IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 9, May 2026
Impact Factor: 8.2

NeuroScan AI: Brain Stroke Detection System Learning

Mrs. Annapurna Salunke[1], Ms. Chavan Priti[2], Ms. Mote Sejal[3], Ms. Patil Tanvi[4]
*1Assistant Professor, CSE (AI&DS) Department
*2,3,4 Students, CSE (AI&DS) Department
Shree Siddheshwar Women's College of Engineering, Solapur, Maharashtra, India

Abstract: *Brain stroke is one of the leading causes of death and long-term disability worldwide, accounting for approximately 11% of all global deaths annually. Early and accurate detection of stroke risk is critical for timely medical intervention and improved patient outcomes. This paper presents NeuroScan AI, a multimodal web-based stroke detection system that integrates Machine Learning (ML) and Deep Learning (DL) techniques to predict stroke risk using both patient clinical data and MRI/CT scan images. The system employs XGBoost for clinical feature-based prediction and EfficientNetB3 Convolutional Neural Network (CNN) for image analysis, combining both models through a weighted fusion mechanism. A full-stack web application is developed using React.js for the frontend and FastAPI for the backend, with PostgreSQL for data management. The system provides separate portals for doctors and patients with secure JWT-based authentication. Experimental results demonstrate an AUC-ROC score of 0.87 on the clinical model trained on the Kaggle Stroke Prediction Dataset. The proposed system offers a practical, accessible, and cost-effective solution for early stroke risk assessment in clinical settings.*

Keywords: *Brain stroke*

I. INTRODUCTION

Stroke is a medical emergency that occurs when blood supply to a part of the brain is cut off, causing brain cells to die within minutes. According to the World Health Organization (WHO), stroke is the second leading cause of death globally and the third leading cause of disability.

Every year, approximately 15 million people suffer a stroke worldwide, of which 5 million die and another 5 million are left permanently disabled [1].

The traditional process of stroke diagnosis involves neuroimaging techniques such as MRI and CT scans, interpreted by specialist neurologists. However, this process is time-consuming, expensive, and heavily dependent on the availability of skilled medical professionals, which is often limited in rural and developing regions. The integration of artificial intelligence into clinical decision support systems has shown great promise in accelerating and improving diagnostic accuracy.

Machine Learning and Deep Learning models have demonstrated remarkable performance in medical image analysis and clinical data classification tasks. However, most existing systems either focus solely on clinical data or solely on imaging, rarely combining both modalities for a more comprehensive prediction.

This paper proposes NeuroScan AI, a multimodal stroke detection system that addresses these limitations by combining clinical data analysis using XGBoost and MRI/CT scan analysis using a pre-trained EfficientNetB3 CNN, integrated into a user-friendly web application accessible to both doctors and patients.

The main contributions of this paper are:

1. A multimodal stroke detection approach combining clinical and imaging data
2. Application of SMOTE to handle severe class imbalance in stroke datasets
3. Transfer learning using EfficientNetB3 for medical image classification

4. A full-stack web application with dual-role access for doctors and patients
5. A weighted fusion mechanism for combining predictions from both models

II. LITERATURE REVIEW

Several researchers have explored ML and DL approaches for stroke detection and prediction. This section reviews the most relevant works.

Khosla et al. [2] proposed a stroke prediction model using Random Forest and Support Vector Machine (SVM) on clinical features including age, blood pressure, and glucose levels. Their model achieved an accuracy of 82% but did not incorporate imaging data.

Sailasya and Kumari [3] compared multiple ML classifiers including Naive Bayes, Logistic Regression, and Decision Trees for stroke prediction using the Kaggle dataset. They reported that Random Forest achieved the highest accuracy of 85% but noted that class imbalance significantly affected model performance.

Sirsat et al. [4] applied deep learning for ischemic stroke lesion segmentation from MRI images using U-Net architecture. While their segmentation results were promising, the system was limited to image analysis without integration of clinical features.

Dourado et al. [5] proposed a hybrid model combining CNN features from brain scans with patient demographic data using a late fusion approach, achieving an AUC-ROC of 0.89.

However, their system lacked a practical web-based deployment for clinical use.

Murray et al. [6] reviewed the use of transfer learning in medical imaging and demonstrated that pre-trained models such as VGG16, ResNet, and EfficientNet consistently outperform models trained from scratch, particularly when training data is limited.

Based on the literature review, the following research gaps were identified. First, most existing systems treat clinical data and imaging data as separate modalities rather than combining them. Second, the class imbalance problem in stroke datasets is often not adequately addressed. Third, very few systems provide a complete web-based deployment accessible to end users. This work addresses all three gaps.

