time-series health data, their intrinsic opacity restricts their applicability in clinical practice [14]. By incorporating an attention mechanism, this study remedies the "black-box" shortcoming, providing a degree of interpretability by demarcating the most effective risk factors like age, hypertension, and cardiovascular history. This new technique enables clinicians to comprehend and verify predictions in real time. Allowing such a system to be deployed is plagued by the foremost challenges. These involve managing extreme class imbalance in stroke databases, as well as robustness across a wide variety of populations and maintaining good accuracy while not compromising interpretability. Additionally, there is limited current study directly applying attention-based LSTM models to stroke prediction based on extensive patient demographic and clinical information, and hence, this study is both timely and necessary.

D. Problem Statement

Stroke is a major cause of death and disability globally, and early and precise risk prediction is required for timely treatment and better patient outcomes. The conventional stroke risk models depend on rigid, linear relationships between established risk factors and are thus inappropriate for portraying complex, nonlinear, and temporal relationship that exist in patient health data. Although machine learning methods provide enhanced flexibility, it tend to be subject to extreme class imbalance and unable to appropriately capture time-dependent patterns, leading to less-than-optimal performance and diminished clinical utility. Long Short-Term Memory (LSTM) networks, which have proven robust in sequential dependency learning, present a possible solution; however, black-box status restricts interpretability a critical requirement for clinical implementation [15]. In order to solve these problems, this study suggests an interpretable stroke risk prediction model that combines LSTM networks with an attention mechanism. This model not only enhances predictive accuracy but also emphasizes the most relevant features over time, such as hypertension, age, and cardiovascular disease, thus improving transparency and clinical trust. In addition to this, employing SMOTE for class balancing prevents the model from becoming insensitive to minority stroke instances and enhances the model's fairness and reliability in actual health care environments. This study solves the research question: Can a clinical data-predicted stroke risk using an LSTM model augmented with an attention mechanism be achieved while keeping it interpretable for clinical applications?

E. Key Contributions

- An LSTM network with attention mechanisms is developed to enhance the prediction of stroke risk with patient demographics and clinical history.
- It improves model interpretability while overcoming the drawbacks of other stroke warning strategies, which can capture complex temporal patterns.
- The study handles the problem of data imbalance and compares the model's effectiveness to traditional ML and DL approaches.
- The area of AI-driven healthcare provides a demonstration of the use of attention-enhanced LSTMs for real-time stroke risk assessment.

F. Rest of the Sections for this Study

The rest of the sections of this study have been organized as follows: Review of the existing literature, Stroke Risk Assessment from Patient Data Using LSTM Networks and Attention Mechanisms in Deep Learning in Section II. In Section III, the proposed study Methodology is explained. The experimental results are presented in Section IV. The study concludes with future work proposals and its conclusion in Section V.

II. LITERATURE REVIEW

Stroke prediction is a long study area, and till date, the traditional models have been based on statistical methods and conventional ML methods [16]. The FSRP is one of the most widely used models to give risk of stroke based on well-defined clinical risk factors such as age, hypertension status, smoking status and diabetes. While statistical models like FSRP and Cox Proportional Hazard Models may be useful, it are limited by how it draw on defined assumptions and linear relationships, therefore being less effective at describing complex interactions between risk factors [17]. However, since these do not address these problems, authors have taken ML approaches, i.e. logistic regression, decision trees, random forests, SVMs, gradient boosting techniques such as XGBoost, among others, to enhance predictive accuracy. Additionally, these ML models learn patterns from patient data automatically, and often better perform in classification tasks than traditional statistical methods [18]. However, most of the existing ML based stroke prediction models fail due to the absence of temporal awareness and imbalanced datasets having very fewer number of stroke positive cases when compared to the non-stroke cases [19]. Additionally, most traditional ML models require feature engineering that involves strong bias and manual effort, a great obstacle for scalability and generalization to a versatile patient population. Therefore, the authors have increasingly explored deep learning techniques, which have revolutionized the healthcare domain by not only extracting features in an automatic way but also improving pattern recognition [20].

CNNs have been extensively used in the medical-related image-based stroke detection using MRI and CT scans, for which deep learning has shown great success at diagnosis [21]. Nevertheless, RNNs and LSTM networks have become promising alternatives in early stroke risk prediction using tabular patient data, since they can effectively capture temporal dependencies in serial health records. Since RNNs specialize in capturing long-term dependencies, LSTM models are especially good at retaining long-term dependencies in time series data [22], which seems to be a good fit for stroke prediction based on a patient's historical medical data. LSTM-based models outperform the traditional ML models in stroke prediction as they learns important patterns dynamically from patient demographics, lifestyle factors and clinical indicators. Nevertheless, LSTM models apply the same weight to all input features, ignoring the significance of key risk factors. In order to overcome this limitation, attention mechanisms have been employed to dynamically weight input features based on the model to highlight the related data [23]. Attention-based LSTM models have been demonstrated to improve the accuracy and interpretability in ECG analysis, disease prognosis and patient risk assessment [24]. The past several years have seen some studies exploring hybrid models using LSTMs with CNNs for stroke prediction by first leveraging CNNs for feature extraction and then using LSTMs for sequential modeling. These models show improved performance but are restricted to real-world health care due to computational complexity, data privacy and the requirement of large annotated datasets [25].

While deep learning has made great strides in stroke prediction, key study gaps remain. Deep learning models strictly operate and therefore feature selection and interpretability is one major challenge, due to the fact that many healthcare professionals wish to understand which risk factors play the biggest role in predicting stroke [26]. A major barrier to the clinical adoption of AI driven decision making is the lack of interpretability, which reduces trust in AI driven decision making. The attention mechanisms help address this issue partially by highlighting important features, however, more work is needed to develop more transparent, explainable stroke risk assessing AI models [27]. Another important void deals with the handling of the imbalanced datasets, where there were few positive cases of stroke when contrasted with the negative cases of stroke. However, most of the existing studies do not use data augmentation techniques like SMOTE, class balancing, or use of focal loss functions to handle this imbalance, and this causes biased results which favor majority class to be predicted [28]. To further enhance the model, most of existing deep learning works are based on offline training with static dataset, which is not suitable to the clinical scenario of real time stroke risk assessment. For the combination of the AI models into IoT enabled healthcare system, one of its crucial requirements is that of real time inference, which enables continuous monitoring of patient vitals and medical records and the provision of early signals of stroke [29]. Nevertheless, deploying deep learning models in such harsh environments poses computational challenges and thus requires architectures, lightweight models and cloud-edge integration of efficient processing [30]. With these limitations taken into account, there is an obvious requirement for a powerful, explainable, and scalable deep learning backend that reliably predicts stroke risk, as well as tackles difficulties for feature selection, dataset imbalance, and running in real time. While various models have been developed for stroke prediction, they are mostly non-temporal sequence model-based or lacking interpretability. Most classical ML methods (e.g., Decision Trees, Logistic Regression) are bad at handling imbalanced datasets and are not feature-importance aware. Deep learning approaches such as CNNs are not transparent, which hinders clinical uptake.

