BRAIN TUMOR SEGMENTATION AND CLASSIFICATION FROM MRI USING A U-NET, GAN AND EFFICIENTNET HYBRID MODEL

[1]Ms. L. Shalini, [2]Pratik Bairagi, [3]Affan Ahmad, [4]Ashutosh Kumar, [5]Aman Mandal
[1]Assistant Professor, [2345]Students
Department of Computer Science and Engineering
Bharath Institute of Science and Technology,
Bharath Institute of Higher Education and Research (BIHER), Chennai, India

***Abstract*— Accurately interpreting MRI scans plays a crucial role in diagnosing brain tumors. However, doing this manually can be both time-consuming and prone to variation between different readers. To address these challenges, this study introduces an automated deep learning framework that improves MRI image quality, precisely identifies tumor regions at the pixel level, and classifies tumor types.The system combines three key components: a Generative Adversarial Network (GAN) for enhancing image quality, a U-Net model for segmentation, and an EfficientNet- based classifier for tumor identification. The overall workflow includes preprocessing MRI scans, enhancing contrast, segmenting tumor regions, refining masks, estimating tumor volume, and finally determining the tumor type.When tested on the BraTS 2020 and BraTS 2021 benchmark datasets, the proposed approach achieved a Dice Score of 92% and an Intersection over Union (IoU) of 85%. These results outperform several existing baseline methods, while also offering faster and more consistent support for clinical decision-making.**

Keyword: Brain Tumor, MRI, U-Net, Generative Adversarial Network, EfficientNet, Deep Learning, Medical Image Segmentation, BraTS 2020, BraTS 2021

I. INTRODUCTION

Cerebral tumors represent one of the most dangerous neurological conditions, with approximately 250,000 new diagnoses recorded globally each year. Prompt and precise identification is critical for patient outcomes, as the probability of survival correlates strongly with how early treatment begins. Magnetic Resonance Imaging (MRI) is the established non-invasive standard for evaluating brain tumor extent, offering high-resolution visualization of soft tissue anatomy.

Despite its diagnostic superiority over alternative imaging techniques, MRI interpretation by specialist radiologists is labor-intensive and susceptible to inter-reader inconsistency. A single patient scan may comprise hundreds of individual slices, each requiring independent scrutiny—a cognitively burdensome and error-prone process. This context motivates the construction of automated computer-aided diagnosis (CAD) frameworks capable of reliably detecting and delineating tumor boundaries.

Conventional image processing approaches—including threshold-based methods, edge detectors, region-growing algorithms, and morphological transforms—prove inadequate for brain tumor delineation owing to the pronounced heterogeneity of tumors. These techniques cannot generalize across varying scanner hardware, imaging protocols, or patient-specific anatomy. Deep learning has since emerged as the prevailing paradigm for tackling these limitations.

Convolutional architectures, particularly U-Net, have achieved leading performance in biomedical segmentation by exploiting encoder-decoder topologies augmented with skip connections that retain fine spatial detail. However, U-Net performance deteriorates when input scan quality is low, annotated training samples are scarce, or tumor margins are obscured by poor MRI contrast.Generative Adversarial Networks (GANs) furnish a complementary capability: they learn to synthesize realistic images and enhance existing scans through adversarial optimization. Conditioning the GAN on suboptimal MRI inputs allows the generator to produce high-contrast, low-noise outputs that directly improve downstream segmentation.

This study introduces a three-stage hybrid pipeline that (1) leverages a GAN to eleva Brain tumors are among the most serious neurological conditions, with around 250,000 new cases reported worldwide each year. Early and accurate detection plays a vital role in improving patient outcomes, as survival rates are closely linked to how quickly treatment begins. Magnetic Resonance Imaging (MRI) remains the preferred non-invasive technique for assessing brain tumors, thanks to its ability to produce detailed images of soft tissues. Even so, analyzing MRI scans is far from straightforward. Radiologists often have to review hundreds of slices for a single patient, carefully examining each one. This process takes time and, understandably, can lead to differences in interpretation between specialists. Because of this, there is a growing need for automated computer-aided diagnosis (CAD) systems that can consistently detect tumors and outline their boundaries with precision.

