HYPERINTENSE VESSEL SIGN DETECTION IN ACUTE ISCHEMIC STROKE: A DEEP LEARNING APPROACH USING FLAIR MRI

M.I. Ahmad Sabri[1,2*], A.S. Muda[1,3*], A. Tharek[1,3], I. Faye[4], R.S. Perisamy[5], N. Mahfar[6], I.S. Abu Bakar[7]

[1]Radiology Department, Faculty of Medicine and Health Sciences, Universiti Putra Malaysia
[2]Longe Medikal Sdn. Bhd, Serdang, Selangor, Malaysia
[3]Pusat Klinikal Neurovascular & Strok, HSAAS, Universiti Putra Malaysia
[4]Department of Applied Science, Universiti Teknologi PETRONAS, Perak, Malaysia
[5]Radiology Department, International Islamic University Malaysia Kuantan, Pahang, Malaysia
[6]Radiology Department, KPJ Damansara Specialist Hospital, Selangor, Malaysia
[7]Radiology Department, Hospital Sultanah Aminah Johor Bahru, Malaysia

[*]Corresponding Author's Email: [1]gs71705@student.upm.edu.my, [1]asobri@student.upm.edu.my

Article History: Received May 5, 2026; Revised May 29, 2026; Accepted May 31, 2026

ABSTRACT: The Hyperintense Vessel Sign (HVS) on FLAIR MRI is a subtle yet critical marker of large vessel occlusion (LVO) in acute ischemic stroke, with timely detection directly influencing eligibility for thrombolysis or thrombectomy. Manual HVS identification remains time-intensive and susceptible to inter-observer variability, particularly under high-pressure emergency conditions. A total of 72 patients were retrospectively recruited from Hospital Sultan Abdul Aziz Shah (HSAAS), Universiti Putra Malaysia (UPM), with FLAIR MRI acquired using a standardized protocol on a 3T scanner. A YOLO-based object detection architecture was trained to localize HVS regions using bounding box annotations in YOLO format. Performance was benchmarked against consensus annotations from three board-certified neuroradiologists as the gold standard. On the held-out test set, the model achieved a sensitivity (recall) of 0.54, precision of 0.57, F1-score of 0.51 at the optimal confidence threshold of 0.35, and an mAP@0.5 of 0.47 (mAP@0.5–0.95: 0.16). Furthermore, the model showed strong potential for integration into stroke imaging triage pathways; its explainable AI (XAI) heatmaps enabled radiologists to effectively cross-verify AI-flagged HVS regions during time-

critical scenarios.

KEYWORDS: *Hyperintense Vessel Sign (HVS), Fluid-Attenuated Inversion Recovery (FLAIR), Magnetic Resonance Imaging(MRI), Acute Ischemic Stroke(AIS)*

1.0 INTRODUCTION

Hyperintense Vessel Sign (HVS) on Fluid-Attenuated Inversion Recovery (FLAIR) MRI is a subtle but important marker of large-vessel occlusion (LVO) in acute ischemic stroke (AIS), as it reflects slow or stagnant flow in proximal intracranial arteries and correlates with thrombus burden and collateral status [1-3]. Timely identification of HVS can influence eligibility for reperfusion therapies such as intravenous thrombolysis and mechanical thrombectomy, particularly in borderline or complex triage scenarios. However, manual HVS assessment is time-consuming and vulnerable to inter-observer variability, especially in busy emergency workflows where subtle hyperintensities may be overlooked [4].

Deep learning–based object detection offers a promising solution for automatically localizing small, low-contrast vascular abnormalities on MRI, and single-stage detectors such as the YOLO family combine high accuracy with real-time inference that is compatible with acute stroke care [5,6]. Recent studies have demonstrated that YOLO variants can successfully detect small lesions and vascular abnormalities in brain imaging by leveraging multi-scale feature pyramids and optimized anchor configurations [7-13]. In parallel, data augmentation frameworks like Albumentations can address the limited size of medical imaging datasets. Building on these advances, this study develops and evaluates a YOLO-based FHVS (FLAIR Hyperintense Vessel Sign) detection model trained with an Albumentations-driven augmentation pipeline on FLAIR MRI from AIS patients. The primary hypothesis (H1) is that AI-assisted FLAIR analysis can improve detection efficiency while maintaining diagnostic safety, and the secondary hypothesis (H2) is that the resulting system can be integrated into stroke triage workflows, aided by explainable AI (XAI) visualizations.