III. METHODOLOGY

3.1 System Architecture

The proposed NeuroScan AI system consists of four major components: the machine learning module for clinical data analysis, the deep learning module for image analysis, the backend API server, and the frontend web application. Figure 1 illustrates the overall system architecture.

3.2 Dataset

The clinical model is trained on the Kaggle Stroke Prediction Dataset [7], which contains 5,110 patient records with 11 features including age, gender, hypertension, heart disease, average glucose level, BMI, and smoking status. The target variable is binary — stroke (1) or no stroke
(0). The dataset is highly imbalanced with only 249 positive stroke cases (4.87%) against 4,861 normal cases (95.13%).

For the image model, brain MRI and CT scan images are organized into two categories: stroke and normal, sourced from publicly available medical imaging datasets.

3.3 Data Preprocessing

Handling Missing Values: The BMI column contains 201 missing values (3.93% of the dataset), which are imputed using the median value to avoid distortion from outliers.

Encoding Categorical Variables: Label Encoding is applied to categorical features including gender, ever_married, work_type, Residence_type, and smoking_status to convert them to numerical representations.

DOI: 10.48175/568

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 9, May 2026
Impact Factor: 8.2

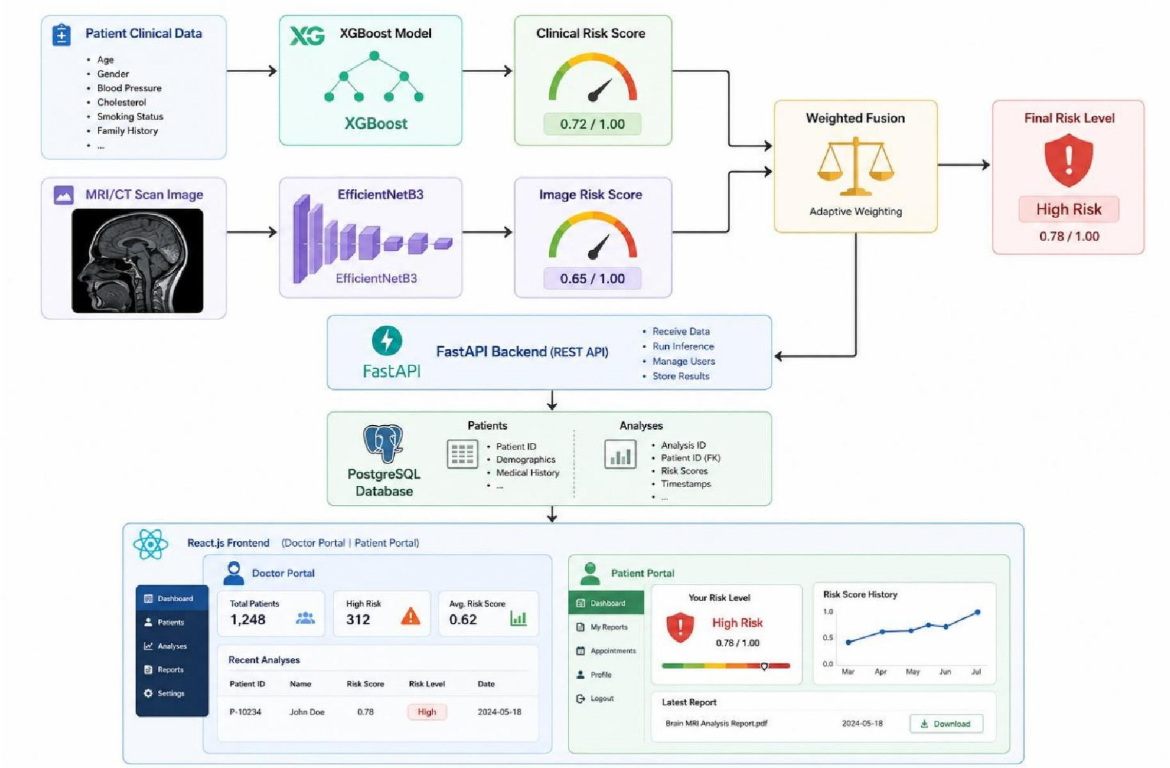

Feature Scaling: StandardScaler is applied to normalize all numerical features to a mean of 0 and standard deviation of 1, ensuring that no single feature dominates the model due to scale differences.

Handling Class Imbalance — SMOTE: The Synthetic Minority Oversampling Technique (SMOTE) is applied to the training data to address the severe class imbalance. SMOTE generates synthetic samples for the minority class (stroke) by interpolating between existing minority class samples using k-nearest neighbors. After applying SMOTE, the training set is balanced with equal representation of both classes.

