Real-Time Stroke Detection in MRI Imaging Using YOLOv8x

Basma Emhamed Ali Dihuom[1*] , Zahrah Mohammed Bileid Jabeer[2*]

[1*]*University of Jafara , Libya*

[2*]*Higher Institute of Industrial Technology / Alnigela*

{*basmadihom@gmail.com*}

Abstract.

This study explores the potential of advanced deep learning techniques, specifically the YOLOv8 architecture, for enhancing stroke detection within MRI imaging. Utilizing the well-established ISLES dataset, we aimed to improve detection accuracy through rigorous preprocessing and data augmentation techniques. By implementing the YOLOv8x model variant, known for its high precision, we conducted a series of experiments to assess the impact of data augmentation on model performance. The experiments were structured around training the model for 60 epochs under different conditions, including original and augmented datasets. The performance was evaluated based on the mean Average Precision (mAP), with the augmented dataset experiment yielding the highest mAP of 96.21%. These results demonstrate the effectiveness of combining advanced YOLO architectures with extensive data augmentation in improving the accuracy and reliability of stroke detection in MRI images. The study underscores the significant potential of integrating AI-driven methods into diagnostic workflows, contributing to early and accurate stroke diagnosis.

Keywords: Stroke Detection, YOLOv8x, Deep Learning, Computer-Aided Diagnosis.

I. Introduction

As a leading cause of severe disability worldwide, ischemic stroke remains a critical challenge to public health, marked by its high incidence, recurrence, and significant social and economic impacts. The burden of this condition is particularly evident among low-income populations and is increasingly affecting younger demographics [1,2]. In the realm of clinical diagnostics, magnetic resonance imaging (MRI) plays a pivotal role, especially in differentiating ischemic stroke from other types such as hemorrhagic stroke, where computed tomography (CT) scans may fall short.

The reliance on manual interpretation of medical images by healthcare professionals introduces substantial delays—often exceeding 24 hours—before a diagnostic conclusion is reached. This delay can critically hinder the timely administration of effective treatments, potentially compromising patient recovery. Furthermore, the demanding workload can lead to fatigue among radiologists, raising the risk of diagnostic errors, including misdiagnoses and missed diagnoses, which directly affect patient outcomes.

To address these issues, the development of automated and intelligent systems for MRI analysis has emerged as a crucial innovation. Traditional machine learning approaches in this field involve labor-intensive processes such as image denoising, segmentation, and manual feature setting, all of which are time-consuming and prone to low accuracy [3]. In contrast, deep learning represents a transformative direction in machine learning, mimicking the hierarchical processing capabilities of the human brain to discern underlying data structures and extract features autonomously, thereby enhancing learning efficacy [4].

Leveraging advancements in GPU technology and expansive datasets, deep learning, particularly through Convolutional Neural Networks (CNNs), has shown remarkable utility in extracting sophisticated features from extensive data samples for more accurate classification, detection, and segmentation. This has substantially improved the feasibility of intelligent MRI interpretation.

Our research has utilized the ISLES 2015 dataset, which includes detailed MRI images and clinical reports, all expertly annotated [5]. We have deployed and assessed deep learning models, notably Faster R-CNN and YOLOv8, to automate the detection of ischemic lesions in MRI scans. Our analysis has provided statistically significant insights into lesion characteristics and potential diagnosis distributions, thus offering crucial tools and data to enhance the prevention and management of ischemic stroke, ultimately benefiting patient health and recovery.

II.Literature Review

The deep learning-based medical image lesion detection and auxiliary diagnosis system significantly enhances lesion feature extraction from medical images, integrating seamlessly with clinical practices to alleviate the workload of physicians. In computer-based lesion detection methods, features of specific body parts or organs are analyzed using supervised learning or traditional image processing techniques such as filtering and mathematical morphology. These supervised learning methods rely on training data that must be carefully labeled by skilled physicians using comprehensive pathological images [6].