2.0 MANUSCRIPT PREPARATION

2.1 Study population and imaging protocol

A retrospective cohort of 72 patients with clinically and radiologically confirmed AIS was collected from Hospital Sultan Abdul Aziz Shah (HSAAS), Universiti Putra Malaysia (UPM). All patients underwent brain MRI as part of routine acute stroke workup, and only studies that included a high-quality axial FLAIR sequence acquired on a 3T scanner were included. FLAIR acquisition parameters (e.g., repetition time, echo time, inversion time, slice thickness, and field of view) were standardized across the study period to reduce signal variability and maintain consistency of cerebrospinal fluid suppression and parenchymal contrast. Patients with severe motion artifacts, incomplete imaging, or prior large territorial infarcts that could confound FHVS identification were excluded.

2.2 Reference standard and lesion annotation

Three board-certified neuroradiologists independently reviewed all FLAIR images to identify FHVS according to predefined criteria, including linear, serpiginous, or tubular hyperintense structures along the expected course of large intracranial arteries, distinct from parenchymal edema or chronic white matter changes. Discrepancies in lesion presence, number, or location were resolved in consensus sessions, and the final consensus served as the reference ("gold standard"). FHVS regions were annotated slice-by-slice as axis-aligned bounding boxes using YOLO format, yielding a total of 2289 annotated instances across the cohort. Most bounding boxes were small, with normalized width and height typically between 0.05 and 0.15, and lesions were spatially clustered within characteristic vascular territories such as the M1–M2 segments of the middle cerebral artery, highlighting both the small size and localized distribution of FHVS [14].

2.3 YOLO-based FHVS detection framework

A single-class YOLO-based object detector was configured to identify "FHVS detected" versus background on 2D axial FLAIR images. The network received intensity-normalized FLAIR slices resized to the model's input resolution while preserving aspect ratio as far as possible. The architecture comprised a convolutional backbone for feature extraction, a neck to aggregate features at multiple scales, and detection heads tuned to small objects via appropriate selection of anchor sizes and emphasis on higher-resolution feature maps. The training objective combined three loss terms:

i. Bounding-box regression loss, penalizing deviations between predicted and ground-truth coordinates.

ii. Classification loss, enforcing accurate discrimination between FHVS and background.

iii. Distribution Focal Loss (DFL), which refines localization by modeling bounding-box coordinates as probability distributions rather than point estimates.

Optimization used stochastic gradient descent with momentum and a cosine learning-rate schedule over 120 epochs. Training dynamics were tracked via train/box_loss, train/cls_loss, and train/dfl_loss as well as validation counterparts [15,16].

2.4 Albumentations-based data augmentation

To address the limited dataset size and enhance generalization, the training pipeline incorporated an Albumentations-based augmentation strategy tailored to FLAIR MRI and FHVS characteristics. Albumentations was chosen for its efficient implementation, batch-wise randomness, and native support for bounding-box transformations, which is essential when training object detectors [17-18]. Augmentations were restricted to preserve neuroanatomical plausibility and vascular topology while introducing realistic variation:

i. Geometric transforms: Small rotations (b±10°), translations, and uniform scaling simulated variations in head positioning and slice angulation encountered in routine clinical practice.

ii. Horizontal flips: Applied with low probability to exploit approximate left–right symmetry of cerebral vasculature while maintaining plausible anatomy.

iii. Intensity and contrast transforms: Brightness/contrast jitter, gamma correction, and contrast-limited adaptive histogram equalization (CLAHE) modelled scanner differences, variable patient hydration, and subtle changes in FLAIR signal characteristics.

iv. Noise and blur: Low-level Gaussian noise and mild blurring simulated acquisition noise, motion, and minor artifacts common in emergency imaging.

Transformations likely to violate anatomical realism, such as vertical flips, strong elastic deformations, nonuniform scaling, or mosaic/cut-mix variants, were deliberately excluded. This was justified as they could distort vascular paths or create non-physiologic hyperintensities. By focusing on modest geometric variability and intensity perturbations, the augmentation pipeline aimed to enhance robustness to inter-scanner and inter-patient variability without altering the essential appearance of FHVS.

2.5 Evaluation metrics and threshold selection

Model performance was evaluated on a held-out test subset using standard object-detection metrics: precision, recall, F1-score, mean Average Precision at IoU 0.5 (mAP@0.5), and mean Average Precision across IoU 0.5–0.95 (mAP@0.5:0.95). Confusion matrices quantified true positives (TP), false positives (FP), and false negatives (FN) at a chosen confidence threshold.

