SCIENTIFIC OASIS

Journal of Soft Computing and Decision Analytics

Journal homepage: www.jscda-journal.org
ISSN: 3009-3481

Advancing Ischemic Stroke Detection Through an In-depth Evaluation of YOLOv10 Models on Diffusion-Weighted Imaging Data

Bilal Bayram[1], Suat Ince[2], Serhat Kilicarslan[3], Erkan Veziroglu[4], Omer Celik[5], Ishak Pacal[6,7,*]

[1] Department of Neurology, University of Health Sciences, Van Education and Research Hospital, 65000 Van, Turkey
[2] Department of Radiology, University of Health Sciences, Van Education and Research Hospital, 65000 Van, Turkey
[3] Department of Software Engineering, Faculty of Engineering, Bandirma Onyedi Eylül University, Balıkesir 10010, Türkiye
[4] Department of Computer Technologies, Simav Vocational School, Kutahya Dumlupinar University, 43500, Simav, Kutahya, Turkey
[5] Department of Healthcare Management, Igdir University, Igdir University, 76002 Igdir, Turkey
[6] Department of Computer Engineering, Faculty of Engineering, Igdir University, 76002 Igdir, Turkey
[7] Department of Electronics and Information Technologies, Faculty of Architecture and Engineering, Nakhchivan State University, Nakhchivan, AZ 7012, Azerbaijan

ARTICLE INFO

Article history:
Received 17 November 2025
Received in revised form 30 December 2025
Accepted 24 January 2026
Available online 30 January 2026

Keywords:
Ischemic stroke detection, Deep learning, YOLOv10, Medical imaging analysis

ABSTRACT

Rapid and accurate detection of acute ischemic stroke (AIS) is critical for minimizing irreversible tissue damage and improving patient outcomes. While deep learning has revolutionized medical imaging, a gap remains in balancing high diagnostic precision with the computational efficiency required for real-time clinical deployment. This study presents a systematic evaluation of the YOLOv10 architecture, utilizing its advanced dual-label assignment and non-maximum suppression (NMS)-free inference, to address this challenge using Diffusion-Weighted Imaging (DWI) data. Utilizing the multi-center ISLES 2022 dataset, five YOLOv10 variants (Nano to Extra-Large) were trained on 1,652 preprocessed DWI images using a rigorous experimental design that included transfer learning and extensive data augmentation to handle anatomical asymmetries and data imbalance. The results indicate that the YOLOv10l (Large) variant emerged as the optimal model for high-stakes diagnostics, achieving a superior precision of 0.933, recall of 0.783, and mAP50 of 0.887, significantly outperforming lighter variants in complex lesion delineation. Conversely, the YOLOv10n (Nano) variant demonstrated ultra-fast processing speeds (0.8 ms), highlighting its potential for resource-constrained environments such as mobile stroke units. These findings confirm that YOLOv10 offers a versatile framework for stroke detection, where YOLOv10l provides the robust accuracy necessary for hospital-grade decision support, and YOLOv10n offers a viable solution for real-time triage, underscoring the necessity of model selection based on specific clinical constraints.

[*] *Corresponding author.*
E-mail address: ishakpacal@igdir.edu.tr

https://doi.org/10.31181/jscda41202680

1. Introduction

Stroke remains one of the foremost causes of mortality and disability worldwide [1]. Within this context, ischemic stroke, typically precipitated by vascular occlusions, constitutes the majority of stroke cases [2]. Diminished cerebral blood flow can result in permanent tissue damage and neurological deficits, underscoring the paramount importance of swift detection and intervention [3]. Identifying stroke at the earliest stage and accurately determining lesion location and extent significantly heightens the likelihood of favorable patient outcomes [4]. Contemporary neuroimaging techniques, capable of detecting even subtle alterations in brain tissue, play a decisive role in stroke diagnosis [5]. Magnetic resonance imaging (MRI) excels due to its high contrast and resolution, mainly through specialized approaches such as diffusion-weighted imaging and apparent diffusion coefficient mapping [6]. Meanwhile, computed tomography (CT), complemented by perfusion-derived metrics, provides additional information regarding lesion size and perfusion abnormalities [7]. The combined use of these multi-modal methods affords a more comprehensive clinical perspective than any single modality alone [7].

Despite technological advancements, manual lesion delineation remains resource-intensive and vulnerable to inter-observer variability [8]. Common challenges in clinical practices such as pronounced brain asymmetries, heterogeneous lesion sizes, and limited data availability further exacerbate this issue [4,9]. Moreover, ischemic stroke often presents initially through minute yet highly consequential changes, elevating the risk of human oversight [10]. These factors underscore the need for automated, efficient, and robust diagnostic tools to alleviate clinicians' workloads [8,10].

In recent years, deep-learning methodologies have facilitated significant progress in medical image analysis [8,10,11]. By exploiting neural network architectures' depth, raw pixel data can be transformed into higher-level, semantically meaningful features [12–16]. This capability has enabled efficient and precise solutions for tasks such as anatomical structure segmentation and pathological lesion detection, substantially reducing the need for human intervention and expediting clinical workflows [17–19]. For instance, encoder-decoder networks like U-Net leverage skip connections to preserve multi-scale information, thereby retaining critical details about lesion boundaries and morphology [20].

Parallel to these developments, the You Only Look Once (YOLO) family has emerged as a leading approach for real-time object detection in medical imaging, including the identification of stroke lesions [21,22]. The most recent version, YOLOv10, incorporates deeper layers, refined optimization strategies, and increased computational efficiency, rendering it suitable for rapid detection scenarios [21]. Precise identification of small or irregularly shaped stroke lesions necessitates careful management of data imbalance and clinical variability [23]. When YOLO-based strategies are paired with meticulous hyperparameter tuning, unbiased data preprocessing, and systematic model improvements, they hold considerable promise as a fast and accurate option for clinical decision support [24].

Numerous studies in the field of medical image analysis have reported that deep learning–based methods significantly surpass traditional image processing and machine learning techniques [25]. We examine several representative methods, outlining their underlying principles, current limitations, and roles in the context of stroke lesion analysis. Chen et al. developed TE-YOLOv5, an advanced method for delineating stroke lesion edges in DWI images based on YOLOv5. The approach utilizes an Aggregate Pool (AP) module to refine high-level lesion features and a Reverse Attention (RA) module to map edge details accurately. TE-YOLOv5 demonstrated superior performance, achieving a precision of 81.5%, a recall of 75.8%, and an mAP@0.5 of 80.7%. Ablation experiments highlighted the critical contributions of the RA and AP modules, with performance dropping to 78.1% mAP@0.5 when both were excluded [26]. Elhanashi et al. [24] introduced a novel approach leveraging YOLOv8

models to analyze extensive datasets featuring images of individuals with and without stroke, focusing on facial paralysis patterns. The model is trained to detect complex features associated with strokes, significantly improving diagnostic precision. To prioritize privacy while retaining performance, the methodology incorporates federated learning, a decentralized framework that allows training on distributed data without exposing sensitive patient information. Utilizing NVIDIA's cutting-edge GPU technology, the proposed system enables real-time analysis, potentially revolutionizing stroke detection and advancing the quality of care in neurology. Chen et al. [23] developed MSA-YOLOv5, a Multi-Scale Attention-enhanced version of YOLOv5, designed to improve the detection of small and low-resolution lesions. The model integrates a Multi-Scale Swin Transformer Prediction Head (MS-STPH) for better accuracy and computational efficiency alongside a Second-Order Channel Attention (SOCA) module to enhance feature discrimination. It achieved impressive results, with a 79.0% mAP0.5 on an AIS dataset, outperforming single-stage models and rivaling two-stage models with fewer parameters, and an 80.0% mAP0.5 on the ISLES 2022 dataset. Visualizations showed its ability to focus precisely on small lesion areas [23]. Zhang et al. developed MDANet, a novel network for stroke segmentation that leverages multi-modal data.

