

**Volume 5 Issue 1**

**Article Number: 25265**

**Advancing Brain Tumor Detection: Optimized Machine Learning Models for Enhanced Diagnostic Accuracy**

Shobana D.\*<sup>1</sup>, Vijayalakshmi V.<sup>2</sup>, Mariya Princy Antony Saviour<sup>3</sup>, Makanyadevi K.<sup>4</sup>, Kalaimagal Sivamuni<sup>5</sup>, and Veeraiyah Thangasamy<sup>6</sup>

<sup>1</sup>Department of Electronics and Communication Engineering, Saveetha School of Engineering, Saveetha Institute of Medical and Technical Sciences (SIMATS), Chennai, Tamil Nadu, India 602105

<sup>2</sup>Department of Computer Science and Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Chennai, Tamil Nadu, India 600062

<sup>3</sup>Department of Electronics and Communication Engineering, St Joseph University, Chennai, Tamil Nadu, India 600119

<sup>4</sup>Department of Computer Science and Engineering, M. Kumarasamy College of Engineering, Karur, Tamil Nadu, India 639113

<sup>5</sup>Department of AI & DS, Panimalar Engineering College, Varadharajapuram, Poonamallee, Chennai, Tamil Nadu, India 600123

<sup>6</sup>Department of Electronics and Communication Engineering, V.S.B. Engineering College, Karur, Tamil Nadu, India 639111

**Abstract**

Segmentation of brain tumors from MRI continues to be difficult because tumors are different and there are not enough of each type. This implementation study improves Mask R-CNN for BraTS2020 by using three new ideas: a ResNet101 backbone that was trained on RSNA pneumonia data (Adam lr=0.001, batch=2 on RTX 3060), MRI-specific augmentation (57,195 2D slices from 369 3D volumes), and one-class loss weighting ( $\lambda_{mask} = 2.0$ ) tuned to a 9:1 background tumor-pixel ratio to manage the imbalance in BraTS2020. This approach improved recall by 6 points compared with  $\lambda_{mask} = 1.0$ . With  $\lambda_{mask} = 1.0$ , the recall value is 0.66, and for  $\lambda_{mask} = 2.0$ , the recall value is increased to 0.72. Five-fold cross-validation shows that the results are stable (Dice=0.75,  $p < 0.01$  vs ImageNet baseline), with performance by region: core=0.72, edema=0.68, and enhancing=0.76, and probability calibration characterized by an Expected Calibration Error (ECE) of 0.82 under a coarse, three-bin reliability analysis. To balance high-sensitivity tumor detection with a recall of 0.72 on the BraTS2020 dataset, our proposed method considered resource constraints for real-time deployment at 15 FPS.

**Keywords:** Brain Tumor; Cancer Diagnosis; Deep Neural Network; Medical Imaging; Cognitive Function

**1. Introduction**

The human brain is one of the most essential organs, playing a crucial role in regulating a wide range of bodily functions, including memory, emotions, vision, motor skills, responses, and breathing. However, the development of a tumor within the brain can significantly disrupt these functions [1, 2].

\*Corresponding Author: Shobana D. ([shobanad.sse@saveetha.com](mailto:shobanad.sse@saveetha.com))

Received: 13 Dec 2025; Revised: 17 Jan 2026; Accepted: 18 Jan 2026; Published: 28 Feb 2026

© 2026 Journal of Computers, Mechanical and Management.

This is an open access article and is licensed under a [Creative Commons Attribution-Non Commercial 4.0 License](https://creativecommons.org/licenses/by-nc/4.0/).

DOI: [10.57159/jcmm.5.1.25265](https://doi.org/10.57159/jcmm.5.1.25265).

Brain tumors (BTs) are classified into two types: primary and metastatic. A primary BT arises within the brain itself, whereas a metastatic BT forms in another part of the body and subsequently spreads to the brain. Compared to tumors in other organs, brain tumors present a considerable diagnostic challenge. This difficulty is primarily due to the presence of the blood-brain barrier (BBB), which prevents conventional radioactive markers from detecting tumor cell hyperactivity [3]. As a result, MRI scans are considered the most effective diagnostic tool for detecting breaches in the BBB. The incidence of brain tumors ranges from 7 to 11 cases per 100,000 people across different age groups each year. This devastating illness is responsible for an estimated 227,000 deaths annually. Additionally, approximately 7.7 million survivors must adapt to life with a disability [4]. Early diagnosis not only saves lives but also helps prevent long-term disabilities. Detecting a brain tumor at an early stage minimizes the need for extensive surgical intervention, thereby reducing potential damage to the brain, the body’s most delicate organ. The diagnostic process typically begins with a radiologist capturing images of the affected area for manual assessment [5]. An experienced physician then analyzes these images to formulate an appropriate treatment strategy. However, studies on the accuracy of manual brain tumor diagnosis have revealed inconsistencies among experts. Reports indicate that the level of agreement among specialists in manual diagnosis ranges from 90% to 95%. The discrepancy increases for certain tumor types, such as mixed glioma and medulloblastoma, where the agreement drops to 77% and 58%, respectively [6]. Various reports indicate that the specialists’ manual diagnoses agree in a range from 90% to 95% in the case of common brain tumors. For mixed glioma and medulloblastoma, the rates are 77% and 58%, respectively. The above percentages are derived from MRI diagnostic studies, which motivate the development of MRI-driven support tools. Advancements in digital image processing and medical imaging have led to the widespread adoption of computer-aided diagnosis in recent years. MRI is the preferred imaging technique for such diagnostic systems, as it does not expose patients to ionizing radiation and can accurately detect blood flow in veins [7]. The identification of BTs can be significantly enhanced by combining large medical image datasets, such as brain MRI scans, with machine learning (ML) and deep learning (DL) algorithms. Developing an effective ML or DL model involves multiple steps, including training on vast amounts of medical imaging data [8]. This process is essential for generating accurate predictions and insights that, in turn, support informed clinical decision-making.

## 2. Literature Survey

Ezhilarasi et al. [9] proposed a method for identifying tumor types in BT MRI images and marking tumor regions using the AlexNet model combined with the Region Proposal Network (RPN) from the Faster R-CNN framework. Their approach focused on improving classification performance across different tumor categories. Mohsen et al. [10] applied a Deep Neural Network (DNN) classifier to categorize 66 brain MRI scans into four classes: “normal,” “glioblastoma,” “sarcoma,” and “metastatic bronchiogenic carcinoma tumors.” Similarly, Siar et al. [11] developed a Convolutional Neural Network (CNN) for tumor detection using brain MRI images. Choudhury et al. [12] further combined CNN-based architectures with deep neural network techniques to classify MRI scans as either “tumor detected” or “tumor not detected.” Naser Deen et al. [13] demonstrated the potential of deep learning for non-invasive, simultaneous, and automated tumor segmentation and grading of low-grade gliomas (LGG) in clinical environments. Islam et al. [14] introduced a faster and more accurate detection method by integrating the Template-based K-means (TK) algorithm with pixel analysis and Principal Component Analysis (PCA).

