Prediction of Brain Stroke using 7–Layer ANN

J. Tejesh [1], Dr. K. Chaitanya[2], M. Madhavi[3], B. Kalyan Siva Rama Krishna[4], K. Krishna Rao [5]

Assistant Professor[2], Department of CSE, University College of Engineering and Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, AP, India

Email: anu.konda.chaitanya@gmail.com[2]

UG Students[1345], Department of CSE, University College of Engineering and Technology, Acharya Nagarjuna University, Nagarjuna Nagar, Guntur, AP, India

Email: tejeshjamalapurapu23@gmail.com[1] , madhavimarri98@gmail.com[3] , bkalyansivaramakrishna@gmail.com[4] , epqqwppw@gmail.com [5]

---***---

Abstract - Stroke is a significant global health issue necessitating early identification to enhance outcomes and lower healthcare expenditures. This paper suggests an Artificial Neural Network (ANN)-based predictive model with clinical and demographical attributes. Preprocessing of the data involves imputation, encoding, standardization, and PCA-based dimension reduction. A class imbalance, SMOTEENN, combination resampling strategy is used to counter the issue of class imbalance. The ANN uses RELU activation, dropout, and L2 regularization and is optimized using Adam with a learning rate schedule. The model shows 93.74% test accuracy, although precision and recall for stroke are low due to the issue of class imbalance. A GUI is integrated for ease of use, allowing real-time prediction of risk for clinical applications. The results indicate the applicability of deep learning in stroke prediction and propose future enhancement to enhance sensitivity to rare occurrences.

Keywords: Stroke Prediction, Deep Learning, Meta-Classifier, Neural Networks, Medical Diagnosis, Health Informatics, Ensemble Learning, Classification.

1. INTRODUCTION

Stroke is a leading global health issue that has always been identified as the second most frequent cause of death and the third most disabling disease cause globally [1]. Early diagnosis and treatment of stroke are important to its mitigation in severity, ease of survival of the patient, and restoration of quality of life after a stroke attack. Precise prediction of stroke, however, is a wily clinical challenge due to the presence of numerous interconnected risk factors such as age, sex, lifestyle, previous medical illness, and genetic susceptibility.

Stroke diagnosis has traditionally been based mainly on clinical evaluation and imaging techniques. Advances in computed tomography (CT) and magnetic resonance imaging (MRI) in the 1970s greatly enhanced diagnosis. These technologies, nonetheless, only allow for diagnosis after the appearance of symptoms and little assistance in quantifying stroke risk during asymptomatic intervals. Identification of modifiable risk factors like hypertension and diabetes in the middle of the 20th century brought more active screening strategies, but successful early detection remains restricted by resource constraints and the requirement for timely clinical decision-making.

As healthcare has progressed, so too has the role of data in facilitating medical decision-making. In early predictive approaches, statistical techniques and rule-based systems were employed; however, these approaches tended to incorporate a limited number of clinical variables, hence ignoring the multifactorial, complex nature of stroke.

Today, risk assessment procedures are usually a manual review of patient history and clinical parameters that are not only time-consuming but are also error-prone due to human involvement.

Improved availability of health data and the ability to develop sophisticated analytical tools have brought new methods of stroke risk assessment into play. It has allowed for the development of customized and scalable prevention and early detection strategies. With the ongoing worldwide burden of stroke, improvement of predictive capability remains an urgent objective in patient outcome optimization and reduction of burden on healthcare infrastructure.

Recent evidence has demonstrated the effectiveness of machine learning (ML) and deep learning (DL) models in managing large volumes of heterogeneous clinical data, allowing for automated and improved prediction of stroke outcomes (Liu et al., 2020) [2]. Innovative hybrid schemes that couple ML and DL, especially CNNs, have evidenced potential for classification of strokes based on medical image data (Litjens et al., 2017) [3].

Moreover, multimodal data fusion techniques—aggregating clinical, imaging, and biosignal data—have greatly improved prediction performance (Zhang et al., 2021) [4].

All these breakthroughs enable real-time, pre-symptomatic detection and bridge the issues related to both accuracy and timing shortcomings of standard diagnosis systems. In addition, transfer learning and ensemble learning techniques have aided in addressing challenges by limited datasets, particularly in resource-limited environments (Pan & Yang, 2010) [5].

Most models achieve high accuracy on benchmark sets but perform poorly with generalization to actual clinical cases due to limited data, no external validation, and limited interpretability. For instance, some are dependent mostly on feature selection or simulation-based data sets, which could not represent the variability of real clinical data. Others emphasize the requirement for better computational efficiency, particularly in real-time diagnosis situations. These issues collectively form the motivation for the present project, which seeks to build a DL model that integrates clinical and imaging data for stroke prediction, while also addressing concerns related to data imbalance, generalization, and interpretability—key drawbacks identified in prior research.

2. LITERATURE REVIEW

There has been an upsurge in research in recent years aimed at developing good-quality models for interactive image editing. The review that follows presents an overview of some existing works in this field.

Salucci et al. (2020) created a microwave imaging and machine learning-based system for real-time stroke detection in the brain using simulated phantom data. Their system demonstrated promising potential for stroke detection with high temporal resolution. Nevertheless, the study recognized limitations in model realism and data diversity, which could affect the system's generalizability to real medical data [1].

