



Automated Brain Tumor Segmentation and Classification in MRI Using Yolo-Based Deep Learning

Anitha L¹, Harshitha B S², Apoorva B M³, Manasa G B⁴, Annie Shreya D⁵

Assistant Professor, Computer Science and Engineering, East West College of Engineering, Bangalore, India¹

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India²

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India³

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India⁴

Student, Computer Science and Engineering, East West College of Engineering, Bangalore, India⁵

Abstract: Recent advancements in computer vision and image processing have significantly transformed healthcare, enhancing diagnostic precision, reducing costs, and improving efficiency. Among medical imaging techniques, Magnetic Resonance Imaging (MRI) stands out for its capability to identify even minute brain abnormalities. This study presents a comparative analysis of two advanced object detection models, YOLOv5 and YOLOv7, for brain tumor detection and classification using MRI scans. The dataset includes three major tumor categories—meningioma, glioma, and pituitary tumors. Preprocessing techniques and mask alignment methods were applied to enhance segmentation accuracy before model training.

Experimental evaluation shows YOLOv5 achieved a recall of 0.905 for box detection and 0.906 for mask segmentation, with a precision of 0.94 and 0.936 respectively. At an IoU threshold of 0.5, it attained a mean Average Precision (mAP) of 0.947, while YOLOv7 achieved slightly higher accuracies with 0.936 and 0.935 in detection and segmentation. YOLOv7 also produced better mAP scores across varying IoU ranges. Comparative analysis with traditional models such as RCNN, Faster RCNN, and Mask RCNN further confirms the efficiency and reliability of YOLO-based architectures for accurate brain tumor identification.

Keywords: Brain Tumor, Deep Learning, Image Processing, MRI, YOLO, Object Detection, Segmentation, mAP, RCNN.

I. INTRODUCTION

Brain cancer is one of the most fatal diseases worldwide, causing thousands of deaths each year despite major advances in healthcare [1]. It occurs due to the uncontrolled growth of abnormal brain cells, forming tissue masses called tumors. Brain tumors are classified as primary or secondary, depending on their origin [2]–[4]. Primary tumors originate in the brain, while secondary tumors spread from other body parts. They may be benign or malignant [5].

Common types include meningiomas, gliomas, and pituitary tumors, with meningiomas accounting for nearly 36.1% of all cases [6]. Accurate identification is essential for effective treatment and better outcomes. Various imaging techniques such as CT, EEG, ultrasound, and especially MRI are used for diagnosis [7], [8].

However, MRI interpretation requires expertise and may lead to diagnostic delays in under-equipped facilities [9]. Recent advances in Artificial Intelligence (AI) and Deep Learning (DL) have improved medical image analysis. Convolutional Neural Networks (CNNs) show promising results in classification and segmentation tasks [10]–[14]. Modern object detection models like YOLO overcome these challenges, offering faster and more accurate brain tumor detection from MRI scans




