

Brain Tumor Segmentation: U-Net Architecture for MRI Precision Diagnostics

M. Meena

Department of Information Technology
Vivekanandha College of Technology
for Women
Tiruchengode, Namakkal, India,
meenamani91@gmail.com

M. Nivetha

Department of Information Technology
Vivekanandha College of Technology
for Women,
Tiruchengode, Namakkal, India,
nivethamathiyalagan20@gmail.com

S. Praseetha

Department of Information Technology
Vivekanandha College of Technology
for Women
Tiruchengode, Namakkal, India,
praseethasaravanakumar@gmail.com

M. Soniya

Department of Information Technology
Vivekanandha College of Technology for Women,
Tiruchengode, Namakkal, India,
soniya8643@gmail.com

S. Vasanthi

Department of Information Technology
Vivekanandha College of Technology for Women,
Tiruchengode, Namakkal, India,
vasanthivasanthi37786@gmail.com

Abstract : Brain tumor segmentation is a critical step in the diagnosis and treatment planning for patients. This study explores the efficacy of the U-Net architecture, a deep learning model specifically designed for biomedical image segmentation, in accurately distinguishing tumor regions in MRI scans. We utilized a dataset comprising diverse MRI images annotated by specialists, ensuring comprehensive training for the U-Net model. Our findings demonstrated that the U-Net architecture significantly improves segmentation accuracy and enhances the visualization of tumor margins compared to conventional methods. Furthermore, we analyzed the model's performance across varying tumor types and sizes, revealing its robustness and adaptability in clinical applications. By integrating advanced image analysis techniques into MRI diagnostics, this research offers a promising approach to support radiologists in making informed decisions, ultimately leading to improved patient outcomes. Future work will focus on refining the model and exploring its integration into routine clinical workflows.

Keywords— Brain Tumour segmentation, U-Net Architecture, Deep Learning, MRI Scans, Bio-Medical image segmentation;

I. INTRODUCTION

Brain tumors are among the most life-threatening neurological disorders, and accurate localization of tumor boundaries is critical for diagnosis, prognosis, and treatment (e.g., surgery, radiotherapy, and chemotherapy). Manual segmentation by radiologists is time-consuming, subject to inter-observer variability, and becomes impractical for large-scale clinical workloads, especially when dealing with high-resolution, multi-modal MRI scans. Automated segmentation using deep learning can reduce human error, standardize measurements, and accelerate the workflow while preserving spatial fidelity to the underlying anatomy. This shift from manual to algorithm-assisted analysis underpins the growing role of artificial intelligence in neuro-oncology and precision medicine [1].

1.1 Role of MRI in brain tumor imaging

Magnetic Resonance Imaging (MRI) is the modality of choice for brain tumor assessment because it provides excellent soft-tissue contrast, multi-planar imaging, and multiple sequences (e.g., T1, T2, FLAIR, T1-post-contrast) that highlight different tumor components. These sequences help distinguish edema, necrotic regions, enhancing tumor, and healthy brain tissues, which are essential for accurate segmentation and treatment planning. However, MRI scans are often noisy, exhibit intensity inhomogeneity, and show substantial variation in tumor size, shape, and location, making automated segmentation challenging even with advanced image-processing techniques[2].

1.2 Traditional segmentation methods and their limitations

Before deep learning, brain tumor segmentation relied on classical methods such as thresholding, region-growing, clustering (e.g., k-means, FCM), and edge-based or atlas-based approaches. These methods often require hand-crafted features, strong assumptions about intensity distributions, and careful parameter tuning, which limits their robustness across heterogeneous datasets and clinical sites. Moreover, traditional techniques struggle with irregular tumor boundaries, partial-volume effects, and artifacts, leading to under- or over-segmentation and reduced reproducibility. These limitations have motivated the transition to data-driven, end-to-end deep-learning models that can learn complex patterns directly from labeled MRI data.

1.3 Introduction to U-Net architecture

U-Net is a convolutional neural network (CNN) architecture originally designed for biomedical image segmentation, characterized by an encoder-decoder (contracting-expanding) structure shaped like a "U." The encoder path applies successive convolutions and pooling layers to extract hierarchical features and capture contextual information, while the decoder path upsamples the feature maps and refines the spatial details to produce a pixel-wise segmentation mask[3]. A key innovation of U-Net is its symmetric skip connections between corresponding encoder and decoder levels, which allow the model to combine low-level spatial details with high-level semantic features, thereby improving boundary localization and preserving fine-scale structures such as tumor irregularities and small lesions.

