A CNN Approach for Accurate Stroke Diagnosis Using Brain Computed Tomography Imaging

Heshan Chandeepa Pathmakumara
Department of Computer and Data Science
NSBM Green University
Homagama, Sri Lanka
bhcpathmakumara@students.nsbm.ac.lk

Gayan Perera
Department of Computer and Data Science
NSBM Green University
Homagama, Sri Lanka
gayanp@nsbm.ac.lk

2025 International Mobile, Intelligent, and Ubiquitous Computing Conference (MIUCC) | 979-8-3315-3922-1/25/$31.00 ©2025 IEEE | DOI: 10.1109/MIUCC66482.2025.11196792

***Abstract*— Stroke detection from medical imaging, which is essential for early diagnosis and prompt intervention, greatly impacts patient outcomes. In this study, provide a technique that uses Convolutional Neural Networks (CNNs) to detect strokes from brain CT (computed tomography) pictures. A collection of 2,501 CT scans, including both normal and stroke cases, was used to develop and evaluate our model. A complete preprocessing pipeline was implemented, including training, validation, and systematic separation into test sets to enhance the data, scale the images, normalize, and reduce overfitting. The proposed CNN architecture consists of two convolutional blocks for feature extraction, followed by dense layers. According to the experimental results, the validation accuracy of the model was 92.93%, indicating its high capacity to distinguish between normal images and stroke. Additionally, the model's performance was thoroughly evaluated across a range of classes and thresholds using ROC curves, accuracy, recall, and F1-score. These comprehensive measurements demonstrated high sensitivity and specificity, underscoring the model's potential as a reliable clinical decision-support tool. According to the research, deep learning techniques - in particular, CNNs - have great potential for use in medical imaging applications for stroke diagnosis, as they can provide both high diagnostic accuracy and possible interpretability. Future studies will concentrate on growing the dataset and examining intricate structures and explanatory techniques in order to enhance therapeutic application.**

***Keywords*— stroke detection, brain CT imaging, deep learning, medical image analysis**

I. INTRODUCTION

Stroke A major disease known as stroke is brought on by a disruption in blood flow to the brain, which harms brain cells and deprives them of oxygen and nutrients. It affects millions of individuals every year and is a leading cause of mortality and permanent disability globally. Eighty-seven percent of the 800,000 strokes that are thought to happen annually are ischemic, ten percent are intracranial hemorrhagic, and three percent are subarachnoid hemorrhagic [1]. Although incidence rates are declining in high-income countries (HICs), projections indicate that there will be 3.4 million more stroke cases in the United States by 2030, highlighting the increasing strain on healthcare systems [2]. Although cutting-edge treatments such as thrombolytics and endovascular thrombectomy (EVT) have yielded improved outcomes, access is still uneven as EVT is currently unavailable in low-income countries [3].

The best standard for diagnosing stroke in emergency situations is non-contrast computed tomography (CT), due to its speed, accessibility, and capacity to distinguish between various stroke subtypes. However, CT interpretation is largely dependent on the skill of radiologists and is subjective, time-consuming, and prone to variability, especially in areas with limited resources and a shortage of specialists. Convolutional Neural Networks (CNNs), a type of artificial intelligence (AI), provide a revolutionary solution by automating highly accurate image analysis [4]. CNNs have shown success in similar applications, such as MRI-based stroke diagnosis [5], and they are particularly adept at learning hierarchical features from medical images.

Despite remarkable advances in neuroimaging technologies, rapid and accurate diagnosis of stroke remains a serious challenge. Human interpretation errors due to subtle radiological findings, high workload pressures in emergency situations, and persistent geographical differences in access to neuroradiology expertise are three major barriers that remain. Poor scalability still limits traditional diagnostic techniques, especially in areas with limited resources and limited access to specialists. While AI-driven automation offers a promising solution, current CNN-based stroke detection systems exhibit two fundamental limitations: insufficient model interpretation capability and insufficient validation in diverse real-world settings, which hinder clinical adoption.