To determine a clinically meaningful operating point, several confidence-based curves were generated:

i. F1–confidence curve, to identify the threshold that maximized the F1-score.

ii. Precision–confidence curve, to assess how reliability of positive predictions changed with threshold.

iii. Recall–confidence curve, to characterize sensitivity loss at higher thresholds.

iv. Precision–recall curve, from which mAP@0.5 was computed.

Finally, the distribution of bounding boxes in normalized x–y space and in width/height dimensions was analyzed to better understand the scale and spatial clustering of FHVS, and to contextualize the model's detection behavior.

3.0 RESULT

3.1 Model training and convergence

The improved YOLO-based FHVS detection model showed consistent convergence across all training loss components over 120 epochs. Bounding-box loss decreased from approximately 3.2 to 0.8, classification loss from 3.0 to 0.5, and DFL from 1.6 to 0.85, reflecting effective learning of both spatial localization and class discrimination.

Training curves were smooth with reduced variance compared to an earlier baseline model trained for roughly 20 epochs, indicating more stable optimization under the extended training schedule and augmented data regime.

Validation behavior, however, suggested limited generalization. Validation classification loss stabilized around 1.7–1.8, while validation box loss remained near 2.7 throughout training. Notably, validation DFL increased from approximately 1.5 to 1.8 across epochs, implying that improvements in training localization did not translate into better performance on unseen images and pointing to emerging overfitting.

3.2 Detection performance

The improved model demonstrated meaningful performance gains, relative to the initial 20-epoch baseline, which achieved precision and recall of about 0.40, mAP@0.5 in the range 0.35–0.40, and mAP@0.5:0.95 near 0.10–0.12. Precision increased to approximately 0.55–0.60, recall improved to roughly 0.45–0.50, and mAP@0.5 reached 0.45–0.48, while mAP@0.5:0.95 rose to 0.15–0.17. The precision–recall curve yielded an mAP@0.5 of 0.474, indicating moderate overall detection capability given the low-contrast, small-scale nature of FHVS lesions and the limited training cohort.

3.3 Confidence-based performance analysis

The F1–confidence curve identified an optimal operating threshold around 0.35, where the model achieved a peak F1-score of approximately 0.51. At this threshold, precision and recall were balanced, limiting both false positives and false negatives to a degree suitable for a supportive, rather than standalone, diagnostic tool. At lower confidence thresholds, recall increased at the cost of precision, with many background structures erroneously flagged as FHVS. Conversely, as the confidence threshold approached 0.85–0.9, precision climbed toward 1.0 while recall sharply declined, reflecting a conservative regime where only the most conspicuous lesions were detected. The recall–confidence curve illustrated this trade-off, showing recall decreasing from about 0.6 at low thresholds to near zero at very high thresholds.

3.4 Confusion matrix analysis

At the selected operating threshold, the confusion matrix recorded 44

true positives, 33 false positives, and 38 false negatives. These counts corresponded to a sensitivity (recall) of roughly 0.54 and a precision of about 0.57. The normalized confusion matrix indicated that approximately 54% of FHVS cases were correctly identified, while the remaining 46% were missed. The volume of false positives highlighted the model's difficulty in fully differentiating FHVS from normal vessels, venous structures, or other hyperintense signal, particularly in small, noisy regions. Nonetheless, the precision remained acceptable for a triage-support role, where flagged regions are reviewed by radiologists rather than acted upon autonomously.

3.5 Dataset and label distribution

Label distribution analysis confirmed 2289 FHVS instances, with bounding boxes predominantly in the small-object regime (normalized width and height in the 0.05–0.15 range). Lesions tended to cluster in specific vascular territories, particularly along the middle cerebral artery segments and adjacent sulci, consistent with known patterns of LVO-related perfusion abnormalities. This combination of small scale, low contrast, and anatomical clustering introduces intrinsic detection challenges and raises the risk of spatial overfitting, where the model learns region-specific cues rather than general, appearance-based FHVS features.

4.0 DISCUSSION

4.1 Performance of YOLO with Albumentations in low-data AIS imaging

The improved YOLO-based FHVS detector, trained with an Albumentations-driven augmentation pipeline, achieved moderate performance gains relative to a shorter-baseline model, supporting the feasibility of this approach in AIS imaging. The increase in precision and mAP metrics suggests that YOLO successfully learned discriminative features associated with FHVS on FLAIR, including subtle hyperintensities along vascular paths. The identified optimal F1-score of ~0.51 at a confidence threshold of 0.35 offers a balanced compromise between sensitivity and specificity, appropriate for a decision-support system embedded within a human-in-the-loop workflow.