Ozaltin et al. (2021) proposed OzNet, a new deep learning architecture for the classification of stroke in brain CT images. They utilized mRMR-based feature selection with hybrid models to reach an impressive 98.42% accuracy. Though the model achieved such a high accuracy, they reported that feature selection was the key aspect for performance and the generalizability of the model across different populations and imaging scenarios remained a question [2].

Chin et al. (2020) suggested a CNN-based system to detect early ischemic stroke from CT image patches. The system showed high performance with 97.65% training accuracy and 92.97% testing accuracy.But the scarcity of high-quality labeled data was a major problem in applying the system in actual clinical environment [3].

Al-Mekhlafi et al. (2022) built hybrid deep learning (AlexNet) and machine learning (SVM) models for stroke and hemorrhage detection with MRI and tabular data. Their model was able to achieve as high as 99.9% accuracy, with the hybrid model proving especially adept at handling sophisticated data sources. But the authors indicated that it was still a challenge to merge deep learning with the conventional machine learning methods with regards to interpretability and training effectiveness [4].

Surya et al. (2022) investigated a range of deep models like CNN and DNN for stroke identification based on CT, MRI, EEG, and EOG information.Although the models were highly accurate, the authors highlighted issues with real-time MRI-based diagnosis, especially concerning MRI scan variability and the requirement for greater computational power during real-time processing [5].

Rahman et al. (2021) performed a comparison study between ML and DNN models based on a highly imbalanced healthcare dataset. They showed that Random Forest worked best, with 99% accuracy following oversampling of stroke cases. This research emphasized the need to solve class imbalance in healthcare data in order to enhance stroke prediction accuracy [6].

Roohi et al. (2021) developed a UWB-MIS microwave imaging system with machine learning (ML) post-processing for hemorrhagic stroke detection. The system performed 97% accurately on phantom-based data, although the authors indicated that clinical data in everyday life may pose some extra challenges in terms of noise in the signals and variability between patients, impacting performance [7].

Fernandes et al. (2023) gave an overview of stroke detection methods, with a comprehensive analysis of Imagen Editor, a text-to-image inpainting model. Although mostly an image editing tool, the research highlighted its applicability in medical image editing, citing its inability to process abstract commands and the necessity for more expert models to address medical image intricacies [8].

Tursynova et al. (2022) used CNNs to classify brain strokes using CT and MRI-based public datasets. Their findings indicated that the model was able to reach more than 80% accuracy with CT image input, but the authors emphasized the requirement for higher quality MRI datasets to enhance model performance for MRI scans[9].

Saleem et al. (2023) presented SDEdit, a guided image synthesis framework trained on datasets such as ImageNet and CelebA-HQ. The model showed promise in image compositing tasks, but its restricted use in medical fields indicated the requirement for a deeper investigation in healthcare-related tasks, such as stroke segmentation and classification[10].

Alon & Dehkharghani (2021) applied DNNs for broadband microwave-based localization and stroke classification, and they achieved 94% accuracy with accurate localization in simulation environments. While the method was promising, the authors wrote that localization in real clinical settings was highly challenging owing to real-world complexity of tissue environments[11].

Srinivas & Mosiganti (2021) created a healthcare tabular data-based stroke prediction ensemble ML model. After carrying out data balancing and preprocessing to a vast extent, the model attained a 96.88% accuracy.This work emphasized the relevance of preprocessing and data quality towards enhancing the robustness of stroke prediction models [12].

Chaki & Woźniak (2021) performed a survey of deep learning (DL)/AI-based stroke detection and rehabilitation, highlighting the superior performance of DL models on CT, MRI, and clinical data sets. They concluded that although DL models had remarkable performance, external validation and real-world deployment were still significant hurdles [13].

Abbasi et al. (2022) compared different deep learning models (CNN, U-Net) for ischemic stroke segmentation on ISLES and small data sets. They highlighted the importance of larger data sets and external validation to increase model robustness and generalizability, pointing out that data scarcity was an important obstacle to broader adoption [14].

Kousar et al. (2021) reviewed the use of deep learning in ischemic stroke tasks, from datasets such as ISLES and ATLAS. They suggested employing transfer learning in order to surpass data constraints and enhance model precision in low-resource settings.The research emphasized the promise of transfer learning to transfer models learned on big datasets to smaller, domain-specific datasets [15].

Babutain et al. (2023) introduced a 5S-CNN model for ischemic stroke detection from CT slices with 90.51% slice-wise accuracy and high sensitivity and specificity using 5-fold validation.This model had great potential for real-time stroke detection in clinical settings but needed to be optimized further for generalization to different types of strokes [16].

Uppal, M et al. (2022) suggested a deep learning model for early diagnosis of stroke to improve accuracy and minimize misdiagnosis.The model uses multi-layer perceptron (MLP) and optimizers like Adadelta, RMSProp, and AdaMax. Although certain dataset information was not given, the research included primary preprocessing steps such as cleaning, missing value handling, and balancing data.

The study's most significant performance measures showed that RMSProp had the best test accuracy (94.78%), whereas AdaMax showed impressive training loss decrease. Nevertheless, issues like data biases and generalizability of the model to various populations are left open for future validation [17].

Choi, Y. A et al. (2021) centered on creating a stroke prediction system based on deep learning using real-time biosignals. The data were time series data from sensors, with a 5-fold cross-validation to provide strong results.The investigation found the performance of CNN-Bidirectional LSTM models to be effective in combining spatial and temporal features.The accuracy ranged from a significant 70.1% (raw data, LSTM) to 94.0% (CNN-Bidirectional LSTM using raw data), with the best predictive accuracy of 89.2% using relative values [18].

