International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IX **Month of publication:** September 2026

DOI: https://doi.org/10.22214/ijraset.2026.84880

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

Prediction of Brain Stroke Risk Using Machine Learning and Deep Learning: An Hybrid XGBoost-ResNet-18 CDSS

G Santhosh Raj[1], A Vijayan[2]
Department of Information Technology, St. Peter's College of Engineering and Technology, Chennai, India

Abstract: Stroke continues to be one of the major causes of global deaths and neurologic impairment. Conventional clinical scoring system such as CHADS₂ scores static and additive in nature and unable to process high dimensional multi-modal data. This paper proposes a hybrid artificial intelligence (AI) based clinical decision support system (CDSS) for evaluating brain stroke risk based on a combination of structured clinical indicators with medical images. The hybrid model comprises of an XGBoost classifier trained on tabular biomarkers such as hypertension status, average glucose level and body mass index (BMI), and a ResNet-18 deep learning model using transfer learning to recognize stroke indicators from MRI/CT scans. With a view to addressing the black-box nature inherent in medical AI applications, we incorporate Explainable AI (XAI) techniques such as SHAP (SHapley Additive exPlanations) to provide an interpretation at the feature level for each prediction. Our prototype utilizes a decoupled React 18 – Flask architecture with PostgresSQL database back end, JSON Web Token (JWT) based role-based authentication and a forensics traceable audit trail with automatic generation of clinical PDF reports.
Keywords: Brain Stroke Prediction; XGBoost; Deep Learning; ResNet-18; Transfer Learning; Explainable AI; SHAP; Clinical Decision Support System (CDSS); Machine Learning in Healthcare.

I. INTRODUCTION

Introduction Convergence of Computer Science and Medicine has heralded a new age of preventive healthcare. Machine Learning (ML) which was once considered experimental has become a fundamental component of precision medicine to discover complex non-linear patterns in high dimensional biomedical datasets that cannot be discovered through manual inspection alone. In case of brain stroke risk stratification, ML models are used to analyze the structured clinical indicators such as hypertension status, average glucose level, Body Mass Index (BMI), smoking status and cardiovascular conditions to calculate risk probability. ML algorithms unlike conventional additive scoring systems can learn the weighted synergy of multiple risk factors from patient data directly [2].

Just as DL is a specific form of ML that mirrors the hierarchical architecture of the human brain, it has also revolutionized medical imaging. CNNs automatically learn features, progressing from low-level edges and texture to more complex anomalies, which require the analysis of the anatomy. In neurology of stroke, where "Time is Brain," timely and accurate interpretation of MRI/CT scans is crucial. The current work applies Transfer Learning with ResNet-18 [1], where residual "skip connections" make deep and stable training possible even without vanishing gradients and allow for fine-tuning of a pre-trained visual model with clinical imaging data of stroke patients from a relatively small dataset.

On the basis of the two paradigms discussed above, the present paper offers an AI-Powered Brain Stroke Risk Prediction and Clinical Decision Support System that combines clinical history (tabular) with radiological imaging (images) in a single probability of risk. As a design goal, the present work solves the problem of the "black box," limiting trust of clinicians in AI predictions through the integration of SHAP (SHapley Additive exPlanations) [3]. The solution is implemented in the form of a secure, decoupled web application – React 18 and Flask, with a PostgreSQL database, JWT authentication and automated PDF report generation. Therefore, all benefits of the model's diagnostics are presented to the physician in a usable form rather than in a research environment.

II. LITERATURE SURVEY

The evolution of the risk prediction for stroke includes several generations of methodology. Linear risk scoring of clinical factors such as hypertension and age was developed early on by means of clinical scoring systems like $CHADS_2$. However, due to its inherent linearity, this methodology cannot account for synergistic effects among biomarkers.

With the emergence of ML, the potential for predicting with tabular data was greatly increased since models such as SVM and gradient-boosted ensembles outperformed those based on a simple summation of weighted features. The performance champion of the latter class is undoubtedly XGBoost [2], proposed by Chen and Guestrin, which allows for missing values and L1/L2 regularization, making it the top choice in healthcare tabular prediction tasks.

Concurrently, the application of CNNs for analyzing neuroimaging and residual architectures introduced by He et al. [1] made it possible to detect ischemic/hemorrhagic markers in scans with just enough data to transfer a pre-trained visual model, not training a network from scratch. However, one of the common threads in the relevant literature is the "trust gap," caused by the opacity of model reasoning. As a result, various solutions were developed to provide the interpretability of predictions, one of them being the SHAP framework [3], which relies on cooperative game theory and gives a quantifiable value of the contribution of each input feature to the prediction.