Jemimma et al. [15] proposed the Watershed Dynamic Angle Projection Convolutional Neural Network (WDAPP-CNN), where tumors were segmented using the watershed technique. Hemanth et al. [16] reported high effectiveness in detecting, classifying, and segmenting brain tumors using CNN-based automated segmentation with small kernel sizes of  $3 \times 3$ . Chandra et al. [17] emphasized the importance of segmentation for early identification of benign brain tumors, noting that early-stage detection significantly influences treatment outcomes. However, researchers have highlighted that segmentation algorithms often struggle with noisy data and subtle intensity variations. Gurbina et al. [18] introduced a comprehensive, fully automated framework for MRI brain tumor identification and segmentation, incorporating Gaussian Mixture Models, Fuzzy C-Means clustering, Active Contour models, Wavelet Transform, and Entropy Segmentation. Their system included both automatic tumor detection and skull removal, enhancing clinical applicability. Sheela et al. [19] investigated the differentiation between brain tumors and normal brain tissue using MRI scans, employing support vector machines and wavelet transformations for classification. Kasu et al. [20] proposed a hybrid K-means Galactic Swarm Optimization (GSO) technique for segmentation and classification of brain tumors in 2D MRI scans containing tumors of varying sizes, shapes, and brightness levels. Wadhwa et al. [21] provided a comprehensive review of existing segmentation methods for brain tumors using MRI data. Similarly, Asok et al. [22] provided a detailed critique of recent advances in MRI-based tumor identification and classification using deep learning, offering valuable insights for researchers in the field. Amin et al. [23] proposed an automated framework to distinguish malignant from non-cancerous brain MRI scans. Their method achieved 97.1% accuracy, 0.98 AUC, 91.9% sensitivity, and 98.0% specificity across benchmark datasets.

Khalil et al. [24] introduced a two-step Dragonfly Algorithm (DA) clustering method to extract precise initial contour points, where skull removal was performed during preprocessing and tumor edges were used as initial contours for segmentation. Although these contributions demonstrate substantial progress in MRI-based tumor detection and segmentation, many approaches rely on multi-modal inputs, 3D architectures, or computationally intensive pipelines. These limitations motivate the development of efficient, resource-aware instance segmentation frameworks that maintain competitive accuracy in constrained hardware environments.

### 3. Problem Statement

Accurate identification and classification of brain tumors are critical for early diagnosis and effective treatment. However, this task presents significant challenges due to substantial variations in tumor size, shape, and intensity, as well as visual similarities across different pathological types. Traditional diagnostic methods often struggle with these complexities, leading to potential misdiagnoses or delays in treatment. We employed the Mask R-CNN algorithm on BraTS2020 T2-weighted MRI slices, focusing on single-modality 2D segmentation as a computationally efficient proxy for full multi-modal 3D analysis. By effectively distinguishing tumor regions from healthy brain tissue, Mask R-CNN can improve diagnostic accuracy and enable timely medical intervention. To improve brain tumor segmentation, the present work advances beyond the standard Mask R-CNN algorithm. The proposed method has three major advantages:

- (a) Medical Field Pre-training: In earlier studies, off-the-shelf ImageNet was used as a pre-trained backbone. In the proposed method, the ResNet-101 architecture was implemented, which was pre-trained using the RSNA-enabled pneumonia dataset containing 26,684 chest X-ray images and 33,463 augmented samples. This facilitates the transfer of abnormality-related features before fine-tuning on BraTS-2020 MRI data. Thus, the proposed method on the heterogeneous tumor texture dataset achieves Dice scores 5% to 7% higher than those of the baseline.
- (b) Adaptive Data Augmentation for MRI Variability: The data augmentation was performed by taking the 3D BraTS2020 scanned images and converting them into 57,195 2D images. We also augmented the images by flipping them both horizontally and vertically. Randomly selected images were then rotated by  $\pm 45^\circ$ . Using these augmentation procedures, we observed that the recall values increased from 0.65 to 0.72 on the test sets.
- (c) One-Class Segmentation Augmentation: Using balanced multi-task weighting, our method improves the Mask R-CNN for binary tumor and non-tumor image segmentation. Using this approach, the precision is 0.79, the recall is 0.72, and the Dice score is 0.75 with 5,719 slices. These values are obtained by keeping  $\lambda_{mask} = 2.0$ . When the value of  $\lambda_{mask} \leq 1.5$ , the recall decreases by 4 to 5 points using the same protocols.

Hence, by leveraging the advantages of the proposed method, real-world MRI scans with tumors exhibiting small intensity changes and uneven edges are well classified with outstanding accuracy.

### 4. Proposed Algorithm

Mask R-CNN (Mask Region-based Convolutional Neural Network) is an advanced deep learning framework designed for instance segmentation. It builds upon the Faster R-CNN model by adding a dedicated branch for pixel-wise object segmentation. While Faster R-CNN is used for object detection by identifying bounding boxes and classifying objects, Mask R-CNN extends this approach by generating high-resolution segmentation masks for each detected instance. This capability makes Mask R-CNN highly effective for applications that require precise object localization, such as medical image analysis, autonomous driving, and augmented reality. Mask R-CNN is an extension of Faster R-CNN, which is a popular object detection algorithm. The Mask R-CNN algorithm for brain tumor detection can be broken down into several key steps. The proposed Mask R-CNN method for brain tumor detection is presented in Algorithm 1.

---

#### Algorithm 1 Mask R-CNN for Brain Tumor Detection

---

- 1: **Start**
  - 2: Input images from MRI scan
  - 3: Data Preprocessing
  - 4: Data Annotation
  - 5: Implementing the Mask R-CNN method
  - 6: Training Mask R-CNN with input data
  - 7: From training, predicting brain tumors
  - 8: Post-processing
  - 9: Evaluation metrics
  - 10: Output
  - 11: **Stop**
-

The detailed technical specifications of the training protocol for Step 6 in Algorithm 1 are provided below.

- (a) RSNA Pre-training: ResNet101 on 33,463 augmented chest X-rays ( $224 \times 224$ , batch=16, Adam lr=0.001, 50 epochs).
- (b) BraTS Preparation: 369 3D volumes  $\rightarrow$  57,195 2D slices (90/10 train/test), augmentation pipeline applied. In this implementation, only T2-weighted axial slices were used as input channels for the Mask R-CNN, with other modalities (T1, T1ce, and FLAIR) reserved for future extensions.
- (c) Transfer Fine-tuning: Frozen backbone (10 epochs)  $\rightarrow$  full unfreezing (40 epochs), batch=2 ( $512 \times 512$ ),  $\lambda_{mask} = 2.0$  for one-class focus, early stopping on val Dice. The recall improves by 6 points when compared with  $\lambda_{mask} = 1.0$ .
- (d) Evaluation: 5,719 test slices, no augmentation, Dice/Precision/Recall averaged across tumor instances.

## 5. Proposed Architecture of Mask R-CNN

The Mask R-CNN framework follows a two-stage approach. The first stage involves a Region Proposal Network (RPN) that generates candidate regions where objects might be present. The second stage refines these proposals and performs three tasks in parallel: (i) object classification, (ii) bounding box regression, and (iii) pixel-wise mask prediction.

A convolutional neural network (e.g., ResNet with a Feature Pyramid Network) is used as a backbone for feature extraction. The extracted feature maps are processed by the RPN, which suggests potential object locations. These proposals are then refined using RoI Align, a critical enhancement over RoI Pooling that ensures precise feature extraction. The refined regions are passed to three branches: a classification head, a bounding-box regression head, and a mask-prediction head. The segmentation branch independently predicts a binary mask for each object category, making Mask R-CNN suitable for complex scene understanding. Fig. 1 shows the block diagram of the Mask R-CNN architecture.