To address these challenges, our study establishes three main objectives:

1. Develop a CNN model specifically designed to identify strokes in CT scans.
2. Measures including precision, accuracy, recall, F1-score, and area under the receiver operating characteristic curve (ROC) can be used to assess performance.
3. Systematically optimizing computational efficiency through hyperparameter tuning without compromising diagnostic accuracy.

979-8-3315-3922-1/25/$31.00 ©2025 IEEE

This study attempts to strike a compromise between computing efficiency and performance to provide real-time relevance in clinical situations, despite the fact that prior research has demonstrated good accuracy in stroke diagnosis. Using a training dataset of 2501 images, the work is validated by creating a CNN model that strikes a balance between high accuracy and computational efficiency.The model shows that it can maintain high performance while significantly lowering computational complexity, with a validation accuracy of 92.93%. This study also highlights the importance of model interpretability and hyperparameter tuning, both of which are essential for creating reliable and therapeutically useful deep learning models.

The remainder of this article is organized as follows. After this introduction, Section II presents the literature review, Section III describes the methodology, Section IV provides the analysis of the results, Section V concludes the work, and Section VI discusses future directions.

II. Literature Review

Worldwide, stroke is a leading cause of mortality and disability, and prompt identification based on imaging is essential for successful treatment [6]. Although early ischemic changes are often mild, non-contrast CT (NCCT) is the first-line modality in the evaluation of acute stroke due to its speed and availability [7]. In fact, only around 10% of acute infarcts can be seen by human readers on NCCT. Manual interpretation by radiologists can be laborious and unpredictable, which is why automated methods are being used [8]. Recent reviews highlight the potential of sophisticated analysis, particularly AI and deep learning, for stroke imaging tasks such ASPECTS grading, lesion recognition, and infarct segmentation.

Early stroke CAD systems used classical image processing and feature-based machine learning. To identify potential lesion sites in brain CT scans, pipelines typically start with preprocessing (such as noise filtering) followed by segmentation techniques such as edge detection (Sobel, Canny) and adaptive thresholding (Otsu's method, adaptive histogram equalization). Morphological operations (erosion, dilation, opening/closing) are also applied to smooth and partition structures [9]. Hand-crafted characteristics, such as statistical descriptors of intensity (mean, variance, skewness, and kurtosis), histogram distributions, texture measures, and form descriptors, are retrieved after candidate regions have been separated. These produced features are then fed into traditional classifiers, which represent lesion geometry and gray-level distributions. To detect stroke or its risk factors, models including logistic regression, linear discriminant analysis, support vector machines, decision trees, and random forests have all been trained using CT-derived data [8]. One review has demonstrated the usefulness of these traditional pipelines, stating that "the application of sophisticated image processing techniques, reliable feature extraction methods, and state-of-the-art segmentation algorithms is highly beneficial for stroke image analysis".

While these techniques have had some success, they require a lot of manual adjustments. Traditional machine learning techniques rely on constraints defined by experts and hand-crafted features. In practice, this means that the feature-engineering phase is time-consuming and performance can degrade if image artifacts or variability change. According to the research, traditional machine learning algorithms are "limited because appropriate discrimination features must be defined by human developers and manually extracted". As a result, much of the work has moved toward more automated methods.

In recent years, deep learning, especially convolutional neural networks (CNNs), has revolutionized medical image analysis. CNNs are able to learn hierarchical features directly from raw images, bypassing manual feature creation. These models use nonlinear layers and convolutional filters to automatically detect complex patterns. When large annotation datasets are available, CNNs often outperform traditional algorithms and, according to reviews, have demonstrated state-of-the-art performance on visualization tasks. Deep CNNs have been used in stroke imaging for automated ASPECTS grading, stroke vs. control classification, and stroke segmentation [9] [8].