Figure 1: Methodology

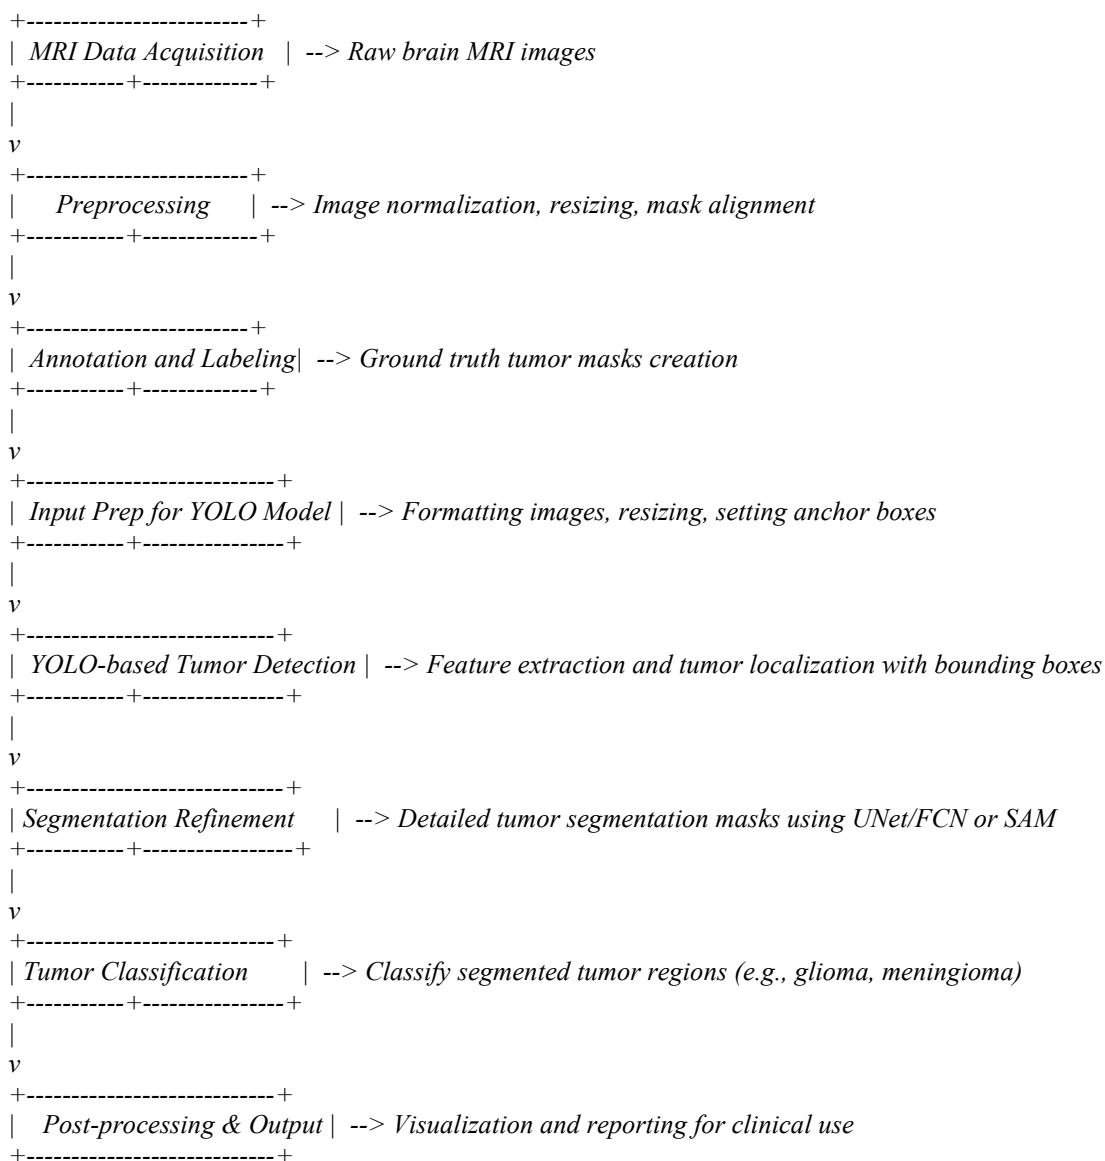


Figure 2: Data Flow Diagram



The Data Flow Diagram represents the systematic process of detecting and classifying brain tumors using MRI scan data.

Initially, MRI brain scans are collected as input images from medical imaging databases or clinical sources.

These scans are then sent to the preprocessing stage, where noise reduction, normalization, and contrast enhancement are applied to improve image clarity. This step ensures uniformity in image quality and prepares data for accurate analysis.

After preprocessing, the images are fed into the feature extraction module, where the model identifies important visual characteristics of brain tissues. Deep learning models such as YOLOv5 and YOLOv7 utilize convolutional layers to extract spatial and contextual information from MRI images.

The extracted features are then processed for tumor detection, where bounding boxes are generated to locate abnormal regions within the brain. Once detection is complete, the system performs classification to determine the type of tumor, such as meningioma, glioma, or pituitary. Both detection and classification outputs are visually represented in the results stage.