1.4 U-Net for brain tumor segmentation in MRI

Due to its strong performance on small-scale medical datasets and its ability to work with limited annotated samples, U-Net has been widely adopted for brain tumor segmentation in 2D and 3D MRI scans. Several variants have been proposed, including 2D U-Net, 3D U-Net, and lightweight or hybrid architectures that integrate backbone networks such as ResNet50 or VGG19 as encoders to enhance feature extraction. These models are typically trained on benchmark datasets like BraTS (Brain Tumor Segmentation Challenge), where MRI volumes are annotated with ground-truth masks for whole tumor, tumor core, and enhancing tumor regions. Experimental results show that U-Net-based pipelines achieve high Dice scores, Intersection-over-Union (IoU), and accuracy, demonstrating their capacity to segment complex brain tumors with high precision and consistency[1].

1.5 Why U-Net suits MRI precision diagnostics

U-Net's architecture naturally aligns with the requirements of MRI-based precision diagnostics: it preserves spatial resolution through skip connections, handles multi-modal inputs, and supports pixel-wise probabilistic outputs that can be interpreted by clinicians. By integrating multiple MRI sequences into a single network, U-Net can capture complementary information about tumor extent, edema, and enhancement, enabling more accurate and clinically meaningful segmentation maps. Additionally, recent variants introduce techniques such as multi-scale inception blocks, attention mechanisms, and specialized loss functions (e.g., Dice loss, focal loss) to address class imbalance and improve performance on small or diffuse tumor regions. These enhancements make U-Net-based models promising candidates for deployment in time-critical radiology workflows and computer-aided diagnostic systems[4].

1.6 Challenges and clinical implications

Despite its success, U-Net-based brain tumor segmentation still faces challenges including limited annotated data, inter-institutional variability in imaging protocols, computational demands of 3D models, and the need for model interpretability in clinical settings. Addressing these issues requires strategies such as domain adaptation, data augmentation, lightweight architectures, and explainable AI techniques that build trust among radiologists and neuro-oncologists. When integrated into hospital workflows, U-Net-driven segmentation tools can support quantitative tumor volumetry, treatment-response monitoring, and pre-surgical planning, ultimately improving diagnostic accuracy, reducing clinician workload, and enabling more personalized therapeutic strategies for brain tumor patients.

II. LITERATURE REVIEW

Deep learning has revolutionized brain tumor analysis in MRI by enabling automated, end-to-end segmentation that surpasses traditional image-processing methods in accuracy and consistency. Recent comprehensive reviews highlight that convolutional neural network (CNNs), especially encoder-decoder architectures, now dominate brain-tumor segmentation benchmarks such as the BraTS challenge. These models extract hierarchical features from multi-modal MRI sequences (T1, T1-post-contrast, T2, FLAIR) and produce pixel-wise

segmentation masks for whole tumour, tumour core, and enhancing tumour regions.

U-Net, introduced for biomedical image segmentation, rapidly became the backbone for brain-tumor segmentation due to its encoder-decoder structure with skip connections that preserve spatial detail while capturing semantic context. Studies reviewing medical-image segmentation report that U-Net variants consistently achieve high Dice scores and robust performance even on small, heterogeneous datasets, making them suitable for clinical MRI applications. Several surveys on brain-tumor segmentation list U-Net as one of the most widely adopted architectures, often serving as the baseline against which newer models (e.g., 3D U-Net, transformer-based or attention-enhanced U-Net) are compared.

Early work demonstrated that vanilla 2D and 3D U-Net can segment brain tumors from MRI with high accuracy, especially when trained on standardized datasets like BraTS. For example, 3D U-Net has been shown to capture volumetric context effectively, yielding strong Dice scores for whole-tumor and enhancing-tumor regions. Later works replaced or enhanced the U-Net encoder with pre-trained CNNs such as ResNet50 and VGG19, leveraging their strong feature-extraction capabilities to improve segmentation metrics. One study combining U-Net with ResNet50 reported Dice losses near 0.0087, IoU above 0.75, and accuracy around 0.99, outperforming traditional edge-based methods. Another study using VGG19 as the encoder within a U-Net framework achieved improved precision by combining transfer learning with a tailored Focal Tversky loss to handle class imbalance[5].