For example, to automatically calculate ASPECTS from NCCT images, Naganuma et al. [8] developed a 3D CNN (the "3D-BHCA" model). They found that the CNN provided better classification results than traditional techniques and that its stroke detection was similar to or superior to that of stroke neurologists. In another multicenter study by Chen et al. , radiologists showed that acute stroke detection in NCCT was greatly improved by a deep learning model. With the help of the DL model, the reader sensitivity increased from approximately 0.254 to 0.333 (and the specificity from 0.896 to 0.915), and the interpretation time was reduced [7]. These cases demonstrate how CNN-based technologies can improve the detection of small ischemic symptoms that are often missed by the unaided eye. In a broader sense, radiological research has also stated that "combining advanced image processing techniques with comprehensive feature analysis results in high sensitivity and specificity in stroke detection".

Even with these developments, deep learning models typically require extensive validation and substantial labeled datasets. To overcome the lack of training opportunities, most reported stroke-DL models use transfer learning or data augmentation [9]. However, it is clear that end-to-end CNN approaches outperform traditional feature-based methods in terms of stroke detection efficiency and accuracy. The study concludes by demonstrating that both CNN-based deep learning models and modern image processing pipelines have improved the ability to detect acute ischemic lesions in brain CT scans, often equaling or surpassing human performance in this regard [7]. This advancement serves as the basis for our ongoing study, which expands the scope of this work to improve automated stroke detection using modern image analysis.

III. Methodology

This section discusses the methodological process for creating and evaluating a CNN model for brain stroke detection using CT scans. Figure 1 shows the steps involved in the process.

The first step is to collect a collection of brain CT scans, comprising both normal and stroke cases. The images are then put through preprocessing steps like scaling, normalization, splitting into training, testing, and validation sets, and augmentation in order to enhance model generalization. A CNN model is subsequently trained using the pre-processed

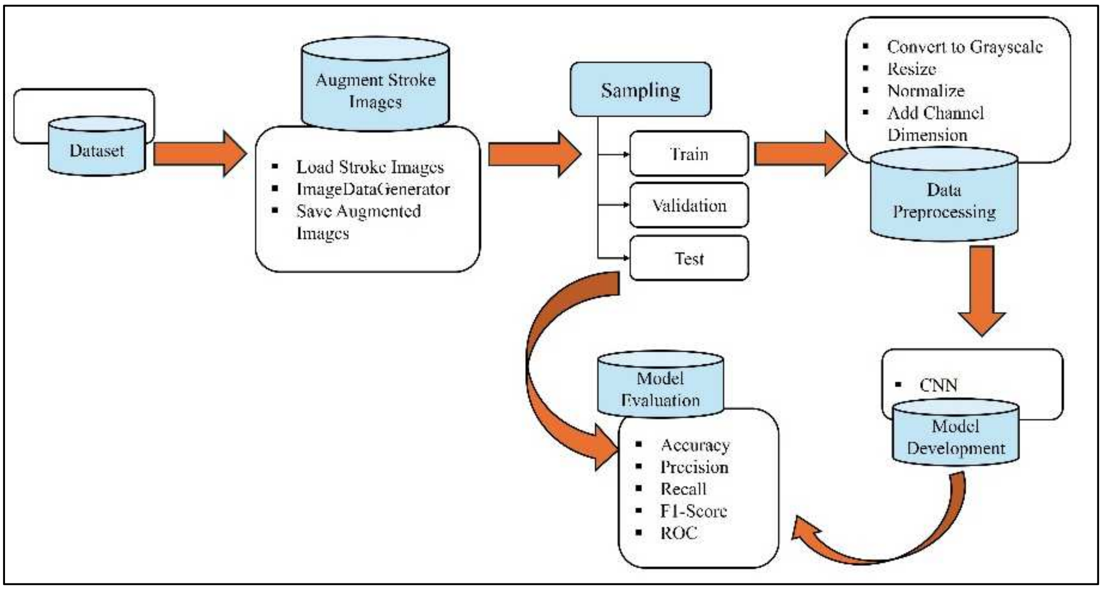

Fig. 1. Processing of Brain CT Images CNN Workflow for Stroke Detection

dataset. The model performance is evaluated using performance metrics. After an ideal model is acquired, the final performance metrics are documented.