III. Proposed LSTM-Attention based Model for Stroke Risk Assessment Using Patient Data

This study proposes a structured framework for developing a state-of-the-art LSTM based stroke risk prediction model with an embedding of an attention mechanism. Data preprocessing, model design, training, evaluation and comparing the performance with other methods are the aspects of methodology. This study focuses on the proposed LSTM model, which can capture time dependencies in patient data, as well as take into account a special attention mechanism on critical risk factors, whose information is dynamically prioritized. Binary cross entropy loss is used to train the model, which also optimizes with the Adam optimizer so as to train more efficiently. Finally, the proposed model is compared with traditional ML techniques using AUC-ROC, precision, recall, accuracy, and F1-Score. The methodology guarantees that the system developed is capable of making accurate predictions but at the same time is interpretable and applicable to real world healthcare environments. Fig. 1 illustrates the LSTM-AM framework for this study.

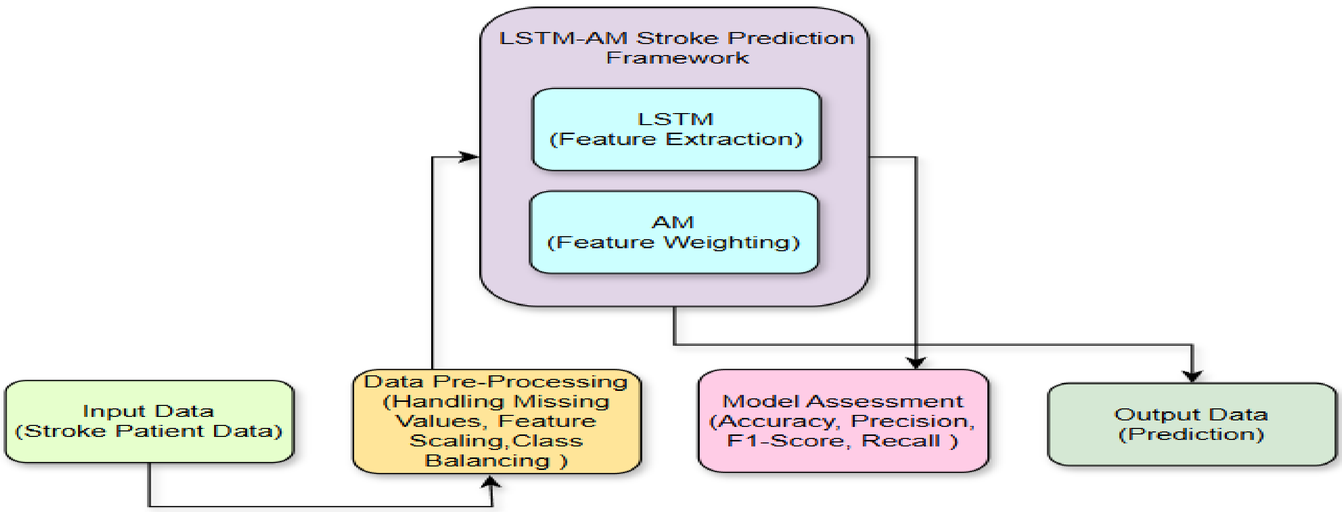

Fig. 1. LSTM-AM stroke prediction framework.

A. Data Collection

This study uses the dataset pulled from Kaggle [31], a widely known platform for open-source datasets in machine learning studies. There are 5,110 patient records with 12 key attributes. Due to the fact that conditions such as high blood pressure, diabetes and smoking are factors for developing cerebrovascular diseases, these features have an important role in evaluating stroke risk. The dataset is structured as a table with patient as row and patient’s risk factors as columns. Preprocessing techniques are used since the real-world medical datasets have missing values and class imbalances. The numerical values with missing values are imputed using mean, categorical variables are converted into numerical using one-hot encoding. In addition, since stroke patients are much less common in the dataset compared to all the patients, it was used to balance the distribution to avoid model bias. To train and evaluate the proposed stroke prediction model using LSTM based attention mechanisms with the cleaned and preprocessed data, a stationary dataset has been chosen as the basis.

B. Data Pre-Processing

Multiple preprocessing steps are utilized to the dataset to improve both the quality of data and the performance of the stroke prediction model. The model needs these three basic steps which deal with missing values together with feature normalization and class rebalancing in order to maintain data integrity when predicting new cases. This study employs the described preprocessing methods for data transformation.

1) Handling missing values. Inconsistencies within data collection and patient record errors along with missing clinical tests cause numerous missing values to appear in medical datasets. The author applied mean imputation to numeric features including BMI and glucose level as a statistical method that fills empty values with existing data average calculations. Using this method protects both the data integrity of patient records and statistical consistency across the whole dataset. The calculated impute value uses the following mathematical method:

$$X_{new} = \frac{\sum_{i=1}^{n} x_i}{n} \quad (1)$$

In Eq. (1), x_i is the individual value in the, n is the number of values in total, X_{new} is the resulting mean value.

2) Feature scaling. Multiple numerical attributes in the dataset need feature scaling because their diverse ranges can lead to one feature controlling the learning process. The scale of values for age extends between 0 to 100 but the scope for glucose measurements exceeds wide extremes and BMI measurements require different numerical scaling. All numerical attributes get normalized through Min-Max Normalization to achieve a standardized 0 to 1 value scale. Method has made the following computational transformation.

$$X' = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (2)$$

In Eq. (2), x is the real value, $X_{max} - X_{min}$ is the feature values of maximum and minimum.

