

# Neuro Vision 3D: Brain Tumor Segmentation and 3D Visualization Using U-Net

Prajwal L

Dept. of Artificial Intelligence and Machine Learning  
Mangalore Institute of Technology & Engineering  
Badaga Mijar, Moodabidri-574225, Karnataka  
lprajwal18@gmail.com

Shreya Shenoy

Dept. of Artificial Intelligence and Machine Learning  
Mangalore Institute of Technology & Engineering  
Badaga Mijar, Moodabidri-574225, Karnataka  
shreyashenoy893@gmail.com

Pourshi V Rai

Dept. of Artificial Intelligence and Machine Learning  
Mangalore Institute of Technology & Engineering  
Badaga Mijar, Moodabidri-574225, Karnataka  
pourshivedavyasarai25@gmail.com

Nandeesh S

Dept. of Artificial Intelligence and Machine Learning  
Mangalore Institute of Technology & Engineering  
Badaga Mijar, Moodabidri-574225, Karnataka  
nandeeshas004@gmail.com

**Abstract**—A key component of neuro-oncology, accurate brain tumor segmentation using magnetic resonance imaging (MRI) has a significant influence on diagnosis, therapy planning, and long-term monitoring. Nevertheless, the current clinical standard of expert manual delineation is time-consuming, prone to human error, and has substantial inter- and intra-observer variability. This crucial work is made more difficult by the heterogeneity of brain cancers, especially gliomas, and differences in imaging procedures. Medical image segmentation can be automated with deep learning approaches, particularly with Convolutional Neural Networks (CNNs) like the U-Net architecture. In order to overcome these difficulties, this research presents *Neuro Vision 3D*, an automated system that uses a 3D U-Net model that works especially well with volumetric medical data, such as MRI. The system processes input using the BraTS 2020 Training Dataset, a well-known benchmark that includes multi-modal MRI images (T1, T1ce, T2, FLAIR) and expert annotations in .nii format via a customized pipeline. Important preprocessing procedures include one-hot encoding of segmentation masks to enable multi-class prediction (Necrotic Core, Edema, Enhancing Tumor), stacking of multiple MRI modalities, spatial cropping to concentrate computational resources on the brain region, and min-max normalization for intensity standardization. The 3D U-Net is trained using the processed data to automatically provide accurate segmentation masks. With its user-friendly interface (developed with Flask and NiiVue for visualization), *Neuro Vision 3D* enables clinicians to upload patient scans and obtain a downloadable .nii.gz segmentation. *Neuro Vision 3D* aims to reduce manual labor and offer useful assistance for clinical decision-making in brain tumor care by automating the segmentation process, resulting in faster, more reliable, and accurate tumor delineation.

**Keywords**—Brain Tumor Segmentation, 3D U-Net, Medical Imaging, BraTS Dataset, Multi-modal MRI, Automatic Segmentation, 3D Visualization

## I. INTRODUCTION

A key component of successful neuro-oncology is the precise segmentation of brain tumors in medical imaging. For accurate diagnosis, strategic treatment planning (including radiation targeting and surgical excision), and longitudinal mon-

itoring of tumor response to therapy, Multi-modal Magnetic Resonance Imaging (MRI) has become the standard imaging modality. It uses sequences like T1-weighted (T1), T1-contrast enhanced (T1ce), T2-weighted (T2), and Fluid Attenuated Inversion Recovery (FLAIR) [6], and offers vital volumetric and spatial data due to its superior soft-tissue contrast and capacity to distinguish between different pathological tissues, including tumor subregions like the enhancing core, necrotic areas, and peritumoral edema.

### A. Motivation

1) *Constant Difficulty*: The clinical workflow mostly depends on the manual delineation of tumor boundaries by skilled radiologists, even with MRI's diagnostic capabilities. There are major ongoing problems with this manual approach. It is incredibly labor-intensive and time-consuming, especially when dealing with large 3D MRI volumes that contain many slices [2]. Furthermore, there is inherent subjectivity in manual segmentation, resulting in significant inter- and intra-observer variability. The heterogeneous nature of brain tumors, particularly gliomas—which feature a variety of sizes, forms, and frequently exhibit infiltrative or ambiguous boundaries—further complicates the task of consistent and precise delineation.