Instead of using pretrained CNN models such as ResNet, EfficientNet, or VGG19, which are originally trained on natural image datasets, we designed a custom CNN from scratch. This approach avoids domain mismatch and ensures the network learns stroke-specific patterns from CT scans directly.

A. Dataset Description

Two sets of brain CT scans "normal" and "stroke"—make up the dataset used in this study. There are 950 stroke and 1551 normal images in the collection, for a total of 2501 images in different sizes and resolutions.

Images in the "normal" category show healthy brain structure, while images in the "stroke" category show cerebral hemorrhage or ischemic areas. The dataset is available at Kaggle [10].

B. Data Preprocessing Pipeline

To minimize computational complexity and preserve uniform input channels, all CT images were first converted to grayscale. Data augmentation was applied to the stroke class to address class imbalance between normal and stroke images, using methods including rotations, shifts, cropping, zooming, and horizontal flips. To establish class balance, augmented images were saved and merged with the original dataset. Visualizations of augmented samples are in Figure 2.

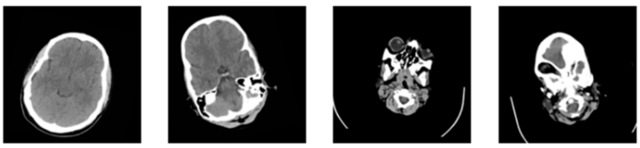

Fig. 2. Visualizations of augmented samples

All images were then normalized to the range [0, 1] and scaled to a fixed resolution of 256 × 256 pixels using bilinear interpolation. To ensure proportional class distributions across all partitions, the dataset was divided into training (80%), validation (10%), and test (10%) subsets.

Two methods were used to load the data. During training, enhancement and batch loading were initially provided by on-the-fly generators using Keras' ImageDataGenerator. The images were then imported into memory, validated, and transformed into TensorFlow tf.data datasets with optimized operations such as grouping, prefetching, and shifting to increase pipeline robustness. Additional enhancements were used throughout the training phase, including brightness, contrast, and random flipping, to improve model generalization.

Standardized inputs, balanced classes, and effective data feeding for model training were all guaranteed by this preprocessing pipeline.

C. Model Development

A convolutional neural network (CNN) architecture was designed to distinguish brain CT images from stroke and normal. Single-channel grayscale images, or 256 × 256 × 1 input images, are supported by the model. Figure 3 shows the model architecture.

The architecture consists of two convolutional blocks. A 3 × 3 kernel Conv2D layer with 32 and 64 filters, respectively, is included in each block. To reduce spatial dimensionality, it employs batch normalization, maximum pooling, and ReLU activation. After smoothing, the collected features are run through a dense layer containing 128 neurons and ReLU activations. One sigmoid-functioning neuron produces the likelihood that the image is in the Stroke class.

Binary cross-entropy loss and Adam optimization (learning rate = 0.001) were used to develop the model. To minimize the possible effects of residual class imbalance, studies were also conducted using a focal loss function. To evaluate the classification performance in a comprehensive way, the evaluation measures included precision, accuracy, recall, and area under the ROC curve.