More recently, the multi-modal fusion of tabular and imaging data became the subject of investigation, allowing for an improvement in the discrimination (increased AUC) of such models compared to unimodal approaches [12]. Moreover, it became obvious that predictive models are clinically useful only if integrated into a secure, user-friendly Clinical Decision Support System (CDSS) built on modern web stacks like React and Flask [6], [7]. The present work occupies the intersection of the mentioned strands of research – a pipeline combining two models (XGBoost [2] + ResNet-18 [1]), an explainability layer (SHAP [3]) and a CDSS, trained and validated on a public stroke dataset [4] with PyTorch [5].

III. EXISTING SYSTEM VS. PROPOSED SYSTEM

Currently, in clinical practice, risk assessment of the occurrence of a stroke is mostly carried out by applying a manual scoring method, for example, $CHADS_2$ and NIHSS. According to this procedure, a specialist manually assigns a point for each binary/categorical risk factor from the list of criteria. Such a procedure is highly subjective, it does not allow modeling non-linear relationships between biomarkers (e.g., the fact that BMI and glucose increase the risk of a stroke together more than separately) and, most importantly, the division of clinical history and imaging into separate processes. The latter leads to a clinically significant delay since each minute of the ischemic period entails a significant neuron loss. The described problem can be solved by using the multi-modal CDSS with a single integrated pipeline that combines structured clinical data with imaging and generates a Risk Probability Index as the output. Each inference is accompanied by the SHAP explanation, and the system itself uses a Role-Based Access Control mechanism and JSON Web Token authentication to be applied in practice, not just in scientific experiments.

IV. SYSTEM ARCHITECTURE AND METHODOLOGY

A. Overall Architecture

The presented platform utilizes the client-server architecture. The front-end is a single-page React 18 web application providing role-based dashboards for doctors and administrators, real-time form validation (React-Hook-Form with Yup) and data visualization in the form of predictions and SHAP charts made with Chart.js library. Backend part is a Python Flask API for orchestrating the requests, processing the input images and preparing the input tabular data for further inference. For storing the patient data, predictions, and audit logs, the backend uses PostgreSQL database supporting ACID transactions and for securing the communication channel, Flask-JWT-Extended library provides issuing and verification of stateless JSON Web Tokens that contain the user role in the payload and are used for access control in the API endpoints.

B. Dataset and Preprocessing

In order to train the tabular classifier, a public stroke prediction clinical dataset [4] was used. It contained the demographic and physiological attributes (age, hypertension, heart disease, mean glucose level, BMI, smoking and many others). All categorical attributes were numerically coded, and all numerical attributes were normalized in order to avoid dominance of the values with wider numeric range. To use MRI/CT images in the imaging classifier, images were rescaled to 224×224 px, normalized against the mean and standard deviation of the training data and converted to PyTorch tensor [5].

C. Clinical Risk Module (XGBoost)

For training the tabular pipeline, XGBoost classifier [2] was selected for its proven effectiveness in imbalanced health care datasets because of the presence of scale_pos_weight parameter that allows to tune the focus towards the minority class. Also, due to feature_importance/gain native implementation, all features' importance is clear and easy to analyze and, thus, it allows to implement the SHAP layer.

Both L1 and L2 regularization prevent the overfitting on a comparatively small number of patient samples. Inference process consists of data vectorization and normalization, running of the pre-trained ensemble of decision trees and transforming the leaf scores sum with the help of a sigmoid function.

D. Imaging Module (ResNet-18 with Transfer Learning)

The imaging pipeline utilizes the ResNet-18 CNN [1], a neural network whose residual architecture represents the best trade-off between representational depth and inference speed in the context of medical imaging, which usually relies on smaller datasets. In order to specialize the network for detection of radiological stroke indicators in MRIs and CTs — often very small differences in tissue density between healthy white matter and an ischemic infarct — rather than to train a whole new one from scratch, we use a pre-trained ResNet-18 CNN, initialized with ImageNet weights and a binary classifier as the last layer ("Normal" vs. "Stroke Indicators Detected"). Only the late, more specialized layers will be fine-tuned on the labeled MRIs/CTs.

E. Explainable AI Module (SHAP)

To address the issue of lack of interpretability, which hinders clinicians' acceptance of ensemble or deep-learning model predictions, the system computes SHAP (SHapley Additive exPlanations) values [3] for each prediction. Based on cooperative game theory, SHAP assigns to each feature of the input a quantitative contribution to the final risk score, represented in the frontend by force plots and summary charts. Thus, a clinician can see, for each individual, which factors — an elevated average glucose level, hypertension or age — were the cause of a high risk classification. It provides additional guarantees of model correctness and sanity-checking of the logic behind the prediction — for instance, pointing at cases where the score is driven by one single anomalous feature.

