International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

A Review of Deep Learning and Artificial Intelligence Approaches for Brain Tumor Segmentation Using 3D MRI Data

Patil Aishwarya Shamkant
Research Scholar, Department of Computer Science, Malwanchal University, Indore
Dr. Manav Thakur
Supervisor, Department of Computer Science, Malwanchal University, Indore

Abstract
Brain tumor segmentation from three-dimensional (3D) magnetic resonance imaging (MRI) is a cornerstone task in computational neuro-oncology, supporting diagnosis, surgical planning, radiotherapy targeting, and longitudinal treatment monitoring. Manual delineation of tumor sub-regions by expert radiologists is accurate but time-consuming, subject to inter- and intra-observer variability, and impractical at scale. Over the past decade, artificial intelligence (AI), and in particular deep learning, has transformed the field by enabling automated, reproducible, and increasingly accurate segmentation of heterogeneous tumor tissue across multi-modal MRI sequences. This review surveys the evolution of deep learning and AI approaches for brain tumor segmentation using volumetric MRI data, with emphasis on convolutional neural network (CNN) architectures, the influential U-Net family and its 3D extensions, encoder-decoder designs, attention mechanisms, and the recent emergence of transformer-based and hybrid models. We examine the role of benchmark datasets, most notably the Brain Tumor Segmentation (BraTS) challenge, in driving methodological progress, and we discuss preprocessing, data augmentation, loss function design, and evaluation metrics such as the Dice similarity coefficient and Hausdorff distance. The review further considers persistent challenges including class imbalance, domain shift across scanners and institutions, limited annotated data, high computational cost of volumetric processing, and the clinical need for uncertainty quantification and interpretability. We organize the literature into four thematic strands: CNN-based volumetric architectures, the U-Net family and encoder-decoder refinements, attention and transformer-based methods, and approaches addressing data efficiency and generalization. We conclude that while deep learning has achieved expert-level performance on curated benchmarks, robust clinical translation requires advances in generalization, trust, computational efficiency, and integration into clinical workflows.

Keywords: brain tumor segmentation; deep learning; 3D MRI; convolutional neural networks; U-Net; transformers; BraTS; medical image analysis

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

1. Introduction

Brain tumors represent a clinically significant and heterogeneous group of neoplasms whose accurate characterization is essential for effective patient management. Gliomas, the most common primary malignant brain tumors in adults, exhibit diffuse infiltration and considerable variability in shape, location, intensity, and internal composition. Magnetic resonance imaging has become the principal modality for non-invasive assessment of these tumors because of its superior soft-tissue contrast and its ability to capture complementary information through multiple sequences. Within this context, the segmentation of tumor regions—delineating the boundaries of the whole tumor, the tumor core, and the enhancing component—forms the quantitative foundation upon which diagnosis, treatment planning, and outcome prediction are built. The transition from manual to automated segmentation, propelled by advances in artificial intelligence, has reshaped both research and clinical practice, motivating the comprehensive review presented here.

1.1 Clinical Importance of Brain Tumor Segmentation

Precise segmentation of brain tumors carries direct clinical consequences across the care continuum. In diagnosis, the volumetric extent and spatial distribution of tumor sub-regions inform grading and contribute to differential interpretation. In neurosurgical planning, segmentation maps guide the definition of resection margins, helping surgeons maximize tumor removal while sparing eloquent brain tissue. In radiation oncology, the gross tumor volume and clinical target volume derived from segmentation directly determine dose distribution, making the accuracy of boundary delineation a determinant of both therapeutic efficacy and toxicity. Beyond initial treatment, longitudinal segmentation enables quantitative monitoring of tumor progression or regression in response to therapy, supporting decisions about treatment continuation, escalation, or change. Manual segmentation by trained neuroradiologists remains the clinical reference standard, yet it is laborious, requiring substantial expert time per case, and it is affected by inter-observer and intra-observer variability that can introduce inconsistency into treatment decisions. These limitations create a compelling clinical rationale for automated, reproducible, and scalable segmentation methods.

1.2 The Role of 3D MRI and Multi-Modal Imaging

Brain MRI is inherently volumetric, acquired as a stack of contiguous two-dimensional slices that together represent a three-dimensional structure. Treating these acquisitions as true 3D volumes, rather than as independent slices, allows segmentation models to exploit spatial context across all three axes and to capture the continuity of tumor tissue through depth. Crucially, brain tumor assessment relies on multiple MRI sequences, each sensitive to different tissue properties. The standard multi-modal protocol used in research and many clinical settings comprises T1-weighted, contrast-enhanced T1-weighted, T2-weighted, and fluid-attenuated inversion recovery

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

sequences. Contrast-enhanced T1 imaging highlights regions where the blood-brain barrier is disrupted, delineating the enhancing tumor; T2 and fluid-attenuated inversion recovery sequences reveal peritumoral edema and infiltrative margins; and native T1 imaging provides anatomical reference. The fusion of these complementary channels supplies the rich, multi-dimensional input that modern deep learning models are designed to exploit, and the volumetric, multi-modal nature of the data fundamentally shapes the architectures discussed throughout this review.