Over the past few decades, deep learning technologies have rapidly advanced and gained extensive application in the medical field, particularly in lesion detection. Convolutional Neural Networks (CNNs) have shown to increase detection accuracy by 13–34% [7]. For instance, Sirinukunwattana et al. employed a Spatially Constrained CNN (SC-CNN) equipped with a Neighboring Ensemble Predictor (NEP) to enhance the accuracy of detecting and classifying colorectal adenocarcinoma cells, surpassing traditional feature classification methods [8]. Similarly, Dou et al. implemented a 3D CNN for automated cerebral hemorrhage detection from MRI images, extracting more representative features and achieving detection accuracies up to 93.16%, which is significantly higher than that of 2D CNNs and manual feature extraction methods [9]. Various CNN architectures have demonstrated remarkable success in diagnosing a range of diseases, illustrating the substantial potential of deep learning in the medical sector. Another key application of deep learning in this domain is lesion recognition, where it excels in extracting valuable information from training data, thereby enhancing the accuracy and efficiency of medical diagnoses. For example, Kooi et al.

utilized CNNs to identify malignant breast lesions, showing the effectiveness of these technologies.

III. Methodology

1) Method Overview

In the methodology section of our study, we detail the process of annotating a dataset consisting of 2,300 MRI images of stroke cases. Utilizing Roboflow, a robust platform for image annotation, we meticulously labeled these images to prepare them for deep learning applications. Following the annotation phase, we trained a YOLOv8 model on this dataset. This cutting-edge object detection model was chosen for its efficiency and accuracy in identifying and classifying features within complex medical images, enabling a more nuanced understanding of stroke imaging characteristics. For a visual representation of our proposed approach, please refer to Figure 1.

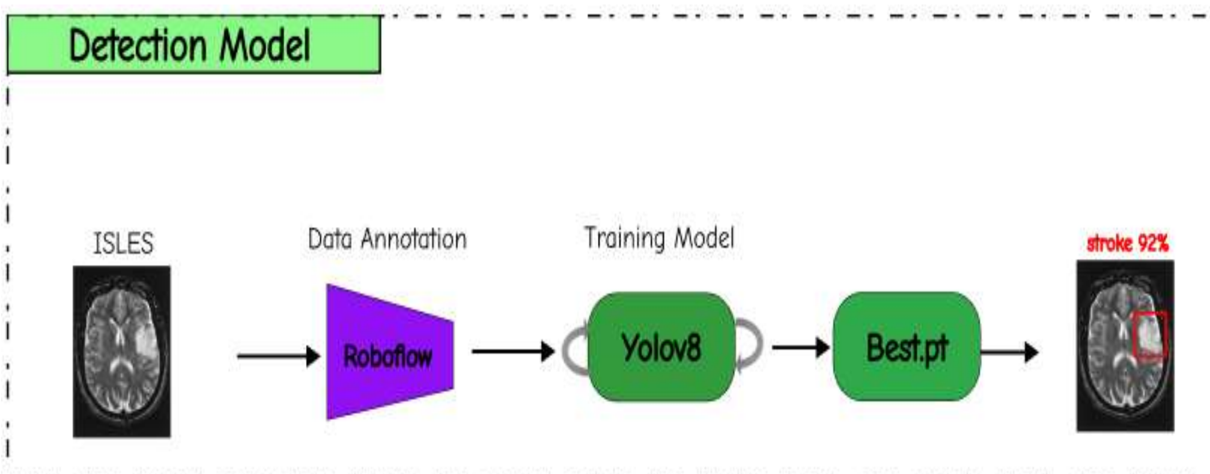

Fig. 1. Diagram of overview method

2) Dataset collection

In our study, we utilized the ISLES 2015 dataset, which comprises both stroke-affected and normal MRI images. For the purposes of our research, we selectively used only the stroke-affected MRI images, amounting to 2,300 in total. These images were meticulously annotated using Roboflow, a comprehensive tool designed for precise image annotation. This selection and annotation process allowed us to focus exclusively on the characteristics relevant to stroke, enhancing the accuracy and relevance of the subsequent YOLOv8 model training.

3) Data Annotation

In the data preparation phase of our study, we chose Roboflow as the annotation tool to handle the intricate task of labeling our dataset, which consists of 2,300 MRI images selected from the ISLES 2015 dataset. This platform was instrumental for its robust functionality in image annotation, allowing for the precise demarcation and labeling of areas affected by stroke.