Earlier image processing methods—such as thresholding, edge detection, region growing, and morphological operations—have struggled to handle the complexity of brain tumors. Tumors can vary widely in shape, size, and appearance, making it difficult for these traditional techniques to perform reliably across different patients, scanners, and imaging conditions. In recent years, deep learning has emerged as a more effective solution to these challenges.

Among deep learning models, convolutional neural networks—especially U-Net—have shown strong performance in medical image segmentation. Its encoder–decoder structure, along with skip connections, helps preserve important spatial details. However, U-Net's effectiveness can drop when the input images are noisy, lack contrast, or when training data is limited. In such cases, accurately identifying tumor boundaries becomes more difficult. This is where

Generative Adversarial Networks (GANs) come into play. GANs are capable of improving image quality by learning to generate more realistic and enhanced versions of input scans. When applied to low-quality MRI images, they can produce clearer, higher-contrast outputs, which in turn make segmentation tasks more reliable. Building on these ideas, this study proposes a three-stage hybrid framework. First, a GAN is used to enhance MRI image quality. Next, a U-Net model performs detailed, pixel-level tumor segmentation. Finally, an EfficientNet-based classifier identifies the type of tumor using transfer learning. The entire approach is evaluated on the publicly available BraTS 2020 and BraTS 2021 datasets, allowing for consistent and reproducible benchmarking.

A. Key Contributions

- A unified three-stage architecture combining GAN-based scan enhancement, U-Net segmentation, and EfficientNet classification within a single end-to-end pipeline.
- Adversarial MRI enhancement that improves image contrast and suppresses acquisition noise prior to segmentation.
- Validation across both BraTS 2020 and BraTS 2021 benchmarks for reproducible cross-dataset performance evaluation.
- A systematic hyperparameter investigation establishing optimal batch sizes (16 for GAN and U-Net; 64 for EfficientNet), learning rate (1×10^{-5}), and epoch count (20) for stable training convergence.
- A GAN-driven synthetic data augmentation strategy for expanding limited annotated medical training corpora.

II. RELATED WORK

Research on brain tumor segmentation has gradually evolved from traditional image processing techniques to advanced deep learning approaches, with developments covering segmentation models, generative methods, and integrated hybrid frameworks.

A. U-Net-Based Segmentation

Ronneberger et al. [1] introduced U-Net, which quickly became a cornerstone in biomedical image segmentation due to its balanced encoder–decoder structure and the use of skip connections to preserve spatial detail. Building on this foundation, later variants such as UNet++ [14] introduced densely connected pathways, allowing for better multi-scale feature learning and improved segmentation accuracy. At the same time, Pereira et al. [13] demonstrated the effectiveness of convolutional neural networks on brain MRI data, showing clear improvements over traditional, non-learning-based approaches. Similarly, Havaei et al. [12] proposed a dual-pathway CNN architecture that captures both local and global contextual features simultaneously, enhancing segmentation performance. To better handle volumetric data, Cicek et al. [8] extended the U-Net framework into a 3D architecture, enabling more effective processing of MRI volumes even when annotations are limited. Despite these advances, such models remain sensitive to variations in image quality, which has led researchers to explore preprocessing and enhancement techniques prior to segmentation

B. GAN-Based Medical Image Enhancement

The Generative Adversarial Network (GAN) framework, introduced by Goodfellow et al. [2], models data through a competitive process between two networks—a generator and a discriminator. While the generator aims to produce realistic outputs, the discriminator works to distinguish between real and generated data, pushing the system to improve over time. Building on this idea, conditional GANs [3, 18] incorporate additional input information to guide the generation process. This makes them particularly well-suited for tasks like MRI enhancement, where outputs need to correspond closely to specific input scans. In medical imaging, GANs have been widely applied to problems such as MRI super-resolution, cross-modality image synthesis, noise reduction, and even the creation of synthetic datasets. By transforming low-quality or low-contrast MRI scans into clearer, more detailed versions, these models help retain important diagnostic features while improving overall image quality—ultimately supporting more accurate performance in downstream tasks like segmentation and classification.