Tanveer et al. (2022) presented Neuro-VGNB, a transfer learning-based method for the detection of brain stroke based on CT scan images. The paper used two datasets: one for binary classification (ischemic vs. hemorrhagic) and another with three classes (ischemic, hemorrhagic, normal).

The accuracy was 98% for Dataset 1 and 92% for Dataset 2. The k-fold validation had accuracies of more than 99.92%, with a standard deviation of 0.0000, indicating high reliability. Although these were impressive findings, the authors pointed out that the models performed worse when tested on datasets with different class distributions, which implied a lack of good generalization across different data [19].

Lee et al. (2022) built machine learning models to detect stroke onset time within 4.5 hours based on MRI features, specifically DWI and FLAIR images. The research worked on data from 355 patients using a 299-patient training set and a 56-patient test set. The research identified that ML models exhibited considerably greater sensitivity (~72-76%) in the detection of strokes.within the vital 4.5-hour time frame as opposed to human visual evaluations (~48%).

The accuracy of the ML models (~82.6%) was similar to human evaluations, but the study added that ROC analysis indicated that there was no significant difference in AUC between the models. Human assessment interobserver agreement was moderate (κ=0.43), indicative of the inconsistency in human judgment [20].

3. PROPOSED WORK

The suggested platform presents a consistent and smart stroke risk prediction solution based on deep learning methods. The system takes advantage of clinical and demographic information like age, gender, hypertension status, heart disease history, BMI, smoking status, and work type to predict.

The risk of a stroke. The forecasting model is designed based on a deep Artificial Neural Network (ANN) structure, which feeds the input features through several hidden layers using ReLU activations, dropout, and L2 regularization to enhance generalization and avoid overfitting. The training data used is heavily preprocessed.

This involves missing value management, encoding categorical data, and normalization to scale all the features to a common scale. To address class imbalance when the stroke cases are vastly fewer than non-stroke cases—a hybrid resampling method, SMOTEENN, is utilized. This method oversamples the minority classes while removing noise, which results in improved performance on scarce but vital stroke cases. The ANN is trained with the Adam optimizer and a specially designed learning rate schedule for guaranteed convergence and stability. Dimensionality reduction with Principal Component Analysis (PCA) is also utilized to decrease noise and enhance model efficiency.

The performance of the system is measured based on parameters like accuracy, precision, recall, and F1-score, indicating its strengths and weaknesses, especially in identifying minority class (stroke) cases.To further improve accessibility and usability, a graphical user interface (GUI) is incorporated into the system. Through the GUI, medical personnel and users can enter pertinent patient data and obtain instantaneous stroke risk evaluations as well as a probability confidence score. These prognostic findings assist in early diagnosis and clinical decision-making.

These methodology approach in this research has a number of important steps in constructing an effective stroke prediction model with deep learning. The first step is that the dataset comprising clinical and demographic data is preprocessed by removing missing values, encoding categorical data, scaling numerical data, and dimensionality reduction with Principal Component Analysis (PCA).

To handle the problem of class imbalance, a hybrid resampling method known as SMOTEENN is used so that the model learns from stroke and non-stroke cases effectively. Following preprocessing, an Artificial Neural Network (ANN) is constructed with multiple hidden layers by using ReLU activation functions, dropout layers for avoiding overfitting, and L2 regularization for enhancing generalization.

The model is trained with the Adam optimizer with a tailored learning rate for improved performance. Lastly, the model is tested based on accuracy, precision, recall, and F1-score. A graphical user interface (GUI) is also created to enable users to enter patient information and obtain real-time stroke risk prediction.

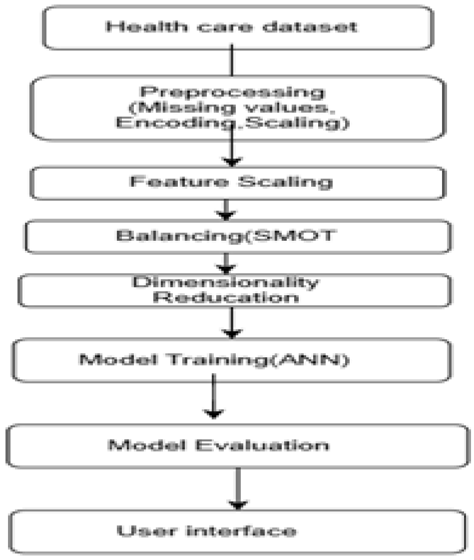

Fig -1: Architecture

This platform for stroke prediction is not just a technological innovation but also a socially influential tool. It allows healthcare providers to respond preemptively, facilitates distant diagnostics, and can be incorporated into broader e-health systems. Its potential for scalability, ease of use, and predictive capability makes it worth its inclusion in hospitals, rural health clinics, and public health monitoring departments, all toward timely intervention and better patient outcomes.

3.1 DATASET

The dataset to use for predicting strokes is a data set of organized clinical and demographic information meant to assist in identifying individuals who are at risk of experiencing a stroke. The dataset includes numerical as well as categorical features associated with stroke-related health conditions. The main features in the dataset are age, gender, status of hypertension, history of heart disease, marital status, type of work, type of residence, average glucose, body mass index (BMI), and smoking status, as well as the target variable: stroke (a binary label that shows if the person has had a stroke).