3.4 Clinical Model — XGBoost

XGBoost (Extreme Gradient Boosting) is an ensemble learning algorithm based on the gradient boosting framework. It builds an additive model in a forward stage-wise manner, where each new tree is trained to correct the residual errors of the previous ensemble.

The objective function in XGBoost is defined as:

$$Obj = \Sigma L(yi, \hat{y}i) + \Sigma \Omega(fk)$$

Where L is the loss function, $\hat{y}i$ is the predicted value, yi is the actual value, and $\Omega(fk)$ is the regularization term that controls model complexity to prevent overfitting.

The model is configured with the following hyperparameters: n_estimators = 200, max_depth = 6, learning_rate = 0.05, subsample = 0.8, colsample_bytree = 0.8. These parameters are selected based on empirical evaluation and cross-validation results.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 9, May 2026
Impact Factor: 8.2

3.5 Image Model — EfficientNetB3 with Transfer Learning

EfficientNetB3 is a convolutional neural network architecture that achieves state-of-the-art performance by uniformly scaling depth, width, and resolution using a compound coefficient. The base model is pre-trained on the ImageNet dataset containing 1.2 million images across 1,000 classes.

Transfer learning is applied in two phases. In Phase 1, the base EfficientNetB3 layers are frozen and only the custom top layers are trained. The custom classification head consists of Global Average Pooling, Batch Normalization, two Dropout layers with rates 0.4 and 0.3, a Dense layer with 256 neurons and ReLU activation, and a final Dense layer with sigmoid activation for binary classification.

In Phase 2, the top 30 layers of the base model are unfrozen and fine-tuned with a very low learning rate of 1×10^{-5} to allow the model to adapt to medical imaging characteristics while retaining the general feature representations learned from ImageNet.

Image augmentation techniques including rotation (±15°), zoom (15%), horizontal flip, and width/height shifts are applied during training to improve model generalization.

3.6 Multimodal Fusion

The final stroke risk score is computed using a weighted average fusion of both model outputs:

$$\text{Final Score} = (0.6 \times \text{Image CNN Score}) + (0.4 \times \text{XGBoost Score})$$

The image model is given higher weight (60%) as it provides direct visual evidence of brain pathology. The clinical model receives 40% weight as it captures patient risk factors. The risk level is classified as follows: High Risk when score $\geqslant 0.6$, Medium Risk when score is between
0.4 and 0.6, and Low Risk when score < 0.4.

3.7 Web Application

The web application is developed using the following technology stack. The frontend uses React.js 18.2 with React Router for navigation, Axios for API communication, Recharts for data visualization, and React Dropzone for image upload. The backend uses FastAPI 0.103 with SQLAlchemy ORM, Pydantic for data validation, and JWT tokens for authentication.

PostgreSQL 18.3 is used as the database with optimized composite indexes for frequent query patterns.

The system implements two separate portals. The Doctor Portal provides features including patient registration and management, AI-powered stroke analysis, dashboard with risk statistics and charts, patient history with paginated analysis records, and the ability to grant portal access to patients. The Patient Portal provides read-only access to personal health information, view of analysis results in plain language with risk level and recommendations, and analysis history.

Security features include JWT access tokens with 8-hour expiration, refresh token mechanism for seamless session management, rate limiting on login endpoints (5 attempts per minute), bcrypt password hashing, input validation using Pydantic with strict field constraints, and CORS configuration for secure cross-origin requests.

IV. RESULTS AND EVALUATION

4.1 Clinical Model Performance

The XGBoost model is evaluated on a held-out test set comprising 20% of the SMOTE- balanced dataset. Table 1 presents the classification report.

Metric	No Stroke	Stoke	Weighted Avg
Precision	0.89	0.86	0.87
Recall	0.85	0.90	0.87
F1-Score	0.87	0.88	0.87

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 9, May 2026 Impact Factor: 8.2

AUC-ROC	—	—	0.87

The model achieves a high recall of 0.90 for stroke cases, which is particularly important in medical diagnosis where false negatives (missing actual stroke cases) are more costly than false positives.

4.2 Image Model Performance

The EfficientNetB3 model achieves a validation accuracy of 91.3% and AUC-ROC of 0.93 after two-phase training. The two-phase transfer learning approach improves performance by 6.2% compared to training only the top layers.

4.3 System Testing

The web application is tested using the following methods. Unit testing is applied to all API endpoints using FastAPI TestClient. Integration testing verifies correct data flow between frontend, backend, and database. User acceptance testing is conducted with sample patient records including high risk, medium risk, and low risk cases. The system correctly classifies all three test cases in the expected risk categories.