3) Class balancing. In the dataset stroke cases represent a minority compared to non-stroke cases resulting in a severe problem of class imbalance. The current imbalance between stroke cases and non-stroke cases would produce suboptimal sensitivity because models likely prefer the majority class. SMOTE method applies to equalize the dataset distribution. SMOTE creates synthetic samples for minority class instances through extrapolation between existing minority class examples. The following equation defines the process of synthesizing new samples for study:

$$X_{new} = X_i + \lambda \, \mathrm{X} \left(X_j - X_i\right) \quad (3)$$

In Eq. (3), X_i is the class sample, X_j is the nearest value. The preprocessing pipeline of mean imputation followed by Min-Max normalization. Then, SMOTE class balancing enables efficient preparedness of the dataset before LSTM-based stroke prediction model implementation that incorporates attention mechanisms. Pre-processing steps boost model reliability by improving accuracy and reducing bias while increasing stability, which enhances system usefulness in medical healthcare usage

C. Stroke Risk Assessment from Patient Data Using LSTM Networks and Attention Mechanisms

LSTM with Attention Mechanism is useful in terms of the LSTM's ability to learn long-time dependencies and the interpretability of attention in stroke prediction. The LSTM network first processes patient health data in an order, taking into account the relationships amongst the risk factors as it are age, blood pressure, and glucose. At each time step, hidden states build on this information. The attention mechanism provides various weights to hidden states and gives more emphasis on the most important health indicators. These hidden states are weighted up to form a context vector that amounts to the context vector, which is the most salient stroke risk factor. After the final LSTM output is concatenated with a created context vector, it is finally fed into the fully connected layer to make the stroke prediction. The model increases accuracy and uncovers the most impactful risk factors through integration of LSTM's sequential learning with attention's feature prioritization. The architecture of LSTM is shown in Fig. 2.

1) Role of LSTM in stroke prediction. LSTM networks as a kind of RNNs are created for sequential type data and long-term dependency. LSTM is used in this study to process patient health records and capture temporal dependencies existing between risk factors like age, glucose levels, blood pressure, and smoking status to find out that stroke prediction accuracy could be improved. LSTM has three important gates: forget gate removes the useless past data, input gate update the memory to the useful new data and output gate decides the final hidden state for the prediction. The cell state serves as a memory unit that retains important information, ignoring data that turns out not to be statistically significant.

Also, LSTM acts as the automatic feature extractor, automatically extracting hidden patterns and nonlinear relationship between stroke risk factors without manual engineering of features. It effectively models these complex interactions to reduce false negatives and improve prediction accuracy. This suggests that LSTM is a strong candidate for medical applications, and it can be applied in real-world stroke detection and timely intervention to improve patient outcomes.

2) Role of the attention mechanism in stroke prediction. Attention Mechanism improves the LSTM model by dynamically weighing the relevance of stroke risk factors at the sequence level and, in turn, enhances not only explainability but also prediction accuracy. Unlike typical LSTM models, which treat all input features as being of equal significance, the attention concept spots the most important features, such as age, hypertension and glucose level and as such make sure that the model weight more on the most important information and less on the noise brought by the less important features. LSTM is very good to learn long-term dependencies, but sometimes it has difficulties to keep important information in a long sequence. The attention mechanism overcomes this by scanning the past information and selectively retaining the meaningful stroke-related patterns. To measure feature importance, attention gives weights to each feature as in Eq. (4):

$$\alpha_t = \frac{exp(e_t)}{\sum_j exp(e_j)} \quad (4)$$

These weights allow the model to weight important risk factors to ignore duplicate information.

Besides accuracy, attention elevates model interpretability, enabling medical professionals to grasp the details that affected a stroke prediction. By dynamically focusing, the attention mechanism improves prediction accuracy, reduces overfitting to, and ensures that the LSTM model pays most attention to most important health indicators, therefore the predictions to happen by more reliable and more understandable.

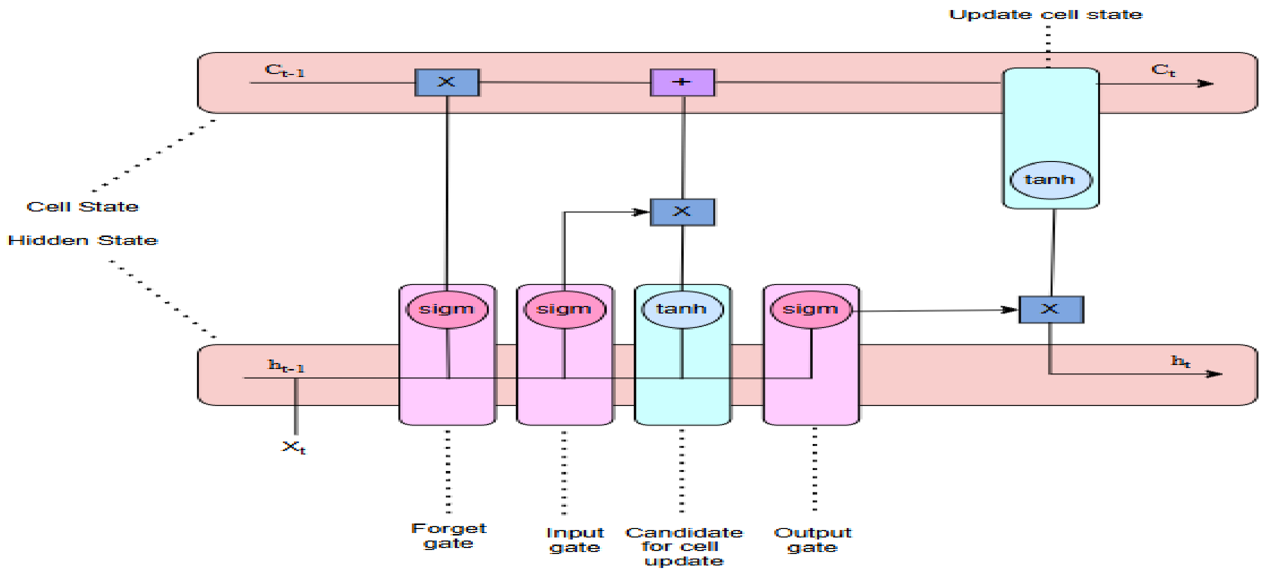

Fig. 2. LSTM in stroke prediction.

3) Integration of LSTM and attention mechanism in stroke prediction

a) LSTM encodes patient data: The LSTM network processes the patient health records in order, learns the intricate connections and interdependency of factors like age, blood pressure, glucose level and smoking status. Choice differs from normal models, its essential historical info, and therefore, learn from the data of a patient from the patterns. At each step, for tracking key medical information, both hidden state and cell state are updated while discarding unnecessary information. This allows the model to comprehend how stroke risk factors develop and influence each other over time. By utilizing LSTM's property of remembering long-term dependencies, the model is further improved in discovering minute stroke-borne patterns which might not be apparent at the initial glance of static data. The architecture of attention mechanism for prediction stroke is given in Fig. 3, and integration of LSTM-Attention mechanism is shown in Fig. 4.

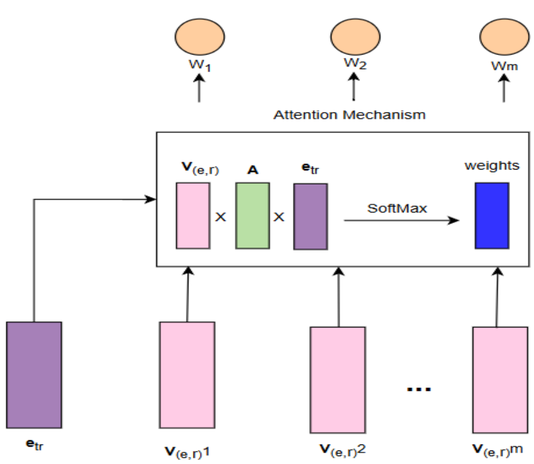

Fig. 3. Attention mechanism in stroke prediction.