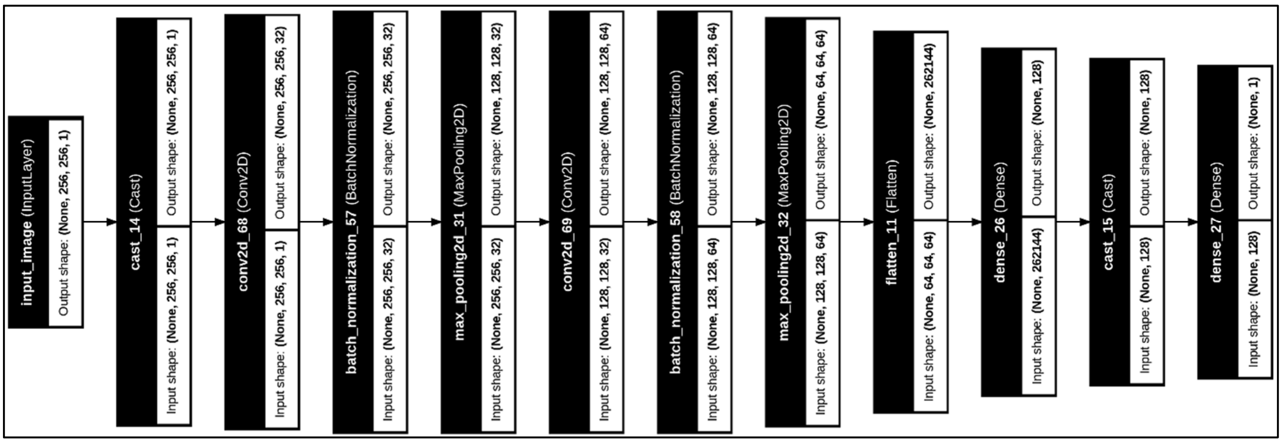

Fig. 3. Architecture of the CNN for Stroke Detection from CT Images.

During model training, early stopping and learning rate reduction techniques were used to improve convergence and avoid overfitting. To ensure effective data loading and augmentation during training, the optimized TensorFlow tf.data pipeline was used.

IV. Results Analysis

This section presents the experimental results obtained from training and evaluating the proposed CNN model on the brain CT dataset. The analysis focuses on key performance metrics, including accuracy, precision, recall, F1-score, and AUC, to assess the model's effectiveness in distinguishing stroke cases from normal CT scans.

To evaluate the efficiency of the model, we employ four widely used classification metrics: precision (1), recall (2), F1-score (3), and accuracy (4). These metrics are defined as follows:

These are calculated as follows:

$$\text{Precision} = \frac{\text{TP}}{\text{TP+FP}} \quad (1)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP+FN}} \quad (2)$$

$$\text{F1} - \text{Score} = \frac{2 \text{ x Precision x Recall}}{\text{Precision+ Recall}} \quad (3)$$

$$Accuracy = \frac{\text{TP+TN}}{\text{TP+FN+TN+FP}} \quad (4)$$

Where TP denotes true positives, FP false positives, and FN false negatives.

The model achieved an F1-score of 0.9308 on the test set, with a precision of 0.9136 and a recall of 0.9487 for the normal class. The stroke class has an F1-score of 0.9276, a recall of 0.9097, and a precision of 0.9463. According to these measurements, the model maintains a high sensitivity for recognizing normal instances while performing marginally better at correctly identifying stroke cases (greater accuracy). In clinical settings, where reducing false positives (unnecessary alarms) and false negatives (missed strokes) is crucial, such balanced performance is crucial.

Initially, training was monitored using learning curves showing the accuracy, loss, and area under the curve (AUC) of the receiver operating characteristic curve over a number of epochs. These graphs (Figure 4) demonstrated a smooth convergence of the model with training and validation losses slowly dropping and plateauing without significant divergence, indicating that the model did not suffer from significant overfitting. Furthermore, the patterns in AUC and validation accuracy closely matched the training data, indicating strong generalization performance.