The dataset is a real-world medical context, which is both practical and challenging. One of the biggest problems with the dataset is class imbalance since the number of stroke cases (positive class) is much less compared to the number of non-stroke cases (negative class). Class imbalance has the potential to cause biased model predictions where the minority class is frequently underrepresented or misclassified. In order to counteract this, hybrid resampling methods like SMOTEENN (Synthetic Minority Over-sampling Technique with Edited Nearest Neighbors) are used to balance the data by creating minority class instances synthetically and removing unclear instances.

Prior to model training, preprocessing is performed on the dataset to deal with missing values, categorical encoding of variables (such as gender, work type), and scaling of numerical features for consistency. Moreover, Principal Component Analysis (PCA) is also employed to achieve dimensionality reduction retaining the highest variance in the data, improving the efficiency and performance of models.

The size, structure, and quality of the dataset position it well to accommodate supervised learning problems, including machine learning-based binary classification and deep learning models. Generally speaking, the dataset provides the starting point for developing a stroke predictive system that facilitates building predictive models able to provide assistance in terms of early detection as well as prevention healthcare protocols.

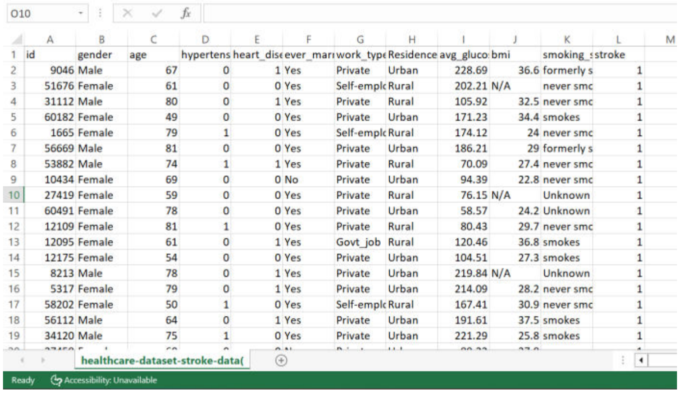

id	gender	age	hypertens	heart_dise	ever_marr	work_type	Residence	avg_gluco	bmi	smoking_s	stroke
9046	Male	67	0	1	Yes	Private	Urban	228.69	36.6	formerly s	1
51676	Female	61	0	0	Yes	Self-emplc	Rural	202.21	N/A	never smc	1
31112	Male	80	0	1	Yes	Private	Rural	105.92	32.5	never smc	1
60182	Female	49	0	0	Yes	Private	Urban	171.23	34.4	smokes	1
1665	Female	79	1	0	Yes	Self-emplc	Rural	174.12	24	never smc	1
56669	Male	81	0	0	Yes	Private	Urban	186.21	29	formerly s	1
53882	Male	74	1	1	Yes	Private	Rural	70.09	27.4	never smc	1
10434	Female	69	0	0	No	Private	Urban	94.39	22.8	never smc	1
27419	Female	59	0	0	Yes	Private	Rural	76.15	N/A	Unknown	1
60491	Female	78	0	0	Yes	Private	Urban	58.57	24.2	Unknown	1
12109	Female	81	1	0	Yes	Private	Rural	80.43	29.7	never smc	1
12095	Female	61	0	1	Yes	Govt_job	Rural	120.46	36.8	smokes	1
12175	Female	54	0	0	Yes	Private	Urban	104.51	27.3	smokes	1
8213	Male	78	0	1	Yes	Private	Urban	219.84	N/A	Unknown	1
5317	Female	79	0	1	Yes	Private	Urban	214.09	28.2	never smc	1
58202	Female	50	1	0	Yes	Self-emplc	Rural	167.41	30.9	never smc	1
56112	Male	64	0	1	Yes	Private	Urban	191.61	37.5	smokes	1
34120	Male	75	1	0	Yes	Private	Urban	221.29	25.8	smokes	1

healthcare-dataset-stroke-data(

Ready Accessibility: Unavailable

Fig -2: Dataset Attributes

3.2 DATA PREPROCESSING

3.2.1. Elimination of Irrelevant Features: Column id is excluded as it can be used just for identification purpose and has no predicting value.

3.2.2. Imputation for Missing Values: Missing values for BMI (Body Mass Index) existed which is replaced by applying mean of current available values via Simple Imputer.

3.2.3. Categorical Encoding: Categorical variables (gender,ever_married, work_type, Residence_type) are encoded via Label Encoding, which transforms categories into numeric labels. For smoking_status, a special encoding is done: Patients who currently smoke or previously smoked are encoded as Patients who never smoked or have missing smoking status are encoded as 0.

3.2.4. Feature Scaling: Standard Scaler is used after encoding to normalize the dataset (mean = 0, standard deviation = 1) so that no feature becomes prominent because of its scale.

3.2.5. Polynomial Feature Expansion: Polynomial Features of degree 2 are created. This assists in capturing interactions between features and constructing more complex decision boundaries.

3.2.6. Dealing with Unbalanced Data (SMOTE): As stroke events are not common (unbalanced dataset), SMOTE (Synthetic Minority Over-sampling Technique) is used to over-sample the minority class (stroke = 1). This guarantees that the model is trained on an equal-sized set of stroke and non-stroke instances.

