ISSN(O): 2320-9801
ISSN(P): 2320-9798

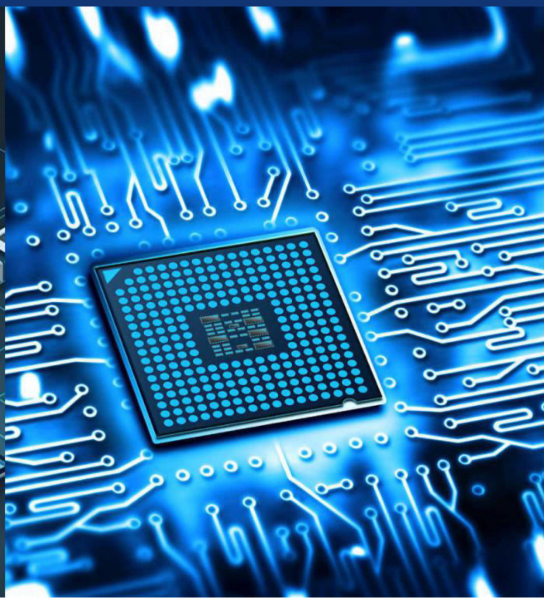

International Journal of Innovative Research in Computer and Communication Engineering

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

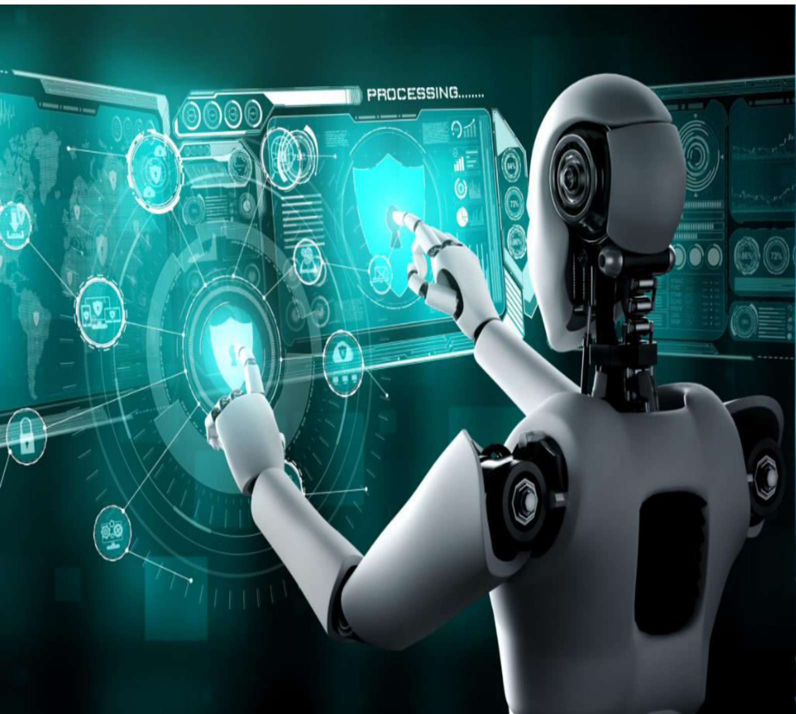

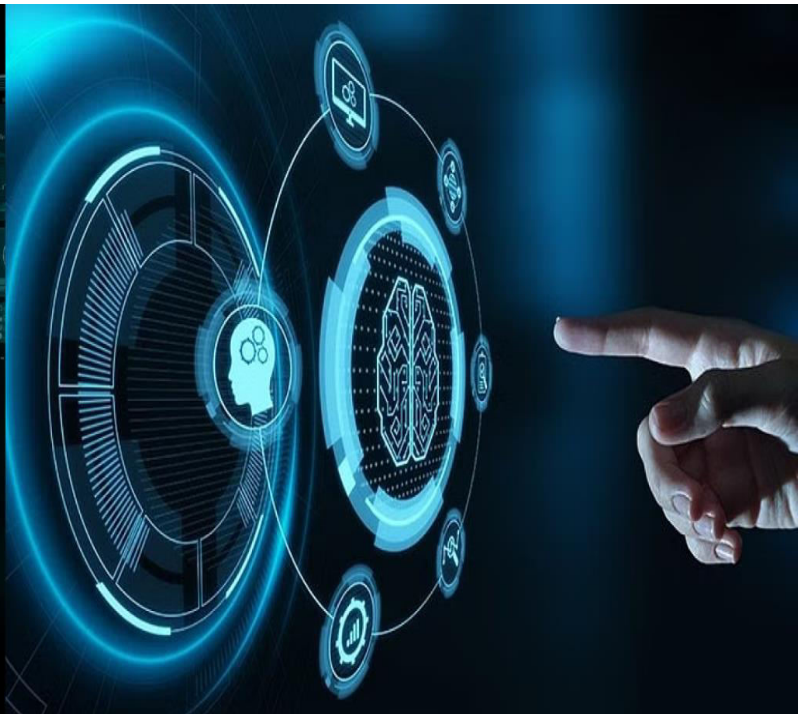

Impact Factor: 8.771

Volume 13, Issue 5, May 2025

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Stroke Detection using Federated Learning and YOLOv8

Prerana S [1], Sahana K M [1], Sahana S Angadi [1], Sangamesh G B[1], Pradeep N[2]

U.G. Student, Department of CS&E, Bapuji Institute of Engineering and Technology, Davanagere, Karnataka, India[1]

Professor, Department of CS&E, Bapuji Institute of Engineering and Technology, Davanagere, Karnataka, India[2]

ABSTRACT: The early diagnosis of brain stroke is crucial to reducing damage and increasing the probability of survival of the patient. Historically, stroke diagnosis has been made by getting the images processed through the interpretation of medical imaging. This method is expensive, complicated, and prone to error. The advancements in machine learning that we have seen today, particularly in real-time image processing, offer a potential solution towards making stroke detection faster and more efficient. The objective of this project, therefore, is to create a real-time brain stroke detection system with Federated Learning in combination with YOLOv8 thus, transforming the private collaborative learning model with the state-of-the-art techniques for object detection. The Federated Learning concept allows a number of health institutions to perform collaborative training on a stroke detection model with their local data, keeping the sensitive patient information to themselves and making it privacy-compliant. The leading real-time object detection algorithm, YOLOv8, is altered in a way to allow the accurate identification of features of brain strokes in medical imagery such as CT scans and MRIs. The infrastructure developed will cover a great storm with rapid and reliable detection of brain-stroked patients without exposing their data. Stroke detection through these means should not only be faster and more accurate but also would help enhance data privacy and create collaborations among medical institutions towards better outcomes and timely interventions to patients.