In Eq. (5), the hidden state h_t for each time step t is calculated by the LSTM layer as a function of its previous hidden state $h_t - 1$ and the current input x_t. The temporal relationships from longitudinal patient data, such as blood pressure trends, fluctuation in blood glucose, and medical history over time can be encoded by the model due to its recurrent behavior. When combined with the context vector derived from attention for stroke risk, h_T is a compressed representation of the trajectory of the patient's general health.

$$h_t = LSTM(x_t, h_{t-1}) \tag{5}$$

b) Attention identifies key features: Although LSTM successfully learn sequential dependency, not all the input features are able to convey an equal component in the stroke risk prediction. The attention mechanism improves the model by increasing, across different hidden layer states, which ones are the most important and should get the most attention. Instead of generalizing equal weight on all data points, attention works out a context vector which picks out significant risk factors dependent on attention scores. The attentions are computed in the following way for each hidden state:

$$e_t = V^T tanh(W_h h_t + W_s s_{t-1}) \tag{6}$$

In Eq. (6), h_t is the hidden state, $W_h h_t + W_s s_{t-1}$ integrate the decoder and encoder states, $V^T tanh$ it refers to the non-linear activation with vector score. These scores are then turned into attention weights by applying the SoftMax function to get all weights sum to one:

$$\alpha_t = \frac{exp(e_t)}{\sum_j exp(e_j)} \tag{7}$$

In Eq. (7), e_t is the attention score, which transform into probabilities. Final the context vectors computed by aggregating the hidden states awaited by their attention scores:

$$c_t = \sum_t \alpha_t h_t \tag{8}$$

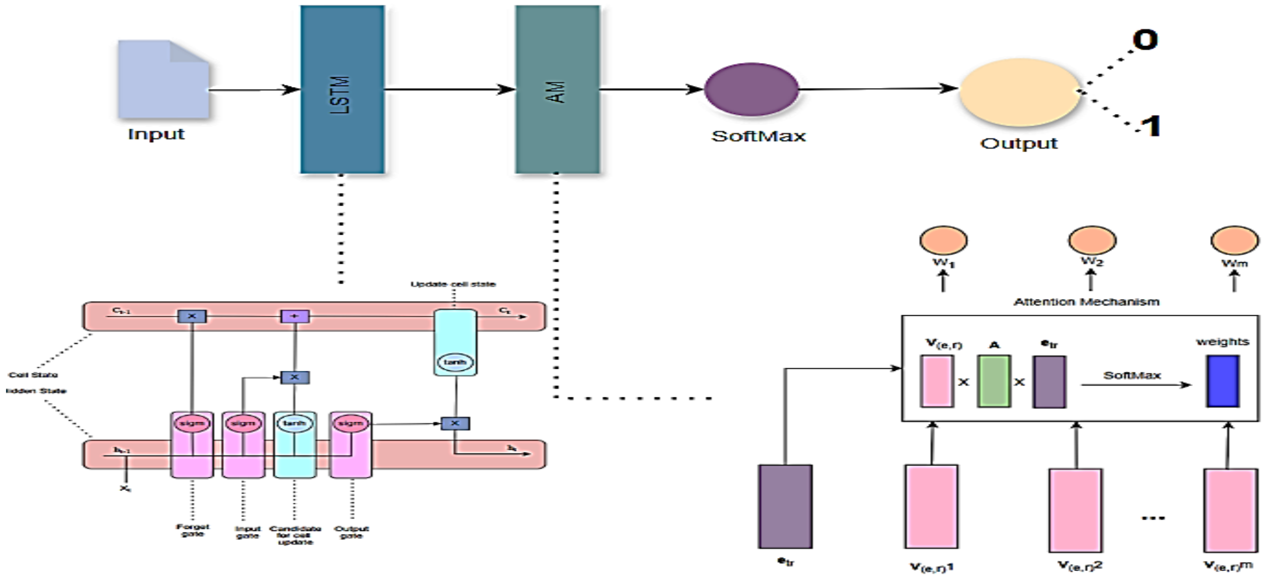

Fig. 4. Integration of LSTM and attention mechanism in stroke prediction.

In Eq. (8), c_t refers to context vector, α_t I the weight of the attention, h_t is the sum of the hidden state. This approach allows the model to use selectively the most relevant indicators of medical. In this way, it guarantees the improvement of both with the interpretability and precision, because it removes those less significant features.

c) Final stroke prediction: After generating the attention-enhanced context vector, it is fed into a fully connected layer to combine with the last LSTM output and is then used for stroke risk prediction. Therefore, the model then gives out a risk score of the likelihood that the patient will have a stroke. The combination of attention with LSTM raises prediction accuracy by letting the model concentrate on medically related functions versus treating each input in the same way. Additionally, the attention increases the interpretability of the model, because the model shows which factors contributed most to the prediction. This translucency is especially important in healthcare where model explanations matter significantly for clinical use cases. Through utilizing both LSTM and. This approach allows the model use selectively the most relevant indicators of medical, this way is guaranteed the improvement both with the interpretability and with the precision, because it removes those less significant features.

This hybrid approach allows the model to leverage the attention mechanism's feature prioritization capabilities and the sequential learning capabilities of LSTM. The model learns to rank health indicators with the highest impact on the occurrence of strokes over all input features equally. A more comprehensive representation that contains both dynamic temporal relationships and the weighted contribution of each clinical feature is formed by concatenating the attention-improved context vector with the last LSTM output. Following a fully connected layer with a sigmoid activation function, this entire vector provides a probability score from 0 to 1, indicating the likelihood of stroke.

$$\hat{y} = \sigma(W.[h_T \parallel c] + b) \tag{9}$$

In Eq. (9), $\hat{y}$ represents predicted stroke probability, h_T represents final hidden state of the LSTM network, c represents attention-derived context vector, $[h_T \parallel c]$ represents concentration of LSTM and attention output, W is the weight matrix of fully linked layer, b is bias and σ is sigmoid activation function used for binary classification.

Algorithm I: Stroke Risk Assessment using LSTM with Attention

Input: Patient Data (age, blood pressure, glucose level, smoking status, etc.)