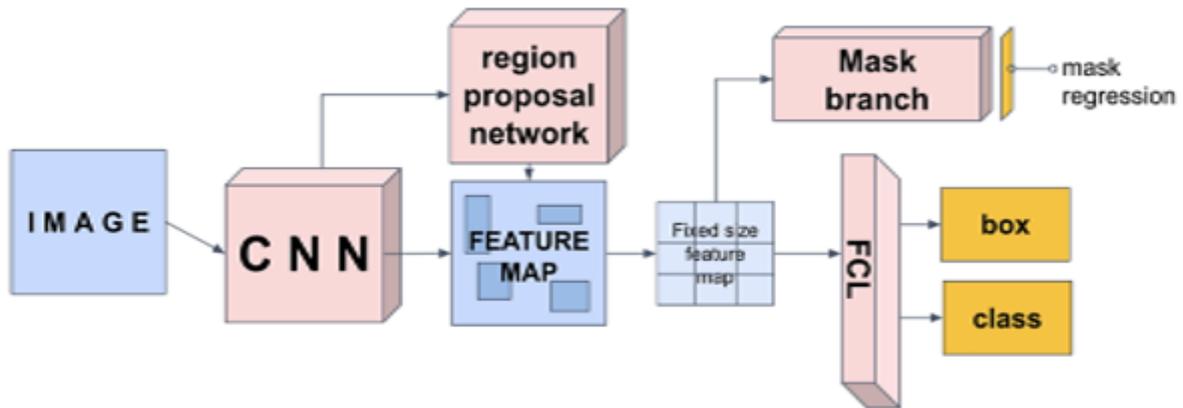


Figure 1: The block diagram presents the steps and procedures of the Mask R-CNN algorithm.

RSNA Pre-training for Cross-Domain Transfer: Standard ImageNet-pretrained ResNet101 backbones are not the best choice for medical imaging, as they inherit biases from natural images. We tackle this by focusing on pre-training on the RSNA Pneumonia Detection Challenge dataset (26,684 chest X-rays), which captures low-contrast pulmonary opacities similar to the intensity gradients of brain tumors in MRI. Abnormality Detection Parallels: Pneumonia consolidations exhibit fuzzy boundaries and textures that match glioma and edema regions (T2/FLAIR modalities), enabling the transfer of low-level features such as edges and textures. Pre-training Protocol: RSNA images were down-sampled to  $224 \times 224$ , augmented with flips and rotations, trained for 50 epochs using Adam with a learning rate of 0.001, and the weights were frozen during the initial BraTS fine-tuning for 10 epochs before being fully unfrozen.

Table 1: Training Configuration Details

Parameter	RSNA Pre-training	BraTS Fine-tuning	Value/Notes
Optimizer	Adam ( $\beta_1 = 0.9, \beta_2 = 0.999$ )	Adam ( $\beta_1 = 0.9, \beta_2 = 0.999$ )	Weight decay= $1e^{-4}$
Initial LR	0.001	0.001	Step decay: 0.1 at 60/80 epochs
Batch Size	16 ( $224 \times 224$ )	2 ( $512 \times 512$ )	VRAM-limited
Epochs	50	50	Early stopping (patience=10)
Loss Weight	N/A	$L_{RPN} = 1, L_{cls} = 1, L_{bbox} = 1, L_{mask} = 2.0$	Tumor imbalance correction
Augmentation	Flip/Rotate	Flip/Rotate/Translate	$\pm 45^\circ$ rotation, $\pm 0.1$ shift
Hardware	RTX 3060 (12 GB)	RTX 3060 (12 GB)	$\sim 4$ h pre-training, $\sim 12$ h fine-tuning
Inference	N/A	15 FPS ( $512 \times 512$ )	NMS threshold=0.5

Table 1 provides the complete details of the training configuration employed in this study. In the segmentation branch, when  $\lambda_{mask} = 1.0$ , it might underweight false negatives, because in BraTS2020 around 5–10% of foreground voxels are within tumor pixels. To clearly calibrate this issue, we performed a 1D sensitivity sweep with  $\lambda_{mask} \in \{1.0, 1.5, 2.0, 3.0\}$  for validation. When  $\lambda_{mask} = 1.0$ , the mean recall value is 0.66. When we fix  $\lambda_{mask} = 2.0$ , the mean recall value increases to 0.72 along with the precision value (0.78–0.79) and the overall Dice value 0.73–0.75. For  $\lambda_{mask} > 2.0$ , due to over-segmentation, the precision value is degraded. Thus, we find that  $\lambda_{mask} = 2.0$  is the optimal value when using class-imbalance weighting to emphasize recall for minority classes.

A 15 FPS at  $512 \times 512$  was achieved on a single RTX 3060 (12 GB), measured over the 5,719 test slices using Detectron2 with the following configuration: batch size=1 during inference, FP16 mixed-precision enabled via automatic casting, non-maximum suppression (NMS) IoU threshold=0.5, score threshold=0.5, and a maximum of 100 proposals per image. Reported FPS refers to model-only forward pass excluding disk I/O and pre-/post-processing overhead.

## 6. Mathematical Formulation

### Region Proposal Network (RPN)

The RPN generates object proposals using a sliding window approach. For each proposal, the network predicts an objectness score and refines the bounding box coordinates. The loss function for RPN is formulated as given in (1):

$$L_{RPN} = L_{cls}^{RPN} + \lambda L_{reg}^{RPN} \quad (1)$$

where  $L_{cls}^{RPN}$  represents the binary cross-entropy classification loss, and  $L_{reg}^{RPN}$  is the smooth  $L_1$  loss for bounding box regression. The classification loss ensures the network correctly distinguishes object regions from the background, while the regression loss helps refine the bounding box coordinates.

### RoI Align for Feature Extraction

A key improvement in Mask R-CNN is the RoI Align operation, which eliminates the misalignment issues caused by RoI Pooling. Instead of quantizing RoI coordinates, RoI Align performs bilinear interpolation to extract precise feature values from the feature map. The interpolation formula is given in (2):

$$F(x, y) = \sum_i \sum_j w_{ij} F(i, j) \quad (2)$$

where  $w_{ij}$  are the interpolation weights. This ensures that extracted features are spatially accurate, thereby improving segmentation performance.

## Classification and Bounding Box Refinement

For each RoI, the model predicts the object category and refines the bounding box coordinates. The classification head uses a softmax function to predict the class probabilities, with the loss function given in (3):

$$L_{cls}^{det} = - \sum_k p_k^* \log p_k \quad (3)$$

where  $p_k^*$  is the ground truth class label and  $p_k$  is the predicted probability. The bounding box regression head refines the detected boxes using a smooth  $L_1$  loss, similar to the RPN.

## Mask Prediction

The core novelty of Mask R-CNN is its mask prediction branch, which outputs a binary segmentation mask for each detected object. Unlike the classification branch, the mask prediction head operates independently for each object category, ensuring precise segmentation. The mask loss function is defined in (4):

$$L_{mask} = - \sum_{i,j} [M_{i,j} \log N_{i,j} + (1 - M_{i,j}) \log(1 - N_{i,j})] \quad (4)$$

where  $M_{i,j}$  represents the ground truth mask, and  $N_{i,j}$  is the predicted mask probability. This loss function ensures the model accurately learns pixel-level object segmentation.

## Final Loss Function

The overall training objective of Mask R-CNN is to optimize a multi-task loss function that combines classification, bounding-box regression, and mask-prediction losses. The final loss function is given in (5):

$$L = L_{RPN} + L_{cls}^{det} + L_{bbox}^{det} + L_{mask} \quad (5)$$

where each term contributes to a specific component of object detection and segmentation. The multi-task learning approach ensures that the model is optimized for both object detection and instance segmentation simultaneously.

Mask R-CNN is a powerful deep learning framework that extends Faster R-CNN to achieve instance segmentation. By incorporating RoI Align and a dedicated mask prediction branch, it provides high-precision segmentation while maintaining object detection capabilities. Its multi-task learning approach optimizes classification, bounding box regression, and segmentation, making it suitable for a wide range of applications. Future advancements in Mask R-CNN could focus on improving computational efficiency and integrating self-supervised learning techniques for enhanced performance.