2) *Limitations of Existing Methods*: CNN-based U-Net topologies are appealing for clinical applications with constrained technology since they are effective and reasonably reliable to train. Although their fixed receptive field restricts long-range spatial dependency extraction, their simple convolutional pipelines allow for faster inference and reduced memory utilization. More sophisticated variants like U-Net++ and attention-enhanced U-Nets boost multi-resolution feature extraction and boundary precision, but bring higher parameter counts and increased sensitivity to hyperparameter adjustment, making training less stable on small medical datasets.

Hybrid Transformer-based models like UNETR and Swin-UNet further improve global context modeling through self-

attention, but result in significantly increased processing demands, quadratic memory scaling with input size, and reliance on large datasets for steady convergence. These architectures require warm-up stages, extensive pretraining, and specific learning-rate schedules, making them unsuitable in everyday clinical settings without access to powerful GPUs and extensive annotated datasets.

3) *Need for Advanced Techniques:* The need for sophisticated, automated methods that can offer accurate, effective, and repeatable brain tumor segmentation is evident and urgent. Deep learning (DL) techniques have demonstrated great potential in meeting these demands in medical image processing.

4) *Potential of the 3D U-Net:* The 3D U-Net was selected for this project due to its proven performance in volumetric medical image segmentation tasks. Its fundamental architecture automatically captures inter-slice spatial context, which is frequently lost in 2D slice-by-slice approaches, making it ideal for 3D MRI data. By combining high-resolution characteristics from the encoder with the upsampled features in the decoder [3], [7], the U-Net's symmetric encoder-decoder structure efficiently learns hierarchical features, with skip connections essential for preserving fine spatial details needed for accurate delineation of complicated tumor boundaries [8].

5) *Impact on Clinical Practice:* The creation of an automated tool based on the 3D U-Net, such as Neuro Vision 3D, has significant potential for clinical impact. It can facilitate more consistent treatment planning and faster diagnostic turnaround times by offering automated segmentation, thereby considerably reducing the manual segmentation workload for clinicians. Additionally, by incorporating an interactive 3D visualization, the structure of the segmented tumor and its relationship to adjacent tissues may be intuitively explored [3]. The ultimate goal is to improve the accuracy and effectiveness of managing brain tumors and ultimately improve patient outcomes.

## B. Project Overview

Neuro Vision 3D, a method for automatically segmenting and visualizing brain tumors from multi-modal MRI data, is presented in this paper. The system accepts multi-modal MRI scans (T1, T1ce, T2, FLAIR) in standard `.nii` format. Pre-processing procedures include intensity normalization, spatial cropping to the brain region, stacking of pertinent modalities into a multi-channel input, and one-hot encoding of annotation masks. A 3D U-Net deep learning model, specifically trained for brain tumor segmentation using the BraTS 2020 dataset, forms the basis of the system.

## II. RELATED WORK

In neuro-oncology, the segmentation of brain tumors from MRI is still a major challenge. Manual tumor delineation is laborious and time-consuming, prone to human error, and suffers from inconsistency across physicians, while MRI offers detailed insights. Accurate segmentation is further complicated by the intrinsic heterogeneity of brain tumors in terms of size,

form, and blurred boundaries. This has prompted a necessary move toward automated techniques.

Deep learning (DL) has become a potent tool in medical image analysis, marking a paradigm shift from conventional image processing methods [5], [8]. Convolutional Neural Networks (CNNs) in particular have demonstrated encouraging results for medical image analysis, excelling at handling vast volumes of unstructured data and learning intricate hierarchical characteristics directly from images, frequently with little human involvement.

The U-Net model has emerged as one of the most popular and successful DL architectures for biomedical image segmentation. Its symmetric encoder-decoder structure allows the encoder (contracting path) to capture contextual information while the decoder (expanding path) enables exact localization. U-Net's key innovation—skip connections that concatenate feature maps from the encoder directly to the decoder—enables the network to integrate deep semantic feature information with shallow, high-resolution features, thereby preserving fine spatial details essential for distinct delineation of tumor shapes and edges.