KEYWORDS: Real-Time Stroke Detection, Federated Learning, YOLOv8, Medical Imaging, Data Privacy, Deep Learning, Object Detection.

I. INTRODUCTION

A brain stroke, or cerebrovascular accident (CVA), is a critical medical emergency that arises when blood flow to a part of the brain is disrupted, depriving brain cells of oxygen and nutrients, ultimately leading to their death. This interruption in blood flow can occur due to two primary types of strokes: ischemic and hemorrhagic. Ischemic strokes are the most common, caused by a blockage in a blood vessel due to a blood clot or fatty deposits (atherosclerosis). Hemorrhagic strokes, on the other hand, occur when a blood vessel in the brain ruptures, leading to bleeding in or around the brain, increasing pressure and causing further damage. The human brain is a highly complex organ responsible for regulating critical functions such as memory, thought, motor skills, emotion, and body processes. Comprising about 60% fat and the rest made up of water, protein, and salts, it contains a network of blood vessels and neurons that support its function. Given the critical role of uninterrupted blood flow, even a brief disruption can result in loss of vital brain functions like speech, movement, or emotional control. Common symptoms of stroke include face drooping, arm weakness, and speech difficulty, summarized in the acronym FAST (Face, Arms, Speech, Time), urging immediate medical attention. Stroke diagnosis is supported by advanced imaging technologies such as CT scans, MRIs, CTA, MRA, and carotid ultrasound, all aimed at accurately identifying the type and severity of the stroke. In recent years, federated learning has emerged as a promising AI model training strategy that allows collaborative model development across multiple organizations without transferring sensitive data. Instead, models are trained locally and updates are aggregated centrally, preserving data privacy. Tools like Flower enable scalable and secure federated learning infrastructure. YOLOv8, the latest in a line of powerful object detection models developed by Ultralytics, enhances stroke detection by identifying key features in medical images using advanced methods such as classification, object detection, and segmentation. It can be trained on labeled data and used to predict stroke presence from CT and MRI scans with high precision, making it a vital tool in real-time, privacy-preserving stroke detection.

