

Neuroimage Analysis for Stroke Detection: A Machine Learning Framework

Vidya D. Argade¹, Diksha Gautam², Diksha Singh³ and Raman Vishwakarma⁴

Computer Engineering, Dhole Patil College of Engineering, Savitribai Phule Pune University, Pune, India

Abstract—This project presents an AI-powered framework for early stroke detection using MRI and CT neuroimages. Leveraging deep learning (DL) models like CNN, the system offers accurate stroke classification and integrates a user-friendly web interface for patients and radiologists. With role-based access, AI chatbot support, and modular design, the framework aims to enhance diagnostic efficiency and accessibility, especially in underserved regions.

Keywords— Stroke Detection, Medical Imaging, MRI, and Deep Learning

I. INTRODUCTION

A stroke is a critical medical condition that occurs when blood flow to the brain is obstructed or reduced, resulting in damage to brain tissue. It continues to be a leading cause of mortality and chronic disability globally. Prompt detection and rapid medical response are essential to improve recovery chances and minimize permanent neurological damage. Traditionally, stroke diagnosis relies on the interpretation of neuroimaging scans, such as CT and MRI, by experienced radiologists. Although this method is generally reliable, it can be time-consuming, influenced by human judgment, and dependent on the availability of skilled professionals. In emergency situations, any delay in diagnosis can significantly impact the patient's prognosis and recovery outcome.

Recent advancements in artificial intelligence (AI) and machine learning (ML) have enabled the automation of medical image analysis. DL approaches, especially Convolutional Neural Networks (CNNs), have shown significant capability in identifying complex patterns in neuroimaging data. These approaches contribute to quicker and more accurate assessments, reduce the diagnostic burden on medical professionals, and help decrease the risk of human error.

Within this framework, designing an intelligent system for stroke detection using ML models represents a major advancement in modern healthcare. Such solutions aim to deliver accurate, reliable, and timely diagnoses, ultimately improving patient care and outcomes.

1.1 Sub Heading

"Neuroimage Analysis for Stroke Detection: A Machine Learning Framework" is an AI-driven system designed to assist in early and accurate stroke diagnosis using CT and MRI brain scans. By integrating deep learning models like CNN, the framework automates image analysis and delivers real-time results through an intuitive user interface. By incorporating features such as an AI-powered chatbot and automated report delivery, the system supports better clinical decisions and improves healthcare accessibility for patients in remote areas, representing a major advancement in smart, scalable, and patient-oriented medical solutions.

II. RELATED WORK

In recent years, machine learning has increasingly been employed in medical diagnostics, particularly in assessing stroke risk and analyzing neuroimaging data. Researchers have examined a diverse array of techniques, from fundamental supervised algorithms to more advanced deep learning models.

Machine learning techniques have demonstrated potential in identifying stroke risk by analyzing patient historical medical data, such as information found in electronic health records (EHRs). Techniques such as Extreme Gradient Boosting (XGBoost) have achieved high predictive accuracy for stroke risk in hypertensive patients by identifying relevant features and patterns in large datasets [1]–[3]. Ensemble models combining various algorithms have further enhanced reliability in detecting individuals at elevated risk [4].

Skull stripping plays a vital role in preparing neuroimages for accurate analysis. According to Buda et al. [5], this process can be addressed using traditional techniques or modern DL approaches. CNNs, in particular, have revealed potential in enhancing the effectiveness and reliability of routine clinical workflows.

Deep reinforcement learning (DRL) has evolved in intricate predictive tasks, including temporal link prediction in dynamic networks. For instance, Kumar et al. [6] demonstrate how DRL can outperform static supervised methods in predicting hidden links in criminal networks, while Grover et al. [7] introduce a generative model for graphs that enhances learning from evolving network structures.

Within stroke diagnosis via imaging, several studies have employed CT and MRI data to predict stroke severity and subtype. Nijjati et al. [8] used CNN-based hybrid structures to predict Modified Rankin Scale (mRS) outcomes with 74% accuracy, while Shree et al. [9] integrated wavelet entropy-based features and probabilistic neural networks for classifying MRI images into normal, stroke, and degenerative disease categories, reporting high classification performance.

Saleem et al. [10] phasiemze the role of patient demographics and clinical attributes—such as age, heart conditions, blood glucose levels, and hypertension—in enhancing prediction models. They propose the use of a perceptron neural network trained on balanced datasets, addressing the limitations of unbalanced data in earlier models.

Advanced deep learning architectures have also been effective in analyzing intracranial hemorrhages. Hussein et al. [11] proposed a framework capable of distinguishing between hemorrhage subtypes (epidural, subdural, and intraparenchymal) with an average accuracy of 96.21%, incorporating lesion quantification to aid emergency clinical decisions.

Hybrid feature extraction techniques that combine Discrete Wavelet Transform (DWT), Gray Level Co-occurrence Matrix (GLCM), and Discrete Cosine Transform (DCT) have proven valuable in classifying CT images. Studies by Huang et al. [12] and Gaurav et al. [13] employed Random Forest classifiers, achieving high classification accuracy by leveraging transformed texture and statistical features.

Preprocessing strategies and neural model design also significantly influence classification outcomes. Ryu et al. [14] demonstrated that integrating CNN and RNN components, along with enhanced

preprocessing pipelines, improved hemorrhage detection performance, reinforcing the utility of deep learning as a clinical decision support tool.

In the current healthcare landscape, stroke diagnosis primarily relies on traditional neuroimaging techniques that involve CT scans, MRI, and ultrasound imaging, interpreted manually by radiologists. These methods, while effective to a certain degree, are often time-consuming, subjective, and prone to human error, especially in emergency settings where time is critical. The existing systems do not typically integrate machine learning or automation, meaning stroke detection heavily depends on the availability and experience of medical professionals. Conventional diagnostic software might highlight anomalies or offer basic image enhancement tools, but it lacks intelligent analysis capabilities to automatically detect strokes with high accuracy. In terms of patient interaction, existing systems are also not user-centric. They do not provide online access for report viewing on neuroimage interpretation. Current technologies inadequately deliver prompt, automated, and accessible stroke detection help, underscoring the pressing necessity for a more sophisticated, integrated, and user-centric solution.

Current models often struggle with limitations that involve inadequate real-time performance, high computational demands, and challenges in handling data from different imaging sources. To address these concerns, the proposed solution utilizes a CNN built upon the VGG16 architecture—an advanced deep learning model originally created by the Visual Geometry Group at Oxford. The architecture is modified by removing its original fully connected layers, transforming it into an efficient feature extractor. A customized binary classification head is then added to differentiate between stroke and non-stroke cases. The terminal output layer comprises a solitary neuron utilizing sigmoid activation, making it ideal for binary classification applications. The model is improved to include both CT and MRI image inputs and is connected with a chatbot, thereby merging high diagnostic precision with an approachable and engaging user interface.

III. PROPOSED WORK

Machine learning—especially deep learning approaches like CNNs—has become increasingly prominent in medical imaging, offering valuable support for early detection of diseases such as stroke.