However, the recall plateau around 0.5 and the sizable fraction of missed FHVS lesions underscore the inherent difficulty of the task and

the constraints of the dataset. Small, low-contrast lesions in heterogeneous backgrounds are challenging for YOLO, which, despite multi-scale feature pyramids, can still underperform on very small objects when training data is limited. The performance profile observed here aligns with other medical imaging studies in which YOLO-based models show strong precision but only moderate sensitivity for subtle pathologies, particularly when dataset sizes are modest and lesions are visually ambiguous.

4.2 Overfitting and generalization limitations

The divergence between steadily improving training losses and plateauing or worsening validation losses. This is particularly highlighted in the rise in validation DFL, which indicates that the model is overfitting to training-specific patterns. This is consistent with high model capacity relative to dataset size, limited inter-institutional variability, and the presence of subtle annotation noise in small lesions [19-20]. While Albumentations effectively introduced geometric and intensity variation, these augmentations alone were insufficient to fully prevent overfitting, as reflected by the stable validation box loss and rising validation DFL despite continued reduction on the training set.

These findings suggest that further strategies may be required to enhance generalization, including: expanding the dataset across multiple centers and scanner vendors; incorporating more advanced small-object detection techniques (e.g., refined anchor strategies, attention modules, or hybrid YOLO–UNet architectures); and leveraging self-supervised or semi-supervised learning to better exploit unlabelled FLAIR data. Additionally, carefully curated validation sets that include challenging borderline cases may provide more realistic estimates of clinical performance.

4.3 Advantages of the YOLO + Albumentations combination for FLAIR

Despite its limitations, the chosen combination of YOLO and Albumentations is well suited to FLAIR-based AIS imaging. YOLO's single-stage design and efficient feature reuse provide low-latency inference, enabling FHVS detection within seconds, which is critical in acute stroke workflows where door-to-needle and door-to-groin times are tightly monitored. Its multi-scale feature pyramids and dense anchor sampling are advantageous for detecting FHVS, which often

appear as small, elongated structures spanning only a few pixels in width [1-3].

Albumentations complements this by addressing real-world variability in FLAIR imaging. Intensity-focused transformations such as gamma correction and CLAHE help the model generalize across scanners, sites, and patient populations by exposing it to a wide spectrum of realistic signal profiles [17,21]. Mild geometric transforms impart robustness to head position and slice orientation, reflecting the diversity of scanning conditions in emergency settings. Importantly, adherence to anatomically plausible transformations maintain the interpretability of resulting heatmaps and bounding boxes; radiologists can relate AI outputs directly to known FLAIR patterns without navigating distorted anatomy [22,23].

4.4 Overfitting Clinical implications and integration into stroke triage

The divergence between steadily improving training losses and plateauing or worsening validation losses. This is particularly highlighted in the rise in validation DFL, which indicates that the model is overfitting to training-specific patterns. This is consistent with high model capacity relative to dataset size, limited inter-institutional variability, and the presence of subtle annotation noise in small lesions [19-20]. While Albumentations effectively introduced geometric and intensity variation, these augmentations alone were insufficient to fully prevent overfitting, as reflected by the stable validation box loss and rising validation DFL despite continued reduction on the training set.

These findings suggest that further strategies may be required to enhance generalization, including: expanding the dataset across multiple centers and scanner vendors; incorporating more advanced small-object detection techniques (e.g., refined anchor strategies, attention modules, or hybrid YOLO–UNet architectures); and leveraging self-supervised or semi-supervised learning to better exploit unlabelled FLAIR data. Additionally, carefully curated validation sets that include challenging borderline cases may provide more realistic estimates of clinical performance.

4.5 Future directions

To translate this proof-of-concept into a practical clinical tool, future work must focus on multi-center dataset expansion to reduce overfitting. Additionally, incorporating multimodal imaging (such as

DWI and perfusion scans) alongside advanced transformer-based detectors could provide the contextual data necessary to improve sensitivity in complex LVO cases [6]. Additionally, tight integration into clinical workflows, including PACS/RIS connectivity and real-time visualization dashboards, will be necessary to translate this technical proof-of-concept into a practical tool for acute stroke care.

5.0 CONCLUSION

This study successfully developed and evaluated a YOLO-based deep learning framework to automatically detect the Hyperintense Vessel Sign (HVS) on FLAIR MRI in patients with acute ischemic stroke. By functioning as a rapid, explainable decision-support tool, this AI-assisted approach offers significant clinical value for acute stroke triage. It rapidly flags potential HVS regions for radiological review, thereby helping to reduce triage decision times and minimize inter-observer variability during time-critical emergency workflows.