Roboflow's intuitive interface and powerful feature set enabled us to systematically identify and annotate stroke features within the MRI images. This included outlining areas of infarction, hemorrhage, and other pathological markers relevant to stroke diagnosis. The annotations involved both bounding boxes to delineate regions of interest and detailed labels describing the type and severity of stroke manifestations observed.

These annotations are critical for training the YOLOv8 model effectively. By providing clear, accurately labeled data, Roboflow helped us train a model that can discern subtle and complex features in MRI scans, which are indicative of stroke. This precision is crucial for developing a reliable automated tool that aids radiologists in diagnosing strokes more swiftly and accurately. The pre-processed, annotated images serve as the foundational data from which the YOLOv8 model learns, ensuring it is well-equipped to handle real-world diagnostic tasks.

4) Detection Model

The YOLO (You Only Look Once) series represents a significant evolution in the field of object detection, leveraging deep learning algorithms to identify and classify various objects within an image in real-time [10], [11], [12]. The architecture of YOLO models can be dissected into three main components: the backbone, neck, and head. The backbone is responsible for feature extraction; it takes the input image and processes it through a series of convolutional layers to create a feature map that contains the essential information needed for object detection. The neck, often composed of additional convolutional and pooling layers, acts as a bridge between the feature-rich backbone and the detection-specific head. It is designed to refine and aggregate these features, making them more suitable for the final detection task. The head of the YOLO architecture then interprets these refined features to predict bounding boxes and class probabilities, effectively pinpointing and classifying objects in the image.

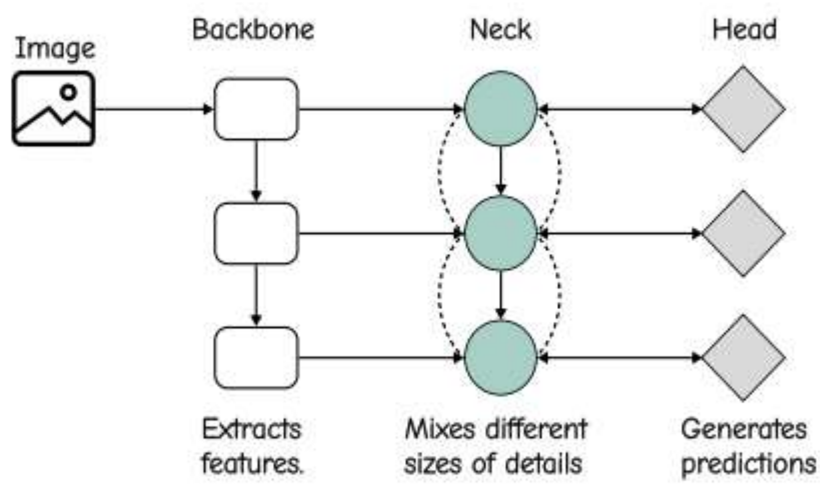

Fig. 2. YOLO Object Detection Architecture

YOLOv8, the latest iteration in the series, has further improved upon its predecessors in terms of speed and accuracy, making it particularly suited for real-time applications [12]. This version introduces enhancements in the backbone's structure for more efficient feature extraction, optimized layers in the neck for better feature integration, and advancements in the head for more accurate object classification and localization. The improvements are geared towards reducing false positives and increasing the precision of the detected objects, which is critical for applications like medical imaging where accuracy is paramount.

Within the YOLOv8 family, there are variants like YOLOv8s, YOLOv8m, and YOLOv8x, designed to cater to different requirements of speed and accuracy. YOLOv8s is optimized for speed, making it ideal for environments where processing time is a critical factor. YOLOv8m offers a balance between speed and accuracy, suitable for general-purpose applications. On the other hand, YOLOv8x, although slower compared to its counterparts, provides the highest precision among the variants. This makes YOLOv8x an excellent choice for scenarios where the accuracy of detection is crucial and can afford slightly longer processing times.