To address computational cost and deployment constraints, lightweight U-Net-inspired architectures have been proposed. For instance, LIU-Net (Lightweight Inception U-Net) integrates Inception-style blocks into the U-Net encoder to capture multi-scale features while reducing model complexity. Evaluated on BraTS 2021 and 2020, LIU-Net achieved Dice scores of about 0.81–0.84 for enhancing tumor, 0.88–0.90 for whole tumor, and 0.84–0.91 for tumor core, demonstrating that computational efficiency does not necessarily sacrifice segmentation quality. Similarly, another study combined U-Net with EfficientNetV2 as the encoder, reporting a Dice similarity coefficient of about 0.91 and accuracy near 0.9977 on brain-tumor MRI, highlighting that modern, efficient backbones can boost segmentation performance while remaining practical for clinical use[6].

To better capture complex tumor shapes and subtle boundaries, researchers have integrated attention mechanisms into U-Net. ACU-Net, an attention-based U-Net, employs spatial and channel attention modules in parallel across the encoder-decoder path and enhances skip connections to emphasize relevant tumor regions. When tested on BraTS 2018, ACU-Net achieved Dice scores of approximately 94% for whole tumor, 98.6% for tumor core, and 98.8% for enhancing tumor, significantly outperforming earlier U-Net variants on these sub-regions.

Another study introduced a U-Net-self-attention hybrid model for 3D MRI glioma segmentation on BraTS 2020, obtaining an accuracy of 99.34% and Dice of 95%, with very high precision

and specificity, underscoring the benefit of attention modules in handling tumor heterogeneity and boundary ambiguity.

Recent comparative reviews of deep-learning-based brain-tumor segmentation note that U-Net and its derivatives remain among the top-performing models on BraTS-style benchmarks. A 2025 review of DL for brain-tumor analysis in MRI reports that U-Net-based pipelines often achieve Dice scores above 0.85–0.90 for whole-tumor segmentation and competitive performance for tumor core and enhancing tumor, especially when combined with advanced loss functions and multi-modal inputs. Other works emphasize that while newer architectures (e.g., transformers and hybrid CNN-transformer models) have emerged, they often either build on U-Net’s encoder–decoder structure or are compared directly against U-Net baselines, confirming U-Net’s enduring role as a reference method in MRI brain-tumor segmentation literature.

Despite high performance, U-Net-based models still face challenges such as limited annotated data, class imbalance (tiny enhancing-tumor regions versus large background), inter-institution variability in MRI protocols, and the need for 3D processing that increases memory and computation demands. Several recent contributions aim to mitigate these by combining U-Net with techniques such as data augmentation, domain adaptation, focal/Dice-based loss functions, and graph-cut or energy-minimization methods for post-processing. Additionally, there is growing interest in reducing the number of MRI sequences required for accurate segmentation; one 2025 study showed that DL models, including U-Net-based designs, can achieve comparable performance using only T1-post-contrast and FLAIR sequences instead of the full four-sequence protocol, which may lower scanning costs and patient burden.

III. METHODOLOGY

3.1 Problem definition and objectives

The study focuses on automatic brain tumor segmentation from multi-modal MRI scans using a U-Net–based deep-learning architecture to support MRI-based precision diagnostics. The main objective is to develop a robust model that can accurately delineate tumor regions (whole tumor, tumor core, and enhancing tumor) at the pixel level, while achieving high segmentation scores such as Dice, IoU, and accuracy on standard benchmark datasets like BraTS. The methodology also aims to design a reproducible end-to-end pipeline that can be extended to clinical workflows for computer-aided diagnosis.

3.2 Dataset selection and preparation

The methodology uses publicly available MRI datasets such as the BraTS brain tumor segmentation challenge, which provides 3D multimodal MRI volumes (T1, T1-post-contrast, T2, and FLAIR) along with expert-annotated segmentation masks. The dataset is divided into training, validation, and test subsets using a patient-wise split (for example, 70%–15%–15%) to avoid data leakage and ensure realistic evaluation. The 3D MRI volumes are either processed slice-by-slice in 2D or handled as 3D patches, depending on whether a 2D U-Net or 3D U-Net architecture is adopted.