Since 2D U-Net models process images slice by slice and fail to fully capture the complete 3D spatial context present in volumetric data like MRI, the 3D U-Net architecture was developed. The 3D U-Net may learn from inter-slice information and gain a better understanding of complicated 3D tumor structures by employing 3D convolutions and pooling to assess the entire 3D volume at once, making it an ideal choice for brain tumor segmentation.

## III. METHODOLOGY

The technique for this project adheres to a thorough workflow encompassing initial data gathering and preprocessing, model training, and culminating in final display, as depicted in Fig. 1.

### A. Dataset

The BraTS 2020 Training Dataset, obtained from Kaggle, was used to train and validate the model. This dataset, which includes scans from roughly 369 patients, is a recognized benchmark for brain tumor segmentation.

**Data Format:** All imaging data is delivered in NIfTI (`.nii`) format, which is common for 3D medical imaging.

Each patient scan comprises four different MRI modalities: Fluid Attenuated Inversion Recovery (FLAIR), T1-weighted (T1), T1-contrast enhanced (T1ce), and T2-weighted (T2). The dataset is completely supervised — an expert-annotated ground-truth segmentation mask (`seg.nii`) is included with every set of patient scans. Each voxel in these 3D volumetric masks is labeled with an integer value corresponding to a discrete tissue class, as shown in Fig. 2.

### B. Preprocessing Pipeline

To prepare the raw `.nii` files for the deep learning model, a specific preprocessing pipeline was applied.

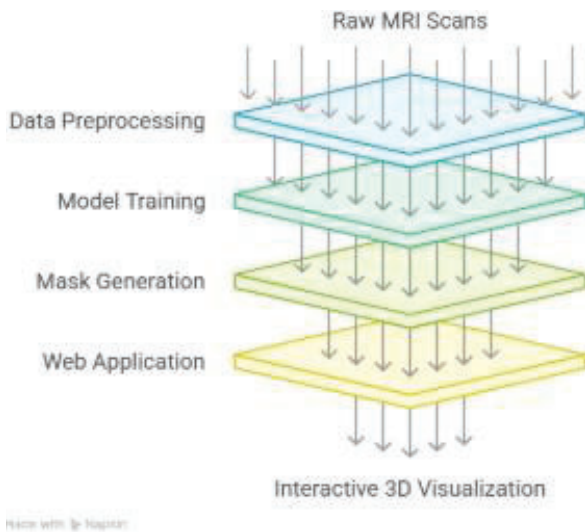


Fig. 1. Neuro Vision 3D Workflow




Fig. 2. MMRI modalities and ground-truth segmentation mask

1) *Normalization*: All voxel intensity values for each MRI scan were scaled to a uniform range of [0, 1] using the `MinMaxScaler` from the `sklearn.preprocessing` library. Different MRI scans possess widely varying intensity ranges; normalization ensures that all scans have a consistent intensity distribution, preventing the model from being biased by high-intensity values and allowing for more stable and uniform training.

2) *Cropping*: The 3D volumes were cropped to a standardized size of  $128 \times 128 \times 128$  voxels to focus on the brain region and remove large areas of empty black background. This step significantly reduces memory usage and speeds up training by focusing the model's computational effort only on the relevant anatomical structures where tumors exist.

3) *Stacking Modalities*: To create a single multi-channel 3D input volume, the preprocessed volumes from the various MRI modalities (FLAIR, T1ce, and T2) were stacked together — analogous to combining the R, G, B channels in a 2D image. Every MRI modality reveals unique tissue properties; stacking them provides the model with a richer, more complete set of data that improves its ability to differentiate between different tumor subregions and healthy tissue.

4) *One-Hot Mask Encoding*: The ground-truth segmentation masks were transformed from single integer labels (0, 1, 2, 3) per voxel into a 4-channel binary format. Every voxel is converted to a four-element vector, with the class-corresponding index set to 1 and all other indexes set to 0 (e.g., Edema  $\rightarrow [0, 1, 0, 0]$ ). This conversion is required for multi-class semantic segmentation, enabling the model to produce a probability for each of the four classes separately per voxel.