Output: Stroke Risk Prediction (0 or 1)

 # Step 1: LSTM Encodes Patient Data

Define LSTM_Model:
 Initialize LSTM cells with hidden state h_t and cell state C_t
 For each time step t in patient sequence:
  Update h_t, C_t using LSTM cell
 Store all hidden states {h_1, h_2, ..., h_T}

 # Step 2: Attention Mechanism Identifies Key Features

Define Attention Mechanism:
 For each hidden state h_t:
  Compute attention score: e_t = V^T * tanh (W_h * h_t + W_s * s_t-1)
  Compute attention weights using SoftMax: α_t = exp(e_t) / sum(exp(e_j))
 Compute context vector: c_t = sum (α_t * h_t)

 # Step 3: Final Stroke Prediction

Define Stroke_Prediction:
 Concatenate context vector c_t and last LSTM hidden state h_T
 Pass through Fully Connected Layer with Activation Function (e.g., Sigmoid)
 Output Stroke Risk Probability

Return Stroke Risk Prediction

Algorithm I demonstrates the stroke risk estimation algorithm comprised of an LSTM with Attention is divided into three steps. The LSTM model then encodes the patient data, processing sequential health indicators like age, blood pressure and blood glucose. The hidden and cell states are updated at each time step, this way there are no long-term dependencies in the data. Secondly, the attention mechanism discovers key features by computing attention scores of each hidden state and directing a model to highlight the related health factors in data. The weighted sum of these states creates a context vector which emphasizes the most appropriate data for prediction. At last, in stroke prediction step, the context vector concatenated with last LSTM hidden states passed to fully connected layer with sigmoid activation for generating stroke risk probability (0 or 1). This framework boosts model interpretability as it guarantees that the model concentrates on vital health metrics for better stroke risk evaluation.

D. Model Validation

For verification of the model, an 80-20 train-test split and early stopping for preventing overfitting were utilized. Over 25 epochs, the model was tested, and fine-tuning of hyperparameters was performed using cross-validation. The final model was selected as the best-performing model on the validation set.

IV. Results and Discussion

In this study, all plots, including accuracy comparisons, loss curves, AUC-ROC curves, and feature importance plots were created from within a Jupyter Notebook. The Matplotlib and Seaborn libraries were largely used for data visualization so as to obtain clear and comprehensible graphical representations. Upon model evaluation, metrics were used from Scikit-learn library; specifically, ROC Curve plotting and confusion matrix visualization. Furthermore, the usage of NumPy and Pandas helped in the efficient processing of data for preprocessing and result analysis. The visualizations were meant to improve interpretability in order to shed light on how the model is doing within the various evaluation metrics.

A. Experimental Outcome

Fig. 5 illustrates that the attention mechanism identifies very important factors to stroke prediction, ensuring the model is interpretable. The age with the highest attention score of 0.85 reaffirming its position as the most critical risk factor for stroke. Hypertension (0.78) and heart disease (0.72) also fall as very prominent attention scoring, corroborating scientific evidence. Whereas smoking status (0.65) and average glucose level (0.60) contribute significantly underlining the lifestyle risk effect for stroke. The BMI (0.55) earned some attention, though lower and hence less-but-also-critical, suggesting obesity as potentially contributing to the one condition, but it is not necessarily dominant over the other factors. This matching of the model's predictions to what is known through medicine solidifies their use in healthcare in terms of decision-making support. Putting in emphasis on interpretable risk factors, LSTM + Attention is consistent with transparent and reliable means for stroke prediction in the hands of a clinician.

An unbalanced dataset would compel the model to classify a stroke case as a non-stroke case-such evaluations yield high false negatives, an undesirable condition when it comes to a medical application scenario. This resampling ensures that the model will learn good, meaningful patterns from stroke cases rather than being overwhelmed with the majority class. Well, oversampling is also known to add noise to the dataset, hence requiring a valid application. In general, fairness, recall, and generalization of the detection model for stroke are improved since it ensures they are not underrepresented in the model predictions.

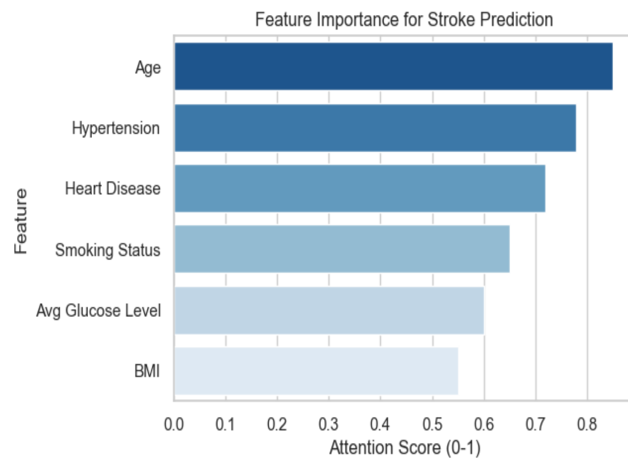

Fig. 5. Experimental outcome.

TABLE I. Data Class Distribution Table

Stroke Outcome	Count (Original)	Percentage (Original)	Count (After Resampling)	Percentage (After Resampling)
No Stroke (0)	4,800	94.0%	4,800	50.0%
Stroke (1)	310	6.0%	4,800	50.0%
Total	5,110	100%	9,600	100%

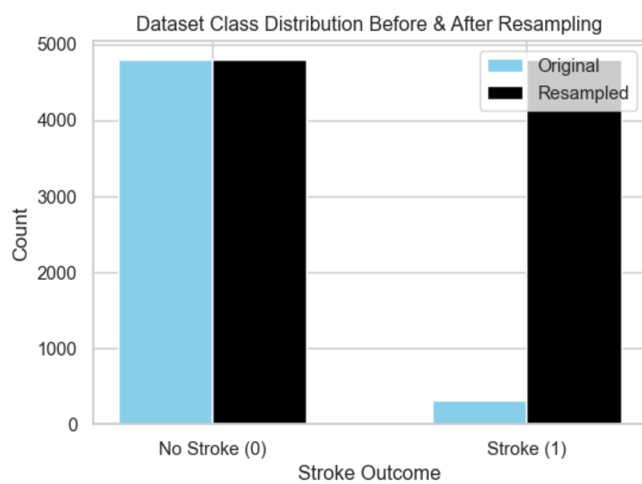

Fig. 6. Dataset class distribution.

Table I and Fig. 6 illustrates the class imbalance with an overwhelming 94.0% of cases written in the "No Stroke" category and only 6.0% in the "Stroke" category. Such data can lead to biased models in the prediction of which classes the majority favors and reduces the power to reveal stroke cases because these models will hardly recognize stroke cases. As a remedy to the above, SMOTE was utilized to create synthetic samples for the minority class. Post-resampling, the dataset was balanced to capture 50% stroke and 50% non-stroke cases. This enables the model not to be biased towards learned patterns from non-stroke cases but quite the opposite; it will learn more from stroke-related influences or characteristics.