The network integrates a difference-aware module, which highlights potential lesion regions by comparing features from different modalities, and a graph convolution fusion block for combining these features globally and locally. A similarity loss function was introduced to enhance its accuracy, further improving the detection of lesion areas. Tested on ISLES 2018 and 2022 datasets, MDANet achieved impressive dice scores of 58.34 and 70.44, surpassing many state-of-the-art methods [27]. Wu et al. proposed FRPNet, an advanced model for segmenting stroke lesions featuring a symmetric encoder-decoder design. It integrates the Twin Attention Gate (TAG) module to extract global and local features and the Multi-Dimension Attention Pooling (MAP) module to reduce feature loss during encoding. TAG employs a dual-directional self-attention mechanism, while MAP focuses on retaining critical details through multidimensional pooling. Experimental results reveal FRPNet's superior performance, achieving higher DSC and lower HD metrics compared to existing methods on comprehensive ischemic stroke datasets [28].

Li et al. introduced a novel approach to effectively combine information from CTP maps for precise segmentation of acute ischemic stroke (AIS) lesions. Their method, called WMHCA-Net, features a U-shaped design with multiple paths for encoding and decoding. To enhance integration, windowed cross-attention mechanisms are applied during encoding, while a specialized module improves information recovery during decoding. The addition of a segmentation optimization module ensures refined results, achieving highly accurate AIS lesion identification to support clinical decision-making [29]. Kaya and Onal introduce innovative automated techniques for classifying and segmenting stroke-related lesions in non-contrast brain CT images. Their method focuses on accurately identifying both hemorrhagic and ischemic lesions, utilizing the U-Net model for precise segmentation. Testing on real-world data demonstrates a classification precision of 95.06%. Additionally, segmentation achieves IoU scores of 92.01% for hemorrhagic lesions and 82.22% for ischemic lesions [7].Chen et al. introduced an advanced computer-aided system designed to automate the segmentation and classification of MRI scans, identifying carotid plaque and assessing stroke risk.

Utilizing pre-trained models, they refined and optimized them to address the specific requirements of this task. The study demonstrated that the YOLO V3 model delivered the highest performance with an accuracy of 94.81%, followed by RCNN at 92.53% and MobileNet at 90.23%. This innovative framework aims to minimize diagnostic inaccuracies stemming from suboptimal image quality and subjective interpretation [30, 31]. Uzun and Okuyar evaluated YOLOv7, YOLOv8, and YOLOv9 models for detecting ischemic and hemorrhagic strokes in brain CT images. They

compared these YOLO models' segmentation performance with U-Net variants and Mask-RCNN to develop an AI-assisted system for physicians. Using 6951 anonymized CT slices (2019–2020) provided by the Turkish Ministry of Health, the YOLOv9-Seg demonstrated superior performance in stroke detection and segmentation [32]. Sinha et al. introduced EnigmaNet, a novel deep-learning model designed for ischemic stroke lesion segmentation in FLAIR and DWI imaging modalities. The architecture integrates innovative Genesis-k blocks and a dual attention mechanism, paired with a customized loss function, Weighted Focal-Tversky-Dice (wFTD) Loss, to boost performance. Tested on the ISLES-2015 dataset, EnigmaNet achieved impressive results, including Dice scores of 0.8965 and 0.8423 for FLAIR and DWI images, respectively. The model demonstrated substantial advancements over existing methods, surpassing U-Net-sharp, FCN-8, and Attention U-Net by 32%, 41%, and 10% in Dice score improvements [33]. Li et al. proposed SrSNet, a two-stage framework designed to identify abnormalities in brain images by leveraging the symmetry between brain hemispheres. The initial stage divides brain scans into spatially informed patches, processed by a Triplet Multi-scale Symmetric Transformer (TMSFormer) to generate preliminary segmentation maps. The refinement stage enhances these outputs using a 2D ResUNet for improved accuracy. Experimental results on ischemic stroke lesion datasets showed that SrSNet outperforms existing advanced methods [34].

A critical review of the contemporary literature indicates a paradigm shift in stroke detection, driven by increasingly innovative deep learning methodologies. Recent studies have largely concentrated on overcoming persistent diagnostic challenges, such as the detection of minute, low-resolution lesions, the precise delineation of infarct boundaries, and the effective fusion of multi-modal imaging data. To enhance segmentation accuracy and classification precision, state-of-the-art approaches have integrated advanced techniques, including attention mechanisms, multi-scale feature pyramids, and domain-specific loss functions. Furthermore, the field is witnessing a pivot towards decentralized learning frameworks, which facilitate large-scale data analysis while preserving patient privacy.

Despite these advancements, a trade-off often persists between computational efficiency and diagnostic sensitivity in clinical settings. Addressing this, the present study introduces a comprehensive evaluation of the YOLOv10 architecture for automated, reliable, and rapid stroke lesion detection. The training and validation phases utilize high-resolution Diffusion-Weighted Imaging (DWI) data from the multi-center ISLES 2022 dataset [35], with ground truth verified by expert radiologists. This approach leverages the architectural efficiency of YOLOv10, specifically its dual-label assignment and NMS-free inference, to mitigate the challenges posed by anatomical asymmetries and class imbalance. By systematically analyzing the YOLOv10 framework across its full spectrum of configurations (Nano through Extra-Large), this work aims to identify the optimal model that balances performance with speed. Ultimately, this research seeks to deploy a robust diagnostic tool capable of alleviating clinician workload and expediting time-sensitive treatment planning to improve patient survival rates [36].

2. Materials and Methods

2.1 Study Design

This research was conducted as a quantitative, analytical study to systematically evaluate the performance of different YOLOv10 model variants for a specific medical imaging task. The primary objective was to identify which YOLOv10 architecture nano, small, medium, large, or extra-large provides the optimal balance between diagnostic accuracy, computational speed, and model complexity using DWI data.

A rigorous data preprocessing phase was implemented. In consultation with medical experts, lesions smaller than 5x5 pixels were excluded from the dataset. This exclusion criterion was applied to minimize the inclusion of noise artifacts and partial volume effects that often mimic micro-lesions but lack immediate clinical significance for acute intervention. Furthermore, removing these ambiguous outliers reduces label noise, ensuring that the model focuses on learning robust features of definitive ischemic events, thereby enhancing training stability and reducing false positives. The curated dataset (Total: 1,652) was then partitioned into a training set (1,321 images; 80%) and a test set (331 images; 20%). To improve model robustness and prevent overfitting, extensive data augmentation techniques were applied to the training data.

The methodology followed an experimental design where the independent variable was the specific YOLOv10 model variant being tested. Five distinct models were evaluated: YOLOv10n, YOLOv10s, YOLOv10m, YOLOv10l, and YOLOv10x. To ensure a fair and controlled comparison, all models were trained using identical hyperparameter settings and on the same high-performance computing environment. The models' performance was assessed using a comprehensive set of dependent variables, including precision, recall, mean Average Precision (mAP50), inference time, and the number of model parameters. The analysis plan involved training each model with transfer learning and subsequently evaluating its performance on the unseen test set to identify the trade-offs between detection accuracy and computational cost, thereby determining the most suitable model for potential clinical application.

2.2 YOLOv10: Real-Time End-to-End Object Detection

The "You Only Look Once" (YOLO) series has undergone several iterations, continually evolving to address many object detection tasks. YOLOv10 represents the latest milestone in this progression, introducing dual-label assignment that removes the need for non-maximum suppression (NMS) while maintaining robust performance. Early YOLO versions frequently relied on a one-to-many assignment strategy often termed Task-Aware Label Assignment (TAL) which generates a rich set of positive samples yet requires NMS in post-processing. By contrast, YOLOv10 combines one-to-many and one-to-one assignments, thus enhancing training supervision and enabling an NMS-free inference pipeline, as seen in Figure 1 [21,37].