CNN architectures have demonstrated strong performance, often achieving over 90% accuracy when classifying CT and MRI images. Nonetheless, these models can struggle with overfitting, particularly when training on limited data. This issue can be mitigated through transfer learning methods, such as adapting pre-trained models like VGG16 to suit medical imaging tasks.

As illustrated in Figure 1, the proposed stroke detection system harnesses DL, specifically a CNN built upon the VGG16 architecture, to streamline and improve the analysis of neuroimages such as CT and MRI scans. The system is designed to identify and pinpoint stroke-affected regions within the images, thereby increasing diagnostic accuracy and minimizing the need for manual assessment. The process begins with several preprocessing steps that involve converting images to grayscale, applying histogram equalization, and using CLAHE (Contrast Limited Adaptive Histogram Equalization) to enhance important image features. After preprocessing, the VGG16 model is employed as a feature extractor, with its original dense layers removed. A custom binary classification head is added to differentiate between normal and stroke cases. The training process employs binary cross-entropy as the loss function and utilizes the Adam optimizer. To ensure the model generalizes well, techniques like data augmentation, class weighting, and early stopping are incorporated to handle imbalances and prevent overfitting.

A significant advancement of the proposed system lies in its dual-interface design: one for medical professionals (radiologists) and another for patients. The admin dashboard enables healthcare providers to review diagnostic results, give feedback to the patient, while the patient interface allows users to access reports and interact with a chatbot. This chatbot serves as an informative assistant, addressing basic stroke-related queries, offering guidance on symptoms, and explaining test procedures, though it does not perform medical diagnoses. Additionally, the system is backed by a secure SQL database for managing patient records and reports, with automated email notifications sent to patients once their results are ready, ensuring timely and convenient access to medical information.

The integration of these features, including a user-friendly interface for both healthcare professionals and patients, automated email notifications, and an AI-powered chatbot, sets this system apart from

existing models. Furthermore, the optimized use of the VGG16 architecture for feature extraction, along with advanced training techniques like dropout and class weighting, helps improve accuracy while minimizing overfitting. This innovative framework not only empowers healthcare providers with a powerful diagnostic tool but also offers patients an accessible and efficient way to manage their health.

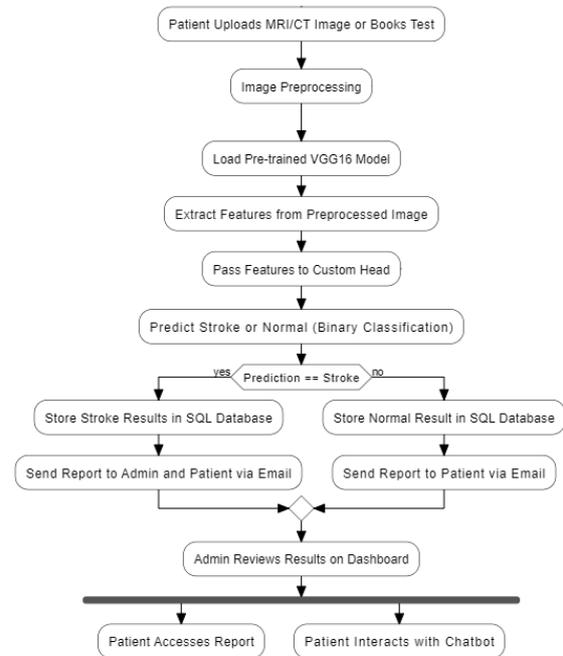


Figure 1: Proposed System

3.1 Data Processing

In the context of machine learning-based image classification, data processing serves as the foundational step to ensure that input data is clean, consistent, and ready for modeling. In Figure 2, the process begins with the dataset undergoing preprocessing, which may include removing noise, correcting artifacts, or handling missing or irrelevant data. This is followed by normalization, which scales data into a consistent range, essential for models that rely on numerical stability. The processed data is then split into training and testing sets through cross-validation, ensuring robust model evaluation. During the training phase, feature extraction identifies key patterns or attributes from the images, and feature selection retains the most relevant ones to reduce dimensionality and improve learning efficiency. A classifier is then trained on this curated data, resulting in a trained module that encapsulates learned knowledge. This module is later applied to the testing data, where the same sequence of feature extraction and selection occurs. The classifier then makes predictions based on previously learned patterns. The

output is a predicted result that undergoes further analysis to assess performance, accuracy, and clinical relevance. This structured data processing pipeline

ensures that the entire system, from raw data input to final result interpretation, is optimized for reliable and interpretable outcomes.

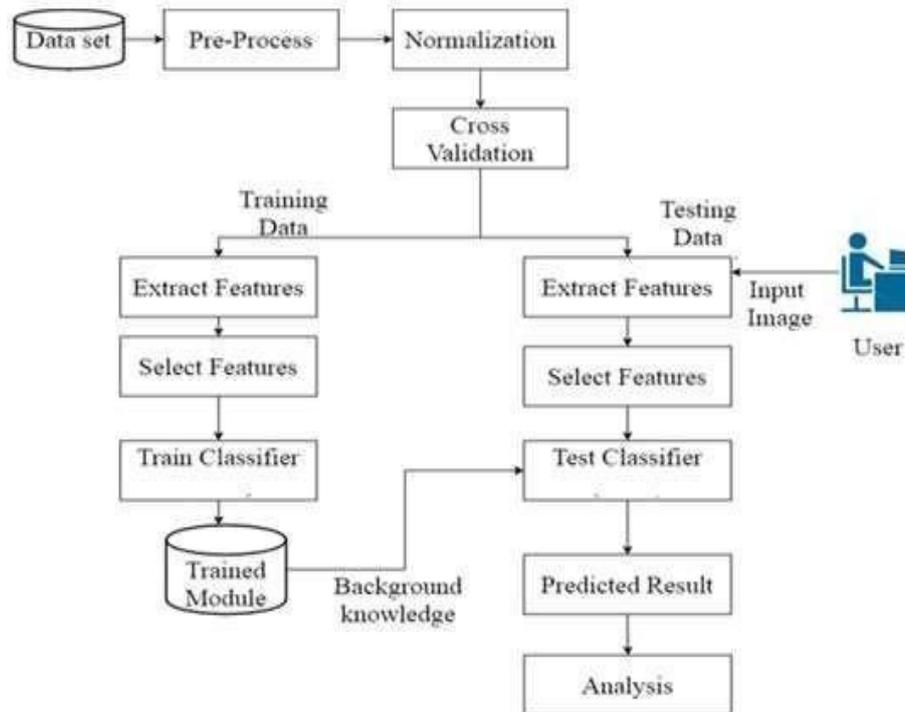


Figure 2: Data Processing

3.2 VGG16 Model

In Figure 3, a deep convolutional neural network originally trained on the large-scale ImageNet dataset. VGG16 is widely recognized for its depth and simplicity, consisting of 13 convolutional layers and 3 fully connected layers arranged in a very uniform architecture. The key advantage of using VGG16 lies in its ability to extract rich hierarchical features from images, which can be transferred and reused for different computer vision tasks beyond the original classification categories in ImageNet.