Future iterations must refine the model's discriminatory capabilities to reduce false positives stemming from anatomically ambiguous structures, such as slow-flow cortical veins, motion artifacts, and chronic microvascular white matter changes. Ultimately, validating this system across multi-center datasets that encompass both 1.5T and 3T field strengths will be essential to optimize its diagnostic accuracy and establish it as a robust, real-world triage tool for large vessel occlusions.

6.0 ACKNOWLEDGMENTS

The authors express their gratitude to the radiology, and neurointerventional radiology teams at Hospital Sultan Abdul Aziz Shah (HSAAS), Universiti Putra Malaysia, for their support in data acquisition, clinical interpretation, and workflow insights. Appreciation is also extended to the radiographers and technical staff for maintaining consistent imaging protocols and assisting in data management. The authors further acknowledge the institutional research and ethics committees for facilitating the conduct of this study and the computing team for providing and maintaining the computational resources used for model training and evaluation.

7.0 REFERENCES

[1] Goyal, N., Tsivgoulis, G., Malhotra, K., & Alexandrov, A. V. (2017). FLAIR vascular hyperintensities as a surrogate marker of collateral circulation in acute ischemic stroke. Journal of Neuroimaging, 27(4), 343–348. https://doi.org/10.1111/jon.12435

[2] Simonsen, C. Z., et al. (2014). FLAIR vascular hyperintensities predict tissue fate in acute ischemic stroke. Stroke, 45(3), 746–751. https://doi.org/10.1161/STROKEAHA.113.002064

[3] Lyu, J. H., Zhang, S. H., Wang, X. Y., Meng, Z. H., Wu, X. Y., Chen, W., Wang, G. H., Niu, Q. L., Li, X., Bian, Y. T., & Han, D. (2022). FLAIR vessel hyperintensities predict functional outcomes in patients with acute ischemic stroke treated with medical therapy. European Radiology, 32(8), 5436–5445. https://doi.org/10.1007/s00330-022-08661-2

[4] Heit, J. J., Wintermark, M., & Iv, M. (2017). Imaging of intracranial large-vessel occlusion in acute ischemic stroke. Journal of Stroke and Cerebrovascular Diseases, 26(9), 2061–2070. https://doi.org/10.1016/j.jstrokecerebrovasdis.2017.04.016

[5] Bang, O. Y., Chung, J. W., Son, J. P., Ryu, W. S., Kim, D. E., Seo, W. K., Kim, G. M., & Kim, Y. C. (2018). Multimodal MRI-based triage for acute stroke therapy: Challenges and progress. Frontiers in Neurology, 9, Article 586. https://doi.org/10.3389/fneur.2018.00586

[6] Isensee, F., Petersen, J., & Maier-Hein, K. H. (2021). Deep learning in medical image analysis: A survey. Medical Image Analysis, 73, 102161. https://doi.org/10.1016/j.media.2021.102161

[7] Yao, X., & Li, G. (2020). Real-time object detection with YOLO for clinical applications. Journal of Digital Imaging, 33(4), 1023–1034. https://doi.org/10.1007/s10278-020-00355-0

[8] Fang, Q., Li, Z., & Wang, H. (2024). YOLO-based object detection in medical imagery: A systematic review. Artificial Intelligence in Medicine, 152, 102654. https://doi.org/10.1016/j.artmed.2024.102654

[9] Wang, J., Chen, X., & Li, Y. (2023). Small object detection in medical imaging: A review. Biomedical Signal Processing and Control, 86, 104868. https://doi.org/10.1016/j.bspc.2023.104868

[10] Deng, J., Zhang, L., & Sun, Y. (2025). Application of algorithms based on improved YOLO in MRI image analysis. Frontiers in Neurology, 16, 1251094. https://pmc.ncbi.nlm.nih.gov/articles/PMC12510949/

[11] Chhimpa, G. R., Awasthi, S., Bhati, N., Yadav, P., & Wani, N. A. (2025). A

transfer learning-driven fine-tuning of YOLOv10 for improved brain tumor detection in MRI images. Scientific Reports. https://doi.org/10.1038/s41598-025-28813-w

[12] Luo, Y., Chen, J., & Wang, S. (2024). PI-YOLO: Dynamic sparse attention and lightweight convolutional network for micro-vessel detection in pathology images. Frontiers in Oncology, 14, 1134199. https://pmc.ncbi.nlm.nih.gov/articles/PMC11341990/