Following training, the test dataset was used to evaluate the model's predictive performance. With a test accuracy of 92.93%, the model demonstrated its ability to accurately classify CT images. Furthermore, the model demonstrated exceptional discriminative power and the capacity to

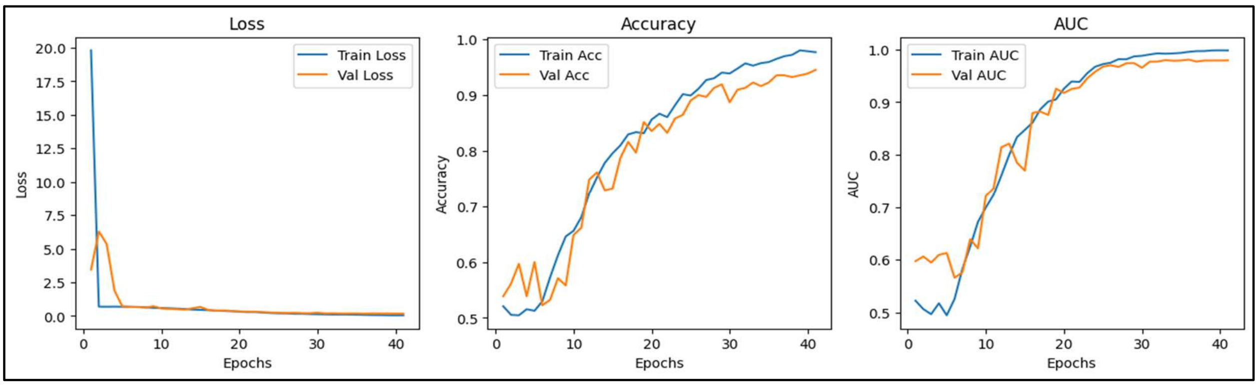

Fig. 4. Learning curves

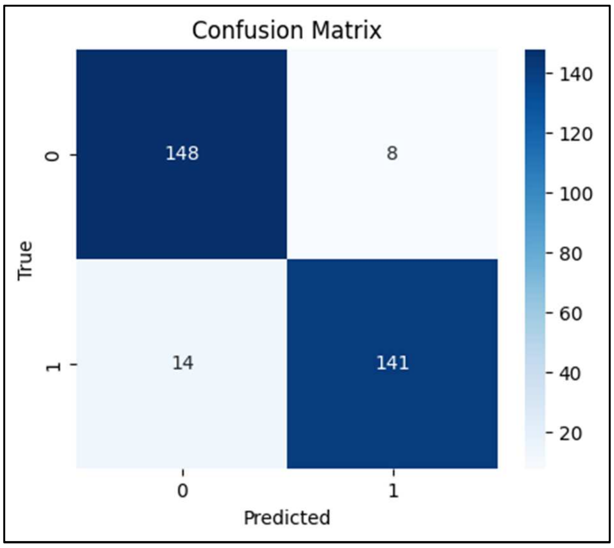

Fig. 5. Confusion matrix

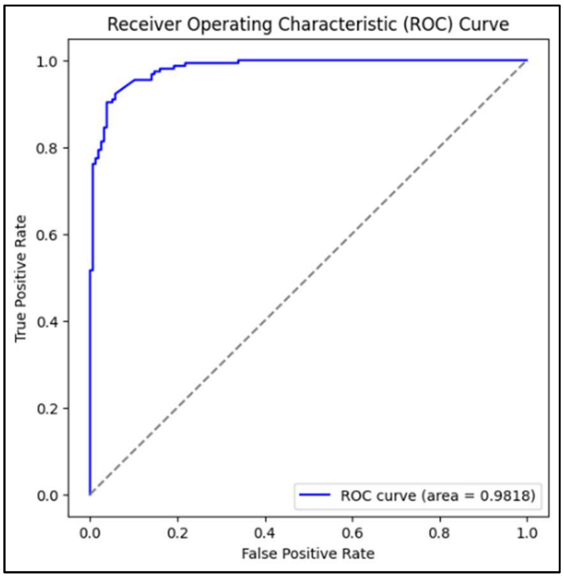

Fig. 6. Receiver Operating Characteristic (ROC) Curve

effectively differentiate between normal and stroke images, as evidenced by its AUC of 0.9818 on the test dataset. High AUC values are important in medical applications because they show that the model can balance sensitivity and specificity across multiple constraints, which is important because the cost of false negatives (i.e., missed strokes) can have substantial clinical effects.