For the purpose of stroke classification in neuroimages, the top layers of the pre-trained VGG16 model, which include the dense classification layers, are removed. This leaves behind the convolutional base, which acts as a robust feature extractor. The convolutional layers are retained in their pre-trained state and are not updated during training. This “frozen” configuration ensures that the valuable general features learned from the large ImageNet dataset, such as edge detectors and texture patterns, are preserved and reused effectively in the new task. Freezing these layers also significantly reduces training time and mitigates the risk of overfitting, especially in domains like medical imaging, where labeled data may be limited.

To adapt the model for binary stroke classification, a custom classification head is added on top of the frozen base. This custom head is designed to learn task-specific patterns from the features extracted by the base model. The first layer in this custom head is a Flatten layer, which transforms the multi-dimensional feature maps from the final convolutional layer into a one-dimensional vector. This flattened representation serves as the input to a fully connected Dense layer consisting of 256 neurons with ReLU (Rectified Linear Unit) activation, allowing the model to learn complex, non-linear combinations of the input features.

To further prevent overfitting, especially given the relatively small size of many medical datasets, a Dropout layer is included after the dense layer. Dropout works by randomly disabling a fraction of the neurons during training, forcing the network to develop redundant and generalized feature representations. Finally, the last layer in the custom head is a Dense layer with a single neuron and a sigmoid activation function. This layer outputs a probability value between 0 and 1, representing the likelihood of the input image being associated with a stroke. A probability threshold (commonly set at 0.5)

is then applied to determine the final class label: stroke (1) or normal (0).

This architecture leverages the strength of VGG16’s generalized visual knowledge while enabling specialized learning through a lightweight, trainable head. It provides an efficient and effective framework for applying deep learning to medical imaging tasks such as stroke detection, where domain-specific labeled data may be limited but diagnostic precision is critical.

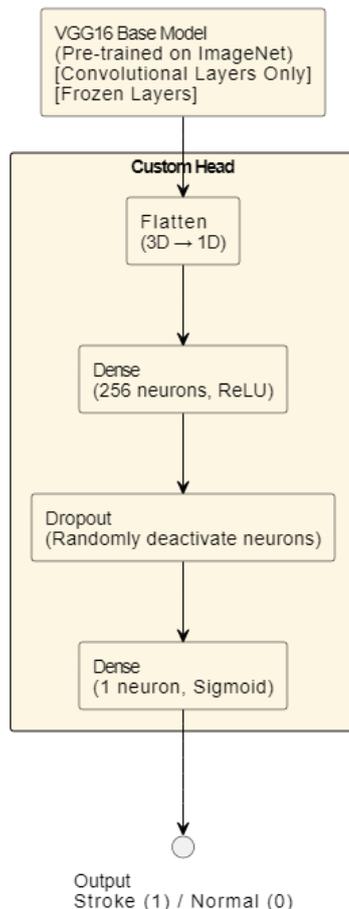


Figure 3 VGG16 Model Transfer Learning Architecture

3.3 Algorithms

- Image Preprocessing & Enhancement:
 - Grayscale conversion
 - Histogram equalization
 - Contrast Limited Adaptive Histogram Equalization (CLAHE)
 - Gamma correction
 - Resizing to 224x224 pixels and normalization
- Transfer Learning using VGG16:
 - The foundational model employed is VGG16, pre-trained on ImageNet, excluding the upper layers. All layers in the

VGG16 architecture are immobilized, permitting the utilization of pre-trained features without modification.

- A custom head is added, which begins with a Flatten layer followed by a Dense layer containing 256 neurons with ReLU activation.
- To reduce overfitting, a Dropout layer is incorporated. The final layer is a Dense layer featuring a single neuron and sigmoid activity, intended for binary classification.
- Loss Function:
 - Binary Cross-Entropy: Used because this is a two-class classification problem.
- Optimizer:
 - Optimizer Used: Adam, an adaptive learning rate optimization algorithm.
 - Learning Rate: Set to 0.0001 for stable and efficient training.
- Model Training Enhancements:
 - Employing Data Augmentation with ImageDataGenerator to mitigate overfitting and enhance generalization.
 - Early Stopping to halt training if the model stops improving.
 - Model Checkpoint to save the best-performing model on validation accuracy.
 - Class Weighting to handle class imbalance between Stroke and Normal images.

IV. RESULTS AND DISCUSSION

The developed machine learning framework for stroke detection using neuroimages has shown highly effective results in classifying brain scans as stroke or non-stroke with strong accuracy. By leveraging the VGG16 convolutional neural network for deep feature extraction, the system efficiently processes pre-processed MRI and CT images to identify potential stroke indicators. The integration of advanced preprocessing techniques like CLAHE, normalization, and gamma correction further enhances image clarity and model performance. Testing reports confirm that the model delivers consistent and reliable predictions across diverse input data. Additionally, the user-friendly interface allows patients to upload scans and receive results, while radiologists can review and manage cases with ease. The chatbot integration and email reporting features also improve user engagement and accessibility. Overall, the system demonstrates great

potential for assisting in early stroke diagnosis and reducing the workload on healthcare professionals.

systems, contributing to timely treatment and improved patient outcomes.

4.1 OUTCOMES

The model supports healthcare professionals by offering fast, automated analysis and shows strong potential for integration into clinical decision-making

4.2 SCREENSHOTS

The Screenshots attached below show the working of the model.