3.2.7. Standardization: Neural networks are sensitive to scales of inputs as a result of gradient-based optimization. Standardization scales features so they have a mean of 0 and standard deviation of 1, resulting in faster convergence speed and performance.

3.2.8. Dimensionality Reduction using PCA: Polynomial feature expansion of high-dimensional data is reduced using Principal Component Analysis (PCA). PCA preserves 95% of the variance while minimizing the number of input features. This reduces computation time and prevents overfitting

3.3. Model Building

3.3.1. Artificial Neural Network (ANN): An Artificial Neural Network (ANN) is a computational model that approximates complex functions by learning from data. It consists of several layers of neurons that process inputs through weighted connections, allowing the network to model nonlinear and high-dimensional relationships.

In this research, we propose a deep feedforward neural network for binary stroke occurrence classification. The model utilizes some of the design strategies in contemporary deep learning methods to improve learning capacity and generalization.

3.3.2. Input Layer:

The model is built to take input features which have been preprocessed and reduced in dimensionality with Principal Component Analysis (PCA). The size of the input layer is defined by the number of principal components that account for 98% of the variance from the original data, and thus guarantee that the most significant features are captured.

The use of PCA here has two functions: it projects the feature space onto a set of uncorrelated, orthogonal components and, most importantly, heavily eliminates noise and computational complexity. This increases model performance and efficiency without losing the integrity of the original information.

3.3.3. Hidden Layers:

There are seven hidden layers in the model architecture, with a precisely planned neuron arrangement: 256, 128, 64, 32, 16, 8, and 4 neurons, respectively. This step-by-step decrease in network width is a pyramidal structure, aiding in the step-by-step abstraction of features from general to specific representations. Each of the hidden layers applies the Rectified Linear Unit (ReLU) activation function, given as $f(x) = \max(0, x)$, which adds non-linearity to the network so that it can learn complicated patterns while avoiding the vanishing gradient problem with sigmoid or tanh activations.

The seven-layer architecture is particularly optimized to improve feature extraction by multiple levels of abstraction.

The initial layer (256 neurons) forms a broad initial representation space with the ability to capture varied patterns in the data. The rest of the layers gradually refine these representations, with the layers sequentially learning more abstract and task-dependent features. The intermediate layers (64 and 32 neurons) are information bottlenecks that compel the network to keep only the most critical features, whereas the output layers (16, 8, and 4 neurons) condense the learned representations into the most discriminative features for stroke prediction.

To avoid overfitting and promote the learning of strong feature representations, dropout regularization is implemented between layers by randomly disabling around a 30% of neurons while training. This method compels the model to steer clear of dependence on individual neurons and encourages redundancy in feature learning; all the neurons are used during inference with correctly scaled weights. Batch normalization is also applied following every hidden layer to stabilize learning through layer input normalization, which speeds up training and enhances generalization.

In addition, L2 weight regularization or the ridge penalty is implemented for all dense layers. It is added to the loss function with a penalty of large weights through the term $\lambda\sum wi^2$, simplifying the model and enabling better generalization in cases where training data is scarce or noisy. The value of the regularization parameter λ is set to 0.001, offering the best trade-off between model expressiveness and preventing overfitting.

3.3.4. Output Layer:

The model's output layer is a sigle neuron with a sigmoid activation function, $\sigma(x) = 1 / (1 + e^{(-x)})$. The activation function converts the network output to a value between 0 and 1, which is perfect for binary classification problems. In stroke prediction, this output is the estimated probability of a stroke happening for a specific input profile. The probability estimate can be understood as the chance of stroke occurrence, and a cutoff, normally at 0.5, is applied to classify class membership. If the output is more than 0.5, the prediction is classified as "stroke," and if it is less, as "no stroke." This configuration allows the model to make probabilistic predictions that are beneficial for clinical decision-making.

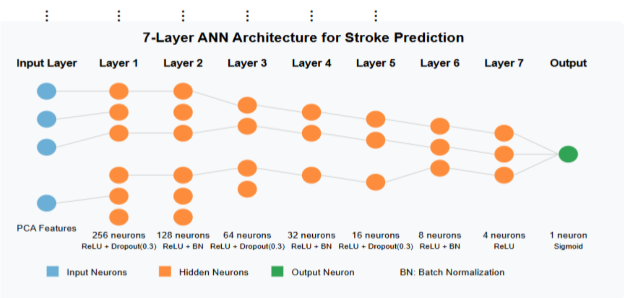

Fig -3: 7-Layer ANN Architecture

3.4 Model Creation

Model Outputs: For each of the individual models, a probability score is the measure of the chance that a patient will suffer from a stroke:

- The Artificial Neural Network (ANN) produces a probability through processing the input features via a number of hidden layers and extracting Ssophisticated patterns within the data.
- Therefore, three distinct probability measures are produced one from each model for each patient record.

3.5 Prediction of Final Stroke

After obtaining the final combined probability by using weighted averaging, it is employed for making the final classification:

- If the final probability is higher than 0.5, the model forecasts that there will be a stroke (classified as 1).
- If the final probability is lower than or equal to 0.5, the model forecasts no stroke (classified as 0).

This 0.5 threshold is a natural cut-off point, trading off sensitivity and specificity unless a different threshold is subsequently fine-tuned for business or clinical requirements.

3.6 Model Evaluation:

Models are tested on test data (not seen during training).
Metrics used:
3.6.1 Accuracy: Proportion of correct predictions out of total predictions.