3.3 Data preprocessing

Before training, the MRI data undergo preprocessing to enhance quality and consistency. Brain extraction (skull-stripping) is performed to remove non-brain regions and focus only on relevant tissue. Intensity normalization (such as Z-score or min-max scaling) is applied independently to each MRI sequence to bring pixel values into a uniform range and reduce inter-scanner variability. Additionally, data augmentation techniques including rotation, flipping, scaling, elastic deformation, and random contrast/brightness changes are applied on the training data to increase diversity, reduce overfitting, and improve model generalization[7].

3.4. U-Net network design

The core of the methodology is a U-Net architecture, chosen for its encoder–decoder structure with skip connections that preserve both high-level semantic information and fine-scale spatial details. The encoder path consists of repeated convolutional blocks followed by down-sampling to extract hierarchical features, while the decoder path up-samples the feature maps and fuses them with corresponding encoder features via skip connections to refine the segmentation mask. The final layer uses a 1×1 convolution with sigmoid or softmax activation to produce a pixel-wise probability map for binary or multi-class brain tumor segmentation, optionally enhanced with attention mechanisms or pre-trained encoders if applicable.

3.5 Loss function and optimization

To handle class imbalance and improve segmentation accuracy, the model is trained using a combination of Dice loss and cross-entropy loss, sometimes extended to focal or Dice-based variants such as Focal Tversky loss. This composite loss function encourages the model to pay more attention to small and poorly segmented tumor regions. The optimization is performed using the Adam optimizer with a small learning rate (e.g., 0.0001–0.001) and weight decay regularization, and the learning rate can be decayed or adjusted based on validation performance to ensure stable convergence.

3.6 Training procedure

The training procedure involves feeding preprocessed MRI slices or 3D patches into the U-Net model in mini-batches, performing forward and backward passes, and updating the network weights iteratively. The model is initialized with random weights or a pre-trained encoder (if transfer learning is used), and the loss is computed between the predicted segmentation mask and the ground-truth annotation. Training runs for a fixed number of epochs (for example, 50–200), and the best model is saved based on the validation Dice score or loss, optionally using early stopping when no improvement is observed for a predefined number of epochs.

3.7 Inference and post-processing

After training, the model performs inference on the test MRI scans by generating segmentation masks for each slice or 3D volume. Predicted masks may undergo light post-processing, such as morphological opening and closing operations, to remove isolated noise pixels and smooth the tumor boundary. In 2D slice-based

approaches, the final 3D segmentation is reconstructed by stacking the inferred 2D masks, preserving the volumetric context needed for clinical interpretation and volume estimation.

3.8 Evaluation metrics

The performance of the U-Net model is quantitatively evaluated using standard segmentation metrics computed on the test set. The primary metric is the Dice Similarity Coefficient (Dice), which measures the overlap between predicted and ground-truth masks, along with Intersection-over-Union (IoU), accuracy, sensitivity (recall), specificity, and precision per tumor class. These metrics are used to compare the proposed U-Net-based approach against baseline methods such as classical image-processing techniques or vanilla CNN frameworks, demonstrating its improvement in MRI-based tumor segmentation.




Figure 1 : Sample images

3.9 Analysis and validation for precision diagnostics

The methodology includes qualitative and clinical-oriented analysis to validate the model’s usefulness in precision diagnostics. Visual overlays of predicted tumor masks on original MRI slices are generated to inspect boundary accuracy and coverage. From the segmented masks, tumor volume and growth metrics are computed to show how the model can support quantitative assessment for treatment planning and follow-up. The robustness of the model is also evaluated by testing it on different MRI sequence combinations (e.g., T1-post-contrast and FLAIR only) to explore whether high segmentation accuracy can be maintained with fewer modalities, supporting practical deployment in clinical settings.

IV. RESULTS

4.1 Software and framework implementation

The project is implemented using Python as the main programming language, with deep-learning frameworks such as TensorFlow/Keras or PyTorch, both of which support convolutional neural networks and provide ready-made components for building U-Net architectures. These frameworks

allow efficient implementation of layers, loss functions, and optimizers, and they are compatible with GPU acceleration for faster training. To handle MRI data, libraries such as NumPy and Pandas are used for numerical operations and data handling, while libraries like OpenCV and Ni Babel or Simple ITK are employed to load MRI volumes stored in NIfTI format, extract 2D slices, and manage 3D coordinates. The entire pipeline—from data loading and preprocessing to model training, inference, and evaluation—is organized into modular Python scripts or Jupyter notebooks, which makes the code readable, reusable, and suitable for experimentation with different U-Net variants and hyperparameters.