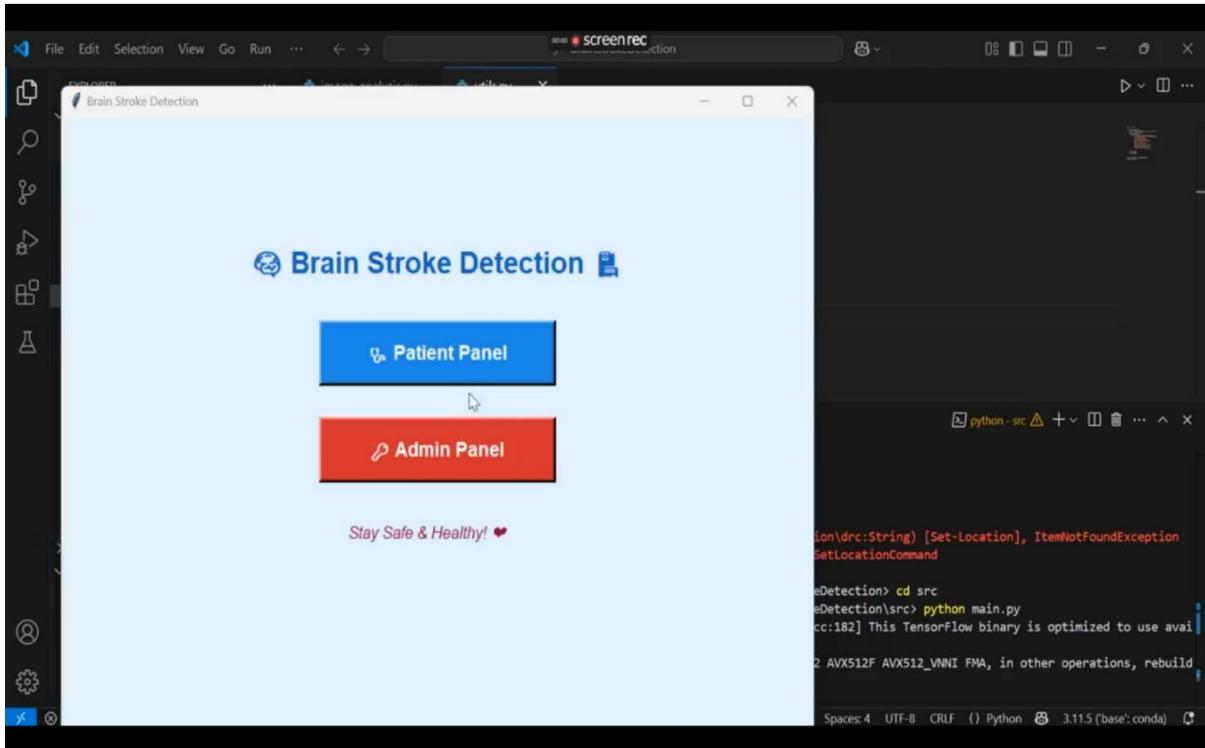


Figure 4.2.1: Home Page

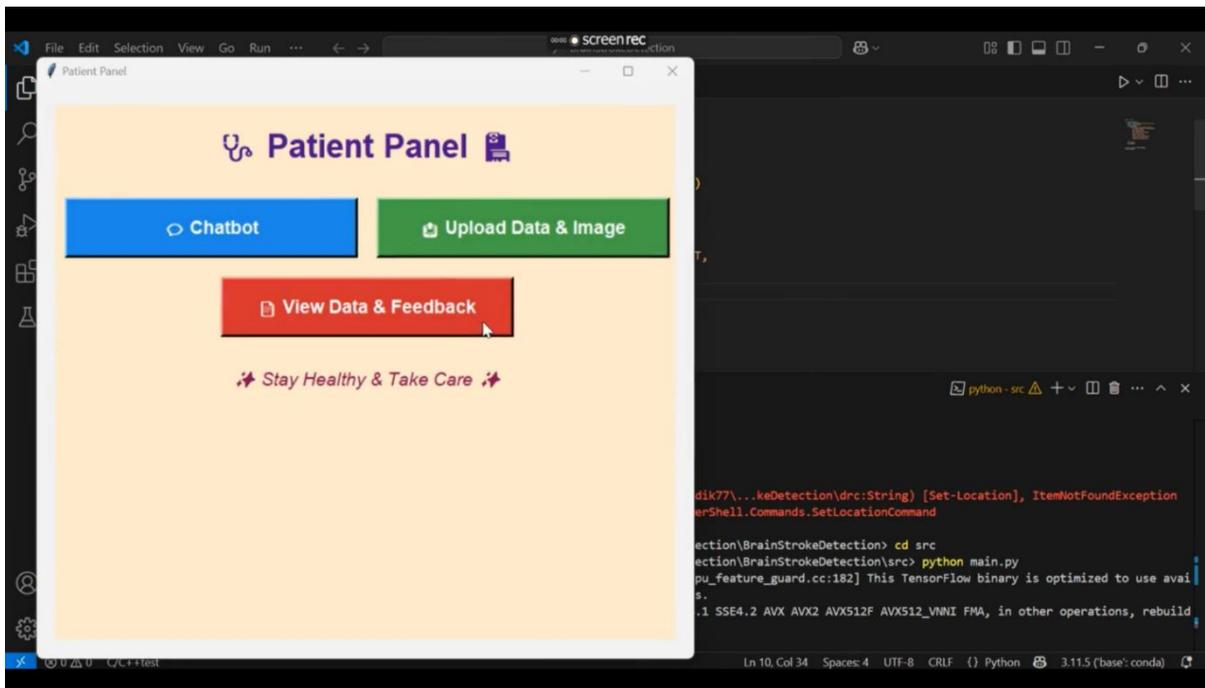


Figure 4.2.2: Patient Panel

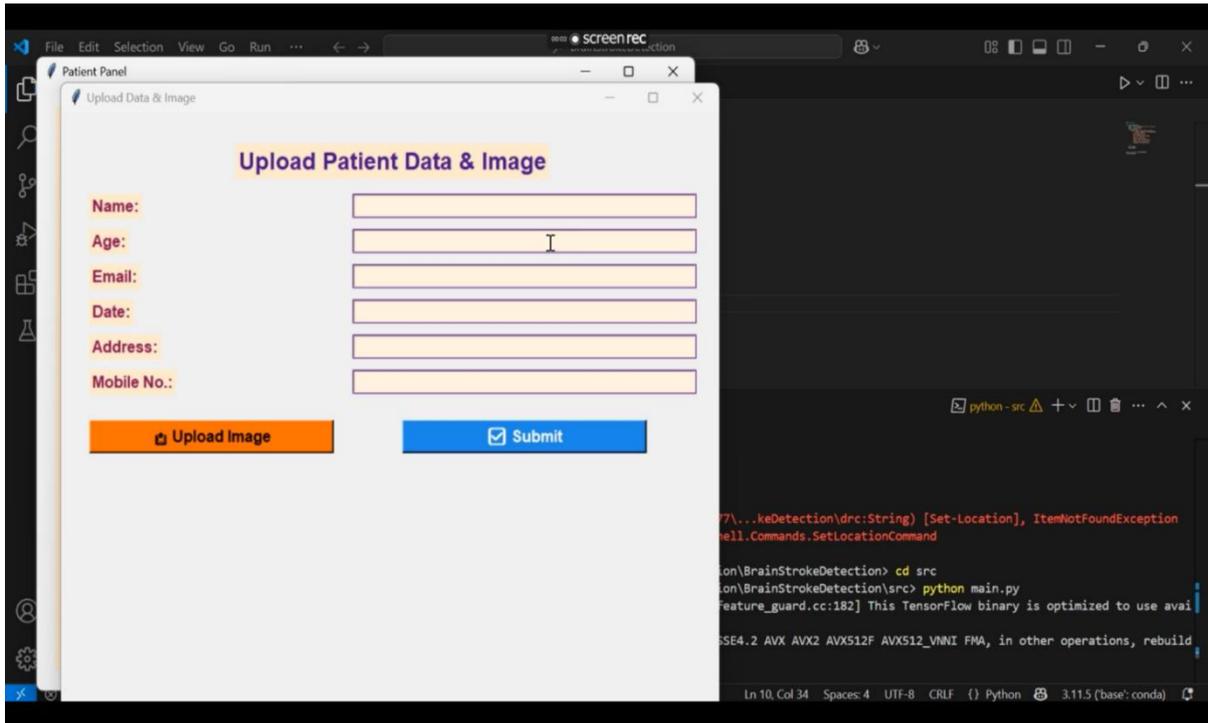


Figure 4.2.3: User Register

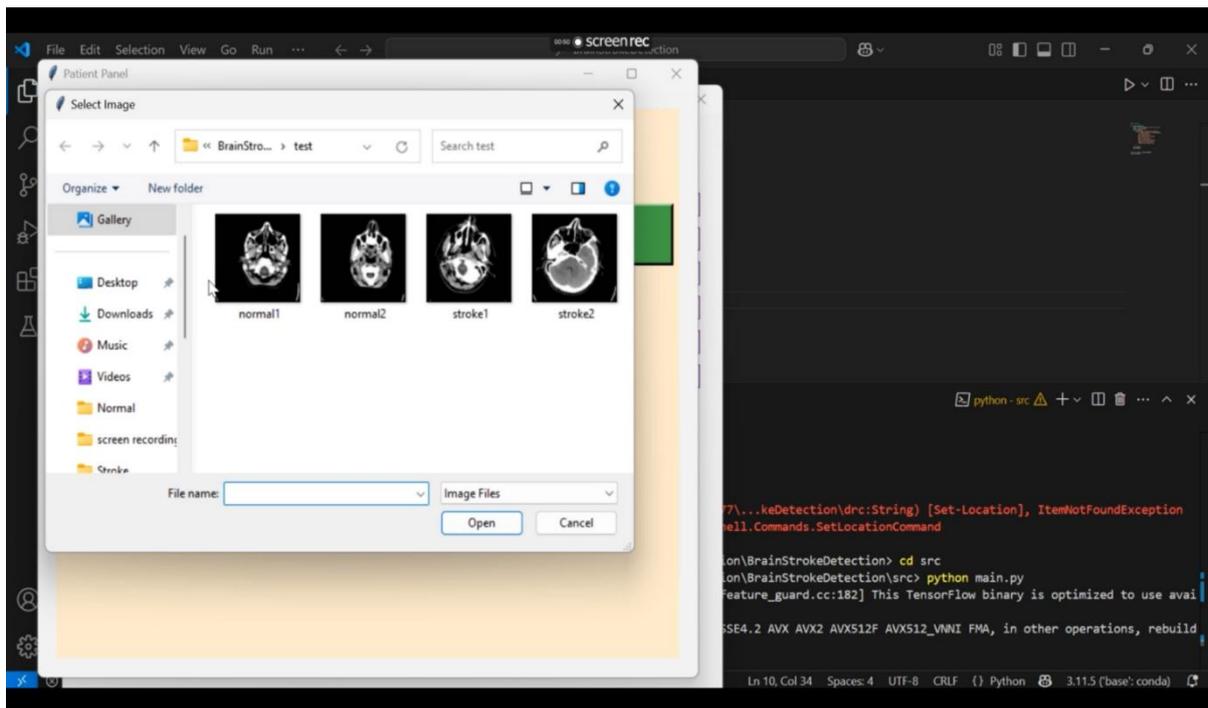


Figure 4.2.4: Upload Image

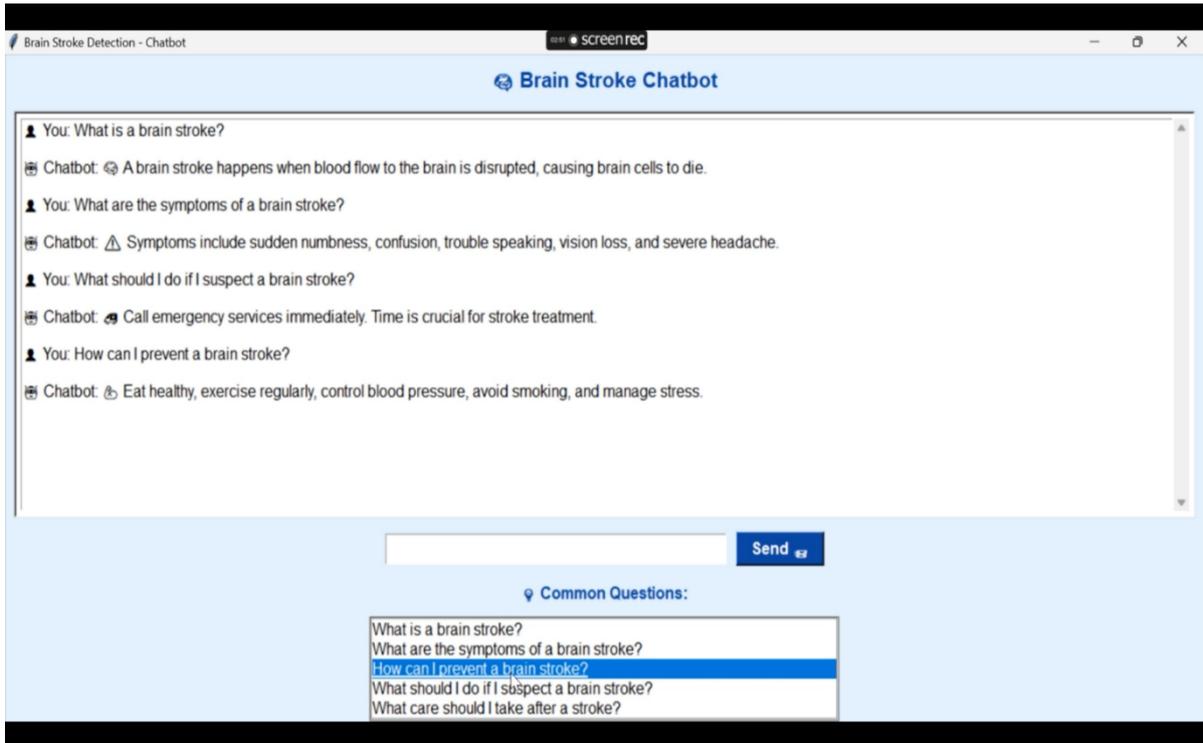


Figure 4.2.5: Chatbot

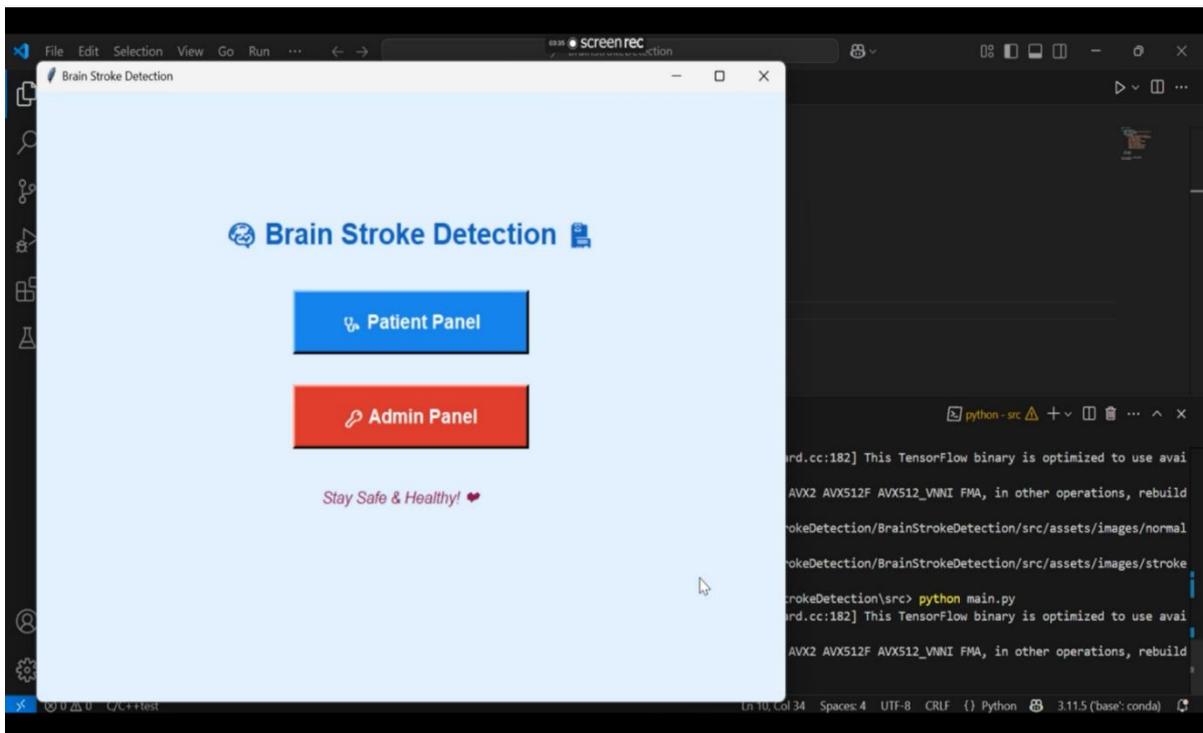


Figure 4.2.6: Home Page

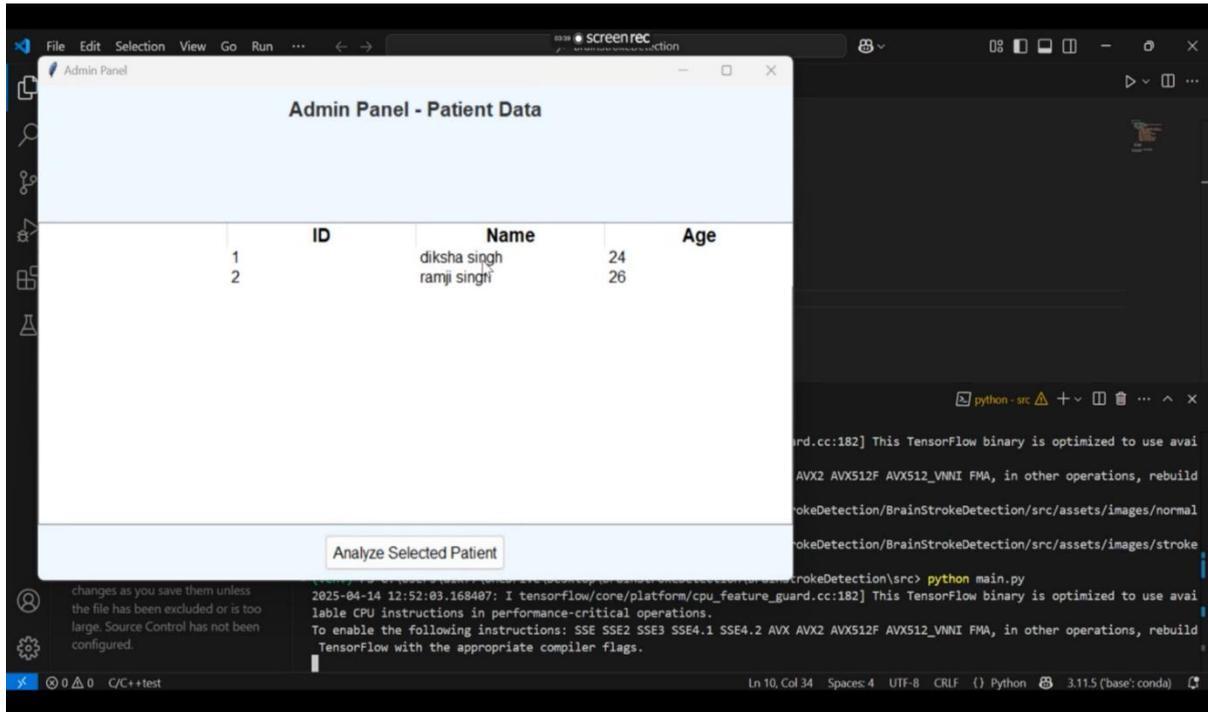


Figure 4.2.7: Admin Panel

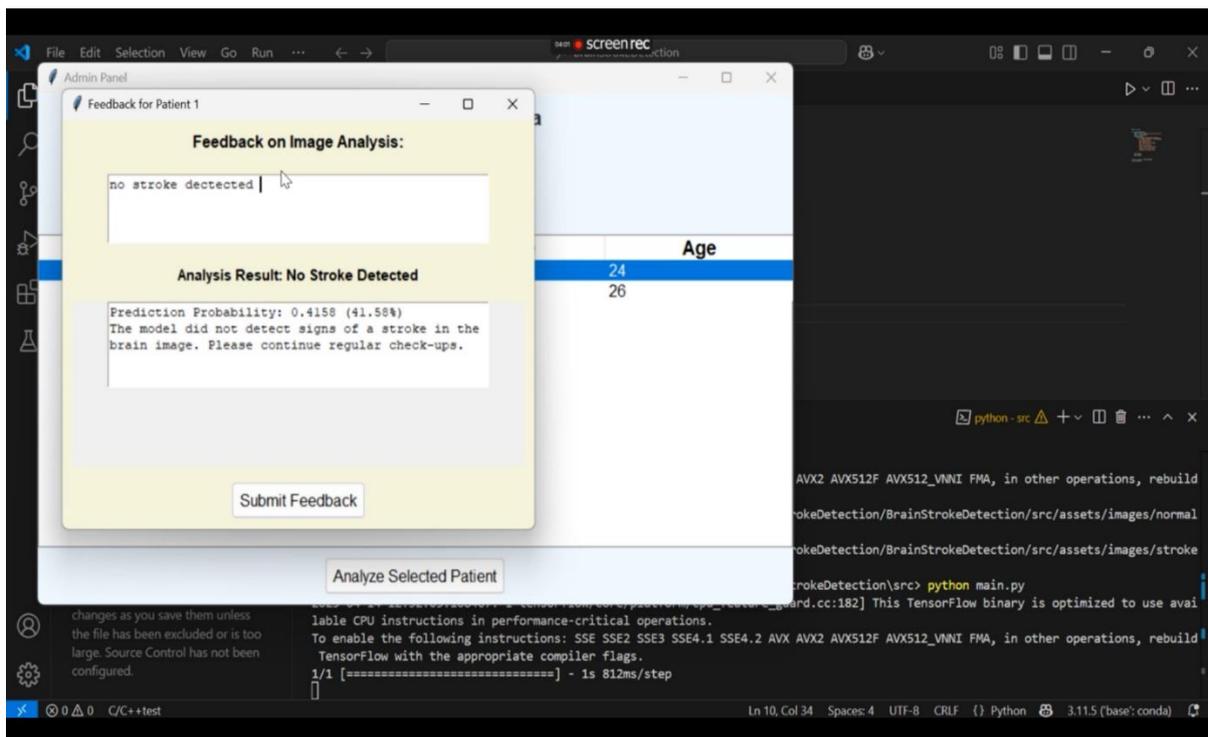


Figure 4.2.8: Feedback On Image Analysis

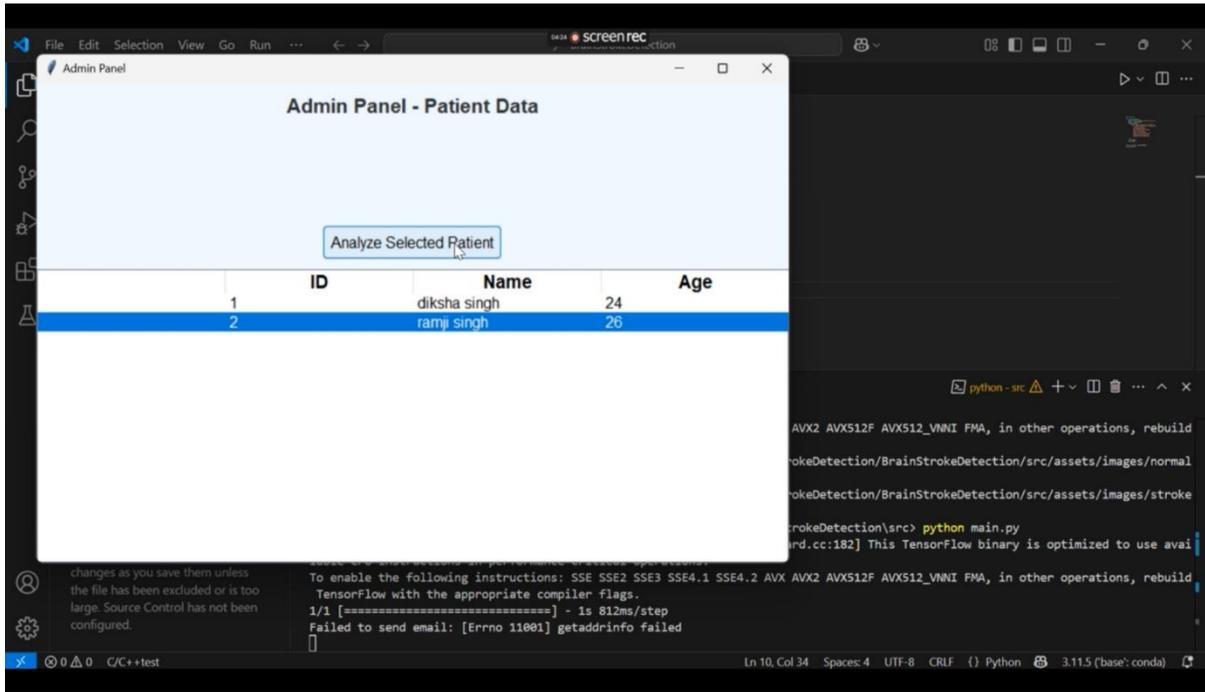


Figure 4.2.9: Analyse Selected Patient

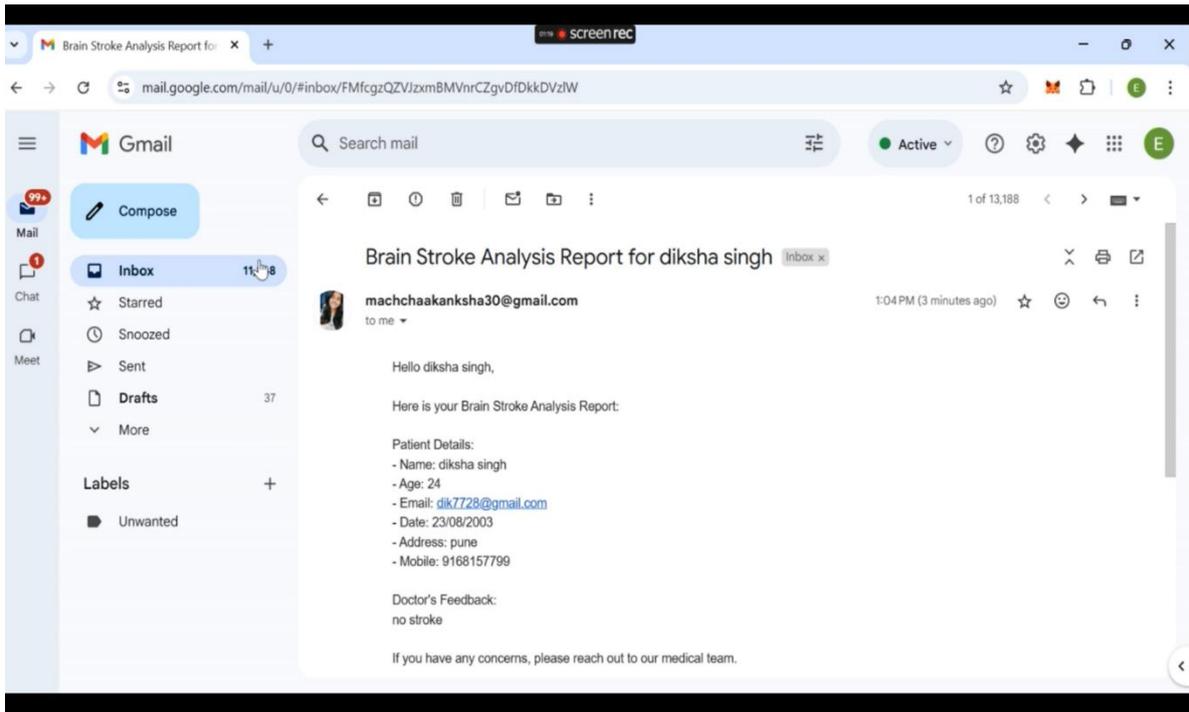


Figure 4.2.10: Report Received Through Mail

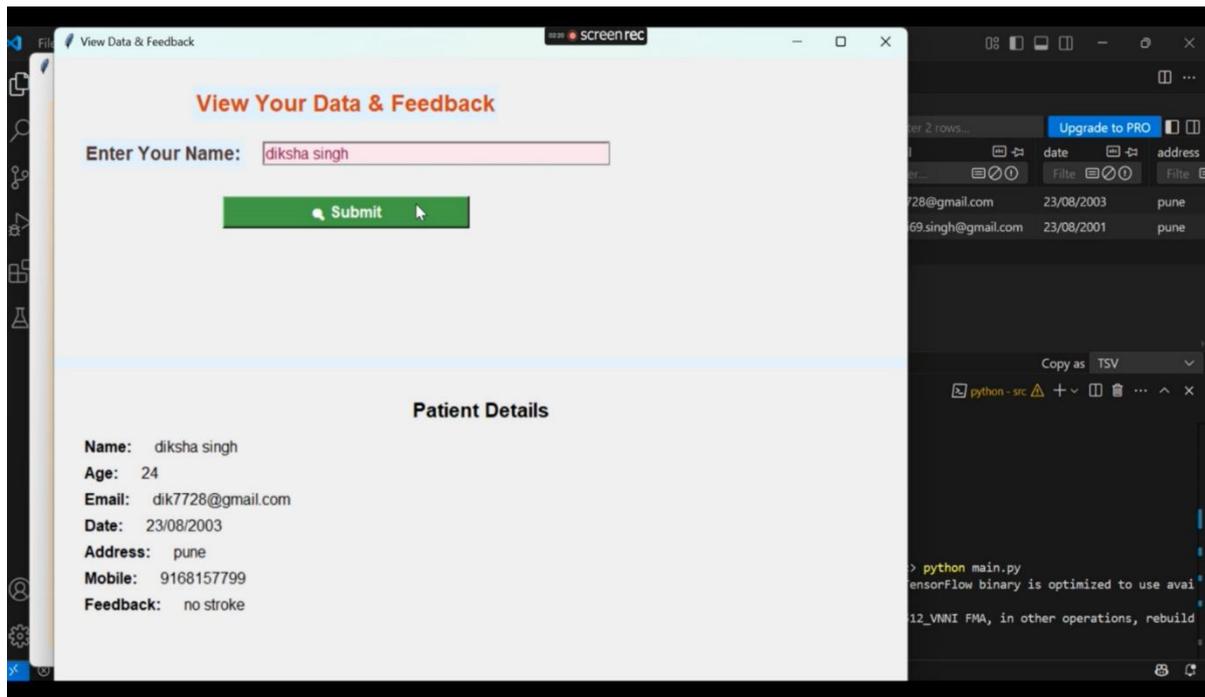


Figure 4.2.11: View Data & Feedback

V. CONCLUSION

In conclusion, the proposed stroke detection framework effectively illustrates the application of CNNs, specifically VGG16, in medical imaging for early stroke diagnosis. The system provides a reliable and automated solution for analyzing CT and MRI scans, supporting healthcare professionals with fast and accurate predictions. Its user-friendly design, combined with essential features like image upload, report generation, and chatbot assistance, makes it accessible for both patients and doctors. By