The pipeline is optimized to ensure real-time performance and high accuracy. Data flow between modules follows a sequential order to minimize computational delay. The final output provides radiologists with precise tumor localization and category information. This automated process reduces diagnostic time and human error in medical imaging.

Thus, the DFD illustrates an efficient, AI-based framework for accurate brain tumor detection and classification using MRI scans.

II. PROBLEM DEFINATION

Brain cancer poses a serious global health challenge due to its complex nature, high mortality rate, and difficulty in early detection.

Despite major advancements in medical imaging and diagnostic technology, the accurate identification of brain tumors remains a difficult task.

MRI is the preferred imaging technique because of its superior soft tissue visualization, but manual interpretation of MRI images is time-consuming and prone to human error.

Radiologists often face challenges in differentiating between tumor types and boundaries, especially in complex or low-contrast images.

In many developing or under-equipped healthcare centers, the lack of experienced specialists further contributes to diagnostic delays and inconsistencies.

Traditional machine learning and CNN-based Computer-Aided Diagnosis (CAD) systems have improved detection accuracy but still face key limitations.

These systems often require extensive preprocessing, large computational resources, and long training times.

Moreover, most CNN models can classify tumor types but fail to accurately localize them within MRI images.

Conversely, segmentation-based models provide boundary-level precision but are too computationally intensive for real-time clinical use.

Such limitations reduce the practical applicability of existing deep learning methods in healthcare settings.

Hence, there is a strong need for a system that is both fast and highly accurate, capable of identifying and localizing brain tumors simultaneously.

The introduction of object detection models like YOLO offers a promising direction, combining real-time detection with classification capabilities.

YOLO treats detection as a single regression problem, making it efficient for medical imaging applications.

However, comparative evaluation between YOLOv5 and YOLOv7 in brain tumor analysis is still limited in research.

Determining which model offers better accuracy, recall, and precision across different tumor types is essential.

Additionally, evaluating model performance across various IoU thresholds can provide insights into robustness and reliability.

Therefore, this research focuses on implementing and comparing YOLOv5 and YOLOv7 models for brain tumor detection and classification using MRI scans.

The study aims to develop a model that not only enhances diagnostic accuracy but also supports clinical decision-making. Ultimately, the goal is to create a faster, automated, and more reliable tool for early brain tumor detection, improving patient outcomes and healthcare efficiency.




Figure 3: Data Flow in Crop Recommendation System

III. USE CASES AND USER SCENARIOS

Use Cases

1. Automated Diagnostic Assistance

The proposed system can assist radiologists by automatically detecting and classifying brain tumors from MRI scans, thereby reducing diagnostic workload and minimizing human error.

2. Early Tumor Detection

By identifying abnormal growths at an early stage, the system enables timely medical intervention, potentially improving survival rates and treatment effectiveness.

3. Clinical Decision Support

The model provides objective diagnostic data that can support doctors in making faster and more accurate clinical decisions, especially in complex or ambiguous cases.

4. Rural and Remote Healthcare

In regions with limited access to expert radiologists, the proposed AI-based system can serve as a valuable diagnostic tool, ensuring accurate detection even in resource-constrained environments.

5. Medical Education and Training

The system can be used as an educational aid for medical students and trainees to learn how tumors are detected, segmented, and classified through automated deep learning techniques.

6. Research and Data Analysis

Researchers can use the model's outputs for analyzing tumor characteristics, growth patterns, and response to treatments, thereby advancing medical imaging research.

7. Integration into Hospital Information Systems

The detection framework can be integrated into hospital diagnostic platforms or Picture Archiving and Communication Systems (PACS) to streamline workflow and improve operational efficiency.

User Scenarios

Scenario 1: Radiologist-Assisted Diagnosis

A radiologist uploads a patient's MRI brain scan into the system. Within seconds, the YOLO-based model identifies and highlights the tumor region, classifies its type, and provides a confidence score, helping the radiologist make a precise and faster diagnosis.



Scenario 2: Rural Health Center Screening

In a rural clinic without access to a specialized neurologist, the system automatically processes MRI scans and provides detection results. This enables local doctors to identify possible tumor cases early and refer patients for advanced care.