II. LITERATURE SURVEY

Elhanashi, A. et al.[1] proposed TeleStroke, a real-time stroke detection system designed for edge devices with federated learning (FL) and YOLOv8. Their primary motivation was to overcome the conventional bottlenecks: centralized learning involves uploading sensitive medical images to a shared server, endangering patient privacy and potentially leading to data breaches. They suggested a decentralized deep learning model in which local stroke detection at client nodes (such as hospitals or edge devices) is carried out using YOLOv8 trained locally on facial imagery data available in the vicinity signifying stroke symptoms, e.g., facial paralysis. The trained weights, rather than the original data, were sent to a central server on which Federated Averaging (FedAvg) averaged out the parameters to construct a more generalized global model. The process iteratively refined detection performance while maintaining data privacy compliance (GDPR/HIPAA). Their use of YOLOv8 was novel due to its real-time inference capability and its adaptation for medical imaging, which involved tuning the model to distinguish subtle signs of stroke on facial expressions. This work lays a solid foundation for privacy-preserving stroke detection in real-time environments, directly aligning with your paper's architecture.

Mansour, A., et al.[2], reported in WISE 2024 a federated deep learning architecture to predict stroke using structured data sets (non-image) with outstanding 98% classification accuracy. Unlike most DL applications needing centrally available data, their framework made cross-institutional collaboration achievable via FL as model training became possible over dispersed electronic health records (EHR) and sensor data sets without revealing raw data. They measured performance against metrics such as precision, recall, F1-score, and confusion matrix analysis to establish the robustness and extensibility of FL in the treatment of real-world medical variability. Their results were notable, as they showed enhanced model convergence speed with strategic client selection and optimized communication rounds. Your utilization of FL is supported by their work as it shows that FL is superior to centralized ML/DL models in stroke diagnosis both performance and privacy-wise.

Sathyanarayana, L.L.[3] his paper in 2024 suggests a telestroke system that combines AI, telemedicine, mobile apps, and cloud infrastructure to provide fast stroke diagnosis and treatment access—notably in remote areas where specialist access is limited. The system applies AI analysis of real-time patient inputs (such as imaging data, vitals, symptoms) gathered from remote devices.

Major innovations are the employment of edge-AI inference models that can execute on mobile applications and cloud-connected diagnostic platforms. The article indicates that this multi-modal integration results in improved synchronization between healthcare professionals, which translates to shorter diagnosis-to-intervention time. Although it is more concerned with systems integration than DL models such as YOLO, it stresses the importance of light-weight, real-time, remote diagnosis systems, which your YOLOv8-based design with FL also facilitates.

Pulaparthi, et al.[4]. Their 2023 work involved developing a CNN-based stroke classification and segmentation model that was trained on CT and MRI scans. Their main contributions were the development of an automatic image preprocessing system, training CNN classifiers to classify between stroke and non-stroke conditions, and implementing an easy-to-use interface where doctors can upload scans through mobile or desktop applications.

Their architecture included convolutional layers, ReLU activations, pooling, and softmax classifiers for a binary classification output. Despite the system being centralized, they resolved practical issues such as data imbalance and overfitting through augmentation and dropout layers. Their conclusion indicated that CNNs were capable of minimizing hemorrhage progression risk significantly by allowing early and accurate detection. This is consistent with your application of YOLOv8 (CNN-based detector) for real-time classification, as your FL integration provides the solutions to the privacy concerns they could not address.

Gowri, P.; N. S ,Sivapriya, G.; G. K, P and M, S[5]. The authors contributed a federated machine learning stroke prediction model combining structured datasets—the two from the web and the one from the hospital. The authors employed federated logistic regression, trained using PySyft, mimicking hospital nodes as clients. The authors' solution averaged model weights after each round across five iterations and ten epochs, yielding a 98.8% accuracy rate, performing better than centralized Random Forest and Logistic Regression models.They also tackled data preprocessing issues, including missing values and heterogeneity in EHR data. Their focus on privacy, scalability, and real-time performance in health systems aligns with the architectural choices in your project. Further, it indicates that FL is not only appropriate for image data (such as YOLOv8), but also for structured hospital datasets.