By examining the confusion matrix, the distribution of accurate and inaccurate predictions across classes was also shown (Figure 5). Out of 156 normal images, 148 were accurately labeled as strokes, while the remaining 8 were incorrectly diagnosed, according to the matrix. On the other hand, 141 out of 155 stroke images were correctly identified, while 14 were misclassified as normal. Although the misclassifications were not that high, it is important to investigate these cases further, as false negative results in stroke detection can have serious clinical consequences.

A receiver operating characteristic (ROC) curve was constructed to evaluate the model's capacity for class discrimination in more detail (Figure 6). The ROC curve shows the trade-off between the false positive rate (1-specificity) and the true positive rate (sensitivity) at different decision thresholds. The model retains exceptional sensitivity and specificity across different thresholds, as demonstrated by the large AUC value of 0. 9818 in the test set. This further supports the model's feasibility for implementation in clinical decision support systems.

Overall, the evaluation findings show that the proposed CNN model distinguishes stroke from normal brain CT images with high accuracy, precision, recall, and AUC. According to these results, deep learning-based CT scan analysis can be of great help in the accurate and timely detection of strokes, which can accelerate diagnosis and improve patient outcomes. Additional validation on external datasets, semantic assessments to clarify model choices, and integration into clinical processes to evaluate real-world performance and practical use can all be part of future research.

Recent research using CT images for stroke diagnosis is compiled in Table I, which also highlights different deep learning techniques and dataset sizes. Our approach performs competitively when compared to this research, highlighting its potential for accurate and effective stroke classification.

Table I. COMPARISON OF EXISTING STUDIES

Study	**Year**	**Dataset Size**	**Model**	**Accuracy**
[11]	2024	5772 CT	CNN	72.00
[12]	2025	Not Specified	JRSN	90.20
[13]	2023	TVGH	CNN	80
This Study	2025	2501 images	CNN	92.93

Note: The studies summarized in Table I used different datasets and experimental settings; therefore, the accuracy values are not directly comparable but serve as indicative references.

V. COCNCLUSION

In this paper, we thoroughly evaluated our classification model, focusing on both overall and class-specific performance. In addition to assessing overall accuracy, we examined ROC curves to see how sensitivity and false positive rates trade off at different thresholds. In addition, we calculated F1-scores, precision, and recall for each class, providing a deeper understanding of the model's performance in distinguishing between positive and negative events.

Recall indicated the model's ability to identify true positive samples, while precision indicated the percentage of positive cases that were correctly detected among all cases that were predicted to be positive. A balanced statistic that accounts for both false positives and false negatives is the F1-score. It is computed as the precision and recall harmonic means. These measures are essential, especially in situations where class imbalance can affect performance or when certain types of errors have more severe consequences.

Our results show that the model achieves strong performance in correctly classifying the data, with certain classes achieving higher accuracy and recall than others. These differences highlight areas where the model can be further improved, perhaps through additional data collection, feature engineering, or fine-tuning of model hyperparameters.

Overall, the proposed CNN-based approach demonstrates promising results for accurate stroke classification using brain CT scans. The model exhibits great promise as a clinical decision-support tool and achieves high diagnostic accuracy by fusing efficient preprocessing with a thoughtfully planned architecture. These findings demonstrate how deep learning techniques might enhance automated stroke diagnosis and provide useful data for future studies in medical imaging.

VI. Future Work

In order to rectify the observed class imbalances and enhance the models' resilience, we intend to explore more sophisticated approaches in subsequent research, such as ensemble methods or cost-sensitive learning. Additionally, we aim to explore model explanation approaches to better interpret the decision-making process, which is essential for deploying such models in practical applications.