C. Hybrid Approaches and Transfer Learning

Integrated pipelines that combine image enhancement with segmentation have consistently shown better performance than models relying on a single component. For instance, Myronenko [23] achieved leading results in the BraTS 2018 challenge using a U-Net architecture regularized with a variational autoencoder, highlighting the benefits of combining complementary techniques. Meanwhile, EfficientNet [15] introduced a more systematic way of scaling neural networks by jointly adjusting depth, width, and input resolution. This balanced approach significantly improved efficiency while maintaining high accuracy.In medical imaging tasks, applying transfer learning with EfficientNet models pretrained on ImageNet has proven especially effective. Compared to training models from scratch, this strategy tends to deliver better performance, particularly in situations where annotated data is limited.

D. Summary of Key Observations

- U-Net achieves precise voxel-level segmentation but is highly sensitive to input scan fidelity.
- GANs improve image clarity and can expand training datasets through synthetic generation.
- EfficientNet with transfer learning attains high classification accuracy even under limited annotation budgets.
- Hybrid pipelines consistently surpass individual components in accuracy and robustness.

III. PROBLEM STATEMENT

Automatically segmenting and classifying brain tumors from MRI scans remains a complex and challenging task. Tumors can differ widely from one patient to another in terms of shape, size, location, and internal composition, making it difficult for rule-based methods to perform reliably. On top of that, MRI data itself often comes with imperfections. Factors like scanner noise, patient movement during imaging, and uneven intensity across the scan can all reduce image quality and make accurate segmentation harder. Another key limitation is the availability of data. Unlike natural image datasets, medical datasets are relatively small and depend on expert annotations, which are time-consuming and costly to obtain. This limits the diversity of training data and adds another layer of difficulty when developing robust deep learning models.

Standalone U-Net models tend to struggle when dealing with low-contrast or noisy MRI inputs, often producing less accurate segmentation results under such conditions. In addition, many existing approaches treat enhancement, segmentation, and classification as separate tasks rather than combining them into a unified, optimized pipeline. This separation between stages can lead to error propagation, where mistakes in one step negatively affect the next. As a result, the overall system performance may fall short of what could be achieved with a more integrated, end-to-end approach.

A. Core Challenges

- Low-fidelity MRI inputs with noise, reduced contrast, and indistinct tumor margins.
- High variability in tumor morphology, scale, and anatomical position across patient cohorts.
- Insufficient volumes of expert-annotated training data for deep model supervision.
- Difficulty in segmenting irregular or diffuse tumor boundaries.
- Lack of a unified end-to-end pipeline integrating enhancement, segmentation, and classification.
- Over-reliance on manual radiological interpretation, introducing subjectivity and inter-observer divergence.

IV. PROPOSED METHODOLOGY

The proposed approach introduces a hybrid deep learning framework that combines a U-Net segmentation model with a Generative Adversarial Network (GAN) to improve the accuracy and reliability of brain tumor detection in MRI scans. Rather than treating each step separately, the system is built as an end-to-end pipeline. It takes raw MRI images as input, enhances their quality, performs segmentation to identify tumor regions, and produces refined detection results. This integrated design helps streamline the entire process while improving overall performance and consistency. The primary goal of the proposed system is to address the shortcomings of traditional approaches by enhancing MRI image quality, expanding the effective diversity of training data, and enabling more accurate detection of tumor boundaries.