Victor, N.; et al.[6] This work introduced the FL-PSO framework—a federated learning setup enriched with Particle Swarm Optimization (PSO) for the tuning of hyperparameters like learning rate, batch size, and number of local epochs. They have utilized this tuned model to forecast brain stroke risk based on real-time sensor readings of smart healthcare devices.

PSO was employed to determine the optimal configuration that optimized accuracy while reducing communication cost between server and client nodes. Their findings showed enhanced model convergence and overfitting reduction. The FL-PSO architecture is evidence that adaptive optimization in FL can result in performance improvement—something that may be investigated in your research to optimize YOLOv8 models for heterogeneous datasets.

Zhao, Ruihui; et al.[7]. presented a federated prediction model that runs on cloud infrastructure for preventing stroke. Their model accommodates asynchronous client connections, variable local training, and applies federated averaging to train models on hospital data without sending sensitive patient information. They observed an improvement of 10–20% in accuracy in small hospitals that earlier lacked sufficient training samples.

Also, they created a mobile app interface that displays stroke risk predictions visually, enabling real-time decision-making. Their research proves that FL can significantly scale stroke prediction models to heterogeneous institutions while ensuring fairness and utility across clients with different data volumes. This makes the federated environment of your model justified.

Srivastava, Utkarsh; et al[8]. In this paper published in 2020, the authors addressed intracranial hemorrhage (ICH) detection with a Time-Distributed Convolutional Neural Network (TDCNN) trained from CT scans. The architecture had both spatial and sequential dependencies on image slices for better localization of hemorrhagic areas.
They noted more than 92% detection accuracy but saw the necessity of a federated learning upgrade to facilitate secure learning among hospitals. Their proposed FL extension would permit parameter model sharing while maintaining scans local. This aligns directly with your application where stroke detection (also hemorrhage involved) is done through YOLOv8 but in federated mode for privacy.

Tan, K.; Marvell, et al[9]. conducted a systematic review of literature (SLR) on the detection of early ischemic stroke using deep learning. Upon considering various methods, they found that CNNs trained with CT and MRI data provided the highest accuracy. Their review placed high value on data curation, augmentation, and bias minimization to reach generalizable models.

They also emphasized the need for interpretable models in medicine, proposing explainable AI systems such as Grad-CAM or SHAP. Your use of YOLOv8 (a CNN-based detector) follows their results, and you might improve your model further by adding interpretability mechanisms.

Rangel, E. et al.[10]. In this 2024 IEEE EMBC paper, the authors created a federated stroke segmentation model that learned from annotated datasets across several institutions. Their approach overcomes an important limitation—most hospitals do not have annotated stroke data for training DL models. Their system enables collaborative training between data-rich and data-poor institutions through federated segmentation methods.

They tested their models on different institutions and found that federated collaboration not only enhanced segmentation performance but also identified domain shift issues. This paper is very important for your project as it shows how FL can be scaled up to image segmentation and assist institutions with fewer resources, which can be an extension of your stroke detection system in the future.

III. PROPOSED SYSTEM

Current stroke detection systems using CNNs analyze CT scans to identify strokes, but they come with several challenges. These systems rely on centralized data collection, which raises privacy concerns and makes sharing patient data between hospitals difficult. They also require powerful hardware, making real-time detection slow and impractical in resource-limited settings. Variations in scan quality, scanner types, and patient demographics can affect accuracy, leading to inconsistent results. Additionally, CNNs need large, well-annotated datasets, which are often hard to obtain. Another major challenge is that their decision-making process is complex and difficult for doctors to interpret, making AI-driven diagnoses harder to trust in critical situations.

Current stroke detection systems using CNNs rely on centralized data collection, raising serious privacy concerns and making data sharing between hospitals difficult. This limits model improvement while also increasing the risk of data breaches. Additionally, high computational demands and variations in CT scan quality affect real-time performance and accuracy. These challenges make existing systems less practical for widespread clinical use.