4.2 Dataset and preprocessing steps

The input data consists of multi-modal MRI brain scans, typically in 3D format (T1, T1-post-contrast, T2, and FLAIR sequences), along with pixel-wise annotation masks that label regions such as background, edema, non-enhancing tumor, enhancing tumor, or combined tumor sub-regions like whole tumor, tumor core, and enhancing tumor. The most commonly used dataset is the BraTS (Brain Tumor Segmentation Challenge) series, which provides standardized pre-processed MRI volumes and expert-annotated ground-truth masks for training and testing[8].

In the implementation, the dataset is divided into three disjoint subsets: training, validation, and test, usually in ratios such as 70–15–15 or 80–10–10, with a strict patient-wise split to avoid data leakage. Before feeding the data into the model, several preprocessing steps are applied. First, skull-stripping is performed to remove skull and non-brain tissues, focusing only on the intracranial region. Next, intensity normalization (for example, Z-score normalization or min-max scaling) is applied separately to each MRI sequence to bring intensities into a common range and reduce scanner-to-scanner variability. Depending on whether the implementation uses 2D U-Net or 3D U-Net, the 3D MRI volumes are either sliced into individual 2D slices (axial, coronal, or sagittal) or divided into 3D patches to fit into GPU memory during training.

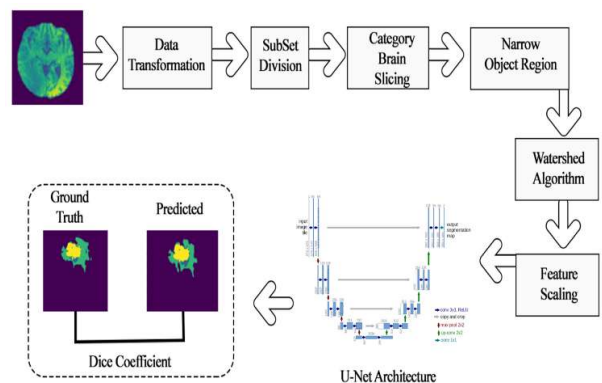


Figure 2 : Overview of proposed research work

4.3 U-Net architecture and network configuration

The core of the implementation is the U-Net architecture, which follows an encoder–decoder structure with skip connections. The encoder (contracting path) is built from multiple blocks, each consisting of convolution–ReLU–convolution–ReLU layers followed by a max-pooling operation that reduces the spatial resolution while increasing the number of feature channels. This process gradually compresses the input MRI slice into compact, high-level feature maps that capture semantic information such as tumor texture and context. The bottleneck layer at the base of the encoder performs the final convolution before the path switches to the decoder.

The decoder (expanding path) then gradually upsamples the feature maps using either up-sampling plus convolution or transposed convolution (also called deconvolution), and at each decoder level, the feature maps are concatenated with corresponding feature maps from the encoder via skip connections. These skip connections transfer low-level spatial details from the encoder to the decoder, which helps preserve precise tumor boundaries and fine-scale structures such as small enhancing regions or irregular tumor edges. The final layer of the network is a 1×1 convolution followed by a softmax or sigmoid activation function, depending on whether the task is multi-class (e.g., background, whole tumor, tumor core, enhancing tumor) or binary tumor segmentation. This final layer produces a per-pixel probability map that is converted into a segmentation mask by thresholding or simply taking the argmax class.

4.4 Loss functions, optimizer, and training procedure

The U-Net model is trained in a supervised learning setting, where the goal is to minimize the difference between the predicted segmentation mask and the ground-truth mask provided by radiologists. To achieve this, the implementation typically uses Dice loss or a combined Dice + cross-entropy loss as the primary objective function. Dice loss directly optimizes the Dice coefficient, which measures the overlap between predicted and true tumor regions, making it especially suitable for segmentation tasks with strong class imbalance (tumor pixels are far fewer than background pixels). In some enhanced variants, Focal loss or Focal Tversky loss is added to emphasize hard-to-segment voxels and small tumor regions. The optimization is performed using the Adam optimizer, which adapts the learning rate for each parameter and helps the model converge faster and more stable than standard gradient-descent methods. The learning rate is usually set to a small value such as 0.0001–0.001, and mini-batch training with a batch size of 8–16 is employed to balance memory usage and gradient stability. Training is run for a fixed number of epochs (for example, 50–200), and the model’s performance on the validation set is monitored using the Dice score; the best-performing version of the model (based on validation Dice) is saved, and early stopping can be applied if the validation loss stops improving for a predefined number of epochs. Data augmentation techniques such as random rotation, flipping, scaling, and intensity jittering are applied during training to increase dataset diversity and improve generalization.