A. Overall System Architecture

The proposed system is designed as a multi-stage architecture, where each component works in sequence to process MRI scans and produce the final output. It begins with preprocessing, followed by image enhancement using a GAN, then tumor segmentation with a U-Net model, and concludes with post-processing to refine the results.

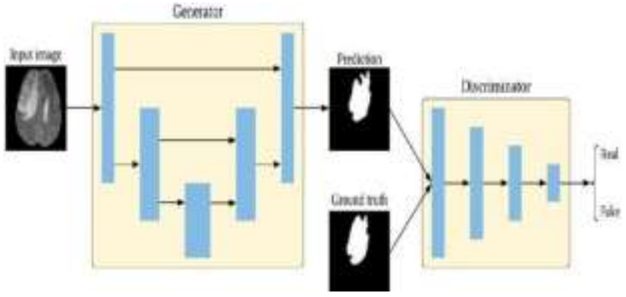

Fig. 1. Proposed Hybrid Model Architecture

V. MATHEMATICAL FORMULATION

A. Dice Coefficient

$$Dice(P, G) = 2|P \cap G| / (|P| + |G|) = 2TP / (2TP + FP + FN)$$

B. Intersection over Union (IoU)

$$IoU = TP / (TP + FP + FN)$$

C. GAN Adversarial Objective

$$\min\_G \max\_D V(D,G) = E[\log D(x)] + E[\log(1 - D(G(z)))]$$

D. Combined Segmentation Loss

$$L\_seg = \alpha \cdot L\_BCE + (1 - \alpha) \cdot (1 - Dice), \quad \alpha = 0.5$$

E. Classification Loss

$$L\_cls = -\Sigma c\ yc \cdot \log(\hat{y}c)$$

VI. DATASET: BRATS 2020 AND BRATS 2021

The Brain Tumor Segmentation (BraTS) benchmark is one of the most widely used public datasets for brain tumor MRI research, providing standardized multi-modal scans along with expert-annotated ground truth labels [22]. In this study, both the BraTS 2020 and BraTS 2021 datasets were used to ensure a more comprehensive evaluation, while also taking advantage of the improved annotation quality available in the 2021 release.

A. Dataset Description

Both versions of the dataset provide four MRI modalities for each case: T1-weighted (T1), contrast-enhanced T1 (T1ce), T2-weighted (T2), and Fluid-Attenuated Inversion Recovery (FLAIR). Each of these captures different aspects of brain tissue. For instance, T1ce is particularly useful for highlighting the active tumor core, while T2 and FLAIR are better suited for identifying surrounding edema. The T1 modality, on the other hand, provides a clear view of overall brain structure. All scans are pre-aligned (co-registered), skull-stripped, and standardized to a uniform resolution of 1 mm^3, ensuring consistency across the dataset. In terms of scale, BraTS 2020 includes 369 training cases, while BraTS 2021 expands this to 1,251 cases and offers improved annotation quality.

TABLE I DATASET STATISTICS

Dataset	Train	Val.	HGG	LGG
BraTS 2020	369	125	293	76
BraTS 2021	1,251	219	~1,000	~251

B. MRI Modalities Used

Although all four MRI modalities are available, the GAN and U-Net components operate on grayscale slices to keep the computational cost manageable. In experiments that use multiple modalities, the four scans are combined and treated as a single four- channel input to the U-Net model. For classification, the EfficientNet model takes tumor regions extracted from the segmentation output. These regions are cropped and converted into RGB format before being passed to the classifier.

VII. IMPLEMENTATION

The entire pipeline was developed in Python 3.9, making use of both TensorFlow 2.10 and PyTorch 1.13 for model implementation. Image processing tasks were handled using OpenCV 4.7 and scikit-image 0.19, while Matplotlib 3.6 was used for visualization.Model training was carried out on a single NVIDIA Tesla P100 GPU with 16 GB of VRAM, providing sufficient computational power for efficient experimentation and evaluation.