The main objectives of our proposal are:

- To ensure data privacy by implementing federated learning for decentralized stroke detection without sharing patient data.
- To enhance real-time detection by utilizing YOLOv8 for fast and accurate stroke identification in CT scans.
- To improve model efficiency and generalization by reducing computational overhead and adapting to variations in scan quality.

The new system swaps out CNNs for federated learning and YOLOv8. This change helps keep data private, boosts efficiency, and allows for quick stroke detection. With federated learning, models can be trained without sharing sensitive data. This makes it safer and helps the model work better with different types of CT scans. YOLOv8 is built to detect strokes quickly and accurately while using less computing power. Overall, this method tackles privacy issues, speeds up detection time, and makes it easier to use for doctors in real settings.

III. METHODOLOGY

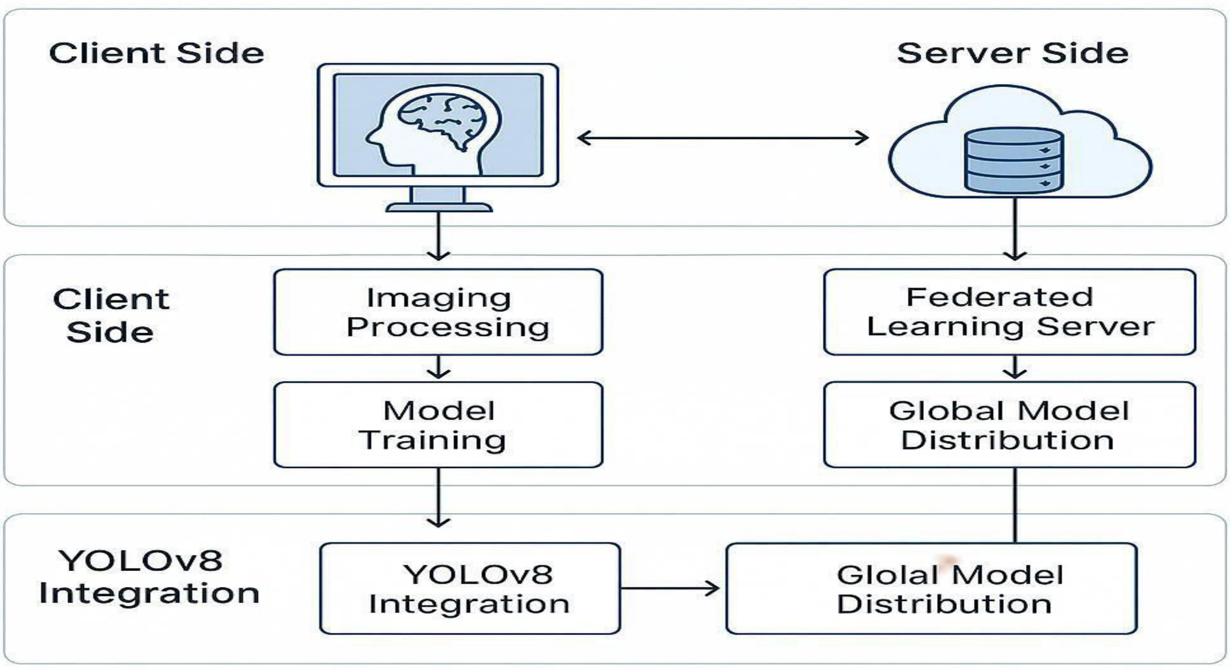

Fig (a) Methodology Diagram

This part describes the approach taken to design and implement a real-time stroke detection system integrating Federated Learning (FL) and YOLOv8. The envisioned system is to provide high accuracy, real-time inference, and rigorous data privacy adherence by dividing the training process among multiple clients (e.g., hospitals), without revealing sensitive patient information. The system consists of three primary components: the client side, the server side, and the YOLOv8 integration layer.

3.1 Client Side (Local Institutional Nodes)

Every client in the federated system is a healthcare institution or a diagnostic center. The clients have local datasets of CT scan or MRI images of the brain. In compliance with privacy laws like HIPAA or GDPR, these datasets are never sent outside; rather, all the training and processing are done locally.