### C. 3D U-Net Architecture

We used a 3D U-Net model that processes 3D volumes directly and captures inter-slice spatial context, making it highly effective for volumetric segmentation tasks [11]. The architecture comprises an encoder, a bottleneck, and a decoder with skip connections.

**Input Layer**: The preprocessed stacked MRI volumes with shape  $(3, 128, 128, 128)$  (3 input channels,  $128 \times 128 \times 128$  spatial dimensions) are accepted by the model.

**Contracting Path (Encoder)**: By gradually downsampling the input, the encoder extracts hierarchical characteristics and records context. There are four major blocks, each consisting of two sets of  $3 \times 3 \times 3$  Conv3D, BatchNorm3D, and ReLU activation followed by a Dropout layer for regularization. A  $2 \times 2 \times 2$  MaxPool3D layer is applied after each block to downsample the feature map. The channel progression grows with depth:  $16 \rightarrow 32 \rightarrow 64 \rightarrow 128$  channels.

**Bottleneck**: The network's deepest layer links the encoder and decoder. It comprises Conv3D, BatchNorm, and ReLU layers with 256 channels, capturing the input volume's most abstract, high-level properties and global context.

**Expansive Path (Decoder)**: The decoder gradually upsamples features to recreate the segmentation map, mirroring the encoder path. Each block starts with a  $2 \times 2 \times 2$  ConvTranspose3D layer to upsample the feature map, followed by concatenation with the corresponding encoder feature map via a skip connection. This critical phase restores high-resolution spatial information lost during encoding. A Dropout layer follows after the combined feature map has been processed by two sets of Conv3D, BatchNorm3D, and ReLU. The channel progression diminishes with increasing depth:  $128 \rightarrow 64 \rightarrow 32 \rightarrow 16$  channels.

**Output Layer**: A final  $1 \times 1 \times 1$  Conv3D layer converts the 16-channel feature map to 4 output channels corresponding to the one-hot encoded classes (Background, Edema, Necrotic Core, Enhancing Tumor). The final output shape is  $(4, 128, 128, 128)$ .

#### D. Training Details

1) *Framework*: The PyTorch deep learning framework was used for model implementation and training.

2) *Loss Function*: A combination of Dice Loss (to address class imbalance) and Focal Loss (to concentrate on difficult-to-classify voxels) was employed.

3) *Optimizer*: The AdamW optimizer was used with an initial learning rate schedule.

4) *Training Configuration*: The model was trained for 100 epochs with a predetermined batch size while monitoring validation loss and Dice score to avoid overfitting.

5) *Hardware*: A GPU was used to accelerate training.

#### E. Visualization and System Integration

The trained model was integrated into a full-stack web application to produce a practical tool for clinicians.

**Backend**: A Flask backend server manages the logic, offering routes to handle file uploads via POST requests and serve processed files. After a user uploads a MRI scan, the backend applies the complete preprocessing and model inference pipeline, then stores the segmentation mask as a `.nii.gz` file using the Nibabel library. File URLs for the initial scan and the generated mask are then returned to the frontend as a JSON response.

**Frontend**: HTML templates rendered by Flask construct the user interface, providing a form for clinicians to upload patient MRI scans. After retrieving the file URLs from the backend's JSON response, it renders the MRI volume with the anticipated tumor mask superimposed in an interactive, in-browser 3D viewer using Niivue, a JavaScript-based 3D visualization toolkit. This eliminates the need for specialist external software and offers instantaneous, interactive feedback.

### IV. RESULTS AND DISCUSSION

#### A. Evaluation Metrics

We employed two well-known metrics that calculate agreement between the predicted mask  $A$  and the expert-annotated ground truth mask  $B$ .

1) *Intersection over Union (IoU)*: This measures overlap as the area of intersection over the area of union, and is a stringent measure of segmentation performance:

$$IoU = \frac{|A \cap B|}{|A \cup B|} \quad (1)$$

2) *Dice Similarity Coefficient (DSC)*: The Dice score is the most common metric in medical segmentation and effectively measures overlap, especially in cases with class imbalance:

$$Dice = \frac{2 \times |A \cap B|}{|A| + |B|} \quad (2)$$

We also monitor standard voxel-wise **Accuracy** (percentage of correctly classified voxels) and **Loss** (a combination of Dice Loss and Focal Loss) to assess model training and convergence.