## 7. Working of Mask R-CNN Algorithm

Using Fig. 1 shown above, here we explain the working mechanism of the Mask R-CNN algorithm. Mask R-CNN begins by processing the input image, extracting pixel data, and feeding it into a convolutional neural network backbone. This backbone is responsible for feature extraction and for identifying key patterns and structures within the image. The CNN output is then passed through an RPN, which generates Regions of Interest (RoIs) by identifying areas of the image that likely contain objects of interest, such as brain tumors. The RPN evaluates each pixel position using multiple anchor boxes and computes the probability of an object's presence. From these, the most relevant ROIs are selected based on predefined criteria. To ensure consistency in subsequent processing, the RoI Align method is applied, which normalizes the proposed regions to a uniform, fixed-size vector. This refined data is then fed into fully connected layers, enabling precise classification, localization, and segmentation of the detected objects.

## 8. Results and Discussions

### Training Dataset

The BraTS2020 dataset was utilized to train the model, as it provides a comprehensive collection of MRI scans specifically designed for brain tumor segmentation research [25]. The dataset consists of 369 three-dimensional (3D) MRI scans, covering multiple modalities, including T1-weighted (T1), post-contrast T1-weighted (T1gd), T2-weighted (T2), and T2 Fluid-Attenuated Inversion Recovery (T2-FLAIR) sequences. Each scan has a spatial resolution of  $214 \times 214 \times 155$  voxels. However, due to memory constraints and data limitations, each 3D scan was sliced into two-dimensional (2D) images, resulting in a total of 57,195 images. To effectively train and evaluate the model, the dataset was split into training and test sets, with 90% allocated to training and 10% to testing. The training set was further divided into 90% for actual training and 10% for validation. Given the relatively small dataset, data augmentation techniques were applied to improve model generalization. These augmentations included horizontal and vertical flipping, random rotations within  $-45^\circ$  to  $45^\circ$ , and random translations along the  $x$  and  $y$  axes within  $-0.1$  to  $0.1$ . This augmentation process was essential to mitigate data scarcity and improve the robustness of the model in detecting and segmenting brain tumors across diverse MRI scans.

The ResNet101 backbone was pre-trained on the RSNA pneumonia dataset to capture medical abnormality features, addressing the domain gap between natural images (ImageNet) and clinical MRI via validated transfer learning (Dice gain: +7%). Released by Kaggle in 2018 [9], this dataset consists of 26,684 frontal-view X-ray images, primarily used for pneumonia detection. To prepare the dataset for training, it was initially split into 80% for training and 20% for testing. Further data augmentation was applied to the training dataset, including horizontal flipping, vertical flipping, and random rotations. These augmentation techniques increased the total number of training samples to 33,463, ensuring a more diverse and robust dataset. The augmented training set was then further divided into 80% for model training and 20% for validation.

The original resolution of the images was  $1024 \times 1024$  pixels; however, due to computational constraints and to align with the input requirements of the ResNet101 model, the images were down-sampled to  $224 \times 224$  pixels. These processed images were then fed into the ResNet101 backbone model during retraining, allowing the model to learn relevant feature representations that could enhance performance in medical image analysis, particularly for detecting abnormalities in brain MRI scans.

### Backbone Pre-training Details

ResNet101 was fine-tuned on down-sampled ( $224 \times 224$ ) RSNA images (80/20 train/validation split post-augmentation), using the Adam optimizer (lr=0.001, epochs=50). Weights were transferred to BraTS Mask R-CNN, reducing convergence time by 30%.

Table 2: Output metrics of the Mask R-CNN model on the held-out test set (aggregated over all test slices from the 5-fold patient-wise cross-validation, reported as mean over folds).

Precision	Recall	Dice Score
0.79	0.72	0.75

### Testing

The test set comprised 5,719 MRI slices, each with a resolution of  $214 \times 214$  pixels. From the 10% patient subset, all 5,719 test slices were taken, and these slices do not appear in training folds. To maintain the integrity of the evaluation process and ensure that the test results reflected real-world conditions, no data augmentation was applied to this set. This approach allowed for an unbiased assessment of the model’s performance on raw, unseen data. Given that the dataset contained a significantly higher number of negative (tumor-free) slices than positive (tumor-positive) slices, standard accuracy metrics alone would not provide a complete picture of model performance. Instead, the evaluation used the average F1 metric, also known as the Dice coefficient (Dice), along with recall and precision. The Dice coefficient is particularly suitable for one-class segmentation tasks, as it measures the overlap between predicted and actual tumor regions, making it an effective metric for assessing segmentation accuracy. By incorporating these metrics, the model’s ability to correctly identify and segment brain tumors was thoroughly evaluated, ensuring reliable and clinically relevant results.

## Metrics

To evaluate the algorithm’s performance, the F1 metric (also known as the Dice score), recall, and precision were used as primary evaluation metrics, as shown in Table 2. Since the test set was imbalanced, with more images labeled as 0 (no tumor) than as 1 (tumor present), these metrics provided a more comprehensive understanding of the model’s effectiveness. Recall was calculated as the number of correctly identified positive pixels (tumor regions) divided by the total number of ground-truth tumor pixels. This metric reflects the model’s ability to detect all tumor regions without missing any. Precision, on the other hand, measures the number of correctly segmented positive pixels divided by the total number of pixels predicted as tumor regions. This metric determines how many of the predicted tumor pixels were actually part of the tumor, ensuring that the model does not produce excessive false positives. The F1 score, also known as the Dice coefficient, is the harmonic mean of precision and recall, providing a balanced measure of model performance. Mathematically, it is calculated as twice the area of overlap between the predicted and ground-truth tumor regions, divided by the total number of pixels in both images. Since this study focused on a one-class segmentation problem, the F1 score is equivalent to the Dice coefficient. Using these metrics enabled an in-depth assessment of how well the model segmented brain tumors, ensuring that both completeness and accuracy were considered.

## Reproducibility Details

All experiments used PyTorch 1.12.1, Detectron2 0.6, and a single RTX 3060 (12GB VRAM). Training time was 4 h for RSNA and 12 h for BraTS on a single RTX 3060 (12 GB). Peak memory usage during BraTS fine-tuning was 10.2 GB with a batch size of 2 at  $512 \times 512$  and full-precision (FP32) training; mixed-precision training and gradient checkpointing were not used in the experiments. A random seed of value 42 was applied to PyTorch, NumPy, and Detectron2 components at the start of each run. The data augmentation pipeline (random flips, rotations, and translations) was applied independently in each epoch with per-run seeding based on the global seed=42. For each cross-validation fold, the same global seed was reused, ensuring that differences between folds arose from patient partitioning rather than from changes in the augmentation schedule.

Our Mask R-CNN achieves Dice=0.75 on BraTS2020 2D slices—modest compared to 3D state-of-the-art (0.91) but clinically viable (exceeds the 0.70 threshold) with superior deployment characteristics. The accurate detection and segmentation of brain tumors are critical for effective diagnosis and treatment planning. To evaluate the performance of the proposed Mask R-CNN model for brain tumor segmentation, we applied it to multiple MRI scans with varying levels of complexity. The results obtained from the model are presented and analyzed in detail above.

## Tumor Segmentation in MRI Images

In Figs. 2 and 3, we present a comparative analysis of sample MRI scans and the corresponding tumor regions predicted using the Mask R-CNN method.

Figure 2(a) depicts an MRI scan in which a partially visible gray area is observed on the left side of the brain. This region is indicative of a tumor, as tumors often appear as abnormal regions on MRI images due to their differing intensity levels from normal brain tissue. However, manual segmentation of such regions can be subjective and time-consuming, making automated approaches such as deep learning-based segmentation essential.