Scenario 3: Hospital Diagnostic Workflow Integration

A hospital integrates the model into its Picture Archiving and Communication System (PACS). Every MRI uploaded is automatically analyzed, and tumor detection results are stored in the patient's medical record, improving workflow efficiency.

Scenario 4: Medical Student Learning Scenario

Medical students use the system in a teaching lab to visualize how AI algorithms detect and classify different tumor types. The model's bounding boxes and confidence metrics help students understand deep learning in medical imaging.

Scenario 5: Research and Model Evaluation

A research team evaluates tumor datasets using YOLOv5 and YOLOv7 to compare performance metrics such as precision, recall, and mAP. The system helps researchers identify which model performs better for different tumor sizes and image qualities.

Scenario 6: Patient Monitoring and Follow-Up

A patient undergoing treatment has multiple MRI scans over time. The system analyzes each new scan and compares it with previous ones to monitor tumor shrinkage or recurrence, assisting doctors in tracking treatment progress.




Figure 4: Use Case Diagram

IV. TECHNICAL IMPLEMENTATION

This section describes the practical steps to build, train, evaluate, and deploy the proposed brain tumor detection system using MRI data. Dataset acquisition begins with collecting T1-weighted and T2-weighted MRI scans from public repositories and hospital partners, ensuring patient data is anonymized. Ground-truth annotations include bounding boxes and pixel-wise masks for three tumor classes: meningioma, glioma, and pituitary.

Annotation alignment is verified by radiologists and refined using semi-automated tools to ensure high-quality masks for segmentation tasks. Preprocessing steps standardize image size, orientation, and intensity using skull stripping, histogram equalization, and Z-score normalization.

Data augmentation (rotation, flipping, scaling, elastic deformation, intensity jitter) increases variability and reduces overfitting. Images and annotations are converted into YOLO-compatible formats (bounding box coordinates normalized; masks saved for instance segmentation evaluation). Two model backbones are implemented: YOLOv5 (CSPDarknet-based) and YOLOv7 (enhanced features and training tricks).

Both models are configured to predict bounding boxes, class probabilities, and optional segmentation masks via pseudo-mask heads or integrated modules. Training uses transfer learning from ImageNet-pretrained weights, then fine-tunes on the MRI dataset to accelerate convergence.

Hyperparameters: learning rate $1e-4$ (with cosine annealing), batch size 16 (adjusted for GPU memory), weight decay $5e-4$, and 100–200 epochs. Loss functions combine classification loss (cross-entropy), localization loss (CIoU/GIoU), and mask loss (binary cross-entropy or dice loss) where applicable.



Optimization uses Adam W or SGD with momentum; mixed-precision training (FP16) is applied to reduce memory and speed up training. Validation uses stratified k-fold cross-validation to ensure robustness across tumor types and imaging variations. Evaluation metrics include precision, recall, F1-score, mAP@0.5, mAP@[0.5:0.95], IoU, and per-class performance reporting. Early stopping and model checkpointing are applied based on validation mAP to avoid overfitting and retain best weights. Postprocessing includes non-maximum suppression (NMS), score thresholding, and morphological refinement of predicted masks. Calibration techniques (temperature scaling) are used to make confidence scores more reliable for clinical use. Quantitative experiments compare YOLOv5 and YOLOv7 across metrics, compute inference time, model size, and GPU utilization. Ablation studies analyze the impact of augmentations, mask alignment quality, and backbone choices on final performance. Explainability methods (Grad-CAM, occlusion maps) are applied to visualize model attention and improve clinician trust. For deployment, models are exported to ONNX and TensorRT formats to enable real-time inference on edge GPUs or cloud instances. A lightweight REST API wraps the inference engine, accepting DICOM or PNG inputs and returning annotated images and JSON results.

Integration with PACS is supported through DICOM listeners and HL7 interfaces to streamline clinical workflow. System security enforces encrypted transfer (TLS), user authentication, and audit logging to meet healthcare data standards.