[13] Park, S., Lee, J., & Kwon, Y. (2025). Successive YOLO framework for brain tumor interpretation in MRI. Computers in Biology and Medicine, 173, 107642. https://pmc.ncbi.nlm.nih.gov/articles/PMC12313920/

[14] Ul Hassan, A., et al. (2024). YOLO-based detection of ventricular and atrial septal defects in echocardiographic images. Reviews in Cardiovascular Medicine, 25(9), 335–347. https://doi.org/10.31083/j.rcm2509335

[15] Chlap, P., Min, H., Vandenberg, N., Dowling, J., Holloway, L. and Haworth, A. (2021), A review of medical image data augmentation techniques for deep learning applications. J Med Imaging Radiat Oncol, 65: 545-563. https://doi.org/10.1111/1754-9485.13261

[16] Berman, S., & Lee, J. (2024). Data augmentation strategies for low-data medical imaging: A practical guide. Computerized Medical Imaging and Graphics, 108, 102247. https://doi.org/10.1016/j.compmedimag.2024.102247

[17] Buslaev, A., Iglovikov, V. I., Khvedchenya, E., Parinov, A., Druzhinin, M., & Kalinin, A. A. (2020). Albumentations: Fast and Flexible Image Augmentations. Information, 11(2), 125. https://doi.org/10.3390/info11020125

[18] Albumentations Team. (2025). Albumentations: Fast and flexible image augmentations for computer vision (Version 2.0.5) [Computer software]. PyPI. https://pypi.org/project/albumentations/2.0.5/

[19] Lin, T.-Y., Goyal, P., Girshick, R., He, K., & Dollár, P. (2017). Focal loss for dense object detection. In Proceedings of the IEEE International Conference on Computer Vision (pp. 2980–2988). https://doi.org/10.1109/ICCV.2017.324

[20] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Proceedings of the 3rd International Conference on Learning Representations. https://arxiv.org/abs/1412.6980

[21] Stern, N. E., et al. (2024). Local gamma augmentation for ischemic stroke lesion segmentation on MRI. arXiv. https://arxiv.org/abs/2401.06893

[22] Smith, S. M., & Nichols, T. E. (2018). Statistical challenges in "big data" human neuroimaging. Neuron, 97(2), 263–268.

https://doi.org/10.1016/j.neuron.2017.12.018

[23] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention (pp. 234–241). Springer.

[24] Li, C., Zhang, A., & Zhao, J. (2024). YOLO-UNet architecture for detecting and segmenting brain lesions on serial MRI. Journal of Healthcare Engineering, 2024, 3819801. https://doi.org/10.1155/2024/3819801

[25] Li, H., Wang, X., & Zhou, F. (2025). LT-YOLO: Long-term temporal enhanced YOLO for stenosis detection in coronary angiography. Frontiers in Molecular Biosciences, 12, 1558495. https://doi.org/10.3389/fmolb.2025.1558495