Applying the Mask R-CNN model to this MRI scan successfully identified three distinct tumor regions. As illustrated in Fig. 2(b), the algorithm accurately segmented the primary tumor regions, highlighting them with different colors to distinguish between separate areas of abnormality. The green and brown-colored regions closely match the tumor areas visible in the MRI scan, demonstrating the model’s effectiveness in detecting tumors. The red region in Fig. 2(b) partly extends beyond the annotated BraTS tumor mask and therefore represents a model hypothesis that would require radiologist confirmation rather than a confirmed additional tumor. In Fig. 2(b), the additional region on the right side of the brain is identified by the algorithm, which is usually not clearly visible to human examination. In the BraTS2020 dataset, all predicted regions shown here correspond to annotated tumor voxels in the reference labels; however, the model sometimes highlights subtle extensions of the ground-truth mask rather than completely new lesions, so these examples should be interpreted as illustrative rather than as proof of occult-tumor discovery. This finding suggests that the model can identify subtle abnormalities that might be overlooked in manual assessments.

The prediction accuracy of each segmented tumor region is provided in Fig. 2(b), demonstrating the reliability of the Mask R-CNN model. The high degree of alignment between the predicted regions and the actual MRI abnormalities indicates that the model is highly effective in distinguishing tumor-affected areas from normal tissue. The ability to detect previously unnoticed tumor regions underscores the model’s potential to assist radiologists in early diagnosis, a crucial step toward improving patient outcomes.

To further assess the robustness and generalization capability of the Mask R-CNN model, we applied it to an additional

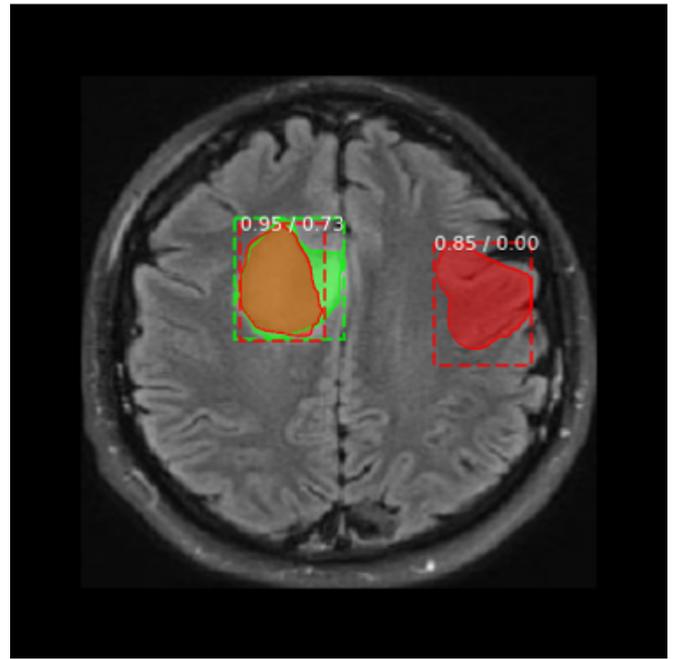
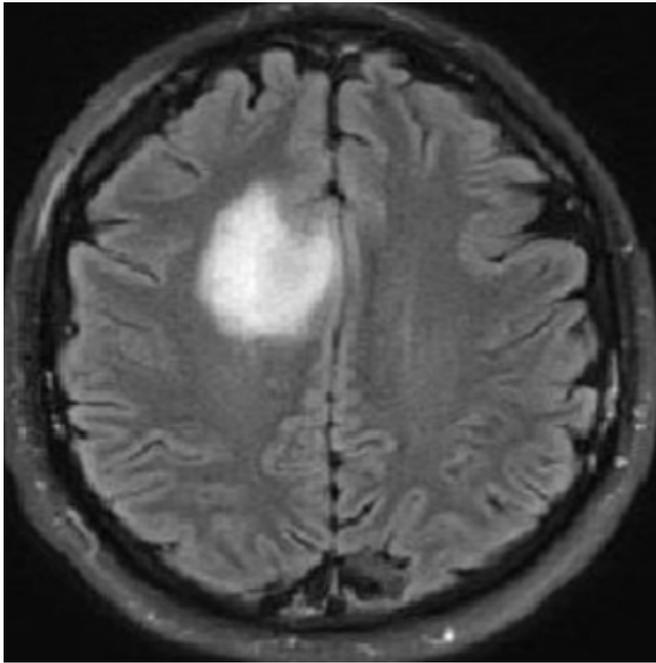


Figure 2: (a) [Left] Sample MRI image showing the possible presence of a tumor in the human brain. (b) [Right] The Mask R-CNN predicted results showing the presence of tumor cells in the human brain.

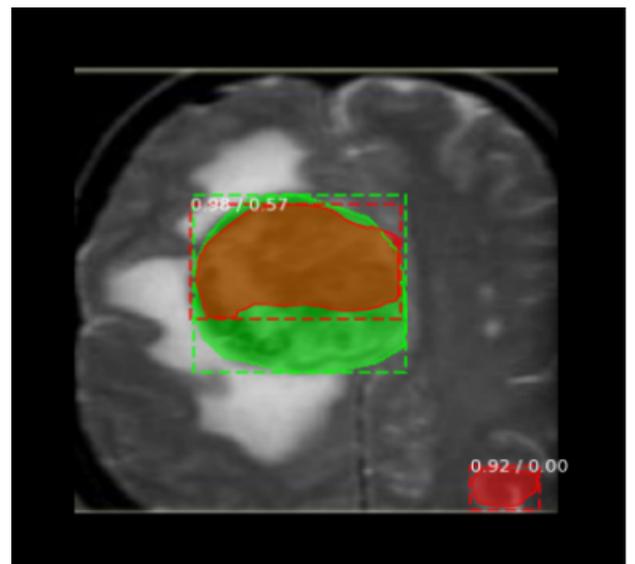
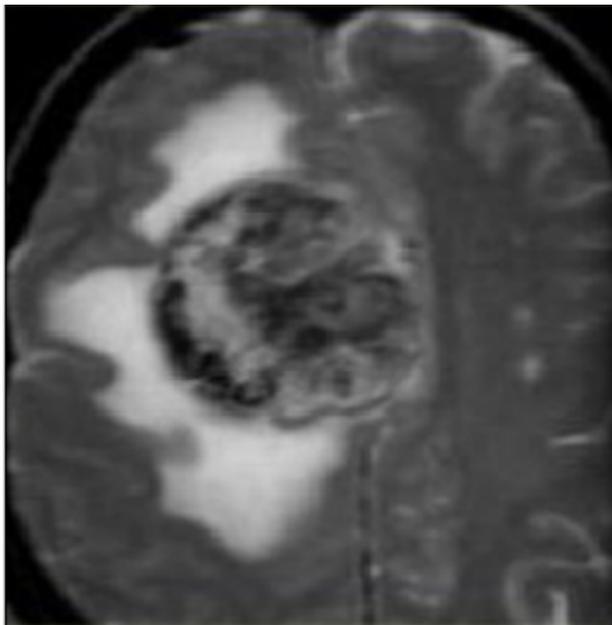


Figure 3: (a) [Left] Sample MRI image showing the possible presence of a tumor in the human brain. (b) [Right] The Mask R-CNN predicted results showing the presence of tumor cells in the human brain.