As seen in Figure 1, this approach holds particular value for brain stroke detection, where timely and precise identification of ischemic lesions is paramount for effective clinical intervention. Like its predecessors, YOLOv10 employs separate loss components for classification, bounding box regression, and abjectness scores, often encapsulated in a single objective in (1).

$$L = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{reg} + \lambda_2 \mathcal{L}_{obj} \quad (1)$$

where $\mathcal{L}_{cls}$ denotes the classification loss, $\mathcal{L}_{reg}$ is the bounding box regression term (for instance, Smooth L1 or GIoU), and $\mathcal{L}_{obj}$ measures objectness or confidence. The weighting factors λ_1 and λ_2 control the relative importance of localization and confidence, respectively. In stroke detection, this structure allows the network to integrate better class-specific lesion cues, such as ischemic regions, with precise bounding box predictions. Conventional YOLO-based pipelines typically employ one-to-many assignment, matching each ground truth lesion with multiple candidate predictions. Although this redundancy yields more supervision, NMS must reconcile overlapping bounding boxes during inference. In contrast, one-to-one assignment pairs each ground truth with exactly one prediction, forgoing NMS at inference but offering fewer positive samples. YOLOv10 reconciles these two paradigms by retaining a one-to-many head for denser supervision and a one-to-one head for streamlined, NMS-free inference particularly important for stroke detection, where real-time localization of small, irregular lesions is essential.

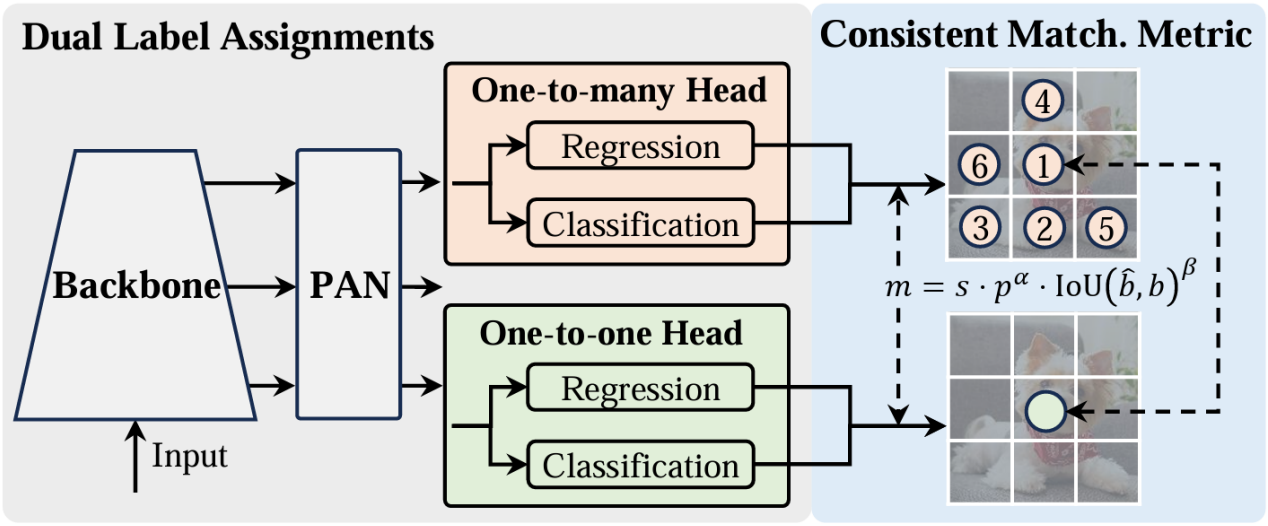

Fig. 1. The YOLOv10 architecture. The model uses a backbone and a Path Aggregation Network (PAN) to process features. Its core innovation is a Dual Label Assignment strategy, which employs two parallel heads: a one-to-many head for enriched training supervision and a one-to-one head for efficient, NMS-free inference. Both heads are optimized using a Consistent Matching Metric, ensuring cohesive training.

The dual-label assignment scheme is at the heart of YOLOv10's design, which simultaneously trains two detection heads. Both heads share the same backbone and neck but differ in assigning ground truth bounding boxes to predictions. Critically, they utilize a consistent matching metric given by (2).

$$m(\alpha, \beta) = s \cdot p^{\alpha} \cdot \mathrm{IoU}(\hat{b}, b)^{\beta} \quad (2)$$

where p is the classification probability, $\hat{b}$ represents the predicted bounding box, b represents the ground truth box, and $\mathrm{IoU}(\hat{b}, b)$ is the Intersection-over-Union between them and depicted in (3).

$$\mathrm{IoU}(\hat{b}, \mathrm{b}) = \frac{|\hat{\mathrm{b}} \cap b|}{|\hat{\mathrm{b}} \cup b|} \quad (3)$$

Here, s indicates whether the anchor point lies within the ground truth lesion region. The exponents α and β regulate the contributions of the classification and localization tasks, respectively. By synchronizing these hyperparameters, YOLOv10 ensures that the top-ranked matches for the one-to-many head often coincide with those of the one-to-one head [38]. During inference for ischemic stroke detection, the network relies solely on the one-to-one branch, eliminating overlapping predictions and obviating NMS. By uniting dense supervision from the one-to-many head with an NMS-free one-to-one output, YOLOv10 offers a streamlined pipeline with reduced computational overhead, which is particularly advantageous in clinical settings. The method excels at localizing small, irregular ischemic lesions common in stroke patients, facilitating more accurate and faster diagnosis. Consequently, YOLOv10 for stroke detection can enhance clinical workflows, improve patient triage, and expedite potentially life-saving interventions.

2.3 Variants of YOLOv10 for Stroke Lesion Detection

YOLOv10 offers several architectural variants: Nano, Small, Medium, and Large, each providing a unique balance between computational efficiency and detection accuracy, making them suitable for diverse clinical scenarios involving brain stroke detection. Nano, an ultra-lightweight design, is

optimal for resource-constrained or portable systems, enabling real-time screening of ischemic lesions at the bedside. Small expands the network depth and channel capacity slightly, delivering enhanced performance while retaining a low computational footprint for moderate-use clinical environments. Medium further augments representational power, capturing subtle or borderline stroke lesions without excessively burdening typical hospital-grade GPU setups. Large, at the upper end of this spectrum, features an extensive backbone and head configuration that excels in delineating intricate lesion characteristics from high-resolution scans. Although it demands considerable computational resources, it offers the most comprehensive detection capabilities, making it ideal for specialized research facilities or complex multi-modal imaging protocols. By selecting an appropriate variant, clinicians and researchers can tailor YOLOv10's performance to their specific hardware constraints and precision requirements, ensuring robust and efficient stroke lesion identification.

2.4 Dataset: ISLES 2022

The ISLES 2022 dataset represents a meticulously curated and expertly annotated collection of MRI scans designed to identify and segment ischemic stroke lesions in their acute to subacute stages [35]. Comprising 400 MRI scans sourced from multiple vendors, the dataset showcases substantial diversity in lesion characteristics, including size, quantity, and spatial distribution. It is divided into 250 publicly available training cases and 150 test cases reserved exclusively for validation purposes. Organized according to the Brain Imaging Data Structure (BIDS) standard, each case includes multi-modal MRI data, such as Diffusion-Weighted Imaging (DWI), Apparent Diffusion Coefficient (ADC), Fluid-Attenuated Inversion Recovery (FLAIR) images, along with ground truth labels, visual snapshots, and accompanying JSON metadata when available.