#### B. Quantitative Results

The 3D U-Net model was trained to 100 epochs, with performance monitored on both the training set and a held-out validation set. Curves of training and validation Loss, Accuracy, and IoU are shown in Figs. 3, 4, and 5.




Fig. 3. Model Loss




Fig. 4. Model Accuracy

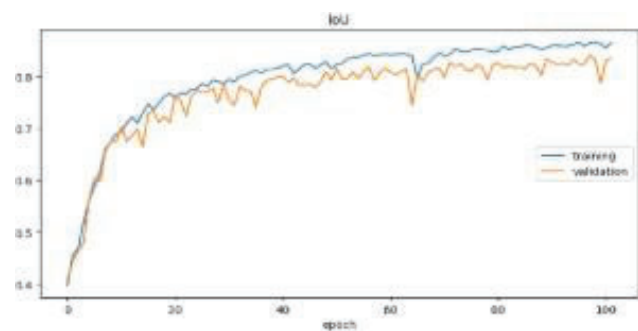


Fig. 5. Intersection over Union (IoU)

The training curves exhibit robust convergence. The Model Loss decreases quickly for both training and validation sets and levels off at small values. The final validation loss (0.115) closely follows the training loss (0.089), reflecting good generalization without severe overfitting. Accuracy and IoU curves are consistently increasing and plateauing towards high-performance values. The final performance metrics on the validation set are tabulated in Table I.

TABLE I  
 PEAK AND FINAL PERFORMANCE METRICS ON BRATS 2020 VALIDATION DATASET

Metric	Final (Epoch 100)	Peak
Average IoU	0.835	0.841
Average Accuracy	0.894	0.907
Average Loss	0.115	0.110 (min)

### C. Qualitative Findings

The model's performance is qualitatively validated through visual examination of the segmentation output. Fig. 6 shows representative segmentation results from our test set, displaying for each case: (a) Input MRI (FLAIR modality), (b) Ground Truth Mask, and (c) Neuro Vision 3D Predicted Mask. Colors denote: Red = Necrotic Core, Blue = Enhancing Tumor, Green = Edema.




Fig. 6. Qualitative segmentation results

As illustrated, the model's predictions closely match the ground truth masks, typically capturing the intricate shapes of the tumors. The system also integrates these segmentation masks into the frontend for real-time 3D visualization, as shown in Fig. 7.

### D. Discussion

The quantitative and qualitative results demonstrate the high performance of the Neuro Vision 3D system. A final average validation IoU of 0.835 and accuracy of 0.894 validate the success of the 3D U-Net architecture for this task (Table I). Its direct volume processing enables learning from inter-slice spatial context, while skip connections successfully integrate shallow high-resolution features with deep semantic features, allowing the model to reconstruct exact segmentation boundaries.

There are, however, a few slight variations, especially in the diffuse borders of peritumoral edema—well-known, challenging regions with naturally hazy boundaries.

The main benefits of the Neuro Vision 3D system are:

- 1) **Automation:** A labor-intensive and time-consuming manual task is successfully automated.
- 2) **Speed:** Model inference is substantially faster than manual segmentation.
- 3) **Reproducibility:** The model is consistent and produces deterministic results on the same dataset.