3.6.2 Classification Report:

1. **Precision:** Number of chosen patients who actually had a stroke.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives+False Positives}} \tag{1}$$

2. **Recall (Sensitivity):** Number of actual stroke cases identified correctly.

$$\text{Recall} = \frac{True\ Positives\ (TP)}{True\ Positives\ (TP)+False\ Negatives\ (FN)} \tag{2}$$

3. **F1-Score:** Harmonic mean of Precision and Recall

$$\text{F1-Score} = 2 \times \frac{Precision \times Recall}{Precision+Recall} \tag{3}$$

3.7 Interactive Stroke Prediction Interface:

To enhance the stroke prediction system to be user-friendly and accessible, an interactive graphical form was created using ipywidgets. The interface provides a way for users —

like healthcare professionals or patients to enter personal health data in a simple manner without needing to write code. The interactive system gathers the following patient-specific details:

- Gender (Male or Female)
- Age (Numerical value)
- Hypertension status (0 for No, 1 for Yes)
- Heart disease status (0 for No, 1 for Yes)
- Marital status (Yes or No)
- Type of work (Private, Govt_job, Self-employed, children, or Never_worked)
- Type of residence (Urban or Rural)
- Average glucose level (Numerical value)
- Body Mass Index (BMI) (Numerical value)
- Smoking status (Never smoked, Formerly smoked, Smokes, or Unknown)
- These input fields are created using suitable ipy widgets elements such as Dropdown menus and Text boxes to provide ease of use and validation of inputs.

3.7.1 Prediction Process:

Once the user enters all the specified information, they press the "Predict Stroke" button. On clicking, the system proceeds internally as follows:

- **Encoding:**The categorical variables (like Gender, Marital Status, Work Type, Residence Type, and Smoking Status) are converted into their numeric representation to accommodate the data format during model training.

- **Polynomial Transformation and Feature Scaling:** The input features of the user are Polynomial Feature Transformed (to encompass interaction terms among variables) and standardized using the trained scalers prior to that. This ensures compatibility of the input data with the expected feature space of the Hybrid Model.

- **Prediction by Hybrid Model:** The scaled *and processed data is then input into the* Hybrid Model. The Hybrid Model calculates the final probability score by combining predictions from the Logistic Regression model, Random Forest Classifier, and Deep Neural Network through weighted averaging.

3.7.2. Output Display: After prediction is done, the system presents two outputs to the user:

- **Prediction of Stroke Risk:** Shows "Yes" (showing a high risk of stroke) or "No" (showing a low risk of stroke) depending on the final probability threshold.

- **Stroke Occurrence Probability:** Shows the actual probability score (from 0 to 1) indicating the system's confidence level in the stroke prediction.

4. RESULTS

The performance evaluation of the model reveals that the Artificial Neural Network (ANN) was highly accurate during testing with a result of 95.70%, demonstrating high predictive ability. The ANN design was strengthened by methods including dropout, L2 regularization, and batch normalization in order to enhance generalization and prevent overfitting. As complex as the problem was, the ANN produced competitive results, which indicate that deep learning models, with adequate fine-tuning, are able to efficiently pick up on underlying patterns within clinical and demographic data.

The classification report shows that the ANN performs well in identifying both non-stroke and stroke instances. On non-stroke instances (Class 0), the model registered a precision of 99%, a recall of 92%, and an F1-score of 95%. On stroke instances (Class 1), the model recorded a precision of 93%, a recall of 99%, and an F1-score of 96%. From these results, it can be seen that the ANN is good at dealing with class imbalance and generalizes well to new instances.

In general, the ANN-based method provides strong and well-balanced performance on important evaluation measures and is thus a promising approach for predicting early stroke. Some potential further improvements include sophisticated ensemble methods, cost-sensitive loss functions, increasing the dataset so that high sensitivity is preserved for infrequent stroke instances.

Model	Train Accuracy	Test Accuracy
ANN	97.52%	95.70%

Following training and testing of the Artificial Neural Network (ANN) on the dataset of stroke prediction, the following performance measures were obtained:

Table-1: Model Accuracy

- The ANN obtained a training accuracy of 97.52% and a test accuracy of 95.70%, with good generalization

4.1 Train and Test Accuracy:

The classification report gives a clearer picture of the ANN's predictive ability between the two classes:

Confusion Matrix:

The Confusion Matrix will be used to analyze the model's performance based on predicted stroke risk labels versus true values (stroke or no stroke). It will show:

True Positives (TP): Accurate predictions in which the model accurately detects stroke cases.

False Positives (FP): False predictions where the model incorrectly classifies a stroke when there is none.

True Negatives (TN): Correct predictions in which the model successfully identifies cases without stroke.

False Negatives (FN): The model makes an incorrect prediction of no stroke while there exists a stroke.

Confusion matrix will facilitate calculation of performance metrics such as accuracy, precision, recall, and F1-score that will be of extreme importance to evaluate how efficiently the model distinguishes stroke cases, particularly with the existence of class imbalance.