reducing dependency on manual interpretation and improving diagnostic efficiency, this framework holds strong potential for integration into clinical settings and telemedicine applications, ultimately contributing to better patient outcomes and timely intervention.

REFERENCES

- [1] Zhou, Y., et al. (2021). "Accurate Prediction of Stroke for Hypertensive Patients Based on Medical Big Data and Machine Learning Algorithms: Retrospective Study." *JMIR Medical Informatics*, 9(11), e30277.
- [2] Rehman, A., et al. (2024). "An Ensemble Machine Learning and Data Mining Approach to Enhance Stroke Prediction." *Bioengineering*, 11(7), 672.
- [3] Dai, W., Jiang, Y., Mou, C., & Zhang, C. (2023). "An Integrative Paradigm for Enhanced Stroke Prediction: Synergizing XGBoost and xDeepFM Algorithms." *arXiv preprint arXiv:2310.16430*.
- [4] Alam, T., et al. (2024). "A Stroke Prediction Framework Using Explainable Ensemble Learning." *Computers in Biology and Medicine*, 169, 107683.
- [5] Buda, M., Saha, A., & Mazurowski, M. A. (2020). Conventional and Deep Learning Methods for Skull Stripping in Brain MRI. *Applied Sciences*, 10(5), 1773.
- [6] Kumar, A., Gupta, R., & Soni, P. (2020). Deep reinforcement learning for link prediction in criminal networks. *Journal of Computational Intelligence and Applications*, 15(2), 67-79.