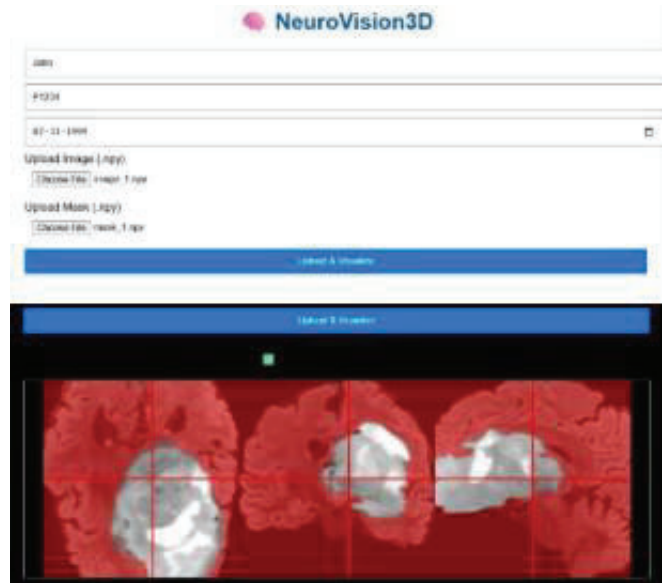


Fig. 7. This screenshot of the Neuro Vision 3D web application frontend shows an interactive 3D visualization of a segmented tumor (multi-colored) superimposed on the patient's MRI scan.

Notwithstanding these benefits, we recognize a number of limitations:

- 1) The model was trained only on the BraTS 2020 dataset. Its effectiveness on “in-the-wild” clinical data from various scanners or hospitals has not yet been confirmed.
- 2) The web application is a prototype and does not currently include user authentication or other security features necessary in a clinical setting.

### V. CONCLUSION AND FUTURE SCOPE

This paper presents Neuro Vision 3D, an automated brain tumor segmentation system developed using a 3D U-Net trained on the BraTS 2020 dataset. The system achieves strong quantitative performance (IoU: 0.835, Accuracy: 0.894) and demonstrates clinically relevant qualitative results. By providing automation, speed, and reproducibility alongside an interactive 3D visualization interface, Neuro Vision 3D has the potential to meaningfully assist radiologists in clinical decision-making and reduce the burden of manual segmentation.

Future directions include:

- Training with larger and more representative datasets to improve generalization to clinical data.
- Examining architectural enhancements, such as lightweight and computationally efficient architectures using depthwise separable convolutions, knowledge distillation, neural architecture search (NAS), and hybrid CNN-Transformer compression techniques to reduce memory and inference loads.
- Adding user authentication and security features for clinical deployment.

- Validation with expert radiologists and broader adaptation to other medical image segmentation applications.
- Employing multimodal MRI synthesis using generative models to handle missing modalities and improve robustness in diverse clinical settings.

#### ACKNOWLEDGMENT

The authors express sincere gratitude to the management of the Mangalore Institute of Technology and Engineering (MITE), Moodabidri, for providing the necessary facilities and a conducive environment to complete this project. They are deeply grateful to the Principal, Dr. Prashanth C M, for his invaluable support and administrative leadership. Heartfelt thanks go to Dr. Sunil Kumar, Head of the Department of Artificial Intelligence and Machine Learning, for his encouragement and departmental resources. The authors also acknowledge their Project Coordinator, Dr. Yogeeshwara Reddy, for his guidance throughout the project duration. Finally, they extend their most profound gratitude to their project guide, Dr. Maryjo M George, for her constant supervision, invaluable technical insights, and unwavering support, which were instrumental in the successful completion of this work.