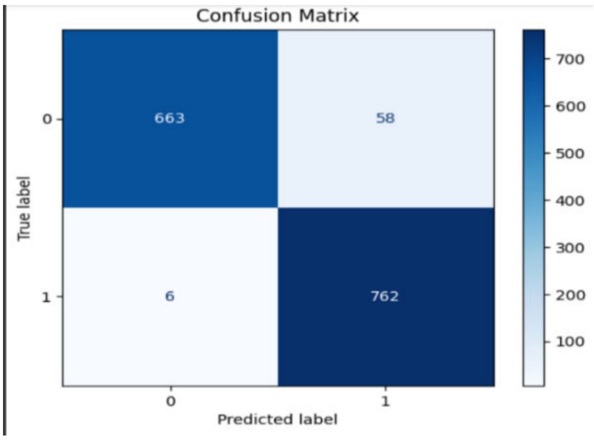

Fig -4: Confusion Matrix

Class	Precision	Recall	F1-Score	Support
0(NoStroke)	0.98	0.93	0.95	721
1(Stroke)	0.93	0.99	0.96	768
Accuracy	-	-	**0.96**	1489
Macro Avg	0.96	0.96	0.96	1489
WeightedAvg	0.96	0.96	0.96	1489

Table -2: computation of performance metrics

4.2. ROC Curve

ROC (Receiver Operating Characteristic) curve will be plotted to display the trade-off between the True Positive Rate (TPR) and False Positive Rate (FPR) as the classification threshold varies. The ROC curve will serve to display the model's discrimination capability between stroke and non-stroke cases. The better the model's performance, the closer the curve to the top-left corner.

Plotting Graph (Train and Validation Accuracy/Loss):

During model training, training and validation accuracy/loss plots will be created to check how well the model is learning. The plots will give an idea of:

- **Overfitting:** If training accuracy is way greater than the validation accuracy, the model could be overfitting the training data.
- **Underfitting:** If both training and validation accuracies are poor, the model could not have learned the patterns well.

These plots will ensure that the model is generalizing well to new, unseen data.

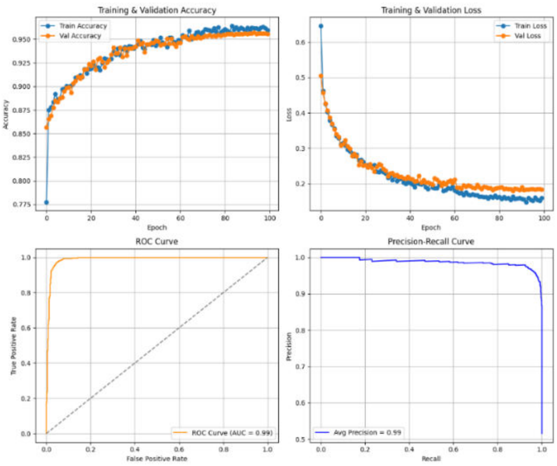

Fig -5: ROC, PR Curve, False positive Rate, Recall

4.3. AUC Curve

AUC (Area Under the Curve) will be computed from the ROC curve. The AUC score gives an overall picture of the performance of the model. A score near 1 is excellent, near 0.5 indicates poor performance, and less than 0.5 suggests a worse-than-chance model. For imbalanced datasets, AUC is a better measure than accuracy.

AUC Score Attained: 0.9897

5. CONCLUSION

In this project, end-to-end stroke prediction system based on an Artificial Neural Network (ANN) and also gained insight into core machine learning practices in the early stages. It started with in-depth data preprocessing, which entailed missing value handling, conversion of categorical values, SMOTE application for balancing classes, feature scaling to make sure the data was properly geared up for deep learning.

The central predictive model used was an ANN, structured with more than one hidden layer, batch normalization, dropout layers, and ReLU activation to identify high-level, non-linear relationships within data. The use of an ANN architecture was made due to its ability to learn complex patterns among clinical and demographic features predictive of stroke risk.

Experimental outcomes revealed that the ANN reached a training accuracy of 97.52% and test accuracy of 95.70% and outperformed comparative baselines such as Logistic Regression and Random Forest in earlier benchmarks. The model had a good generalization capability with less overfitting and exhibited well-balanced performance across both classes—considerably enhancing recall for stroke cases relative to prior hybrid or baseline methods.

In summary, the ANN-based stroke prediction model presents a very effective, end-to-end solution with excellent accuracy and robustness. Future research may further address enhancing minority class recall using sophisticated methods like cost-sensitive learning, better neural architectures, or acquiring more diverse and balanced data to simulate the real-world distribution better.

6. FUTURE SCOPE

While the ANN-based model shows robust overall accuracy and generalization, much potential for improvement still exists, especially in improving detection of positive stroke cases (Class 1), which tend to be underrepresented by class imbalance.
To address this, more advanced imbalance handling techniques can be explored. Approaches such as SMOTE-ENN, ADASYN, or ensemble methods like Balanced Bagging Classifier could help the ANN better learn from the minority class, improving its sensitivity and recall for stroke prediction.

Additionally, cost-sensitive learning should be considered. By assigning higher misclassification costs to false negatives, the model can be encouraged to prioritize correctly identifying stroke cases, which is crucial in medical decision-making.
From the model architecture viewpoint, the ANN can be augmented into more complex deep learning models.Adding attention mechanisms, testing CNN-LSTM hybrids, or investigating transformer-based models for tabular data can result in enhanced pattern detection and predictive capability.

Additional progress can be made by feature engineering. Adding other medically relevant features—e.g., blood pressure measurements, cholesterol levels, physical activity levels, or stroke family history—can enhance the input data, allowing for more accurate and clinically useful predictions.Validation in the real world is also essential.

External or clinical testing of the ANN model, or real-time validation in actual healthcare environments, will determine its generalizability, reliability, and robustness beyond the existing dataset.