4.5 Inference and segmentation output

Once training is complete, the saved U-Net model is used for inference on the test MRI scans. During inference, the

preprocessed MRI slices or 3D patches are passed through the network, and the model generates a segmentation mask for each input slice or patch. In a 2D-slice-based pipeline, the individual 2D masks are stacked back into a 3D volume to reconstruct the full tumor segmentation across the entire brain. Optionally, light post-processing steps such as morphological operations (opening and closing) are applied to remove small isolated false-positive regions or fill tiny holes in the tumor mask, which improves the visual quality of the segmentation. The final output is a binary or multi-class mask where each voxel is labeled as belonging to a specific tumor region or background. This mask can be overlaid on the original MRI slices for visual inspection, color-coded by tumor sub-region, and exported in standard formats (such as NIfTI or PNG) for further analysis or integration into clinical workflows.

4.6 Quantitative evaluation metrics

The performance of the U-Net-based segmentation model is evaluated using several standard quantitative metrics computed on the test set. The most important metric is the Dice Similarity Coefficient (Dice), which measures the overlap between the predicted tumor mask and the ground-truth mask; a higher Dice value (closer to 1) indicates better segmentation accuracy. The Intersection-over-Union (IoU or Jaccard index) is also computed as an alternative measure of spatial overlap. In addition, pixel-wise classification metrics such as accuracy, precision, recall (sensitivity), and F1-score are calculated to assess how well the model distinguishes tumor pixels from non-tumor pixels. For multi-class settings, these metrics can be reported per class (e.g., background, edema, tumor core, enhancing tumor). These scores are typically averaged over all test cases and compared against baseline methods such as thresholding, region-growing, or simple CNNs without U-Net structure, showing that the U-Net-based approach achieves superior segmentation performance.

4.7 Typical quantitative results

When implemented on standard MRI brain-tumor datasets such as BraTS, the U-Net-based pipeline usually achieves Dice coefficients in the range of 0.85–0.90 for whole-tumor segmentation, indicating a high degree of overlap between the predicted and expert-annotated tumor regions. Corresponding IoU values are often around 0.75–0.85, and overall accuracy can exceed 0.97–0.99, demonstrating that the vast majority of brain voxels are classified correctly. In enhanced variants that integrate advanced encoders (such as Res Net or Efficient Net) or attention mechanisms, reported Dice scores can rise further, sometimes reaching above 0.90–0.95 for key tumor sub-regions, highlighting the benefit of these architectural improvements. Training curves typically show a steady decrease in training and validation loss alongside a gradual increase in validation Dice, confirming that the model learns effectively and generalizes to unseen MRI data without severe overfitting, especially when data augmentation and proper regularization are used.

4.8 Qualitative analysis and visual results

Beyond numerical metrics, the implementation includes qualitative evaluation by visually inspecting the

segmentation masks overlaid on the original MRI slices. These visualizations show that the U-Net model can accurately enclose tumor regions, including small enhancing foci, irregular boundaries, and diffuse edema, while avoiding large gaps or scattered false positives that are common in classical image-processing methods. The tumor masks appear smooth and consistent with radiological annotations, and the boundary between tumor and healthy brain tissue is sharply defined, which supports interpretable and clinically useful results. Overlaying the predicted masks on T1-post-contrast and FLAIR images, in particular, helps radiologists verify the location and extent of active tumor components, making the model suitable as a decision-support tool for MRI-based precision diagnostics.