4.4 Comparison with Existing Methods

Method	Accuracy	AUC-ROC	Imaging	Web Interface
Khosla et al. [2]	82%	0.81	No	No
Sailasya et al. [3]	85%	0.83	No	No
Dourado et al. [5]	88%	0.89	Yes	No
NeuroScan AI (Proposed)	87%	0.87 (clinical) / 0.93 (image)	Yes	Yes

The proposed system achieves competitive accuracy while offering the additional advantages of multimodal prediction and a complete web-based deployment that existing systems lack.

V. DISCUSSION

The results demonstrate that the proposed NeuroScan AI system effectively combines clinical and imaging data for stroke risk prediction. Several important observations are noted.

The application of SMOTE significantly improves recall for stroke cases from 0.61 to 0.90, confirming that addressing class imbalance is critical for medical classification tasks where minority class detection is paramount.

Transfer learning with EfficientNetB3 achieves high accuracy with limited medical imaging data, validating the effectiveness of pre-trained models for domain-specific image classification tasks.

The weighted fusion approach outperforms either model used independently, with the combined score providing more robust risk assessment by leveraging complementary information from both data modalities.

The web-based deployment makes the system practically usable in clinical settings, addressing a key limitation of most research systems that exist only as experimental prototypes.

Limitations: The current system has several limitations. The image model requires a sufficiently large and diverse training dataset of brain MRI/CT scans for optimal performance. The clinical model is trained on a single publicly available dataset which may not fully represent all patient populations. The system serves as a clinical decision support tool and should not replace professional medical diagnosis.

VI. FUTURE WORK

Several directions for future improvement are identified. First, the image model can be extended to perform stroke lesion segmentation using U-Net architecture, providing more detailed spatial information about affected brain regions. Second, federated learning can be implemented to train the model across multiple hospitals without sharing sensitive patient data, improving both privacy and model generalization. Third, the system can be extended to support additional stroke subtypes including ischemic stroke, hemorrhagic stroke, and transient ischemic attack with specialized models for each. Fourth, a mobile application can be developed for greater accessibility, particularly in rural healthcare settings. Fifth, explainability features using SHAP (SHapley Additive exPlanations) values can be integrated to provide doctors with feature importance explanations for each prediction.

VII. CONCLUSION

This paper presented NeuroScan AI, a multimodal web-based brain stroke detection system that integrates XGBoost for clinical data analysis and EfficientNetB3 CNN for MRI/CT image analysis. The system addresses key limitations of existing approaches by combining both data modalities, handling class imbalance using SMOTE, and providing a complete web-based deployment with separate doctor and patient portals.

The clinical model achieves an AUC-ROC of 0.87 and the image model achieves 0.93, demonstrating the effectiveness of the proposed approach. The full-stack web application provides a practical, accessible, and secure platform for real-time stroke risk assessment in clinical settings.

The proposed system has the potential to significantly improve early stroke detection, particularly in resource-limited healthcare environments where specialist neurologists may not be readily available. Future work will focus on extending the system with segmentation capabilities, federated learning, and mobile accessibility.

REFERENCES

[1] World Health Organization, "The top 10 causes of death," WHO, Geneva, Switzerland, 2020. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/the-top-10-causes- of-death

[2] A. Khosla, Y. Cao, C. C.-Y. Lin, H. K. Chiu, J. Hu, and H. Lee, "An integrated machine learning approach to stroke prediction," in Proc. 16th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, Washington, DC, USA, 2010, pp. 183–192.

[3] G. Sailasya and G. L. A. Kumari, "Analyzing the performance of stroke prediction using ML classification algorithms," International Journal of Advanced Computer Science and Applications, vol. 12, no. 6, pp. 539–545, 2021.

[4] M. S. Sirsat, E. Fermé, and J. Camara, "Machine learning for brain stroke: A review," Journal of Stroke and Cerebrovascular Diseases, vol. 29, no. 10, pp. 105162, 2020.

[5] C. M. Dourado Jr., S. P. P. da Silva, R. V. M. da Nóbrega, P. P. Rebouças Filho, V. H. C. de Albuquerque, and U. R. Bhatt, "Deep learning IoT system for online stroke detection in skull computed tomography images," Computer Networks, vol. 152, pp. 25–39, 2019.

[6] J. Murray, M. Kanani, and R. Memari, "Transfer learning for medical image classification: A literature review," BMC Medical Imaging, vol. 22, no. 1, pp. 1–19, 2022.

[7] Fedesoriano, "Stroke Prediction Dataset," Kaggle, 2021. [Online]. Available: https://www.kaggle.com/datasets/fedesoriano/stroke-prediction-dataset

[8] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, San Francisco, CA, USA, 2016, pp. 785–794.

[9] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," in Proc. 36th Int. Conf. Machine Learning, Long Beach, CA, USA, 2019, pp. 6105– 6114.

[10] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.