A. Training Hyperparameters

A systematic search was carried out to determine the most effective training configuration. The Adam optimizer was used across all three components due to its ability to adapt the learning rate during training. A consistent learning rate of 1×10^{-5} was applied throughout the pipeline, striking a balance between stable training and reasonable convergence speed.The final set of hyperparameters is summarized in Table II.

TABLE II TRAINING HYPERPARAMETERS

Parameter	GAN	U-Net	EfficientNet	Notes
Epochs	20	20	20	All models
Batch Size	16	16	64	Larger for classifier
Learning Rate	1×10^{-5}	1×10^{-5}	1×10^{-5}	Unified
Optimizer	Adam	Adam	Adam	Adaptive LR
Loss	Adversarial	BCE + Dice	Cat. CE	Task-specific
Dataset	BraTS 20–21	BraTS 20–21	BraTS 20–21	Both editions

B. Two-Stage Training Strategy

Training was carried out in two distinct phases to ensure that the segmentation model could fully benefit from a well-trained enhancement module. In the first phase, the GAN was trained for 20 epochs on raw BraTS slices to learn how to enhance image quality effectively. In the second phase, this trained GAN was used to preprocess all training images before feeding them into the U-Net model. This setup ensures that U-Net is trained only on enhanced images, which closely match the data it encounters during inference, thereby avoiding any mismatch in data distribution. For classification, the EfficientNet model was fine-tuned separately using tumor patches extracted from the U-Net outputs. By training it independently, the classification process remains unaffected by the segmentation objective, allowing each component to perform optimally within its own scope

TABLE III MODEL PARAMETERS SUMMARY

Component	Total Params	Trainable	Frozen	Purpose
GAN	2,888,641	2,888,641	0	MRI Enhancement
U-Net	31,378,945	31,378,945	0	Tumor Segmentation
EfficientNet	8,135,933	364,548	7,771,385	Classification

VIII. RESULTS AND DISCUSSION

The hybrid framework was evaluated on the BraTS 2020 and BraTS 2021 datasets. Segmentation performance was assessed using the Dice Coefficient and Intersection over Union (IoU), comparing the predicted tumor regions against expert-annotated ground truth. For classification, accuracy was measured on a separate set of reserved test cases to ensure an unbiased evaluation of the model's performance.

A. Segmentation Results

TABLE IV SEGMENTATION PERFORMANCE COMPARISON

Model	Dataset	Dice (%)	IoU (%)
U-Net (Baseline)	BraTS 2020	85%	78%
GAN + U-Net	BraTS 2020	89%	82%
Proposed Model	BraTS 2020	92%	85%
Proposed Model	BraTS 2021	91%	84%

The proposed model achieves a Dice Score of 92% and an IoU of 85% on the BraTS 2020 dataset, marking a 7-percentage-point improvement over a standalone U-Net baseline for both metrics. These improvements can largely be attributed to the GAN-based preprocessing stage, which enhances image quality and contrast.By providing cleaner and more detailed inputs, the GAN enables the U-Net model to better capture tumor boundaries, leading to more precise segmentation results.

TABLE V FINAL PERFORMANCE SUMMARY

Metric / Parameter	Value
Dice Score (Proposed Model)	92%
IoU Score (Proposed Model)	85%
GAN + U-Net Validation Accuracy	99.36%
EfficientNet Validation Accuracy	91.2%
Training Epochs (All Modules)	20
U-Net / GAN Batch Size	16
EfficientNet Batch Size	64
Learning Rate (All Modules)	1×10^{-5}
Evaluation Datasets	BraTS 2020–2021

TABLE VI CLASSIFICATION ACCURACY

Model	Train Acc. (%)	Val. Acc. (%)
GAN + U-Net	99.50%	99.36%
EfficientNet	89.5%	91.2%