As the centerpiece of the ISLES 2022 challenge, this dataset has been instrumental in evaluating and benchmarking automated methods for stroke lesion segmentation. While it provides diverse imaging modalities, many studies have prioritized DWI due to its unparalleled sensitivity in detecting ischemic stroke lesions, cementing its role as the clinical gold standard. The FLAIR modality was excluded from our analysis due to image alignment and accuracy challenges, and ADC data were similarly omitted to maintain analytical consistency. These decisions underscore the nuanced considerations required when utilizing multi-modal imaging datasets to develop reliable and effective segmentation models. Figure 2 presents a selection of random DWI images from the ISLES 2022 dataset, showcasing single and multiple lesions and their corresponding masks.

3. Experiments and Results

3.1 Implementation

All computational experiments were executed on a dedicated high-performance workstation running Ubuntu 24.04 LTS. The hardware infrastructure was configured with an Intel Core i9-14900K processor, 64 GB of DDR5 RAM, and an NVIDIA GeForce RTX 4090 GPU (24 GB VRAM) to handle intensive deep learning workloads. The software environment was standardized on Python 3.10 and the PyTorch 2.2 framework, leveraging NVIDIA CUDA 12.1 for GPU-accelerated training and inference. The YOLOv10 variants were implemented using the Ultralytics framework. To enhance feature extraction capabilities and accelerate convergence, transfer learning was employed by initializing the models with weights pre-trained on the MS COCO dataset3. To ensure the reproducibility of the results, a strict hyperparameter protocol was followed. The models were optimized using Stochastic Gradient Descent (SGD) with a momentum of 0.937 and a weight decay of 0.00054. A cosine annealing learning rate scheduler was utilized, commencing with an initial learning rate (lr_0) of 0.01 and decaying to a final learning rate factor (lrf) of 0.015. Training was

conducted with a batch size of 16 for a total of 500 epochs to guarantee model stability and convergence.

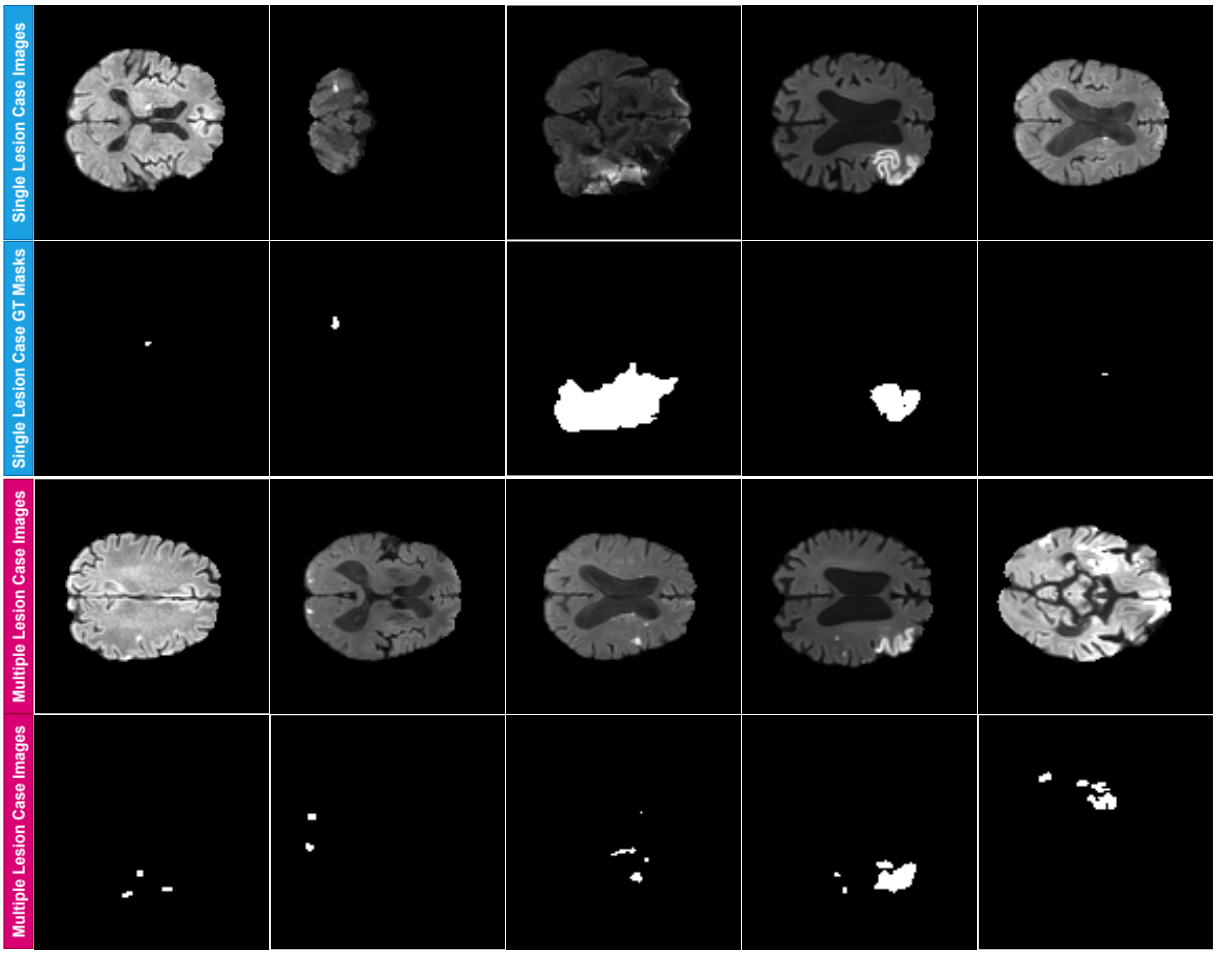

Fig. 2. Examples from the ISLES 2022 dataset illustrate single and multiple lesions on DWI images alongside their corresponding ground truth (GT) masks.

3.2 Metrics

Various metrics are employed to evaluate the performance of YOLO-based models in brain stroke detection to ensure a comprehensive assessment of the model's capabilities. These metrics measure detection accuracy, localization precision, and overall robustness. This section provides an overview of the key metrics used to analyze the effectiveness of the proposed YOLO models for medical imaging tasks. Precision measures the proportion of correctly identified positive cases (True Positives, TP) out of all predicted positives, minimizing false alarms. Recall (Sensitivity) evaluates the proportion of actual positives correctly detected, reducing the risk of missed strokes. F1 Score provides a harmonic mean of precision and recall, balancing the trade-off between false positives and false negatives. Mean Average Precision (mAP) calculates the average precision across multiple IoU thresholds, evaluating the overall detection performance. mAP@0.5: This metric evaluates the average precision at a specific IoU threshold 0.5. It is commonly used as a standard benchmark for assessing object detection models. An IoU of 0.5 indicates that the overlap between the predicted and actual bounding boxes must be at least 50% for the detection to be deemed accurate. Intersection-over-Union (IoU) measures the overlap between predicted and ground truth bounding boxes, ensuring precise localization. Inference Time assesses computational efficiency, reflecting the time required for the model to process an image. A low inference time ensures real-time applicability,

which is especially important for emergency stroke detection systems. These metrics collectively provide a comprehensive evaluation of the YOLO-based stroke detection system, addressing accuracy and practicality as depicted from (4) - (7).

$$Precision = \frac{TP}{TP + FP} \tag{4}$$

$$Recall = \frac{TP}{TP + FN} \tag{5}$$

$$F1\ Score = 2\frac{Precision \cdot Recall}{Precision + Recall} \tag{6}$$

$$mAP@0.5: = \frac{1}{N}\sum_{i=1}^{N} AP_i \tag{7}$$

3.3 Data Processing and Data Augmentation

The ISLES 2022 dataset comprises ADC, DWI, and FLAIR modalities, with DWI being widely recognized as the gold standard for ischemic stroke detection due to its high informational value. Some lesions in the dataset are as small as 5x5 pixels or smaller, which can negatively impact model performance, particularly during training. To address this issue, lesions of these dimensions were excluded through image processing techniques in consultation with radiology and neurology experts. This refinement resulted in a total of 1652 usable images. Of these, 80% (1321 images) were allocated for training and 20% (331 images) for testing, following a splitting methodology similar to that used in COCO detection datasets. To ensure uniformity and reduce unnecessary black borders, all images were resized to 640 × 640 pixels using adaptive scaling.