- [7] Grover, A., & Leskovec, J. (2016). node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 855-864). ACM.
- [8] Nijati, M., Zhang, L., & Wang, H. (2019). Predicting stroke outcomes using CNN-based hybrid structures: An evaluation of the Modified Rankin Scale (mRS). *Journal of Medical Imaging and Health Informatics*, 9(8), 1254-1262.
- [9] Shree, V., Patel, A., & Gupta, S. (2020). Wavelet entropy-based features for

classification of MRI images in stroke and degenerative disease. *Journal of Neuroscience Methods*, 325, 108356.

- [10] Saleem, M., Ali, M., & Khan, F. (2018). Enhancing prediction models for stroke outcomes using patient demographics and clinical attributes: A perceptron neural network approach. *International Journal of Medical Informatics*, 118, 1-9.
- [11] Hussein, A. A., Sharaf, M., & Rahman, M. A. (2020). A deep learning framework for intracranial hemorrhage subtype classification and lesion quantification. *Journal of Neurosurgery*, 132(5), 1351-1359.
- [12] Huang, H., Li, S., & Zhang, Y. (2019). Hybrid feature extraction using DWT, GLCM, and DCT for CT image classification in stroke diagnosis. *Journal of Medical Imaging*, 26(2), 202-210.
- [13] Gaurav, R., Singh, D., & Sharma, P. (2020). Random forest classification for stroke detection using texture features extracted from CT images. *Journal of Digital Imaging*, 33(1), 118-126.