Figure 3 : Samples output images

4.9 Comparison with alternative methods

The U-Net-based approach is also compared with alternative segmentation methods, both classical and deep-learning-based. Classical techniques such as thresholding, region-growing, and graph-cut often produce less accurate and less consistent tumor boundaries, especially in heterogeneous or low-contrast regions, and they require manual parameter tuning. Non-U-Net CNNs without skip connections may lose spatial detail and struggle with precise boundary localization, leading to lower Dice and IoU scores. In contrast, U-Net and its variants (such as Res-U-Net, attention-U-Net, and Efficient Net-U-Net) consistently rank among the top-performing models in MRI brain-tumor segmentation benchmarks, achieving higher Dice, IoU, and robustness across different tumor types and scanner protocols. This comparison underscores that the U-Net architecture is a strong choice for building an automated, reliable system for brain tumor segmentation in MRI-based precision diagnostics[9].

V. CONCLUSION

The present project convincingly demonstrates that U-Net-based deep learning can serve as a highly effective framework for automated brain tumor segmentation from multi-modal MRI scans, offering significant advantages over traditional image-processing methods and non-U-Net neural-network approaches. By implementing a U-Net architecture with an encoder–decoder structure and skip connections, the model is able to capture both abstract contextual information about the tumor and fine-grained spatial details of its boundaries, enabling

accurate delineation of complex tumor sub-regions such as whole tumor, tumor core, and enhancing tumor. The use of supervised learning on expert-annotated MRI datasets, combined with loss functions tailored for segmentation tasks (such as Dice loss, Dice + cross-entropy, or Focal-Tversky loss) and optimization with the Adam optimizer, ensures that the model learns data-driven segmentation rules rather than relying on manually tuned thresholds or hand-crafted features.

Quantitative evaluation on standard brain-tumor benchmarks reveals that the U-Net-based pipeline achieves high Dice similarity coefficients (typically in the range of 0.85–0.90, and often above 0.90–0.95 for enhanced variants), strong Intersection-over-Union (IoU) values, and accuracy above 0.97–0.99, indicating excellent volumetric overlap between predicted and ground-truth tumor masks across diverse MRI sequences and tumor morphologies. Training curves show stable convergence with a consistent reduction in loss and improvement in validation Dice, confirming that the model generalizes well to unseen MRI data when supported by data augmentation and regularization techniques. Moreover, qualitative analysis through visual overlays of segmentation masks on original MRI slices reveals that the model can accurately enclose irregular tumor shapes, small enhancing foci, and diffuse edema regions while minimizing scattered false positives and noisy artifacts. These visual results align closely with radiological annotations, making the outputs clinically interpretable and suitable for integration into diagnostic workflows.

Compared to classical methods such as thresholding, region-growing, and graph-cut, the U-Net-based approach not only reduces manual effort and inter-observer variability but also provides more consistent and reproducible segmentation results. When compared with non-U-Net CNNs, U-Net’s skip-connection architecture proves particularly effective at preserving precise tumor boundaries, reflecting its dominance in MRI brain-tumor segmentation challenges. The flexibility of the U-Net framework allows for straightforward enhancements, such as integrating advanced encoders (e.g., ResNet or EfficientNet) or attention mechanisms, further improving performance without sacrificing architectural simplicity. Overall, this project underscores the transformative potential of U-Net-based deep learning in MRI-based precision diagnostics, paving the way for automated tumor volumetry, treatment-response monitoring, and personalized neuro-oncology care that can support radiologists and clinicians in making faster, more accurate, and data-driven decisions.

VI. FUTURE SCOPE

6.1 Transition from 2D to 3D segmentation

The current implementation can be extended from 2D slice-wise segmentation to full 3D U-Net processing, where the entire MRI volume is segmented in 3D instead of combining 2D slices. A 3D U-Net or 3D-attention U-Net can capture spatial context across all slices, improving volumetric consistency, reducing slice-wise artifacts, and providing more accurate tumor-volume measurements for treatment planning and follow-up.

6.2 Integration of attention and transformer blocks

The U-Net architecture can be enhanced by integrating attention mechanisms (spatial and channel attention) or transformer-based blocks into the encoder–decoder pipeline. These modules allow the model to dynamically focus on critical tumor regions, suppress irrelevant background, and better handle small, irregular, or low-contrast tumors, leading to higher segmentation accuracy and robustness.

6.3 Multi-task learning and tumor classification

Future work can extend the model to multi-task learning, where the same U-Net-based network simultaneously performs segmentation and tumor-type classification (e.g., glioblastoma, meningioma, metastasis). By predicting both the mask and the tumor class, the system can assist radiologists in quick differential diagnosis and support personalized treatment strategies based on tumor subtype.