#### REFERENCES

- [1] K. Lakshmi, S. Amaran, G. Subbulakshmi, S. Padmini, G. P. Joshi, and W. Cho, "Explainable artificial intelligence with UNet based segmentation and Bayesian machine learning for classification of brain tumors using MRI images," *Scientific Reports*, vol. 15, no. 1, p. 690, Jan. 2025. doi: 10.1038/s41598-024-84692-7.
- [2] Y. Pang et al., "Online Self-distillation and Self-modeling for 3D Brain Tumor Segmentation," *IEEE J. Biomed. Health Inform.*, 2025.
- [3] A. Guennich, M. Othmani, and H. Ltifi, "Advanced Brain Tumor Segmentation With a Multiscale CNN and Conditional Random Fields," *IEEE Access*, vol. 11, pp. 1–14, 2023.
- [4] C. M. Umarani, S. G. Gollagi, S. Allagi, K. Sambrekar, and S. B. Ankali, "Advancements in deep learning techniques for brain tumor segmentation: A survey," *Biomed. Signal Process. Control*, vol. 98, p. 106015, 2024.
- [5] M. M. Islam, Z. Wang, M. A. Iqbal, and G. Song, "Brain Tumor Segmentation on MR Images Using Anisotropic Deeply Supervised Convolutional Neural Network," *Neurocomputing*, vol. 570, p. 127063, 2023.
- [6] M. Lee, J. H. Kim, W. Choi, and K. H. Lee, "AI-assisted Segmentation Tool for Brain Tumor MR Image Analysis," *J. Imaging Inform. Med.*, vol. 38, no. 1, pp. 74–83, Feb. 2025. doi: 10.1007/s10278-024-01187-7.
- [7] S. Sangui, T. Iqbal, P. C. Chandra, S. K. Ghosh, and A. Ghosh, "3D MRI Segmentation using U-Net Architecture for the Detection of Brain Tumor," *Procedia Comput. Sci.*, vol. 235, pp. 1203–1211, 2024.
- [8] S.-Y. Lin and C.-L. Lin, "Brain Tumor Segmentation Using U-Net in Conjunction with EfficientNet," *Front. Neurosci.*, vol. 18, p. 10773611, 2024. doi: 10.3389/fnins.2024.10773611.
- [9] Z. U. Abidin, R. A. Naqvi, A. Haider, H. S. Kim, D. Jeong, and S. W. Lee, "Recent deep learning-based brain tumor segmentation models using multi-modality magnetic resonance imaging: a prospective survey," *Front. Bioeng. Biotechnol.*, vol. 12, p. 1392807, 2024.
- [10] H. Byeon et al., "Brain tumor segmentation using neuro-technology enabled intelligence-cascaded U-Net model," *Front. Comput. Neurosci.*, vol. 18, p. 1391025, 2024.
- [11] O. C. İc, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger, "3D U-Net: Learning dense volumetric segmentation from sparse annotation," in *Proc. MICCAI 2016*, pp. 424–432.
- [12] A. S. Farhan, M. Khalid, and U. Manzoor, "XAI-MRI: an ensemble dual-modality approach for 3D brain tumor segmentation using magnetic resonance imaging," *Front. Artif. Intell.*, vol. 8, p. 1525240, 2025.
- [13] A. Guennich, M. Othmani, and H. Ltifi, "Advanced brain tumor segmentation with a multiscale CNN and conditional random fields," *IEEE Access*, vol. 13, pp. 34925–34935, 2025.
- [14] M. M. Islam, Z. Wang, and M. A. Iqbal, "Brain tumor segmentation on MR images using anisotropic deeply supervised convolutional neural network," in *Proc. ICIT 2018*.
- [15] M. Lee, J. H. Kim, W. Choi, and K. H. Lee, "AI-assisted segmentation tool for brain tumor MR image analysis," *J. Imaging Inform. Med.*, vol. 38, pp. 74–83, 2025.
- [16] Y. Pang et al., "Online self-distillation and self-modeling for 3D brain tumor segmentation," *IEEE J. Biomed. Health Inform.*, 2025.
- [17] D. Rastogi et al., "Deep learning-integrated MRI brain tumor analysis: feature extraction, segmentation, and survival prediction using replicator and volumetric networks," *Scientific Reports*, vol. 15, p. 1437, 2025.
- [18] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. MICCAI 2015*, pp. 234–241.
- [19] M. Hassan et al., "Unfolding Explainable AI for Brain Tumor Segmentation," *Neurocomputing*, vol. 599, p. 128058, 2024.
- [20] C. M. Umarani et al., "Advancements in deep learning techniques for brain tumor segmentation: A survey," *Informatics in Medicine Unlocked*, vol. 50, p. 101576, 2024.
- [21] M. Zhou et al., "Generating 3D brain tumor regions in MRI using vector-quantization Generative Adversarial Networks," *Comput. Biol. Med.*, vol. 185, p. 109502, 2025.
- [22] A. H. Nizamani et al., "Advance brain tumor segmentation using feature fusion methods with deep UNet model with CNN for MRI data," *J. King Saud Univ. Comput. Inf. Sci.*, vol. 16, p. 101793, 2023.
- [23] L. Jiang et al., "Multimodal 3D Brain Tumor Segmentation with Adversarial Training and Conditional Random Field," *arXiv preprint*, 2024.