Various data augmentation techniques were applied to mitigate the risk of overfitting and enhance the diversity of the dataset. These methods included adjustments in color attributes (hue, saturation, and brightness), horizontal flipping, mosaic augmentation, random erasing, automatic augmentations (randaugment), and other geometric transformations. These augmentation strategies significantly improved the model's robustness and generalizability, enabling it to adapt effectively to different data variations and deliver highly accurate results. The augmented dataset improved the model's performance while maintaining its reliability in detecting stroke lesions.

3.4 Results

This study utilized the DWI modality from the ISLES 2022 dataset to evaluate the effectiveness of the state-of-the-art YOLOv10 algorithm in identifying ischemic stroke lesions. The analysis encompassed multiple YOLOv10 variants, including nano (YOLOv10n), small (YOLOv10s), medium (YOLOv10m), large (YOLOv10l), and extra-large (YOLOv10x) models. These variants were rigorously assessed and compared using key performance metrics such as precision, recall, F1 score, mean Average Precision (mAP), computational efficiency, inference time, and GPU resource usage. Table 1 provides a detailed overview of the experimental outcomes for each YOLOv10 model variant, highlighting their respective advantages and limitations. The results underscore the nuanced trade-offs between computational requirements and predictive performance across the different model configurations.

The results in Table 1 provide a comprehensive evaluation of the YOLOv10 algorithm's variants nano (YOLOv10n), small (YOLOv10s), medium (YOLOv10m), large (YOLOv10l), and extra-large (YOLOv10x) in detecting ischemic stroke lesions from Diffusion Weighted Imaging (DWI) in the ISLES 2022 dataset. Key performance indicators such as precision, recall, mean Average Precision (mAP50),

inference time, and parameter count were analyzed to elucidate the trade-offs between detection accuracy and computational efficiency, emphasizing their relevance for cerebral vascular occlusion detection.

Table 1
Experimental results of YOLOv10 variants on the ISLES 2022 dataset (RTX 4090, batch size = 16; ms: milliseconds, M: million)

Model	Precision	Recall	mAP50	Inference (ms)	Params(M)
YOLOv10n	0.931	0.764	0.865	0.8	2.69
YOLOv10s	0.917	0.776	0.887	1.4	8.04
YOLOv10m	0.912	0.783	0.873	2.9	16.45
YOLOv10l	0.933	0.783	0.887	4.5	25.72
YOLOv10x	0.918	0.766	0.881	6.7	31.59

Precision and recall, essential metrics for evaluating detection models, indicate a model's ability to minimize false positives and false negatives. Among the variants, YOLOv10l demonstrated the highest precision at 0.933, reflecting its effectiveness in avoiding false positives, while YOLOv10l and YOLOv10m exhibited the highest recall values at 0.783, indicating strong detection rates for true lesions. In contrast, YOLOv10n and YOLOv10x exhibited slightly lower recall values (0.764 and 0.766, respectively), highlighting potential limitations in identifying smaller or less distinct lesions. These findings underline the necessity of selecting a model that aligns with the specific clinical requirement of prioritizing either high recall or precision.

Detection accuracy, measured through mAP50, offers a comprehensive view of overall model performance. The highest mAP50 values of 0.887 were achieved by YOLOv10s and YOLOv10l, demonstrating balanced accuracy across diverse lesion characteristics. YOLOv10m closely followed with an mAP50 of 0.873, achieving robust accuracy with moderate resource demands. While YOLOv10n (0.865) and YOLOv10x (0.881) scored slightly lower, they maintained competitive performance, particularly for scenarios demanding either low computational cost or enhanced feature representation.

Inference time and computational efficiency are critical factors for real-time diagnostic applications. YOLOv10n emerged as the fastest model with an inference time of 0.8 ms, making it an excellent candidate for real-time scenarios such as portable diagnostic devices or emergency settings. On the other hand, YOLOv10x took the most time (6.7 ms) due to its high computational complexity and parameter count, rendering it more suitable for offline or research-focused tasks. YOLOv10m achieved a practical balance with an inference time of 2.9 ms, efficiently supporting real-time and offline diagnostic workflows.

The parameter count of each model reflects its complexity and resource demands. YOLOv10n, with the smallest parameter count of 2.69M, demonstrated exceptional efficiency for deployment in resource-constrained environments. In contrast, YOLOv10x offered superior feature representation with 31.59M parameters but demanded significant computational resources. YOLOv10m, with a parameter count of 16.45M, represented an optimal compromise between complexity and performance, making it adaptable to various clinical environments. These findings demonstrate the efficacy of YOLOv10 models in detecting ischemic stroke lesions from DWI images, highlighting their adaptability to a range of clinical and computational scenarios. The results emphasize the potential of YOLOv10 to enhance stroke detection workflows, enabling improved diagnostic accuracy and timely clinical decision-making. Figure 3 presents the performance graphs of the YOLOv10l model, identified as the best-performing model.

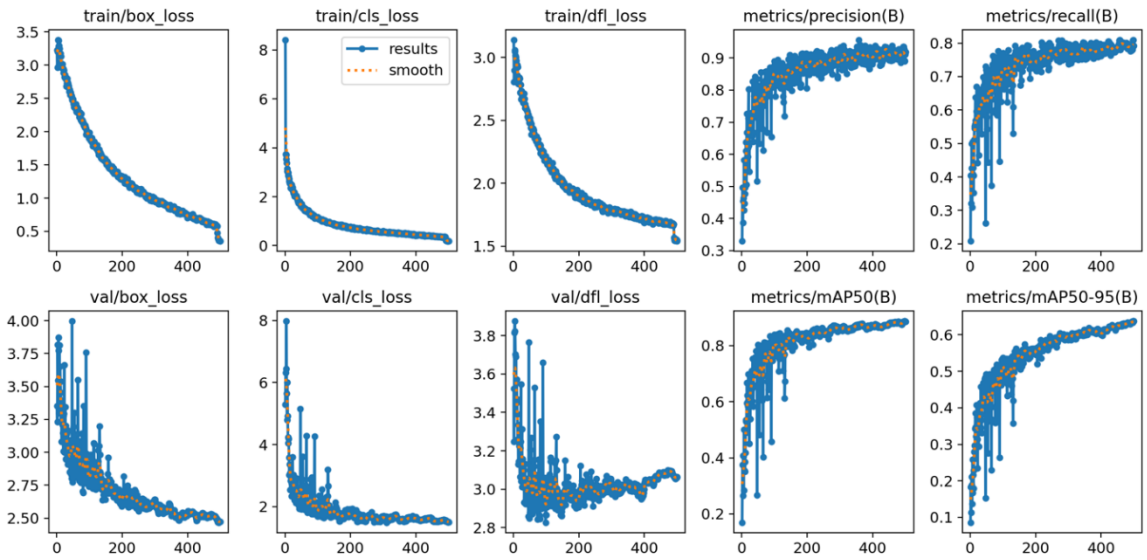

Fig. 3. Training and validation performance of the YOLOv10l model over 500 epochs. The top row shows training progress, while the bottom row shows validation performance on a separate dataset. The plots from left to right display: (1) bounding box loss (box_loss), (2) classification loss (cls_loss), (3) Distribution Focal Loss (dfl_loss), (4) overall precision, and (5) overall recall. The final two validation plots show mean Average Precision at IoU thresholds of 0.50 (mAP50) and 0.50-0.95 (mAP50-95). The x-axis on all plots represents epochs, and the y-axis represents the metric value.