For our system, which requires high precision in detecting stroke-related abnormalities in MRI images, YOLOv8x was chosen as the optimal model. Despite its slower speed, the high precision offered by YOLOv8x is invaluable in ensuring the reliability and effectiveness of stroke detection. This choice underscores our prioritization of detection accuracy over processing speed, given the critical nature of stroke diagnosis where the consequences of false negatives or false positives can be extremely high.

IV. Experimental Results

In this study, we rigorously evaluated the YOLOv8x model's performance in detecting stroke from MRI images using the ISLES dataset. Before commencing our experiments, we partitioned our data into two distinct sets: 80% for training purposes and 20% for testing, ensuring a comprehensive evaluation of the model's performance under various conditions.

The evaluation process was methodically structured into distinct experiments to ascertain the impact of data augmentation on model accuracy. Initially, the model underwent training for 60 epochs on the ISLES dataset both with and without data augmentation. For augmenting the data, we utilized Roboflow, an AI-centric augmentation tool that expanded the dataset by a factor of three, thereby enhancing the diversity and volume of training samples. This augmentation aimed to mimic various real-world variations in MRI images, such as rotations, translations, and scaling, to robustly train the YOLOv8x model against overfitting and to improve its generalization capabilities. In these initial experiments, the model achieved mean Average Precision (mAP) scores of 83.24% without data augmentation and 85.68% with augmentation, indicating a notable improvement in detection performance due to the diversified training samples.

Subsequent evaluations involved additional tests on the augmented and non-augmented ISLES dataset. These tests were crucial in understanding the model's adaptability and performance consistency. The results exhibited mAP scores of 92.87% without augmentation and 94.13% with augmentation, reinforcing the value of augmented datasets in enhancing model accuracy.

Finally, to maximize the potential of our model, we conducted an experiment on the ISLES dataset augmented threefold using Roboflow. This comprehensive approach aimed to leverage the full spectrum of variations and complexities present in MRI images. The augmented dataset led to the highest mAP score of 96.21%, showcasing the substantial impact of extensive data augmentation and dataset diversity on improving the accuracy and reliability of stroke detection using the YOLOv8x model. Figure bellow show results of yolov8x stroke detection model.