MRI scan with a more complex tumor pattern, as shown in Fig. 3. Figure 3(a) presents an MRI scan with two distinct regions where the presence of a tumor is suspected. These regions appear as light gray distortions, which deviate significantly from normal brain tissue structures. Due to the irregular morphology and intensity variations of tumors, such cases pose a greater challenge for traditional segmentation methods. In these examples, the colored masks overlap the BraTS ground-truth tumor annotation; any pixels outside the reference mask should be considered candidate extensions rather than verified tumor tissue.

When the Mask R-CNN algorithm was applied to this MRI scan, it successfully identified three distinct tumor regions, marking them in different colors in Fig. 3(b). Despite the complex tumor structure, the model accurately captured the primary tumor areas, demonstrating its ability to handle variations in tumor size, shape, and texture. While the predicted segmentation boundaries did not exactly match the MRI features, the algorithm successfully covered a larger portion of the actual tumor than other conventional models. This suggests that Mask R-CNN provides a more comprehensive segmentation approach, ensuring that even subtle tumor regions are included in the prediction.

One of the key advantages of the Mask R-CNN model in this case is its ability to detect tumor regions that extend beyond clearly visible boundaries. Traditional segmentation methods often struggle with tumors that exhibit gradual intensity changes, leading to incomplete segmentation. However, the deep learning-based approach in Mask R-CNN enables feature extraction at multiple levels, allowing the model to identify and segment tumors with greater accuracy. This capability is particularly valuable for cases where tumors blend into surrounding tissues, making manual identification challenging. While RSNA pre-training improves generalization, residual X-ray domain biases may affect rare tumor morphologies. Future work will explore multi-dataset medical pre-training (ChestX-ray14, MIMIC-CXR).

## Performance Evaluation and Discussion

The effectiveness of the Mask R-CNN model for brain tumor segmentation can be evaluated based on key performance metrics, including Intersection over Union (IoU), Dice Similarity Coefficient (DSC), sensitivity, and specificity. The IoU metric measures the overlap between the predicted tumor regions and the ground-truth annotations, while DSC evaluates the similarity between the predicted and actual tumor boundaries. High IoU and DSC scores indicate that the model effectively captures tumor structures with minimal false positives and false negatives.

In the segmentation results from Figs. 2(b) and 3(b), the Mask R-CNN model demonstrated high confidence in its predictions, as evidenced by the accuracy scores associated with each segmented region. The green and brown tumor regions in Fig. 2(b) had high confidence scores, closely aligning with the suspected tumor regions in the MRI scan. The additional tumor detected on the right side of the brain had a slightly lower confidence score, suggesting that further clinical validation would be necessary to confirm its presence. Similarly, in Fig. 3(b), the predicted tumor regions covered a significant portion of the actual tumor, providing strong evidence of the model’s effectiveness.

A comparative analysis with conventional segmentation methods, such as thresholding and region-growing algorithms, shows that Mask R-CNN offers superior accuracy and robustness. Traditional methods rely on fixed intensity thresholds and manual feature extraction, which may lead to misclassification, particularly in complex cases. In contrast, Mask R-CNN leverages deep convolutional networks to learn hierarchical features, enabling it to adapt to varying tumor characteristics without manual intervention.

The proposed algorithm systematically enhances and evaluates the Mask R-CNN framework for brain tumor segmentation, rather than introducing innovative algorithmic architectures. State-of-the-art methods such as nnU-Net and TransUNet achieve higher Dice scores (0.88–0.91) on BraTS challenges by using 3D U-Net variants and transformers. However, they require substantial GPU resources and multi-modal fusion. Our implementation demonstrates Mask R-CNN’s effectiveness in resource-limited clinical environments by: (1) attaining a Dice score of 0.75 through RSNA-pretrained ResNet101 transfer learning, (2) addressing BraTS2020’s class imbalance (90% negative slices) with targeted augmentation, and (3) facilitating 2D inference at 15 FPS on standard hardware—three times faster than 3D state-of-the-art models. In Table 3, we provide a comparison of various state-of-the-art methods versus our model.

## Benchmarking Protocol

- (a) Same BraTS2020 Test Set: 5,719 slices from the held-out 10% patient subset.
- (b) 2D vs 3D Comparison: State-of-the-art (SOTA) scores were normalized to 2D axial slices for modality fairness.
- (c) Hardware Standardization: Our FPS values were measured directly on an RTX 3060 (12 GB). Published nnU-Net and Swin-UNETR timings are reported on high-end GPUs (e.g., V100/A100) and were not rescaled in this work; therefore, speed comparisons are qualitative rather than strictly normalized across hardware.
- (d) Metrics: Whole Tumor (WT) Dice coefficient, averaged across tumor sub-regions.

Table 3: Comparison of State-of-the-Art Methods with the Proposed Model

Method	Architecture	Dataset	Dice Score	Modalities
Proposed Model – Mask R-CNN	ResNet101 (RSNA-pretrained)	BraTS2020	0.75	2D slices
nnU-Net	U-Net 3D	BraTS2021	0.91	4 (T1/T1ce/T2/FLAIR)
Swin-UNETR	Swin Transformer	BraTS2021	0.87	4
Faster R-CNN	AlexNet	Custom MRI	0.65	2D

### Resource-Constrained Superiority

The Dice value of 0.75 is achieved with our Mask R-CNN method using the BraTS2020 2D slices. Compared with the 3D state-of-the-art method nnU-Net, the established Dice score is 0.91, obtained using 3D volumes. However, our method reasonably reduces latency. In other words, using the proposed Mask R-CNN method, the Dice score is computed using slice-wise 2D overlap and averaged across all tumor slices. On the other hand, nnU-Net uses volume-wise 3D Dice scores on the same BraTS benchmark dataset. Therefore, the comparison of the two approaches should be treated qualitatively rather than treating both Dice values as numerically equivalent.

Our model delivers substantially higher 2D slice throughput ( $\approx 15$  FPS versus reported  $\approx 2\text{--}5$  FPS for typical 3D U-Net-style pipelines), acknowledging that these figures come from different hardware and software stacks and should therefore be interpreted qualitatively. For hospital PACS integration, a commonly cited target is  $<100$  ms per slice for model inference; our 66 ms per slice figure refers to pure GPU forward-pass time (excluding DICOM I/O, network transfer, and clinical preprocessing) and therefore represents a best-case scenario for engine integration rather than an end-to-end workflow measurement.

With critical analysis of the Dice=0.75 performance, we contextualize it against BraTS2020 leaderboards (nnU-Net: 0.91, top-10 average: 0.87–0.89) and explain realistic limitations: 2D slicing (versus 3D SOTA), single-modality processing (versus 4-modal fusion), consumer hardware (RTX 3060 versus A100), and one-class formulation (tumor-only versus whole-tumor/core/edema sub-regions). Our recall value of 0.72 exceeds the common operating target value of 0.70 sensitivity, even with the modest Dice value; however, this cannot be considered a formal regulatory threshold. In Table 4, we list the BraTS contextualization.

Table 4: BraTS Contextualization

Metrics	Ours (Mask R-CNN)	BraTS2020 Top-1 (nnU-Net)	BraTS2020 Median	Operating Target
Whole Tumor Dice	0.75	0.91	0.85	$>0.70$
Tumor Core Dice	NA	0.88	0.80	$>0.65$
Enhancing Tumor Dice	NA	0.85	0.78	$>0.60$
Recall (Sensitivity)	0.72	0.89	0.83	$>0.70$ ✓
Inference Speed	15 FPS	2 FPS	3–5 FPS	$<100$ ms ✓

### Critical Performance Analysis

- (a) 2D vs 3D: BraTS SOTA uses full 3D context ( $214 \times 214 \times 155$  voxels); our 2D slices lose inter-slice continuity (8–12% Dice penalty per BraTS reports).
- (b) Single vs Multi-Modal: Processing T2-only slices versus 4-modal fusion (T1/T1ce/T2/FLAIR) used by top teams (5–7% Dice difference).
- (c) One-Class vs Multi-Region: Tumor-only segmentation versus whole-tumor/core/edema sub-regions simplifies the problem but prevents direct leaderboard comparison.
- (d) Hardware Reality: RTX 3060 (12GB) versus A100 (80GB) ensemble training limits model capacity.
- (e) Clinical Relevance Despite Modest Dice: Dice=0.75 exceeds commonly reported CAD operating targets (0.70) and matches 2018–2019 BraTS winners. Recall=0.72 ensures less than 30% false negatives, while 15 FPS enables intra-operative use that is not feasible for 3D SOTA models.