As seen in Figure 3, the performance of the YOLOv10l model, as illustrated by the training and validation graphs over 500 epochs, indicates consistent advancements across all major metrics. Training losses, including the box, classification, and distribution focal loss, show a steady reduction throughout the epochs. By the end of the training, the box loss stabilizes around 0.5, and the classification loss approaches zero, reflecting the model's enhanced capability to detect and classify ischemic stroke lesions accurately. Although validation losses follow a similar declining trend, stabilizing at approximately 2.5 for box loss and 3 for classification loss, they remain higher than their training counterparts. This suggests the model effectively generalizes but may face challenges with more ambiguous or smaller lesions in the validation dataset. The disparity between training and validation losses highlights a potential risk of overfitting, which could be mitigated through additional regularization strategies or extended data augmentation techniques.

The performance metrics, evaluated over 500 epochs, further affirm the robustness of YOLOv10l. The model achieves a precision of 0.933, underscoring its ability to minimize false positives, while recall reaches 0.783, signifying its efficiency in detecting true lesions. The mAP50 score of 0.89 demonstrates the model's strong overall detection accuracy, while the mAP50-95 score, peaking at 0.6, suggests room for improvement in detecting smaller or more complex lesions. These results showcase YOLOv10l's effectiveness in stroke lesion detection, particularly for clear and well-defined cases. However, to fully optimize its clinical utility, addressing the limitations in generalization and performance in edge cases is crucial, ensuring the model provides consistent and comprehensive diagnostic support across various lesion complexities. Figure 4 illustrates the performance of the YOLOv10l model, showcasing DWI images with actual stroke lesion locations and their predicted counterparts, along with the associated bounding boxes. The label "0" represents the ground truth bounding boxes for the lesions, while the predicted bounding boxes indicate the model's detected stroke lesions.

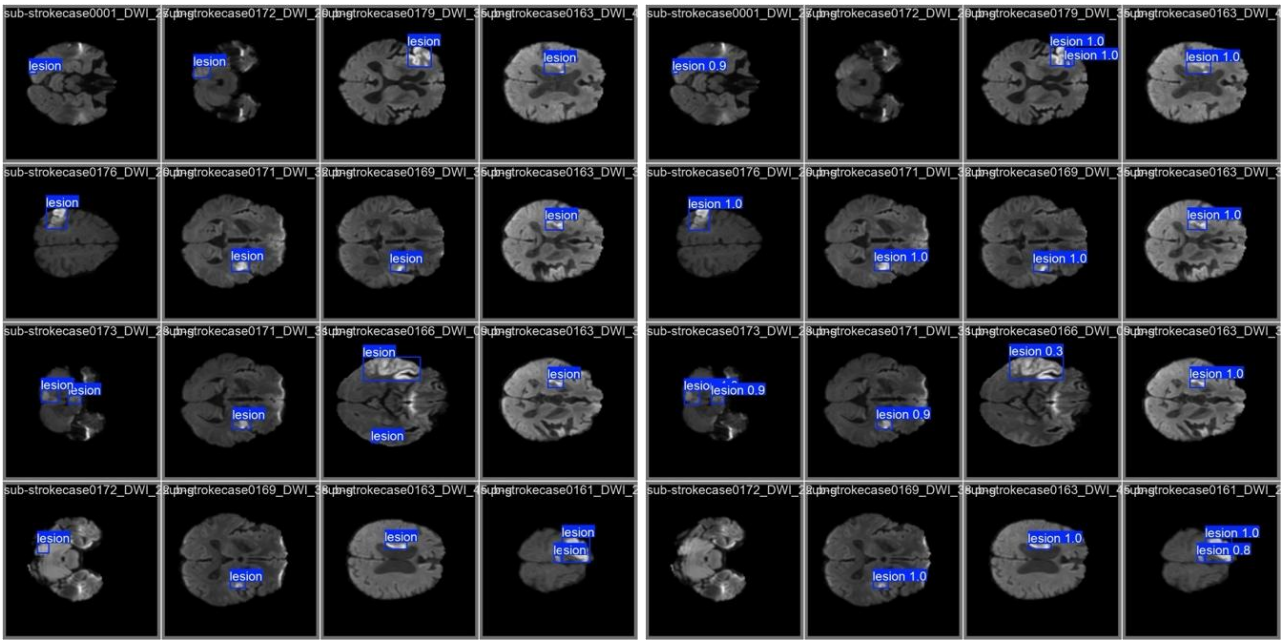

Fig. 4. Qualitative results of the YOLOv10l model on test images from the ISLES 2022 dataset. Each pair of images shows a comparison: (Left Panel) The ground truth (GT) bounding boxes manually drawn by experts are shown overlaid on the DWI scans. (Right Panel) The corresponding bounding boxes predicted by the YOLOv10l model are displayed. Each prediction includes a confidence score, indicating the model's certainty.

As shown in Figure 4, the provided images compare the ground truth and predicted stroke lesion masks. The images on the left display the actual lesion masks, labeled as "lesion," representing the true locations of stroke lesions in the DWI images. In contrast, the images on the right illustrate the predictions made by the YOLOv10l model, which are labeled as "lesion" and include bounding boxes with associated confidence scores. These visualizations highlight the model's ability to detect and localize ischemic stroke lesions accurately.

While the predicted bounding boxes align well with the ground truth masks, some discrepancies are observed, such as differences in confidence scores and minor misalignments. Additionally, in certain cases, the model fails to detect lesions, while in others, it produces false positives or incorrect predictions. These findings emphasize the effectiveness of the YOLOv10l model in stroke lesion detection but also point to areas for improvement, particularly in handling smaller or less distinct lesions and reducing false positive rates.

3.5 Discussion

This study provides a comprehensive benchmarking of YOLOv10-based models for ischemic stroke lesion detection, leveraging the high-quality ISLES 2022 dataset. The results demonstrate that the YOLOv10l variant achieves state-of-the-art performance with a precision of 0.933 and an mAP50 of 0.887. When compared to recent advancements in the field, YOLOv10l outperforms the attention-enhanced TE-YOLOv5, which reported a precision of 81.5% and mAP@0.5 of 80.7%, as well as the MSA-YOLOv5 model, which achieved an mAP0.5 of 80.0% on similar datasets. This superiority suggests that YOLOv10's architectural innovations, specifically the elimination of NMS and the use of dual-label assignment, provide a more effective mechanism for feature localization in medical imaging than traditional attention modules alone.

A critical finding of this research is the quantifiable trade-off between diagnostic accuracy and computational efficiency. While YOLOv10l is ideal for high-precision tasks within hospital

infrastructures, the YOLOv10n variant, with an inference time of just 0.8 ms, presents a groundbreaking opportunity for pre-hospital diagnostics. Such lightweight models could be integrated into mobile stroke units or portable MRI scanners, facilitating rapid "stroke/no-stroke" triage before the patient reaches the hospital.

However, the integration of these models into real-world clinical workflows requires addressing several translational hurdles. Deployment into Picture Archiving and Communication Systems (PACS) necessitates not only high accuracy but also regulatory compliance and interoperability standards. Furthermore, the robustness of these models across heterogeneous hardware, varying magnetic field strengths and vendor-specific imaging protocols, remains a pivotal factor for widespread adoption.