FIGURE LEGEND

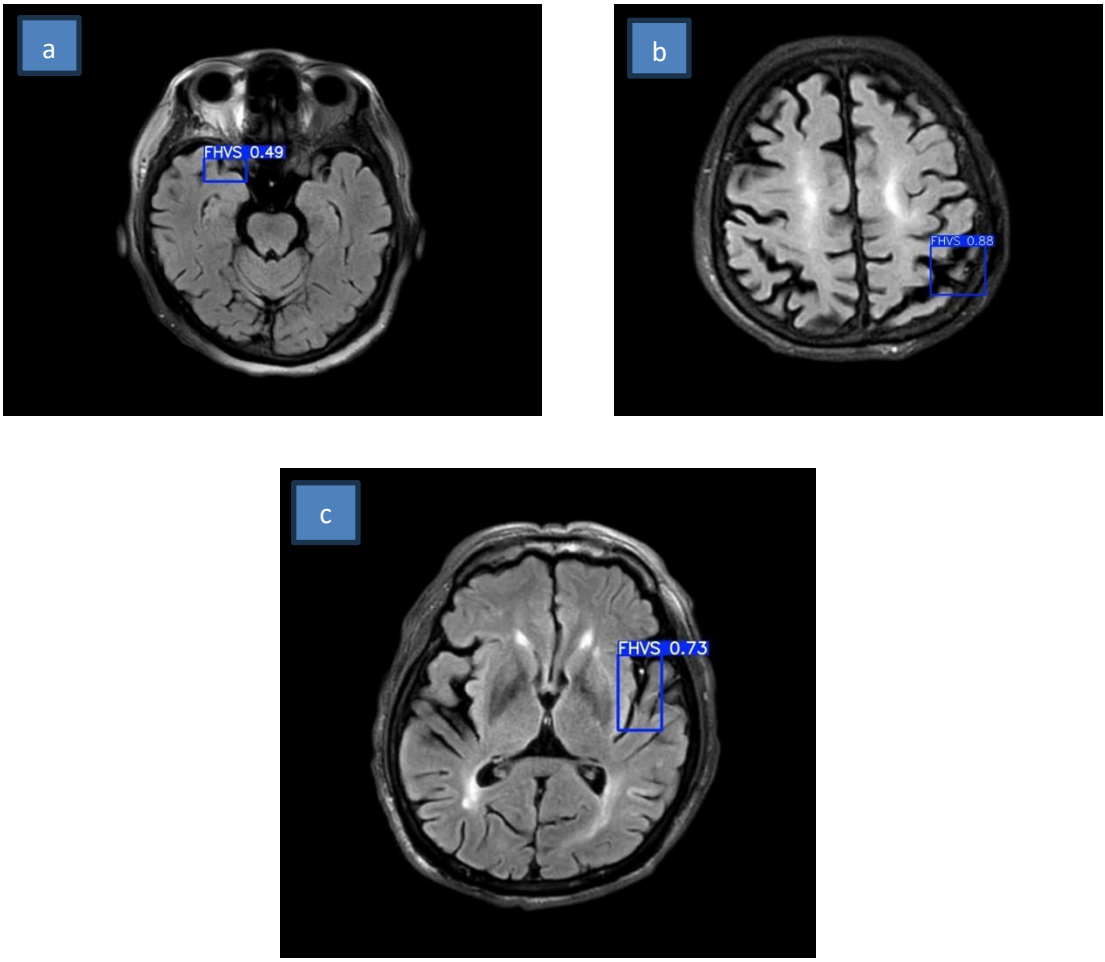

Figure 1 : LVO Detection using FHVS as a visual cue, showing bounding box outputs across three axial slice levels: (a) Low confidence (0.49) at the posterior fossa/basilar region demonstrating the challenge of subtle, elongated HVS (b) High confidence (0.88) at the parietal/MCA territory, clear hyperintense vessel captured precisely (c) Moderate confidence (0.73) at the thalamo-capsular region, elongated HVS morphology partially captured

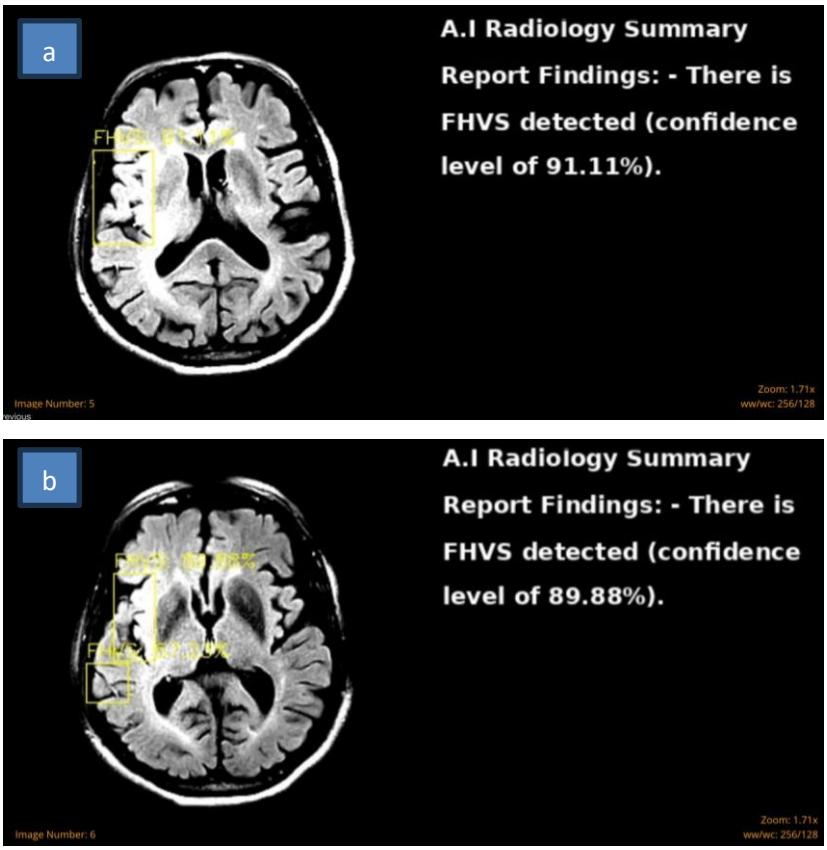

Figure 2 : Panels (a) and (b) demonstrate the full clinical pipeline, bounding box localisation with confidence scores alongside the auto-generated "A.I Radiology Summary Report Findings" text output.