3.1.1 Image Preprocessing

Before model training, CT or MRI images go through various preprocessing operations to stabilize and enhance model performance:

Noise Reduction: Denoising filters are used to clean images and minimize background noise. **Normalization:** Pixel intensity values are normalized to a fixed range to allow smooth training convergence. **Resizing:** All images are resized to 640×640 pixels, compatible with the input of YOLOv8.
Annotation: Each image is annotated in YOLO style, with bounding boxes indicating stroke or non-stroke area.
This preprocessing pipeline converts raw medical images into organized datasets appropriate for object detection tasks.

3.1.2 Local Model Training

After preprocessing, each client starts and fine-tunes a local YOLOv8 model with its own data. YOLOv8 is chosen for its real-time detection performance and high accuracy in object localization and classification.

Key features of this stage are:

- Use of pretrained YOLOv8 weights, enabling faster convergence and less overfitting.
- Binary classification of stroke and non-stroke cases, with stroke-affected areas being boxed in bounding boxes.
- Data privacy is maintained since no raw images are sent to the server—only the model weights are transmitted.
- The output of this stage is a locally optimized model that can detect stroke features in CT scans.

3.2 Server Side (Federated Aggregator)

The central server serves as the coordinator of the federated learning framework. It supports model updates, aggregation, and redistribution without direct access to client data.

3.2.1 Federated Averaging (FedAvg)

The server collects model parameters (weights) from all contributing clients and does federated averaging using the FedAvg algorithm. This process is given by:

$$\theta_{\text{global}} = \sum_{i=1}^{n} \frac{n_i}{N} \theta_i$$

Where:

- θ_i denotes the model weights from client i,
- n_i is the number of samples at client i,
- N is the total number of samples across all clients,
- θ_{global} is the new aggregated model.

This gives a global model that aggregates the learned representation across all institutions, enhancing generalization to diverse patient populations and imaging configurations.

3.3 Global Model Redistribution

After aggregation, the global model is broadcast again to all the clients. All the clients proceed to use the model as a starting point for the next rounds of local training. This process is repeated over and over again until the model converges in accuracy, precision, and recall for all datasets.This way, the process guarantees that: Model performance becomes better with each round. Institutions gain collective knowledge without releasing sensitive information. The global model becomes resilient to non-IID (non-independent and identically distributed) variations in data.

1. **YOLOv8 Integration Layer**

© 2025 IJIRCCE | Volume 13, Issue 5, May 2025| DOI:10.15680/IJIRCCE.2025.1305143

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

YOLOv8 is the main object detection engine of the system. It is used both at the client side for local training and at the global side for aggregation and inference.

3.3.1 Model Architecture and Capabilities

YOLOv8, created by Ultralytics, is a fast deep learning model optimized for real-time object detection. It offers image classification, segmentation, and bounding box regression, which are all suitable for medical imaging applications where precise localization is essential.

Advantages of YOLOv8 in this application are:

- Real-time processing rates adequate for emergency diagnostics.
- High accuracy detection for small and subtle details, like ischemic or hemorrhagic areas.
- Compatibility with resource-limited devices through light variants (e.g., YOLOv8n, YOLOv8s).

3.4 Deployment and Inference

Once multiple rounds of federated updates are performed, the resulting YOLOv8 model is deployed to every client. This model is employed for:

- Real-time stroke detection on newly acquired CT scans.
- Aiding clinical decision-making without needing remote server access.
- Smooth updates through continuous federated rounds so that the model stays up to date with new medical data.

3.5 Iterative Learning Workflow

The entire training pipeline works iteratively as described below:

1. The server starts a global YOLOv8 model.
2. The model is shared with all participating clients.
3. The clients train the model locally on their respective data.
4. Train weights are sent back to the server.
5. The server combines the weights through FedAvg.
6. The new global model is shared with clients.
7. Steps 3–6 are repeated for several communication rounds until convergence.

3.6 Benefits of the Proposed Methodology

Feature Benefit: Federated Learning Provides end-to-end patient data privacy across institutions.
YOLOv8 Architecture: Provides accurate and quick stroke detection from CT images.
Localized Training: Fine-tunes the model for client-specific data distributions.
Global Aggregation: Increases final model's generalization and robustness.
Real-Time Inference: Permits instantaneous clinical decision in emergency cases.