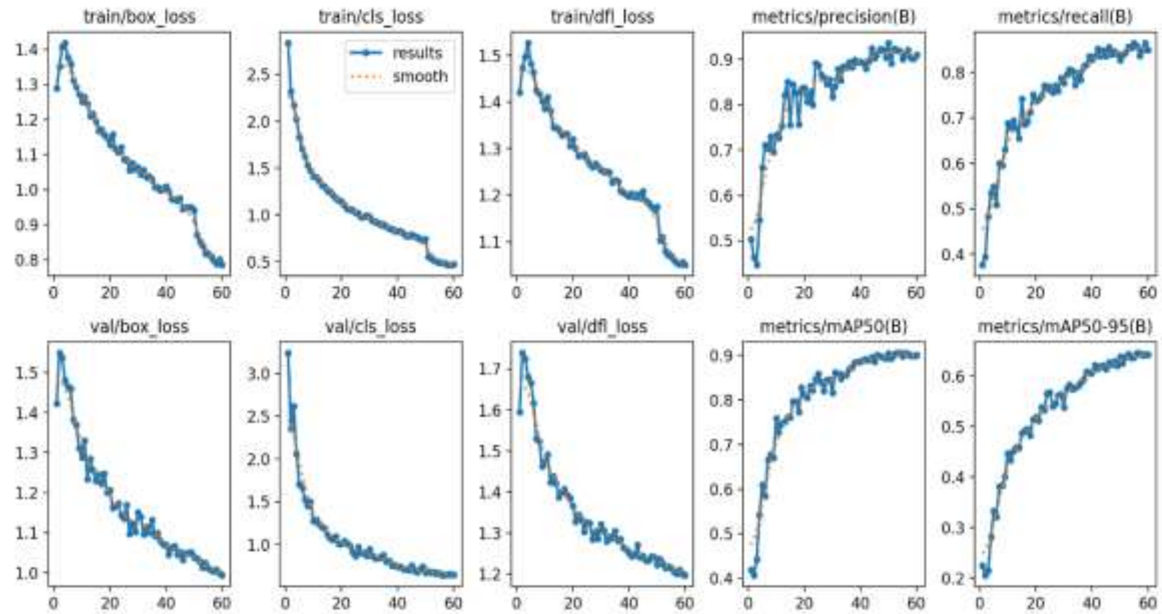

Fig. 3. YOLO Object Detection Architecture

V. Conclusion

Our study validates the significant potential of the YOLOv8x architecture in improving stroke detection through MRI analysis. Utilizing the ISLES dataset, we demonstrated that deep learning models, particularly when enhanced with data augmentation and trained on diverse datasets, significantly increase detection accuracy. The highest mean Average Precision (mAP) achieved was 96.21%, showcasing the model's high precision.

The findings advocate for the incorporation of AI-driven tools like YOLOv8x into clinical settings, despite its slower processing speed, due to its superior accuracy in critical medical diagnostics. Future research should extend to refining these models and exploring their real-world clinical efficacy, aiming to bridge the gap between technological advancement and practical healthcare applications. By enhancing early detection methods, we move closer to improving patient outcomes and advancing the fight against stroke.

VI. References

1. Karthik, R., Menaka, R., Johnson, A., & Anand, S. (2020). Neuroimaging and deep learning for brain stroke detection-A review of recent advancements and future prospects. Computer Methods and Programs in Biomedicine, 197, 105728.
2. Semenov, S. Y., & Corfield, D. R. (2008). Microwave Tomography for Brain Imaging: Feasibility Assessment for Stroke Detection. International Journal of Antennas & Propagation.
3. Rahman, M. A., Hossain, M. F., Riheen, M. A., & Sekhar, P. K. (2020). Early brain stroke detection using flexible monopole antenna. Progress In Electromagnetics Research C, 99, 99-110.
4. Gaidhani, B. R., Rajamenakshi, R. R., & Sonavane, S. (2019, September). Brain stroke detection using convolutional neural network and deep learning models. In 2019 2nd International conference on intelligent communication and computational techniques (ICCT) (pp. 242-249). IEEE.
5. Maier, O., Menze, B. H., Von der Gablentz, J., Häni, L., Heinrich, M. P., Liebrand, M., ... & Reyes, M. (2017). ISLES 2015-A public evaluation benchmark for ischemic stroke lesion segmentation from multispectral MRI. Medical image analysis, 35, 250-269.
6. Xie, Xiaozheng, et al. "A survey on incorporating domain knowledge into deep learning for medical image analysis." Medical Image Analysis 69 (2021): 101985.
7. Sirinukunwattana, K., Raza, S. E. A., Tsang, Y. W., Snead, D. R., Cree, I. A., & Rajpoot, N. M. (2016). Locality sensitive deep learning for detection and classification of nuclei in routine colon cancer histology images. IEEE transactions on medical imaging, 35(5), 1196-1206.
8. Dou, Q., Chen, H., Yu, L., Zhao, L., Qin, J., Wang, D., ... & Heng, P. A. (2016). Automatic detection of cerebral microbleeds from MR images via 3D convolutional neural networks. IEEE transactions on medical imaging, 35(5), 1182-1195.
9. Kooi, T., Litjens, G., Van Ginneken, B., Gubern-Mérida, A., Sánchez, C. I., Mann, R., ... & Karssemeijer, N. (2017). Large scale deep learning for computer aided detection of mammographic lesions. Medical image analysis, 35, 303-312.
10. J. Terven and D. Cordova-Esparza, “A comprehensive review of yolo: From yolov1 to yolov8 and beyond,” arXiv preprint arXiv:2304.00501, 2023.
11. F. M. Talaat and H. ZainEldin, “An improved fire detection approach based on yolo-v8 for smart cities,” Neural Computing and Applications, vol. 35, no. 28, pp. 20 939–20 954, 2023.
12. B. Aldughayfiq, F. Ashfaq, N. Jhanjhi, and M. Humayun, “Yolo-based deep learning model for pressure ulcer detection and classification,” in Healthcare, vol. 11, no. 9. MDPI,