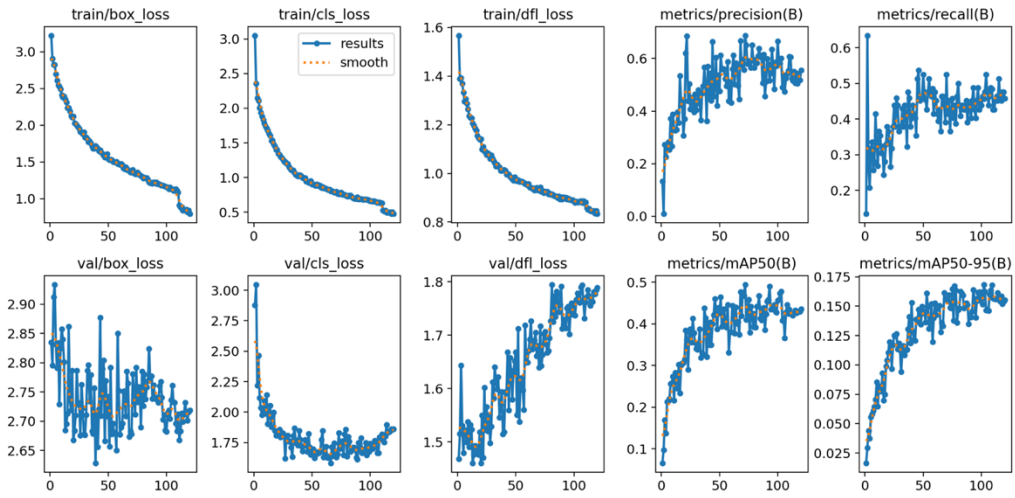

Figure 3 : Panels (a–f) illustrate the training dynamics of the YOLO-based FHVS detector, including box, classification, and DFL losses for the training and validation sets, together with the corresponding precision, recall, and mAP trajectories over 120 epochs.

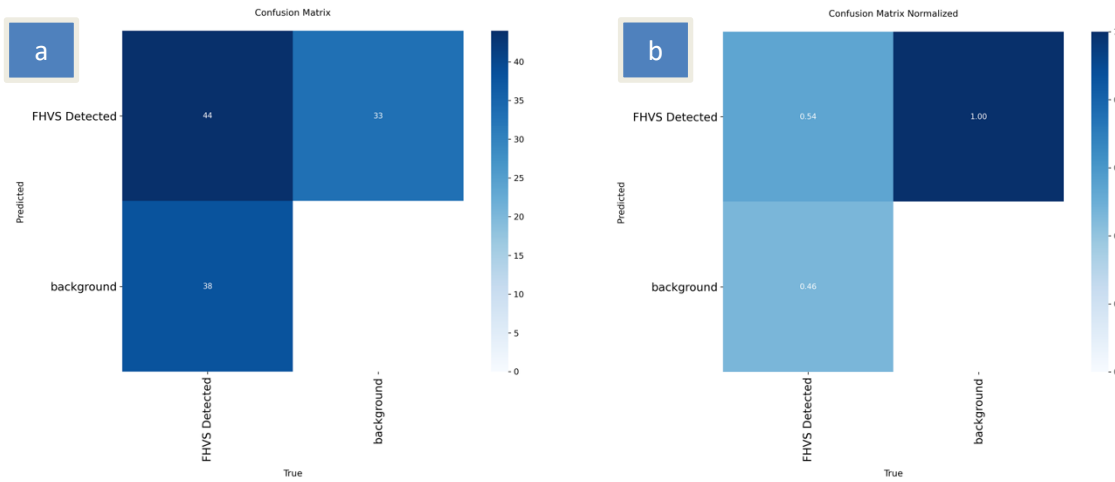

Figure 4. Panels (a) and (b) show the confusion matrices for FHVS detection on the held-out test set, with (a) displaying the absolute counts of true/false positives and negatives and (b) the corresponding normalized proportions for each class.