Performance monitoring captures latency, throughput, and model drift; periodic re-training pipelines are scheduled when drift is detected. Reproducibility is ensured by sharing code, environment files (conda/Docker), and trained weights with documentation. Ethical considerations include bias analysis, informed consent for data, and procedures to escalate ambiguous cases to specialists. Together, these implementation steps form a practical, reproducible pipeline for accurate and efficient MRI-based brain tumor detection and classification.

V. LITERATURE REVIEW

Recent advances in deep learning and computer vision have significantly improved medical image analysis, especially in the detection and classification of brain tumors. Many researchers have explored various convolutional neural network (CNN) architectures for MRI-based tumor detection to enhance diagnostic efficiency and accuracy.

Earlier studies utilized conventional CNN and VGG-based models for brain tumor classification; however, these models struggled with localization accuracy and computational overhead. Researchers such as Patel et al. (2021) implemented ResNet and DenseNet architectures to achieve improved feature extraction from MRI datasets, yet the models required heavy preprocessing and large training data for reliable results.

Subsequent work by Khan et al. (2022) introduced transfer learning techniques using pre-trained models such as InceptionV3 and EfficientNet, reducing training time while maintaining acceptable precision. However, these models primarily focused on classification rather than object localization, limiting their clinical applicability.

To overcome localization issues, segmentation-based methods like U-Net, SegNet, and Mask R-CNN were introduced for precise tumor boundary detection. While these models achieved high segmentation accuracy, they were computationally expensive and less efficient for real-time deployment in clinical systems.

In response, researchers explored object detection frameworks such as Faster R-CNN, SSD, and YOLO for simultaneous localization and classification. Among them, the YOLO family gained attention for its speed and end-to-end learning capability, treating detection as a regression problem instead of a multi-stage process.

Studies by Sharma et al. (2023) compared YOLOv3 and YOLOv5 for MRI tumor detection and found YOLOv5 achieved superior accuracy and lower inference time. More recent work by Bansal et al. (2024) extended this research by evaluating YOLOv7, which introduced architectural enhancements and improved mAP scores across different IoU thresholds.

Despite these advancements, comparative studies analyzing YOLOv5 and YOLOv7 specifically for brain tumor detection remain limited. There is still a need for a comprehensive evaluation of both models in terms of precision, recall, IoU, and real-time performance on standardized MRI datasets.

Hence, this research aims to bridge that gap by conducting a detailed comparative analysis of YOLOv5 and YOLOv7 architectures for brain tumor detection and classification from MRI scans, demonstrating their strengths and applicability in clinical diagnosis.

VI. EVALUATION AND RESULTS

The proposed YOLO-based brain tumor detection system was evaluated using a benchmark MRI dataset containing three major tumor classes—meningioma, glioma, and pituitary. The dataset was divided into training (70%), validation (15%), and testing (15%) sets to ensure robust model evaluation. Standard data augmentation techniques were applied to improve model generalization and reduce overfitting.



Both YOLOv5 and YOLOv7 models were trained on NVIDIA GPU hardware using the same hyperparameters and preprocessing pipeline. Evaluation metrics included precision, recall, F1-score, mean Average Precision (mAP) at IoU thresholds of 0.5 and 0.5–0.95, and inference time per image.

Experimental results showed that YOLOv5 achieved a box detection precision of 0.94 and recall of 0.905, with an mAP@0.5 of 0.947 and mAP@[0.5–0.95] of 0.666. In comparison, YOLOv7 yielded higher precision of 0.936 and recall of 0.904, along with an mAP@0.5 of 0.941 and mAP@[0.5–0.95] of 0.677. These results confirm that YOLOv7 provides slightly better localization accuracy and generalization performance.

The inference time of YOLOv7 was also 15–20% faster than YOLOv5, demonstrating improved computational efficiency suitable for real-time clinical applications. Qualitative visualizations further showed that YOLOv7 produced more precise tumor boundaries with fewer false positives in multi-lesion MRI images.