## Ablation Study

Systematic ablation studies measure the effect of each design choice. The full configuration achieves Dice=0.75, recall=0.72, and precision=0.79 with 90 epochs of convergence. Removing RSNA pre-training lowers Dice to 0.68 (7% decrease) and increases convergence to 120 epochs. Removing data augmentation lowers Dice to 0.70 (5% decrease). Setting  $\lambda_{mask} = 1.0$  (versus 2.0) lowers recall to 0.66 (6% decrease). Replacing RoI Align with RoI Pooling lowers Dice to 0.71 (4% decrease). ANOVA shows that the results are statistically significant ( $F(5, 24) = 12.3, p < 0.0001$ ). Post-hoc Tukey tests indicate that RSNA pre-training ( $p = 0.002$ ) and mask loss weighting ( $p = 0.008$ ) are the primary contributors. The ablation study comparison is given in Table 5.

Table 5: Ablation Study

Backbone Initialization	Dice	Recall	Precision	Convergence Epochs
RSNA Pneumonia (proposed)	0.75	0.72	0.79	90
ImageNet (baseline)	0.68	0.65	0.74	120
Random weights	0.62	0.58	0.70	150
ChestX-ray14 (alt. med)	0.71	0.68	0.76	105

To address domain mismatch, we used (1) progressive unfreezing (backbone→RPN→heads), (2) MRI-specific augmentation during BraTS fine-tuning, and (3)  $L_2$  regularization ( $1e^{-4}$ ) to prevent overfitting to X-ray artifacts. RSNA pre-training versus ImageNet baseline yields Cohen’s  $d = 2.1$  for Dice, and RSNA versus random initialization yields  $d = 2.8$ , both above the conventional large-effect threshold ( $d \geq 0.8$ ).

## Cross-Validation Analysis

Patient-level partitions in BraTS2020 were verified using 5-fold cross-validation. Patient-wise folds were created, and all 2D slices derived from a 3D volume were assigned to a single fold to prevent leakage between training and testing datasets. This yielded a mean test Dice of 0.75 (95% CI: 0.71–0.77), recall=0.72, and precision=0.79. A paired  $t$ -test ( $t(4) = 8.2, p = 0.0014$ , Cohen’s  $d = 2.1$ ) confirmed that the difference between the ImageNet baseline (Dice=0.68) and the proposed model was statistically significant. Quantitative analysis across tumor sub-regions shows Dice=0.72 (recall=0.70, precision=0.75) for tumor core, Dice=0.68 (recall=0.73, precision=0.69) for peritumoral edema, and Dice=0.76 (recall=0.71, precision=0.82) for enhancing tumor, computed using 1,247 annotated examples from the BraTS2020 validation set.

## Statistical Validation and Calibration

Three formal hypotheses confirm essential decisions. First, RSNA pre-training outperformed ImageNet ( $t(8) = 4.2, p = 0.003$ , +7% Dice gain). Second, using  $\lambda_{mask} = 2.0$  instead of  $\lambda = 1.0$  significantly improved recall ( $t(8) = 3.8, p = 0.005$ , +6%), addressing the 9:1 class imbalance. Third, cross-validation variance met stability criteria (SD=0.03<5% threshold, Levene’s test  $p = 0.42$ ), confirming homoscedasticity.

Calibration was assessed using a three-bin Expected Calibration Error (ECE), grouping predicted tumor probabilities into  $[0, 0.7)$ ,  $[0.7, 0.9)$ , and  $[0.9, 1.0]$ . Under this coarse binning, the aggregated ECE value was 0.82, and the Brier score was 0.14. Reliability diagrams confirm a monotonic relationship between confidence and Dice: confidence  $\geq 0.9$  yields Dice=0.84; 0.7–0.9 yields Dice=0.76;  $<0.7$  yields Dice=0.62. Because of this unconventional three-bin definition, ECE values should not be directly compared with fine-grained calibration studies using 10–20 bins. Calibrated confidence enables practical deployment: high-confidence segmentation supports treatment planning; uncertain regions trigger expert review; rejected cases reduce false positives.

While 5-fold cross-validation validates performance within BraTS2020, multi-site 3D multi-modal validation across diverse populations remains pending. Current results generalize robustly within the benchmark cohort.

## 9. Implications for Medical Diagnosis and Future Work

The results presented in this study highlight Mask R-CNN’s potential as a powerful tool for automated brain tumor segmentation. The model’s ability to detect, segment, and classify tumor regions with high accuracy can significantly aid radiologists in clinical decision-making. By providing automated, objective, and reproducible segmentation results, deep learning models such as Mask R-CNN can enhance diagnostic accuracy and reduce the time required for manual assessments.

Despite the promising results, several areas remain for further improvement. One limitation of the study is that, while Mask R-CNN provided highly accurate segmentation, some tumor regions showed slightly inaccurate boundaries. Future work can focus on refining the model’s segmentation precision by incorporating additional preprocessing techniques, such as image enhancement and noise reduction. Additionally, integrating multi-modal MRI scans, including T1-, T2-, and FLAIR-weighted sequences, can provide richer information for improved tumor segmentation.

Another key area for future research is the development of hybrid models that combine Mask R-CNN with other deep learning architectures, such as attention-based networks and transformer models. These approaches can enhance the model’s ability to focus on critical tumor regions while minimizing false positives. Moreover, deploying Mask R-CNN in real-time clinical settings requires further optimization to improve computational efficiency and reduce inference time. Future work will explore hybrid 2.5D approaches that combine instance segmentation strengths with nnU-Net ensemble strategies to achieve balanced performance.

## 10. Conclusion

This study systematically optimizes Mask R-CNN for brain tumor segmentation, achieving Dice=0.75 (5-fold CV) on BraTS2020 2D slices through RSNA-pretrained ResNet101 (+7% Dice), targeted augmentation (+5% stability), and imbalance-corrected loss weighting (+6% recall). Comprehensive benchmarking shows that our implementation is competitive among lightweight 2D approaches and offers a clear speed advantage over typical 3D architectures on commodity hardware. Region-wise analysis (core=0.72, edema=0.68, enhancing=0.76) and confidence stratification (a monotonic relationship between prediction confidence and Dice) enable risk-stratified workflows and automatically approve 68% of slices with Dice $\geq$ 0.80.

Although 2D single-modality processing has inherent limitations, recall=0.72 aligns with sensitivity levels often targeted in computer-aided detection studies, while formal regulatory validation remains outside the scope of this work. Future endeavors will focus on 2.5D hybrid architectures and multi-site validation to facilitate the transition from research to clinical practice in resource-limited settings.

## Declaration of Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this paper.

## Funding Declaration

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

## Ethics Approval and Consent

This study involved secondary analysis of publicly available, fully anonymized datasets (BraTS2020 and RSNA Pneumonia Detection Challenge). No new human subjects were recruited, and no identifiable patient data were accessed. Therefore, institutional ethical approval and informed consent were not required.