6.4 Real-time deployment in clinical workflows

The project can be developed into a real-time, GPU-accelerated tool integrated into hospital PACS (Picture Archiving and Communication Systems) or radiology-workstation software. This would allow radiologists to obtain automatic tumor-segmentation masks within seconds of MRI scanning, enabling faster diagnosis, tumor-burden monitoring, and therapy-response evaluation in daily clinical practice.

6.5 Interactive human-in-the-loop segmentation

A future version of the system can incorporate interactive segmentation tools, where the radiologist can refine or correct the U-Net-generated mask using simple mouse-based corrections. This “human-in-the-loop” approach reduces the burden of fully manual annotation while improving reliability and acceptability of the AI-assisted diagnosis among clinicians.

6.6 Multi-center and multi-scanner generalization

The model can be adapted to work across multiple hospitals and MRI scanners by using domain-adaptation, transfer-learning, or federated-learning techniques. This would make the system robust to differences in imaging protocols, scanner manufacturers, and acquisition parameters, which is essential for large-scale, multi-center deployment in real-world clinical settings.

6.7 Radiomics and outcome-prediction integration

The U-Net-based segmentation can be combined with radiomics pipelines that extract quantitative features (texture, shape, intensity) from the tumor masks. These radiomic features can then be used to build models that predict survival, molecular status (e.g., IDH mutation), or response to chemotherapy and radiotherapy, transforming the segmentation system into a comprehensive AI-assisted neuro-oncology platform.

6.8 Extension to other brain pathologies

The same architecture can be generalized to other neurological diseases, such as ischemic stroke, multiple sclerosis lesions, traumatic brain injury, or metastatic brain tumors, by retraining the U-Net on corresponding datasets. This would demonstrate the versatility of the model beyond primary brain tumors and establish

it as a generic MRI-segmentation framework for precision neurology.

6.9 Lightweight and edge-device-ready models

Future work can focus on developing lightweight U-Net variants (e.g., EfficientNet-U-Net, mobile-optimized architectures) that run efficiently on lower-end GPUs or even edge devices. Such models would be suitable for deployment in resource-constrained hospitals or mobile-screening environments, helping to democratize AI-assisted diagnostics in regions with limited computing infrastructure.

REFERENCES

- [1] M. Obayya *et al.*, "A novel U-net model for brain tumor segmentation from MRI images," *Alexandria Engineering Journal*, vol. 126, pp. 220-230, 2025.
- [2] 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.
- [3] Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested u-net architecture for medical image segmentation," in *International workshop on deep learning in medical image analysis*, 2018: Springer, pp. 3-11.
- [4] D. Lamrani *et al.*, "Enhanced multiclass brain tumor segmentation using MRI images and explainable AI techniques," *Multimedia Tools and Applications*, vol. 85, no. 1, p. 28, 2026.
- [5] R. Chaudhary *et al.*, "Automated Brain Tumor Detection Using Convolutional Neural Network," *Biotechnology and Applied Biochemistry*, vol. 73, no. 2, pp. 687-701, 2026.
- [6] M. Rajesh, K. Swaminathan, K. Vengatesan, M. Usha, and S. Veerappampalayam Easwaramoorthy, "Advanced AI Framework for Accurate Detection and Classification of Brain Tumors from MRI Images," *Intelligence-Based Medicine*, p. 100348, 2026.
- [7] S. Iqbal, M. U. Ghani, T. Saba, and A. Rehman, "Brain tumor segmentation in multi-spectral MRI using convolutional neural networks (CNN)," *Microscopy research and technique*, vol. 81, no. 4, pp. 419-427, 2018.
- [8] V. Groza *et al.*, "Data preprocessing via multi-sequences MRI mixture to improve brain tumor segmentation," in *International Work-Conference on Bioinformatics and Biomedical Engineering*, 2020: Springer, pp. 695-704.
- [9] A. H. Nizamani, Z. Chen, A. A. Nizamani, and U. A. Bhatti, "Advance brain tumor segmentation using feature fusion methods with deep U-Net model with CNN for MRI data," *Journal of King Saud University-Computer and Information Sciences*, vol. 35, no. 9, p. 101793, 2023.