© 2025 IJIRCCE | Volume 13, Issue 5, May 2025| DOI:10.15680/IJIRCCE.2025.1305143

www.ijircce.com | e-ISSN: 2320-9801, p-ISSN: 2320-9798| Impact Factor: 8.771| ESTD Year: 2013|

International Journal of Innovative Research in Computer and Communication Engineering (IJIRCCE)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

IV. RESULTS

The following chapter presents the results derived at the end of the flow.

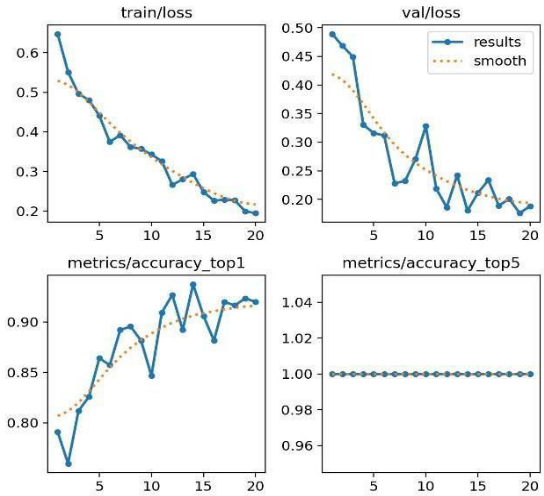

Fig(b): Model Performance Metrics

Fig(c): Training Metrics

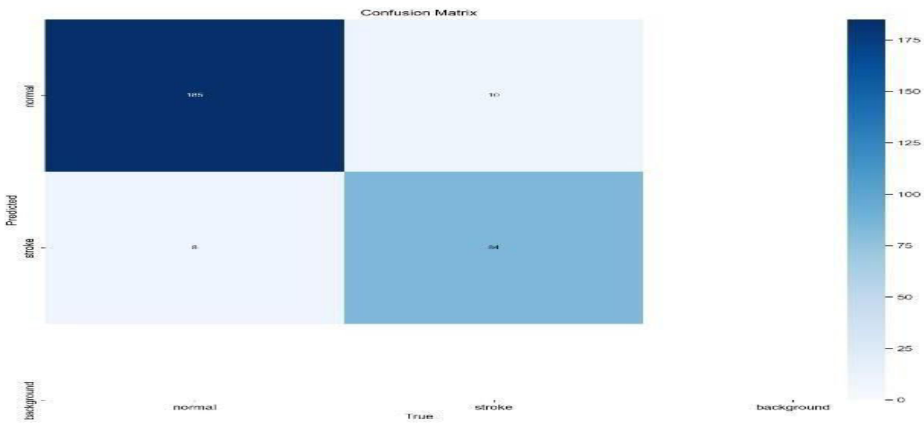

Fig(d): Confusion Matrix

Key Metrics from This: Accuracy = (TP + TN) / Total
= (185 + 84) / (185 + 10 + 8 + 84)
= 269 / 287 ≈ 93.7%
Precision (Stroke) = TP / (TP + FP) = 84 / (84 + 10) ≈ 89.4% Recall (Stroke) = TP / (TP + FN) = 84 / (84 + 8) ≈ 91.3%
F1 Score (Stroke) = Harmonic mean of precision & recall ≈ 90.3%.

V. CONCLUSION AND FUTURE SCOPE

In conclusion, this presents a significant advancement in the field of medical diagnostics by addressing key limitations of traditional centralized stroke detection systems. By integrating Federated Learning, the system ensures data privacy and security, allowing multiple medical institutions to collaboratively train models without exposing sensitive patient data. This decentralized approach not only mitigates privacy concerns but also enhances the robustness and adaptability of the model across diverse datasets. Meanwhile, the adoption of YOLOv8, a cutting-edge object detection algorithm, facilitates rapid and accurate identification of stroke-affected regions in CT scans with minimal computational overhead.

Together, these technologies create a powerful framework that enables timely diagnosis, which is critical in preventing long-term damage and improving patient survival rates. The system is designed to operate efficiently in real- world clinical settings, including resource-constrained environments, making it a scalable and practical solution. Furthermore, its ability to generalize across various scan qualities and medical imaging devices enhances its reliability and usability in heterogeneous healthcare ecosystems. In summary, this project demonstrates how modern AI techniques can be harnessed to deliver intelligent, privacy-conscious, and life-saving healthcare tools that pave the way for future innovations in real-time medical diagnostics.