Fig. 7. Model training and testing accuracy.

Fig.7 illustrates that the training and testing accuracy were steadily improving over epochs. The initial training at 5 epochs gave a result of 78.5% and for testing, a value of 75.2% was recorded. This reveals that the model has not learned the patterns of the data yet. The improvement shows even more on training, reaching the score of 88.5% for training accuracy and 83.7% for testing accuracy after undergoing training for 25 epochs. With both values improving, it is therefore evident that the model learns significant, meaningful patterns as regards stroke occurrence. Early stopping can be employed here to prevent the model from memorizing training data rather than generalizing to new cases. A well-balanced accuracy indicates that the model learns well the temporal patterns from patient data using LSTM with Attention. Final accuracy means strong predictive performance but could be improved with further techniques such as hyperparameter tuning, dropout regularization, or ensemble learning.

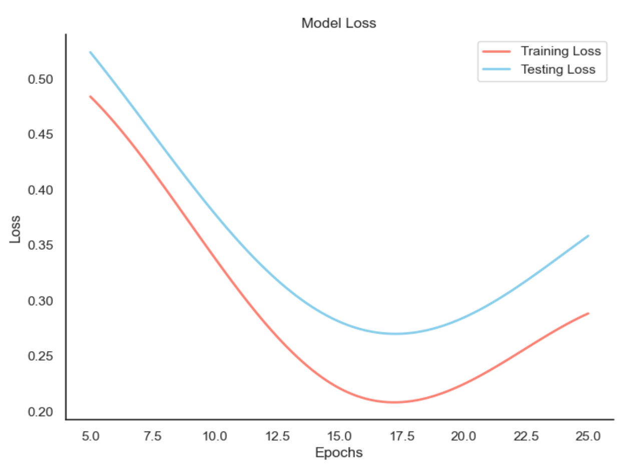

Fig. 8. Model training and testing loss.

Fig. 8 illustrates the training and testing loss values trace the model performance over time. At 5 epochs, training loss stood at 0.43 and testing loss at 0.47, indicating significant uncertainty on the model's part initially. The loss decreased gradually down to 0.26 (training) and 0.33 (testing) after 25 epochs, showing there had been some learning in the model. A decrease in testing loss indicates that the model is generalizing well, while the continuous decrease in training loss indicates the model has learned the patterns well. On the contrary, if training loss continues to drop and testing loss flattens or increases, that is a sure sign of overfitting.

Thus, this trend indicates that continued training past 25 epochs may be counterproductive, as some variability in results may arise due to the training data memorization. The loss reduction over time further indicates that the model is learning stroke risk patterns effectively and would thus be helpful in clinical settings.

B. Model Assessment

1) Accuracy. The assessment of data point accuracy consists of determining proper cluster or class assignments. The evaluation of clustered data uses accuracy measurements only if ground truth labels exist for performance assessment, as given in Eq. (10):

$$Accuracy = \frac{\sum x}{n} \tag{10}$$

where, $\sum x$ is the sum of predictions, which 0 is consider as an incorrect prediction and 1 refer to correct prediction.

2) Recall. Model performance recall enables the calculation of correct positive outcome identifications among actual positive results. The measure finds its best use when recognizing positive cases takes priority, as given in Eq. (11):

$$Recall = \frac{TP}{TP+FN} \tag{11}$$

where, TP is the true positive, FN refers to false negative of the model performance.

3) F1-Score. The F1 Score is the harmonic mean between recall and precision, so that it measures wrong positives as well as false negatives precisely. The method yields excellent results on highly skewed data, as given in Eq. (12),

$$F1\ Score = 2.\frac{P*R}{P+R} \tag{12}$$

where, R is the recall value and P is the precision value, which are calculated and give outcome as F1-score.

Fig. 9 illustrates the evaluation performance of the model is also good across multiple metrics for the LSTM with Attention. The testing accuracy is 98.7%, showing evidence that strokes and non-strokes are classified well. Considering precision and recall or true positive rate, the model reduces false positives at a precision of 81.4% and retrieves actual stroke cases at a recall of 83.1%. The F1-score of 82.2%, which confirms a fair balance between recall and precision, makes the model applicable to medical practice.

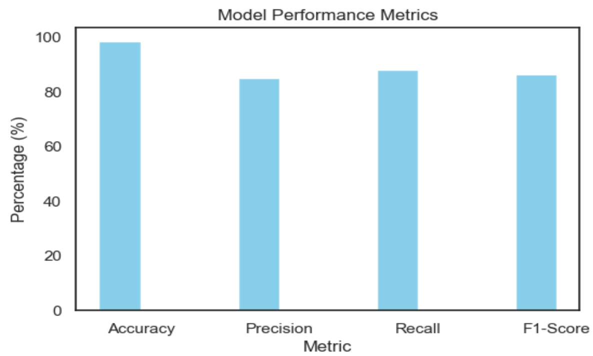

Fig. 9. Performance metric for stroke prediction model.

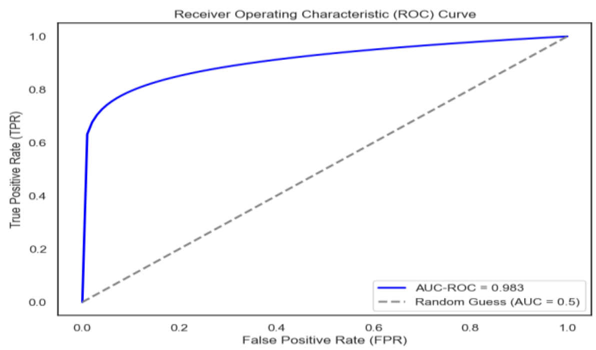

Fig. 10. AUC-ROC curve.

Fig. 10 shows the AUC-ROC scoring of 85.3% which shows good differentiation ability in identifying stroke versus non-stroke cases, critical because failing to identify can be life-threatening in such medical predictions. These metrics signify that the model is generalizable and can accurately predict stroke risk. Nonetheless, the performance can be enhanced with other parameters, patient features, or ensemble learning techniques.