- [14] Ryu, H. J., Lee, S. H., & Kim, J. (2019). Improving hemorrhage detection in CT images through the integration of CNN and RNN. *IEEE Access*, 7, 136601-136611.

Author's biography



Vidya D. Argade an accomplished academican and software developer with 10 years of total experience—6 years in the software industry and 4 years in academia. Currently serving as an Assistant Professor at the Dhole Patil College of Engineering, Pune, the author specializes in full-stack development with expertise in Python, Django, Flask, Tkinter, C#.NET, Streamlit, and SQL Server. Holding an M.E. in Computer Engineering and a B.E. in Information Technology from Sinhgad College of Engineering, Pune University, the author integrates innovative teaching methodologies—discussions, demonstrations, and case studies—to engage and educate students. With strong skills in object-oriented programming, database design, and SQL, the author

brings adaptability, quick learning, and effective communication to both teaching and development roles.



DIKSHA GAUTAM received her B.E. Degree in Computer Engineering from Dhole Patil College of Engineering, Savitribai Phule Pune University, Pune, India. She is particularly passionate about exploring the applications of intelligent algorithms in solving real-world problems, especially in the domains of healthcare and image analysis. Her research interests include Artificial Intelligence, Full-Stack Development, Image Classification and Deep Learning.



Diksha Singh received her Bachelor of Engineering (B.E.) degree in Computer Engineering from Dhole Patil College of Engineering, affiliated with Savitribai Phule Pune University, Pune, India. During her academic journey, she actively engaged in research and development projects, particularly in the fields of artificial intelligence, machine learning, and medical image analysis. Her recent project focused on developing a machine learning framework for stroke detection using neuroimages, showcasing her interest in applying cutting-edge technology to real-world healthcare challenges.



Raman Vishwakarma, and I am a student at Dhole Patil College of Engineering in Pune, where I have received my BE degree in Computer Engineering. I am passionate about data science and machine learning, and I am currently working on a project titled "Neuroimage Analysis for Stroke Detection: A Machine Learning Framework." This project aims to enhance the accuracy of stroke detection using advanced neuroimaging techniques. I am eager to expand my knowledge in emerging technologies and continuously seek opportunities to learn and innovate in the field of data science.