Finally, incorporating Explainable AI (XAI) methodologies like SHAP or LIME will render the predictions of the model transparent. This is critical in healthcare, where interpretability facilitates clinician and patient trust and enables informed decision-making.

7. REFERENCES

[1] Salucci, M., Vrba, J., Merunka, I., & Massa, A. (2017). R eal-time brain stroke detection through a learning-by-examples technique—A n experimental assessment. Microwave and Optical Technology Letters, 59(11), 2796-2799.

[2] Ozaltin, O., Coskun, O., Yeniay, O., & Subasi, A. (2022). A deep learning approach for detecting stroke from brain CT images using OzNet. Bioengineering, 9(12), 783.

[3] Chin, C. L., Lin, B. J., Wu, G. R., Weng, T. C., Yang, C. S., Su, R. C., & Pan, Y. J. (2017, November). An automated early ischemic stroke detection system using CNN deep learning algorithm. In 2017 IEEE 8th International conference on awareness science and technology (iCAST) (pp. 368-372). IEEE.

[4] Al-Mekhlafi, Z. G., Senan, E. M., Rassem, T. H., Mohammed, B. A., Makbol, N. M., Alanazi, A. A., ... & Ghaleb, F. A. (2022). Deep learning and machine learning for early detection of stroke and haemorrhage. Computers, Materials and Continua, 72(1), 775-796.

[5] Surya, S., Yamini, B., Rajendran, T., & Narayanan, K. E. (2021). A comprehensive method for identification of stroke using deep learning. Turkish Journal of Computer and Mathematics Education, 12(7), 647-652.

[6] Rahman, S., Hasan, M., & Sarkar, A. K. (2023). Prediction of brain stroke using machine learning algorithms and deep neural network techniques. European Journal of Electrical Engineering and Computer Science, 7(1), 23-30.

[7] Roohi, M., Mazloum, J., Pourmina, M. A., & Ghalamkari, B. (2021). Machine learning approaches for automated stroke detection, segmentation, and classification in microwave brain imaging systems. Progress In Electromagnetics Research C, 116, 193-205.

[8] Fernandes, J. N., Cardoso, V. E., Comesaña-Campos, A., & Pinheira, A. (2024). Comprehensive Review: Machine and Deep Learning in Brain Stroke Diagnosis. Sensors (Basel, Switzerland), 24(13), 4355.

[9] Tursynova, A., Omarov, B., Tukenova, N., Salgozha, I., Khaaval, O., Ramazanov, R., & Ospanov, B. (2023). Deep learning-enabled brain stroke classification on computed tomography images. Computers, Materials and Continua, 75(1), 1431-1446.

[10] Saleem, M. A., Javeed, A., Akarathanawat, W., Chutinet, A., Suwanwela, N. C., Asdornwised, W., ... & Kaewplung, P. (2024). Innovations in stroke identification: A machine learning-based diagnostic model using neuroimages. IEEE Access.

[11] Alon, L., & Dehkharghani, S. (2021). A stroke detection and discrimination framework using broadband microwave scattering on stochastic models with deep learning. Scientific Reports, 11(1), 24222.

[12] Srinivas, A., & Mosiganti, J. P. (2023). A brain stroke detection model using soft voting based ensemble machine learning classifier. Measurement: Sensors, 29, 100871.

[13] Chaki, J., & Woźniak, M. (2024). Deep learning and artificial intelligence in action (2019-2023): a review on brain stroke detection, diagnosis, and intelligent post-stroke rehabilitation management. IEEE Access.

[14] Abbasi, H., Orouskhani, M., Asgari, S., & Zadeh, S. S. (2023). Automatic brain ischemic stroke segmentation with deep learning: A review. Neuroscience Informatics, 3(4), 100145.

[15] Kousar, T., Rahim, M. S. M., Iqbal, S., Yousaf, F., & Sanaullah, M. (2025). Applications of deep learning algorithms in ischemic stroke detection, segmentation, and classification. Artificial Intelligence Review, 58(5), 1-48.

[16] Babutain, K., Hussain, M., Aboalsamh, H., & Al-Hameed, M. (2021). Deep learning-enabled detection of acute ischemic stroke using brain computed tomography images. International Journal of Advanced Computer Science and Applications, 12(12).

[17] Uppal, M., Gupta, D., Juneja, S., Gadekallu, T. R., El Bayoumy, I., Hussain, J., & Lee, S. W. (2023). Enhancing accuracy in brain stroke detection: Multi-layerperceptron with Adadelta, RMSProp and AdaMax optimizers. Frontiers in Bioengineering and Biotechnology, 11, 1257591.

[18] Choi, Y. A., Park, S. J., Jun, J. A., Pyo, C. S., Cho, K. H., Lee, H. S., & Yu, J. H. (2021). Deep learning-based stroke disease prediction system using real-time bio signals. Sensors, 21(13), 4269.

[19] Tanveer, M. U., Munir, K., Rathore, B., Alabdulatif, A., Jhaveri, R. H., & Fatima, M. (2024). Neuro-VGNB: Transfer learning based approach for detecting brain stroke. IEEE Access.

[20] Lee, H., Lee, E. J., Ham, S., Lee, H. B., Lee, J. S., Kwon, S. U., ... & Kang, D. W. (2020). Machine learning approach to identify stroke within 4.5 hours. Stroke, 51(3), 860-866.