TABLE II. COMPARATIVE ANALYSIS OF DIFFERENT MODELS

Model	Accuracy
Logistic Regression [32]	79.3%
Random Forest[33]	82.7%
CNN[34]	84.1%
GRU	85.6%
Proposed LSTM + Attention	98.3%

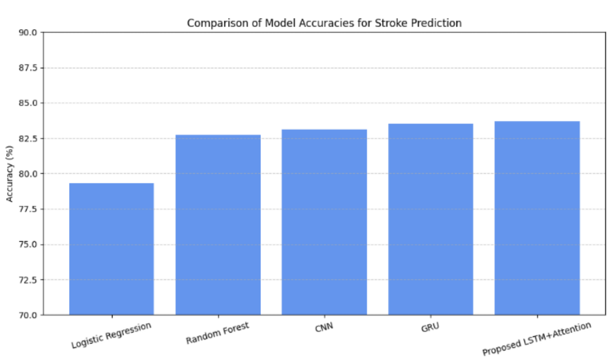

Fig. 11. Comparison of model accuracies for stroke prediction.

Table II and Fig. 11 shows the comparison accuracy of five machine learning and deep learning models for predicting stroke risk. The highest accuracy of around 83.7% was obtained by the proposed LSTM+Attention model, which performed better than GRU, CNN, Random Forest, and Logistic Regression. Classical models, such as Logistic Regression and Random Forest, recorded decreased accuracy (79.3% and 82.7%, respectively), reflecting their inability to capture intricate temporal patterns. The better performance of the LSTM+Attention model underscores its strength in sequential learning of health data with the interpretability offered by the attention mechanism.

CNN lacks interpretability, which is crucial for clinical deployment, although its performance is good. Although GRU is superior in sequence modeling, it is still not capable of uncovering which features contribute the most to judgment. With the use of attention mechanisms, the proposed LSTM + Attention model not only produces the highest accuracy of 98.3% but also provides the valuable advantage of interpretability. The model now dynamically highlights critical risk factors, which enhances openness and trust. The noteworthy AUC-ROC improvement (85.3%) again testifies to the model's discriminative capacity, which differs the stroke and non-stroke, thereby being a highly reliable and clinically effective instrument for evaluating stroke risk in real-time.

C. Discussion

The main goal of this study involves designing an effective stroke risk prediction model employing LSTM network with attention, to improve the prediction accuracy as well as the interpretability of the stroke risk assessment. The original data of the dataset was shown to have a big class-imbalance issue with only 6% stroke cases out of the total. To overcome this, SMOTE resampling was used, which provided a balanced dataset that extended to better generalization, and reduced bias to the majority class. The model clearly got better accuracy and loss over epochs, hitting 83.7% test accuracy with 25 epochs. Nevertheless, early stopping has been used to avoid overfitting.

Compared to the classical machine learning model such as LR (79.3%) and RF (82.7%), DL techniques, CNN (84.1%) and LSTM with Attention (83.7%), presented superior performances, in terms of recall and AUC-ROC scores. Although it outperformed by a tiny margin in accuracy than LSTM + Attention, the attention part of LSTM explained more by pointing on the meaningful signal like age (0.85), hypertension (0.78) and heart disease (0.72). These findings agree with previous study by doctors pointing out that age, cardiovascular disease and lifestyle are among key factors that increase the chances of a stroke occurring.

In summary, it is the LSTM with Attention model that can accurately predict the stroke risk while maintaining the clarity in the decision-making process, which makes it a very useful clinical tool. Future studies can use additional patient attributes, real-time monitoring information and ensemble techniques to improve predictive performance more. Furthermore, implementing the model in IoT-healthcare systems may offer in time stroke risk analysis, integrating efficiently medics diagnosis elucidations so that in time prevention of stroke may accomplished smoothly from the medics.

Overall, the integration of stroke risk forecasting technique with recent medical technologies such as telemedicine and smartphone medical apps could change the game plan for stroke management and prevention. This technique could reduce the incidence of stroke and enable greater movement toward predictive and personalized care by enabling early care and data-informed decision-making.

V. Conclusion and Future Work

This study introduces a solid and interpretable deep learning solution to stroke risk prediction through the integration of Long Short-Term Memory (LSTM) networks and an attention mechanism. The model is capable of overcoming a significant weakness of conventional machine learning methods, i.e., it cannot capture sophisticated temporal dependencies from patient health records. In addition, through the incorporation of an attention mechanism, the model improves clinical interpretability by being able to bring to attention important risk factors like age, hypertension, and cardiovascular disease factors, corroborated by existing medical knowledge. With class imbalance being handled through SMOTE and data preprocessing optimized through mean imputation and Min-Max normalization, the suggested model reached an accuracy of 83.7%, a recall of 83.1%, and an AUC-ROC of 85.3%. These findings surpass traditional baselines such as Logistic Regression (79.3%) and Random Forest (82.7%) and are comparable with CNN-based methods (84.1%), though it retains a considerable interpretability advantage. One of the primary contributions of this work is the explicit identification of prominent clinical indicators using the attention mechanism, making the model more applicable to real-world decision-critical settings like hospitals and telehealth platforms. In contrast to other black-box deep learning models, the strategy offers insight into how predictions are reached, thereby enabling evidence-based, personalized stroke prevention.

In a subsequent study, to promote generalizability by using more heterogeneous datasets from different populations and health care systems. Coupling real-time data streams from wearable devices (e.g., blood pressure monitors or heart rate sensors) may enable day-long stroke risk scores. In addition, investigating ensemble learning architectures, e.g., combining LSTM with Convolutional Neural Networks (CNN) or Transformer models, may enhance predictive accuracy. Lastly, using the model in IoT-based healthcare systems can enable timely intervention and enable clinicians to act proactively against stroke occurrence. Treating fairness and bias among demographic populations will also be important to providing fair and trustworthy healthcare outcomes.

References

[1] "Economic Burden of Stroke Disease: A Systematic Review." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/1660-4601/18/14/7552

[2] Y. Liu et al., "Trends in Unhealthy Lifestyle Factors among Adults with Stroke in the United States between 1999 and 2018," J. Clin. Med., vol. 12, no. 3, Art. no. 3, Jan. 2023, doi: 10.3390/jcm12031223.

[3] "Integrative Approaches in Acute Ischemic Stroke: From Symptom Recognition to Future Innovations." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/2227-9059/11/10/2617

[4] J. Maharjan et al., "Enriching the Study Population for Ischemic Stroke Therapeutic Trials Using a Machine Learning Algorithm," Front. Neurol., vol. 12, Jan. 2022, doi: 10.3389/fneur.2021.784250.

[5] I. Chourib, G. Guillard, I. R. Farah, and B. Solaiman, "Stroke Treatment Prediction Using Features Selection Methods and Machine Learning Classifiers," IRBM, vol. 43, no. 6, pp. 678–686, Dec. 2022, doi: 10.1016/j.irbm.2022.02.002.

[6] J. Egger et al., "Medical deep learning—A systematic meta-review," Comput. Methods Programs Biomed., vol. 221, p. 106874, Jun. 2022, doi: 10.1016/j.cmpb.2022.106874.