Despite the strong overall performance, the detection of small, punctate lesions remains a persistent challenge, a limitation shared by many object detection frameworks in radiology. While our preprocessing strategy excluded extremely small artifacts to stabilize training, the lower recall in lighter variants highlights the need for specialized small-object detection heads or multi-scale attention mechanisms in future architectures to ensure no subtle ischemic event is overlooked.

3.6 Study Limitations and Future Directions

The findings of this study should be interpreted within the context of certain limitations, which serve as a roadmap for future research. First, the evaluation was conducted exclusively on the ISLES 2022 dataset. Although this is a multi-center and diverse dataset, the lack of external validation on independent cohorts from different institutions limits the generalizability of the findings. Future studies must prioritize cross-center validation to assess model robustness against the variability inherent in real-world clinical data, such as differences in scanner vendors and acquisition protocols.

Second, the study utilized a unimodal approach, relying solely on Diffusion-Weighted Imaging. While DWI is the gold standard for acute ischemia, clinical diagnosis often benefits from a multi-modal assessment involving additional sequences like ADC maps and FLAIR images to distinguish between acute and subacute infarcts and to estimate tissue viability. Future research should focus on developing multi-modal fusion architectures that can synthesize information from concurrent MRI sequences, potentially improving the detection of subtle lesions and reducing false positives.

Third, the detection of small and fragmented lesions remains a challenge for current YOLO architectures. Future work should explore the integration of dedicated small-object detection layers, contextual attention mechanisms, or hybrid Transformer-CNN architectures to enhance sensitivity to minute ischemic changes.

Finally, to bridge the gap between AI research and clinical practice, future investigations should incorporate Federated Learning frameworks. This approach would enable the training of robust models across multiple institutions without compromising patient privacy, a critical requirement for medical AI. Additionally, incorporating Explainable AI modules, such as saliency maps or uncertainty quantification, will be essential to build clinician trust and facilitate the regulatory approval necessary for deploying these tools as accredited Clinical Decision Support Systems.

4. Conclusion

This study systematically evaluated the YOLOv10 framework for the automated detection of ischemic stroke lesions, establishing a new benchmark for speed and accuracy in medical object detection. The results identify YOLOv10l as a high-performance candidate for stationary clinical settings, offering superior precision and mean average precision that rival or surpass existing attention-based models. Conversely, the YOLOv10n variant demonstrates exceptional efficiency, making it uniquely specific for real-time, edge-computing applications in emergency medicine. While the study highlights the potential of these models to transform diagnostic workflows, it also

delineates critical avenues for future development, including multi-modal data fusion and privacy-preserving federated learning. Ultimately, integrating these advanced deep learning tools into clinical practice holds the promise of accelerating stroke diagnosis, alleviating radiologist workload, and improving long-term patient outcomes.

List of abbreviations

ADC: Apparent Diffusion Coefficient
AIS: Acute Ischemic Stroke
BIDS: Brain Imaging Data Structure
CT: Computed Tomography
DFL: Distribution Focal Loss
DWI: Diffusion-Weighted Imaging
FLAIR: Fluid-Attenuated Inversion Recovery
GPU: Graphics Processing Unit
IoU: Intersection-over-Union
mAP: mean Average Precision
MRI: Magnetic Resonance Imaging
NMS: Non-Maximum Suppression
PAN: Path Aggregation Network
YOLO: You Only Look Once

Declaration of Generative AI and AI-assisted Technologies in the Writing Process

During the preparation of this work, the authors used large language models (LLMs) in order to improve the readability and language quality of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Acknowledgments

The authors gratefully acknowledge the financial support provided by the Health Institutes of Türkiye (TUSEB) under the call code '2023-C1-YZ' (Project No: 33934). We also extend our appreciation to Igdir University for providing the high-performance computing infrastructure at the Artificial Intelligence and Big Data Application and Research Center, which was utilized for the experimental analysis in this study.

Conflict of Interest

The authors declare no conflicts of interest.

Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

References

[1] Spence, J. D. (2009). Stroke: Atrial fibrillation, stroke prevention therapy and aging. *Nature Reviews Cardiology, 6*(8), 448–450. https://doi.org/10.1038/nrcardio.2009.98

[2] Langlois, J. A., Rutland-Brown, W., & Wald, M. M. (2006). The epidemiology and impact of traumatic brain injury: A brief overview. *Journal of Head Trauma Rehabilitation, 21*(5), 375–378. https://doi.org/10.1097/00001199-200609000-00001

[3] He, Q., Wang, Y., Fang, C., Feng, Z., Yin, M., Huang, J., & Ma, Y., Mo, Z. (2024). Advancing stroke therapy: A deep dive into early phase of ischemic stroke and recanalization. *CNS Neuroscience & Therapeutics, 30*(4), e14634. https://doi.org/10.1111/cns.14634

[4] Feigin, V. L., Barker-Collo, S., Krishnamurthi, R., Theadom, A., & Starkey, N. (2010). Epidemiology of ischaemic stroke and traumatic brain injury. *Best Practice & Research Clinical Anaesthesiology, 24*(4), 485–494. https://doi.org/10.1016/j.bpa.2010.10.006

[5] Lo, E. H., Dalkara, T., & Moskowitz, M. A. (2003). Mechanisms, challenges and opportunities in stroke. *Nature Reviews Neuroscience, 4*(5), 399–414. https://doi.org/10.1038/nrn1106

[6] Campbell, B. C. V. (2024). Hyperacute ischemic stroke care—Current treatment and future directions. *International Journal of Stroke, 19*(7), 718–726. https://doi.org/10.1177/17474930241267353

[7] Majumder, D. (2024). Ischemic stroke: Pathophysiology and evolving treatment approaches. *Neuroscience Insights, 19*. https://doi.org/10.1177/26331055241292600

[8] Liu, Y., Wen, Z., Wang, Y., Zhong, Y., Wang, J., Hu, Y., Zhou, P., & Guo, S. (2024). Artificial intelligence in ischemic stroke images: Current applications and future directions. *Frontiers in Neurology, 15*, 1418060. https://doi.org/10.3389/fneur.2024.1418060

[9] Hill, M. D., & Hachinski, V. (1998). Stroke treatment: Time is brain. *The Lancet, 352*(SIII), 10–14. https://doi.org/10.1016/S0140-6736(98)90088-5

[10] Jiang, B., Pham, N., van Staalduinen, E. K., Liu, Y., Nazari-Farsani, S., Sanaat, A., van Voorst, H., Fettahoglu, A., Kim, D., Ouyang, J., Kumar, A., Srivatsan, A., Hussein, R., Lansberg, M. G., Boada, F., & Zaharchuk, G. (2025). Deep learning applications in imaging of acute ischemic stroke: A systematic review and narrative summary. *Radiology, 315*(2). https://doi.org/10.1148/radiol.240775

[11] Pacal, I. (2024). MaxCerVixT: A novel lightweight vision transformer-based approach for precise cervical cancer detection. Knowledge-Based Systems, 289, 111482. https://doi.org/10.1016/j.knosys.2024.111482

[12] Ganie, S. M., Pramanik, P. K. D., & Zhao, Z. (2024). Improved liver disease prediction from clinical data through an evaluation of ensemble learning approaches. *BMC Medical Informatics and Decision Making, 24*(1), 160. https://doi.org/10.1186/s12911-024-02550-y

[13] Ganie, S. M., Malik, M. B., Aadil, M., Muteeb, G., Farhan, M., & Aatif, M. (2025). Explainable AI based hybrid DRM-Net transfer learning model for breast cancer detection and classification using ultrasound images. *Scientific Reports, 15*(1), 44170. https://doi.org/10.1038/s41598-025-27934-6