In envisioning the future scope of our project, which holds immense potential in transforming stroke detection and broader medical diagnostics. One of the primary extensions is the real-time deployment of this system in hospitals and emergency units, where rapid diagnosis can be critical to saving lives. Integrating the model into user-friendly web and mobile applications will make it accessible to healthcare professionals anytime and anywhere, ensuring swift decision-making in critical scenarios. The system can also be scaled to support other types of brain abnormalities such as tumors, hemorrhages, and traumatic injuries by fine-tuning the detection algorithms. Moreover, incorporating multimodal data like patient medical history, MRI images, and genetic information could significantly improve the model's accuracy and clinical relevance. As the federated learning infrastructure expands, it can connect hospitals across different regions or countries, creating a vast collaborative ecosystem that continuously enhances diagnostic performance while preserving privacy. Additionally, incorporating explainable AI techniques—such as heatmaps or confidence visualizations—will help doctors better understand and trust the model's outputs, leading to higher adoption in clinical environments. With support for multiple languages and adaptive learning features, the system can be customized for global usage, even in rural or underserved areas. Overall, this project lays a strong foundation for developing next-generation, intelligent, and ethical healthcare technologies.

REFERENCES

1. Elhanashi, A., Dini, P., Saponara, S. et al. TeleStroke: real-time stroke detection with federated learning and YOLOv8 on edge devices. J Real-Time Image Proc 21, 121 (2024). 10.1007/s11554-024-01500-1
2. Mansour, A., Besbes, O., Abdellatif, T. (2025). Federated Deep Learning Models for Stroke Prediction. In: Barhamgi, M., Wang, H., Wang, X. (eds) Web Information Systems Engineering – WISE 2024. Lecture Notes in Computer Science, vol 15439. Springer, Singapore.
3. Sathyanarayana, L. L. (2024). Distributed AI-Driven Telestroke Solution for Rapid and Accurate Stroke Diagnosis. Mapana Journal of Sciences.
4. Naga MahaLakshmi Pulaparthi, Madhulika Dabbiru, Charishma Penkey, Dr. Nrusimhadri. Brain Stroke Detection Using Deep Learning. International Journal of Research Publication and Reviews, April 2023.
5. Gowri, P., N. S, Sivapriya, G., G. K, P., S. M. Stroke prediction analysis using federated machine learning. 2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT), Delhi, India, 2023.
6. Victor, N., Bhattacharya, S., Maddikunta, P. K. R., Alotaibi, F. M., Gadekallu, T. R., Jhaveri, R. H. FL-PSO: A Federated Learning approach with Particle Swarm Optimization for Brain Stroke Prediction. 2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW), Bangalore, India, 2023.
7. Zhao, R., Sun, J., Wei, X., Zhao, B., Liu, Y., Li, H., Chen, T., Zhang, X., Gao, D., Tan, B., Yu, H., Jin, Y. (2020).
8. Privacy-Preserving Technology to Help Millions of People: Federated Prediction Model for Stroke Prevention. 10.48550/arXiv.2006.10517
9. Srivastava, U., Upadhyay, D., Sharma, V. (2020). Intracranial Hemorrhage Detection Using Neural Network Based Methods With Federated Learning. 10.13140/RG.2.2.10157.31206
10. Kodi, D. (2024). Automating Software Engineering Workflows: Integrating Scripting and Coding in the Development Lifecycle . Journal of Computational Analysis and Applications (JoCAAA), 33(4), 635–652.
11. Tan, K., Marvell, Y. A., Gunawan, A. A. S. Early Ischemic Stroke Detection Using Deep Learning: A Systematic Literature Review. 2023 International Seminar on Application for Technology of Information and Communication (iSemantic), Semarang, Indonesia, 2023.
12. Rangel, E., Gómez, S., Mantilla, D., Camacho, P., Martínez, F. (2024, July). A federated stroke segmentation to impact limited data institutions. 2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 1–4. IEEE.

INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN COMPUTER & COMMUNICATION ENGINEERING

9940 572 462 6381 907 438 ijircce@gmail.com

www.ijircce.com

Scan to save the contact details