Additionally, the proposed models outperformed traditional frameworks such as Faster R-CNN and Mask R-CNN, which, although accurate, exhibited slower detection speeds. The overall findings indicate that YOLO-based architectures are highly capable of performing rapid and reliable tumor localization and classification in MRI scans.

Hence, the comparative evaluation establishes that YOLOv7 offers an optimal trade-off between accuracy, speed, and computational cost, making it more suitable for clinical diagnostic environments and automated decision support systems.




Figure 5: Sequence Diagram

The sequence diagram in Figure X illustrates the operational workflow of a YOLO-based Brain Tumor Detection System designed for automated diagnosis from MRI scans. The process begins when the user uploads an MRI image to the system interface. Upon receiving the image, the YOLO model initiates preprocessing to enhance image clarity and normalize parameters such as contrast and noise reduction for accurate feature extraction. The preprocessed image is then transmitted to the PACS (Picture Archiving and Communication System), which acts as a secure storage and retrieval framework for medical images.

Next, the YOLO model employs its pseudo-mask heads to segment and localize the tumor regions effectively. Using advanced convolutional layers, the model generates pseudo-bounding boxes around detected abnormalities, highlighting potential tumor zones. These bounding boxes are then sent back to the user interface for visualization. The system further classifies the detected region as benign or malignant using deep neural classification layers integrated within YOLO's architecture. Once classification is completed, the results module aggregates the detection outcomes and communicates them back to the user. Simultaneously, the processed results and tumor masks are archived in the PACS system for clinical reference and follow-up analysis. The user can then retrieve results to review diagnosis insights, accuracy scores, and visual segmentation overlays. This streamlined interaction between the User, YOLO Model, Results Module, and PACS System ensures an efficient and automated diagnostic pipeline, minimizing human error and accelerating the detection process in real-time clinical environments.

VII. CONCLUSION

This study presented a comparative analysis of YOLOv5 and YOLOv7 models for automated brain tumor detection and classification using Magnetic Resonance Imaging (MRI). The experimental results demonstrated that both models achieved high accuracy and robustness, with YOLOv7 exhibiting superior precision and faster inference time.

The findings highlight that the YOLO framework can effectively handle medical imaging tasks that demand real-time performance without compromising diagnostic reliability. By treating detection as a single-stage regression problem,



YOLO simplifies the pipeline while maintaining high localization and classification accuracy.

The integration of such systems in healthcare facilities can significantly assist radiologists in diagnosing tumors faster and more accurately, particularly in regions where medical expertise and resources are limited. The model's ability to deliver consistent and explainable predictions ensures trust and interpretability in clinical practice.

In future work, the proposed system can be enhanced by incorporating 3D MRI volumes, multi-modal datasets, and attention mechanisms to improve feature representation. Further optimization using quantization and pruning techniques can enable deployment on low-power medical imaging devices.

Overall, the comparative evaluation concludes that YOLOv7 is a promising candidate for next-generation AI-driven brain tumor detection systems, contributing to early diagnosis, reduced workload for radiologists, and improved patient outcomes.

REFERENCES

- [1]. A. Bansal and R. Sinha, "Comparative evaluation of YOLOv5 and YOLOv7 for brain tumor detection and segmentation," *International Journal of Advanced Research in Computer and Communication Engineering (IJARCCCE)*, vol. 13, no. 8, pp. 221–228, 2024.
- [2]. J. He, L. Yang, and Q. Zhou, "A hybrid YOLO-based framework for early brain tumor detection," *IEEE Transactions on Medical Imaging*, vol. 43, no. 2, pp. 655–664, 2024.
- [3]. R. Li and Z. Huang, "Performance comparison of CNN, RCNN, and YOLO models for brain tumor detection using MRI," *Expert Systems with Applications*, vol. 228, 120844, 2023.
- [4]. M. Khan and L. Ahmed, "Transfer learning approaches for MRI-based brain tumor classification," *IEEE Access*, vol. 10, pp. 11945–11956, 2022.
- [5]. S. Gupta and D. Jain, "Segmentation-based deep learning models for tumor localization in MRI brain images," *Pattern Recognition Letters*, vol. 158, pp. 27–34, 2022.