[14] Ganie, S. M., & Pramanik, P. K. D. (2025). Interpretable lung cancer risk prediction using ensemble learning and XAI based on lifestyle and demographic data. *Computational Biology and Chemistry, 117*, 108438. https://doi.org/10.1016/j.compbiolchem.2025.108438

[15] Banerjee, T. (2025). Towards automated and reliable lung cancer detection in histopathological images using DY-FSPAN: A feature-summarized pyramidal attention network for explainable AI. *Computational Biology and Chemistry, 118*, 108500. https://doi.org/10.1016/j.compbiolchem.2025.108500

[16] Banerjee, T., & Pacal, I. (2026). Towards accurate and interpretable brain tumor diagnosis: T-FSPANNet with tri-attribute and pyramidal attention-based feature fusion. *Biomedical Signal Processing and Control, 113*, 108852. https://doi.org/10.1016/j.bspc.2025.108852

[17] Pacal, I. (2024). A novel Swin transformer approach utilizing residual multi-layer perceptron for diagnosing brain tumors in MRI images. *International Journal of Machine Learning and Cybernetics*. Advance online publication. https://doi.org/10.1007/s13042-024-02110-w

[18] Pacal, I. (2025). Chaotic learning rate scheduling for improved CNN-based breast cancer ultrasound classification. *Chaos Theory and Applications, 7*(1), 297–306. https://doi.org/10.51537/chaos.1807694

[19] Banerjee, T., & Paçal, İ. (2025). A systematic review of machine learning in heart disease prediction. *Turkish Journal of Biology, 49*(1), 600–634. https://doi.org/10.55730/1300-0152.2766

[20] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. *arXiv*. https://arxiv.org/abs/1505.04597

[21] Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., & Han, J., Ding, G. (2024). YOLOv10: Real-time end-to-end object detection. *arXiv*. https://doi.org/10.48550/arXiv.2405.14458

[22] Redmon, J., Divvala, S., Girshick, A., & Farhadi, R. (2016). You only look once: Unified, real-time object detection. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 779–788. https://doi.org/10.1109/CVPR.2016.91

[23] Chen, S., Duan, J., Zhang, N., Qi, M., Li, J., Wang, H., Wang, R., Ju, R., Duan, Y., & Qi, S. (2023). MSA-YOLOv5: Multi-scale attention-based YOLOv5 for automatic detection of acute ischemic stroke from multi-modality MRI images. *Computers in Biology and Medicine, 165*, 107471. https://doi.org/10.1016/j.compbiomed.2023.107471

[24] Elhanashi, A., Dini, P., Saponara, S., & Zheng, Q. (2024). TeleStroke: Real-time stroke detection with federated learning and YOLOv8 on edge devices. Journal of Real-Time Image Processing, 21(3), 70. https://doi.org/10.1007/s11554-024-01500-1

[25] Ince, S., Kunduracioglu, I., Algarni, A., Bayram, B., & Pacal, I. (2025). Deep learning for cerebral vascular occlusion segmentation: A novel ConvNeXtV2 and GRN-integrated U-Net framework for diffusion-weighted imaging. *Neuroscience, 574*, 42–53. https://doi.org/10.1016/j.neuroscience.2025.04.010

[26] Chen, S., Duan, J., Wang, H., Wang, R., Li, J., Qi, M., Duan, Y., & Qi, S. (2022). Automatic detection of stroke lesion from diffusion-weighted imaging via the improved YOLOv5. *Computers in Biology and Medicine, 150*, 106120. https://doi.org/10.1016/j.compbiomed.2022.106120

[27] Zhang, K., Zhu, Y., Li, H., Zeng, Z., Liu, Y., & Zhang, Y. (2024). MDANet: Multimodal difference aware network for brain stroke segmentation. *Biomedical Signal Processing and Control, 95*, 106383. https://doi.org/10.1016/j.bspc.2024.106383

[28] Wu, Z., Zhang, X., Li, F., Wang, S., & Li, J. (2024). A feature-enhanced network for stroke lesion segmentation from brain MRI images. *Computers in Biology and Medicine, 174*, 108326. https://doi.org/10.1016/j.compbiomed.2024.108326

[29] Li, L., Liu, J., Chen, S., Wang, J., Li, Y., Liao, Q., Zhang, L., Peng, X., & Pu, X. (2025). Segmentation of acute ischemic stroke lesions based on deep feature fusion. *Information Fusion, 114*, 102724. https://doi.org/10.1016/j.inffus.2024.102724

[30] Kaya, B., & Önal, M. (2023). A CNN transfer learning-based approach for segmentation and classification of brain stroke from noncontrast CT images. *International Journal of Imaging Systems and Technology, 33*(5), 1335–1352. https://doi.org/10.1002/ima.22864

[31] Chen, Y. F., Chen, Z. J., Lin, Y. Y., Lin, Z. Q., Chen, C. N., Yang, M. L., Zhang, J. Y., Li, Y. Z., Wang, Y., & Huang, Y. H. (2023). Stroke risk study based on deep learning-based magnetic resonance imaging carotid plaque automatic segmentation algorithm. *Frontiers in Cardiovascular Medicine, 10*, 1101765. https://doi.org/10.3389/fcvm.2023.1101765

[32] Uzun, S., & Okuyar, M. (2024). A new deep learning-based GUI design and implementation for automatic detection of brain strokes with CT images. *The European Physical Journal Special Topics*. Advance online publication. https://doi.org/10.1140/epjs/s11734-024-01423-9

[33] Sinha, S., Bhatt, M., Anand, A., Areeckal, A. S., & Aparna, V. (2024). EnigmaNet: A novel attention-enhanced segmentation framework for ischemic stroke lesion detection in brain MRI. *IEEE Access, 12*, 91480–91498. https://doi.org/10.1109/ACCESS.2024.3422025

[34] Li, T., An, X., Di, Y., Gui, C., Yan, Y., Liu, S., & Ming, D. (2024). SrSNet: Accurate segmentation of stroke lesions by a two-stage segmentation framework with asymmetry information. *Expert Systems with Applications, 254*, 124329. https://doi.org/10.1016/j.eswa.2024.124329

[35] Hernandez Petzsche, M. R., de la Rosa, E., Hanning, U., Wiest, R., Valenzuela, W., Reyes, M., Meyer, M., Liew, S. L., Kofler, F., Ezhov, I., Robben, D., Hutton, A., Friedrich, T., Zarth, T., Bürkle, J., Baran, T. A., Menze, G., Broocks, G., Meyer, L., ... Ikenberg, B., & Wiestler, B., Kirschke, J. S. (2022). ISLES 2022: A multi-center magnetic resonance imaging stroke lesion segmentation dataset. *Scientific Data, 9*(1), 762. https://doi.org/10.1038/s41597-022-01875-5

[36] Global Burden of Cardiovascular Diseases and Risks 2023 Collaborators. (2025). Global, regional, and national burden of cardiovascular diseases and risk factors in 204 countries and territories, 1990-2023. *Journal of the American College of Cardiology*, *86*(22), 2167-2243. https://doi.org/10.1016/j.jacc.2025.08.015

[37] Zaher, H. F., & Roy, J. K. (2025). Decentralized finance and sustainability analysis of global research patterns and emerging themes. *Discover Sustainability*. 7, 72. https://doi.org/10.1007/s43621-025-02311-5

[38] Anikó, N., Tarek, A., Arshad, S., Nóra, G., Miklós, N., Mirzaei, M., ... & Mohammed, S. (2025). Real-time monitoring of ammonia emissions from cereal crops using LoRaWAN-based sensing technology. *Scientific Reports*. 16, 1446. https://doi.org/10.1038/s41598-025-31661-3