F. Security, Audit Logging and Automated Reporting

Being a sensitive data management system, every authenticated request made by a user is handled by Role-Based Access Control based on JWT tokens, where access permissions are divided into "Doctor" and "Admin". The system logs all of the security-relevant actions: authentication events, prediction requests and generation of the report, allowing administrators to have a traceable log immutable audit trail, consistent with the requirements of healthcare data governance, like HIPAA/GDPR [10]. After the prediction is made, the Report Generation Module uses ReportLab [11] to generate a PDF report of the fused risk score, SHAP explanation and the patient's clinical summary.

V. RESULTS AND DISCUSSION

All of the components described above – from clinical data ingestion to PDF report compilation – were implemented and functionally tested end-to-end in the React–Flask–PostgreSQL stack introduced in Section IV. Per-feature gain scores exposed by the XGBoost module were verified to match the known risk factors (age, hypertension, average glucose level and BMI), and SHAP layer was verified to output force plot explanations that could be traced back to these exact features, thus ensuring the transparency goal mentioned in Section I. ResNet-18 module was also verified to process uploaded MRI/CTs in a pipeline that consists of resize, normalization, tensorization and inference steps, returning an imaging-based risk contribution to the fused score.

Discussion. In line with multi-modal fusion research [12], the combination of tabular model and imaging model is likely to deliver a more complete risk picture than each of the modalities separately, because they capture different types of information — long-term physiological history vs acute anatomical state. Explicit incorporation of SHAP-based explanations is supposed to tackle the problem of lack of interpretability outlined in Section II. Quantitative benchmarking of the deployed models — accuracy, precision, recall, F1-score and AUC on the held-out clinical test set — is planned after the collection of larger dataset, and the authors recommend reporting these metrics alongside with this paper, rather than providing illustrative numbers here.

VI. CONCLUSION

In summary, this paper has described an AI-Powered Brain Stroke Risk Prediction and Clinical Decision Support System which integrates an XGBoost classifier of structured clinical data with a ResNet-18 deep neural network for medical image analysis into a single Risk Probability Index. With the incorporation of Explainable AI based on SHAP values, this approach directly tackles the challenge of the "black box" problem, providing a feature-level explanation for each prediction instead of the probability value. From the architectural perspective, the separation of React 18 frontend and Flask backend, the PostgreSQL database, the use of JWT-based Role-Based Access Control and forensic auditing, and the creation of PDF reports prove that the system was designed not just as a research prototype but as a potential candidate for implementation in clinical practice. Thus, this paper demonstrates that the ability to make predictions can be combined with explainability and audibility in one system with the aim to help detect stroke risk earlier and more evidence-based.

VII. FUTURE ENHANCEMENTS

Four main directions of future development have been identified. First, interoperability: adaptation of HL7 FHIR standards for automatic population of patient demographics, lab results and previous images from the Electronic Health Records (EHR) system, such as Epic and Cerner. Second, scalability: generalization of the hybrid pipeline of structured data plus medical image analysis not only for the purpose of detection of stroke risk but for a larger range of neurodiagnostics, such as brain tumors, Alzheimer's disease-related atrophy and multiple sclerosis lesions. Third, accessibility: quantization of the models for execution even in the conditions of poor connectivity using such tools as TensorRT or ONNX. And, fourth, prognosis: adding RNN/LSTM-based time-series forecasting for longitudinal prediction of stroke risk.

VIII. ACKNOWLEDGMENT

We would like to express our gratitude to the Department of Information Technology, St. Peter's College of Engineering and Technology (Anna University), Chennai, for their guidance and facilities.

REFERENCES

[1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.
[2] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785–794.
[3] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.
[4] Kaggle, "Stroke Prediction Dataset," 2021. [Online]. Available: https://www.kaggle.com/datasets/fedesoriano/stroke-prediction-dataset
[5] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, 2019.
[6] M. Grinberg, Flask Web Development: Developing Web Applications with Python. Sebastopol, CA: O'Reilly Media, 2018.
[7] A. Boduch and R. Derks, React and React Native: A Complete Guide to Building Cross-Platform Apps with React. Birmingham, U.K.: Packt Publishing, 2020.
[8] World Health Organization, "Stroke: A global medical emergency," 2022. [Online]. Available: https://www.who.int
[9] M. Winand, SQL Performance Explained. Self-published, 2012.
[10] L. O. Gostin and S. J. Nass, Beyond the HIPAA Privacy Rule: Enhancing Privacy, Improving Health Through Research. Washington, DC: National Academies Press, 2009.
[11] ReportLab Europe Ltd., "ReportLab PDF Library User Guide," 2024. [Online]. Available: https://www.reportlab.com/docs/reportlab-userguide.pdf
[12] S. C. Huang, A. Pareek, S. Seyyedi, I. Banerjee, and M. P. Lungren, "Fusion of medical imaging and electronic health records using deep learning: A systematic review and implementation guidelines," NPJ Digital Medicine, vol. 3, no. 1, pp. 1–9, 2020.