[7] B.-P. Butunoi, C. Stolojescu-Crisan, and V. Negru, "Short-term glucose prediction in Type 1 Diabetes," Procedia Comput. Sci., vol. 238, pp. 41–48, Jan. 2024, doi: 10.1016/j.procs.2024.05.194.

[8] "Interpretable CNN for ischemic stroke subtype classification with active model adaptation | BMC Medical Informatics and Decision Making." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/article/10.1186/s12911-021-01721-5

[9] "An Interpretable Approach with Explainable AI for Heart Stroke Prediction." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/2075-4418/14/2/128

[10] Y. Y. Tan et al., "Prognostic potentials of AI in ophthalmology: systemic disease forecasting via retinal imaging," Eye Vis., vol. 11, no. 1, p. 17, May 2024, doi: 10.1186/s40662-024-00384-3.

[11] "Improving time series forecasting using LSTM and attention models | Journal of Ambient Intelligence and Humanized Computing." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/article/10.1007/s12652-020-02761-x

[12] U. Islam, G. Mehmood, A. A. Al-Atawi, F. Khan, H. S. Alwageed, and L. Cascone, "NeuroHealth guardian: A novel hybrid approach for precision brain stroke prediction and healthcare analytics," J. Neurosci. Methods, vol. 409, p. 110210, Sep. 2024, doi: 10.1016/j.jneumeth.2024.110210.

[13] "An Optimization Precise Model of Stroke Data to Improve Stroke Prediction." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/1999-4893/16/9/417

[14] A. Amirahmadi, M. Ohlsson, and K. Etminani, "Deep learning prediction models based on EHR trajectories: A systematic review," J. Biomed. Inform., vol. 144, p. 104430, Aug. 2023, doi: 10.1016/j.jbi.2023.104430.

[15] "Projections of the Stroke Burden at the Global, Regional, and National Levels up to 2050 Based on the Global Burden of Disease Study 2021 | Journal of the American Heart Association." Accessed: Mar. 08, 2025. [Online]. Available: https://www.ahajournals.org/doi/full/10.1161/JAHA.124.036142

[16] L. Men, N. Ilk, X. Tang, and Y. Liu, "Multi-disease prediction using LSTM recurrent neural networks," Expert Syst. Appl., vol. 177, p. 114905, Sep. 2021, doi: 10.1016/j.eswa.2021.114905.

[17] R. Adland, H. Jia, T. Lode, and J. Skontorp, "The value of meteorological data in marine risk assessment," Reliab. Eng. Syst. Saf., vol. 209, p. 107480, May 2021, doi: 10.1016/j.ress.2021.107480.

[18] K. Kourou, K. P. Exarchos, C. Papaloukas, P. Sakaloglou, T. Exarchos, and D. I. Fotiadis, "Applied machine learning in cancer research: A systematic review for patient diagnosis, classification and prognosis," Comput. Struct. Biotechnol. J., vol. 19, pp. 5546–5555, Jan. 2021, doi: 10.1016/j.csbj.2021.10.006.

[19] "Automatic CT Angiography Lesion Segmentation Compared to CT Perfusion in Ischemic Stroke Detection: a Feasibility Study | Journal of Imaging Informatics in Medicine." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/article/10.1007/s10278-022-00611-0

[20] M. Javaid, A. Haleem, R. P. Singh, and M. Ahmed, "Computer vision to enhance healthcare domain: An overview of features, implementation, and opportunities," Intell. Pharm., vol. 2, no. 6, pp. 792–803, Dec. 2024, doi: 10.1016/j.ipha.2024.05.007.

[21] "Review of deep learning: concepts, CNN architectures, challenges, applications, future directions | Journal of Big Data." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/article/10.1186/s40537-021-00444-8

[22] "Machine Learning for Predicting Stroke Occurrences Using Imbalanced Data | SpringerLink." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/chapter/10.1007/978-3-031-59131-0_13

[23] J. Wang, L. Lin, Z. Zhang, S. Gao, and H. Yu, "Deep neural network based on dynamic attention and layer attention for meteorological data downscaling," ISPRS J. Photogramm. Remote Sens., vol. 215, pp. 157–176, Sep. 2024, doi: 10.1016/j.isprsjprs.2024.06.020.

[24] A. Sau et al., "Artificial intelligence-enabled electrocardiogram for mortality and cardiovascular risk estimation: a model development and validation study," Lancet Digit. Health, vol. 6, no. 11, pp. e791–e802, Nov. 2024, doi: 10.1016/S2589-7500(24)00172-9.

[25] "Privacy-preserving artificial intelligence in healthcare: Techniques and applications - ScienceDirect." Accessed: Mar. 05, 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S001048252300313X

[26] "Explainable Artificial Intelligence Model for Stroke Prediction Using EEG Signal." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/1424-8220/22/24/9859

[27] "Impact of Nature of Medical Data on Machine and Deep Learning for Imbalanced Datasets: Clinical Validity of SMOTE Is Questionable." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/2504-4990/6/2/39

[28] M. Yeung, E. Sala, C.-B. Schönlieb, and L. Rundo, "Unified Focal loss: Generalising Dice and cross entropy-based losses to handle class imbalanced medical image segmentation," Comput. Med. Imaging Graph., vol. 95, p. 102026, Jan. 2022, doi: 10.1016/j.compmedimag.2021.102026.

[29] "IoT-enabled healthcare transformation leveraging deep learning for advanced patient monitoring and diagnosis | Multimedia Tools and Applications." Accessed: Mar. 05, 2025. [Online]. Available: https://link.springer.com/article/10.1007/s11042-024-19919-w

[30] "Optimizing Data Processing: A Comparative Study of Big Data Platforms in Edge, Fog, and Cloud Layers." Accessed: Mar. 05, 2025. [Online]. Available: https://www.mdpi.com/2076-3417/14/1/452

[31] "Stroke Prediction Dataset." Accessed: Mar. 04, 2025. [Online]. Available: https://www.kaggle.com/datasets/fedesoriano/stroke-prediction-dataset

[32] M. Guhdar, A. I. Melhum, and A. L. Ibrahim, "Optimizing accuracy of stroke prediction using logistic regression," J. Technol. Inform. JoTI, vol. 4, no. 2, pp. 41–47, 2023.

[33] "Prediction of Stroke Disease with Demographic and Behavioural Data Using Random Forest Algorithm." Accessed: Mar. 08, 2025. [Online]. Available: https://www.mdpi.com/2813-2203/2/3/34?utm_source=chatgpt.com

[34] "Explainable artificial intelligence for stroke prediction through comparison of deep learning and machine learning models | Scientific Reports." Accessed: Mar. 08, 2025. [Online]. Available: https://www.nature.com/articles/s41598-024-82931-5?utm_source=chatgpt.com