Furthermore, we recognize the need to validate the proposed CNN in more diverse and clinically relevant contexts. Future directions will therefore include:

- Multi-center data collection using varied imaging protocols to improve generalization.
- Integration of explainability tools such as Grad-CAM or SHAP to enhance transparency and clinician trust.

Prospective validation in emergency settings to confirm real-world clinical utility.

VII. References

[1] S. Mainali, M. E. Darsie, and K. S. Smetana, “Machine Learning in Action: Stroke Diagnosis and Outcome Prediction,” *Front. Neurol.*, vol. 12, p. 734345, 2021, doi: 10.3389/fneur.2021.734345.

[2] G. Bathla *et al.*, “Advances in Acute Ischemic Stroke Treatment: Current Status and Future Directions,” *AJNR Am. J. Neuroradiol.*, vol. 44, no. 7, pp. 750–758, Jul. 2023, doi: 10.3174/ajnr.A7872.

[3] X. Zhao, L. Wang, Y. Zhang, X. Han, M. Deveci, and M. Parmar, “A review of convolutional neural networks in computer vision,” *Artif. Intell. Rev.*, vol. 57, no. 4, p. 99, Mar. 2024, doi: 10.1007/s10462-024-10721-6.

[4] G. Altamirano-Gomez and C. Gershenson, “Quaternion Convolutional Neural Networks: Current Advances and Future Directions,” *Adv. Appl. Clifford Algebr.*, vol. 34, no. 4, p. 42, Aug. 2024, doi: 10.1007/s00006-024-01350-x.

[5] V. Abedi *et al.*, “Prediction of Long-Term Stroke Recurrence Using Machine Learning Models,” *J. Clin. Med.*, vol. 10, no. 6, Art. no. 6, Jan. 2021, doi: 10.3390/jcm10061286.

[6] J. E. Soun *et al.*, “Artificial Intelligence and Acute Stroke Imaging,” *AJNR Am. J. Neuroradiol.*, vol. 42, no. 1, pp. 2–11, Jan. 2021, doi: 10.3174/ajnr.A6883.

[7] W. Chen *et al.*, “Improving the diagnosis of acute ischemic stroke on non-contrast CT using deep learning: a multicenter study,” *Insights Imaging*, vol. 13, no. 1, p. 184, Dec. 2022, doi: 10.1186/s13244-022-01331-3.

[8] L. Cui *et al.*, “Deep Learning in Ischemic Stroke Imaging Analysis: A Comprehensive Review,” *BioMed Res. Int.*, vol. 2022, p. 2456550, Nov. 2022, doi: 10.1155/2022/2456550.

[9] A. Hossain, J. Sultana, S. Hossain, and H. Rahman, “Radiomic Feature Extraction Based on Computed Tomography for Stroke Patients,” *J. Med. Phys.*, vol. 50, no. 2, p. 337, Jun. 2025, doi: 10.4103/jmp.jmp_199_24.

[10] “Brain Stroke CT Image Dataset.” Accessed: Jul. 15, 2025. [Online]. Available: https://www.kaggle.com/datasets/afridirahman/brain-stroke-ct-image-dataset

[11] A. Fontanella *et al.*, “Development of a deep learning method to identify acute ischaemic stroke lesions on brain CT,” *Stroke Vasc. Neurol.*, p. svn-2024-003372, Dec. 2024, doi: 10.1136/svn-2024-003372.

[12] “Enhanced Multi-Layer Vision Transformer for Brain Stroke Detection Using Advanced Image Segmentation and Deep Learning Techniques,” *Int. J. Intell. Eng. Syst.*, vol. 18, no. 6, pp. 167–179, Jul. 2025, doi: 10.22266/ijies2025.0731.11.

[13] “StrokeNeXt: an automated stroke classification model using computed tomography and magnetic resonance images | BMC Medical Imaging | Full Text.” Accessed: Jul. 15, 2025. [Online]. Available: https://bmcmedimaging.biomedcentral.com/articles/10.1186/s12880-025-01721-1