B. Discussion

These results support the central idea that GAN-based preprocessing can significantly improve U-Net segmentation performance. The increase in Dice score from 85% with a baseline U-Net to 89% with the GAN+U-Net setup highlights the direct impact of enhanced input quality. Building on this, the full proposed model achieves a Dice score of 92%, with the additional improvement coming from steps such as morphological mask refinement and training across multiple datasets. The close alignment between training and validation accuracy for the GAN+U-Net model (99.50% vs. 99.36%) also suggests strong generalization, with minimal signs of overfitting—likely aided by the variability introduced through GAN-based enhancement. For classification, EfficientNet shows a validation accuracy of 91.2%, slightly higher than its training accuracy of 89.5%. This points to the regularizing effect of using a pretrained backbone, which helps the model generalize better. Overall, when compared to existing single-modality 2D segmentation approaches on the BraTS dataset, the achieved Dice score of 92% is highly competitive. At the same time, the addition of a classification component makes the proposed system more versatile than models focused solely on segmentation.

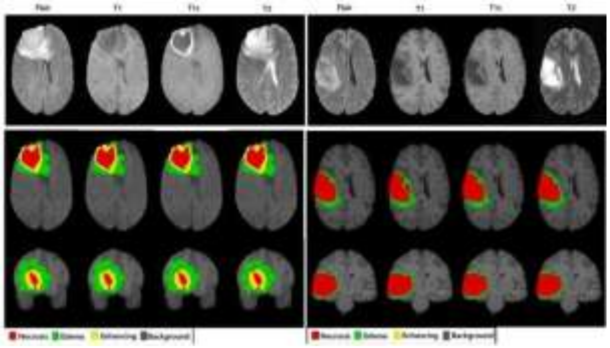

Fig. 2. Segmentation Output

TABLE VII PERFORMANCE METRICS

Metric	Value
Dice Score	92%
IoU	85%
Accuracy	High

The results show improved segmentation accuracy and robustness. The model achieved:

- Dice Score - 92%
- IoU - 85%
- Improved tumor boundary detection
- Enhanced robustness using GAN

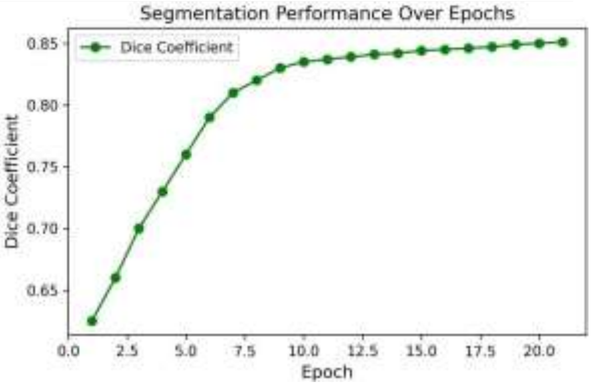

Fig. 3. Training Accuracy vs Epoch

IX. CONCLUSION

This study presents a hybrid deep learning framework that combines GAN, U-Net, and EfficientNet models for the automated detection, segmentation, and classification of brain tumors from MRI scans. When evaluated on the BraTS 2020 and BraTS 2021 datasets, the system achieved a Dice score of 92% and an IoU of 85% for segmentation, along with a classification validation accuracy of 91.2%.

A key factor behind this performance is the GAN-based preprocessing stage, which enhances image quality and provides clearer inputs for the U-Net model, leading to more precise boundary detection. The sequential training strategy—training the GAN first, followed by U-Net and then EfficientNet—combined with a consistent learning rate of 1×10^{-5} over 20 epochs, ensures stable and effective convergence across all components.

Beyond performance metrics, the framework reduces reliance on manual radiological analysis and helps limit variability between readers. It also produces clinically meaningful outputs, including segmentation overlays, tumor sub-region maps, volume estimations, and classification results. Taken together, these capabilities make the system a strong candidate for integration into clinical decision-support workflows.