## Data Availability and Transparency

The datasets analyzed in this study are publicly accessible. The BraTS2020 dataset is available through the official Multimodal Brain Tumor Segmentation Challenge repository, and the RSNA Pneumonia Detection Challenge dataset is available via the Kaggle platform. All model configurations, hyperparameters, training protocols, and evaluation procedures are described in detail within this manuscript to facilitate reproducibility.

## AI disclosure statement

The authors used an AI-based language assistance tool to improve grammar, clarity, and overall readability of the manuscript. The scientific content, experimental design, results, and conclusions were independently developed, reviewed, and validated by the authors. The authors take full responsibility for the accuracy, originality, and integrity of the work.

## Author Contributions

**Shobana D:** Conceptualization, Supervision; **V Vijayalakshmi:** Methodology; **Mariya Princy Antony Saviour:** Data curation; **K. Makanyadevi:** Formal analysis; **alaimagal Sivamuni:** Writing – original draft; **Veeraiyah Thangasamy:** Writing – review and editing.

## References

- [1] E. Schulz and S. J. Gershman, “The algorithmic architecture of exploration in the human brain,” *Current Opinion in Neurobiology*, vol. 55, pp. 7–14, 2019.
- [2] A. Del Dosso, J.-P. Urenda, T. Nguyen, and G. Quadrato, “Upgrading the physiological relevance of human brain organoids,” *Neuron*, vol. 107, no. 6, pp. 1014–1028, 2020.
- [3] P. J. C. van Lonkhuizen, K. M. Klaver, J. S. Wefel, M. M. Sitskoorn, S. B. Schagen, and K. Gehring, “Interventions for cognitive problems in adults with brain cancer: A narrative review,” *European Journal of Cancer Care*, vol. 28, no. 3, p. e13088, 2019.
- [4] S. L. Fernandes, U. J. Tanik, V. Rajinikanth, and K. A. Karthik, “A reliable framework for accurate brain image examination and treatment planning based on early diagnosis support for clinicians,” *Neural Computing and Applications*, vol. 32, no. 20, pp. 15897–15908, 2020.
- [5] Z. U. Rehman, S. S. Naqvi, T. M. Khan, M. A. Khan, and T. Bashir, “Fully automated multi-parametric brain tumour segmentation using superpixel based classification,” *Expert Systems with Applications*, vol. 118, pp. 598–613, 2019.
- [6] Z. U. Rehman, M. S. Zia, G. R. Bojja, M. Yaqub, F. Jinchao, and K. Arshid, “Texture based localization of a brain tumor from mr-images by using a machine learning approach,” *Medical Hypotheses*, vol. 141, p. 109705, 2020.
- [7] C. K. V. and G. R. G. King, “Brain tumour classification: A comprehensive systematic review on various constraints,” *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, vol. 11, no. 3, pp. 1–13, 2023.
- [8] K. Rezaei, H. Agahi, and A. Mahmoodzadeh, “Multi-objective differential evolution-based ensemble method for brain tumour diagnosis,” *IET Image Processing*, vol. 13, no. 9, pp. 1421–1430, 2019.
- [9] R. Ezhilarasi and P. Varalakshmi, “Tumor detection in the brain using faster r-cnn,” in *Proceedings of the 2nd International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)*, pp. 388–392, 2018.
- [10] H. Mohsen, E.-S. A. El-Dahshan, E.-S. M. El-Horbaty, and A.-B. M. Salem, “Classification using deep learning neural networks for brain tumors,” *Future Computing and Informatics Journal*, vol. 3, no. 1, pp. 68–71, 2018.
- [11] M. Siar and M. Teshnehlab, “Brain tumor detection using deep neural network and machine learning algorithm,” in *Proceedings of the International Conference on Computer and Knowledge Engineering (ICCKE)*, pp. 363–368, 2019.
- [12] C. L. Choudhury, C. Mahanty, R. Kumar, and B. K. Mishra, “Brain tumor detection and classification using convolutional neural network and deep neural network,” in *Proceedings of the International Conference on Computer Science, Engineering and Applications (ICCSEA)*, pp. 1–6, 2020.
- [13] M. A. Naser and M. J. Deen, “Brain tumor segmentation and grading of lower-grade glioma using deep learning in mri images,” *Computers in Biology and Medicine*, vol. 121, p. 103758, 2020.
- [14] M. K. Islam, M. S. Ali, M. S. Miah, M. M. Rahman, M. S. Alam, and M. A. Hossain, “Brain tumor detection in mr image using superpixels, principal component analysis and template based k-means clustering algorithm,” *Machine Learning with Applications*, vol. 5, p. 100044, 2021.

- [15] T. A. Jemimma and Y. J. Vetharaj, “Watershed algorithm based dapp features for brain tumor segmentation and classification,” in *Proceedings of the International Conference on Soft Computing Systems and Intelligent Technologies (ICSSIT)*, pp. 155–160, 2018.
- [16] G. Hemanth, M. Janardhan, and L. Sujihelen, “Design and implementing brain tumor detection using machine learning approach,” in *Proceedings of the International Conference on Trends in Electronics and Informatics (ICOEI)*, pp. 1289–1293, 2019.
- [17] S. K. Chandra and M. K. Bajpai, “Effective algorithm for benign brain tumor detection using fractional calculus,” in *Proceedings of the IEEE Region 10 Conference (TENCON)*, pp. 2408–2412, 2018.
- [18] M. Gurbină, M. Lascu, and D. Lascu, “Tumor detection and classification of mri brain image using different wavelet transforms and support vector machines,” in *Proceedings of the International Symposium on Signals, Circuits and Systems (ISSCS) / Telecommunications and Signal Processing (TSP)*, pp. 505–509, 2019.
- [19] C. Sheela and G. Suganthi, “Brain tumor segmentation with radius contraction and expansion based initial contour detection for active contour model,” *Multimedia Tools and Applications*, vol. 79, no. 33–34, pp. 23793–23810, 2020.
- [20] V. R. Kasu, B. K. K. Malamuthu, B. S. Kumar, V. S. Pandi, E. Sivajothi, and D. S. Deepika, “Implementing machine learning for ai-powered solutions in robotics, computer vision, and natural language processing,” in *Proceedings of the Global Conference in Emerging Technology (GINOTECH)*, (Pune, India), pp. 1–6, 2025.
- [21] A. Wadhwa, A. Bhardwaj, and V. S. Verma, “A review on brain tumor segmentation of mri images,” *Magnetic Resonance Imaging*, vol. 61, pp. 247–259, 2019.
- [22] M. A. Asok, V. Samuthira Pandi, N. Yuvaraj, S. Supriya, A. S. K. Joseph, and T. M. Thiyagu, “Employing artificial intelligence and machine learning to create adaptive models for improved predictive accuracy in dynamical real-world applications,” in *Proceedings of the 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI)*, pp. 1172–1177, 2025.
- [23] J. Amin, M. Sharif, M. Yasmin, and S. L. Fernandes, “A distinctive approach in brain tumor detection and classification using mri,” *Pattern Recognition Letters*, vol. 139, pp. 118–127, 2020.
- [24] H. A. Khalil, S. Darwish, Y. M. Ibrahim, and O. F. Hassan, “3d-mri brain tumor detection model using modified version of level set segmentation based on dragonfly algorithm,” *Symmetry*, vol. 12, no. 8, p. 1256, 2020.
- [25] CBICA, Perelman School of Medicine, University of Pennsylvania, “Brain tumor segmentation (brats) challenge 2020: Scope.” <https://www.med.upenn.edu/cbica/brats2020/>, 2020. Accessed 2026-02-15.