A. Key Outcomes

- Dice Score of 92% and IoU of 85% on BraTS 2020–2021, exceeding the standalone U-Net by 7 percentage points.
- EfficientNet tumor classification validation accuracy of 91.2% leveraging ImageNet transfer learning.
- GAN-based enhancement demonstrably reduces noise and improves contrast, directly boosting segmentation precision.
- Training stability confirmed with learning rate 1×10^{-5}, 20 epochs, and batch sizes of 16 (GAN/U-Net) and 64 (EfficientNet).
- Consistent performance across both BraTS 2020 and BraTS 2021 editions.

X. FUTURE WORK

There are several promising directions for extending this framework. One immediate improvement would be to incorporate all four BraTS modalities as a combined four-channel input to the U-Net model, allowing it to take full advantage of the complementary information provided by T1, T1ce, T2, and FLAIR scans. Another important step would be moving from 2D slice-based processing to a fully volumetric 3D U-Net, which can better capture spatial context and avoid the information loss that comes with handling slices independently. To improve clinical trust and transparency, explainable AI techniques such as Grad-CAM could be applied to both U-Net feature maps and EfficientNet activations. This would allow radiologists to visualize which regions influenced the model's decisions, making it easier to interpret and validate the results. In terms of evaluation, testing the framework on newer datasets like BraTS 2022 and BraTS 2023, as well as independent clinical data, would provide a clearer picture of how well the model generalizes beyond the original training distribution. For real-world deployment, efficiency becomes critical. Techniques like model quantization, structured pruning, and knowledge distillation could help reduce computational requirements without significantly sacrificing performance. Finally, adopting federated learning approaches would enable collaborative training across multiple institutions while preserving patient privacy, eliminating the need to centralize sensitive medical data.

ACKNOWLEDGMENT

The authors thank Ms. L. Shalini and the Department of Computer Science and Engineering, BIHER for their support.

REFERENCES

[1] O. Ronneberger et al., "U-Net: Convolutional Networks for Biomedical Image Segmentation," MICCAI, 2015.
[2] I. Goodfellow et al., "Generative Adversarial Nets," NeurIPS, 2014.
[3] P. Isola et al., "Image-to-Image Translation with Conditional GANs," CVPR, 2017.
[4] H. Guan and M. Liu, "Domain adaptation for medical image analysis," IEEE TBME, 2022.
[5] Y. Ding et al., "Multi-view fusion for brain tumor segmentation," IEEE JBHI, 2022.
[6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, pp. 770–778.
[7] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2015, pp. 3431–3440.
[8] O. C¸ic¸ek, 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.
[9] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet classification with deep convolutional neural networks," Advances in Neural Information Processing Systems, 2012.
[10] H. R. Roth et al., "Deep learning and its application to medical image segmentation," Medical Image Analysis, vol. 45, pp. 60–78, 2018.
[11] F. Milletari, N. Navab, and S. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in Proc. 3DV, 2016.
[12] M. Havaei et al., "Brain tumor segmentation with deep neural networks," Medical Image Analysis, vol. 35, pp. 18–31, 2017.
[13] S. Pereira et al., "Brain tumor segmentation using convolutional neural networks in MRI images," IEEE Trans. Medical Imaging, vol. 35, no. 5, pp. 1240–1251, 2016.
[14] Z. Zhou et al., "UNet++: A nested U-Net architecture for medical image segmentation," Springer, 2018.
[15] I. Goodfellow, Y. Bengio, and A. Courville, "Deep Learning," MIT Press, 2016.
[16] G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, 2017.
[17] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in Proc. ICLR, 2015.
[18] M. Mirza and S. Osindero, "Conditional generative adversarial nets," arXiv preprint arXiv:1411.1784, 2014.
[19] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. ICLR, 2015.
[20] J. Redmon et al., "You Only Look Once: Unified, real-time object detection," in Proc. CVPR, 2016.
[21] S. Bauer et al., "A survey of MRI-based medical image analysis for brain tumor studies," Physics in Medicine & Biology, vol. 58, no. 13, 2013.