1.3 Evolution from Traditional Methods to Deep Learning

Before the deep learning era, brain tumor segmentation depended on traditional image-processing and classical machine-learning techniques. Early approaches included intensity thresholding, region growing, and edge-based methods, which were intuitive but fragile in the presence of the noise, intensity inhomogeneity, and overlapping tissue intensities characteristic of brain MRI. More sophisticated pipelines employed hand-crafted features such as texture descriptors and intensity statistics, fed into classifiers like support vector machines and random forests, often combined with atlas-based registration or probabilistic tissue models. While these methods achieved moderate success, their performance was bounded by the expressiveness of manually engineered features and by their limited capacity to model the complex, non-linear relationships present in tumor imaging. The advent of deep learning marked a decisive shift: by learning hierarchical feature representations directly from data, convolutional neural networks removed the dependence on hand-crafted features and delivered substantial gains in accuracy. The introduction of architectures specifically suited to dense prediction, together with the availability of standardized benchmark datasets, catalyzed rapid and sustained progress in automated segmentation.

1.4 Scope and Organization of the Review

This review concentrates on deep learning and artificial intelligence approaches applied to brain tumor segmentation using volumetric MRI data, with particular attention to gliomas and the methodological lineage that the BraTS benchmark has shaped. The objective is to provide a structured synthesis that is useful to both computational researchers entering the field and clinicians seeking to understand the technical landscape. Following this introduction, the literature review is organized into four thematic subsections. The first examines convolutional neural network architectures designed for volumetric segmentation. The second traces the U-Net family and its encoder-decoder refinements, which have become the dominant design paradigm. The third addresses attention mechanisms and transformer-based models that capture long-range dependencies. The fourth surveys methods aimed at data efficiency, robustness, and generalization, including augmentation, transfer learning, and emerging learning paradigms. A

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

conclusion summarizes the principal findings, and a final section outlines promising directions for future research and clinical translation.

2. Literature Review

The literature on deep learning for brain tumor segmentation has expanded dramatically since the mid-2010s, propelled by architectural innovation, the standardization of benchmark datasets, and steady improvements in computational hardware. This section synthesizes that body of work along four thematic lines that together trace the methodological evolution of the field, from foundational convolutional architectures to the most recent transformer-based and data-efficient approaches.

2.1 Convolutional Neural Network Architectures for Volumetric Segmentation

Convolutional neural networks constitute the foundational architecture for automated brain tumor segmentation. Early deep learning approaches adapted two-dimensional CNNs to operate on individual MRI slices, treating segmentation as a patch-wise or pixel-wise classification problem. Havaei and colleagues introduced an influential two-pathway cascaded CNN that processed both local detail and broader contextual information, demonstrating that learned features could substantially outperform traditional hand-crafted pipelines on the BraTS benchmark (Havaei et al., 2017). Although effective, purely two-dimensional models discard the volumetric continuity inherent in MRI, motivating the development of architectures that operate directly on three-dimensional data.

Three-dimensional convolutional networks address this limitation by applying volumetric convolutions that aggregate spatial context along all three axes. Kamnitsas and colleagues proposed DeepMedic, a dual-pathway 3D CNN that combined multi-scale processing with a fully connected conditional random field for post-processing, achieving strong performance on brain lesion segmentation and establishing volumetric convolution as a powerful paradigm (Kamnitsas et al., 2017). The principal trade-off introduced by 3D architectures is computational: volumetric convolutions and the large memory footprint of full MRI volumes constrain batch sizes and input dimensions, frequently necessitating patch-based training in which the volume is processed in overlapping sub-regions. Researchers have explored numerous strategies to mitigate this cost, including efficient convolution designs, cascaded coarse-to-fine pipelines, and hybrid two-and-a-half-dimensional models that process adjacent slices to approximate volumetric context at reduced expense. Pereira and colleagues demonstrated that careful architectural choices, such as the use of small convolutional kernels stacked in depth, could yield both efficiency and accuracy improvements for glioma segmentation (Pereira et al., 2016). Collectively, this line of work established that the capacity of CNNs to learn discriminative, hierarchical representations from multi-modal MRI is central to modern segmentation performance, while also exposing the computational constraints that subsequent architectural innovations would seek to resolve.

2.2 The U-Net Family and Encoder-Decoder Architectures

The U-Net architecture, introduced by Ronneberger and colleagues for biomedical image segmentation, became the single most influential design in the field (Ronneberger et al., 2015). Its symmetric encoder-decoder structure couples a contracting path that captures context with an expanding path that enables precise localization, while skip connections transfer high-resolution features from encoder to decoder to recover spatial detail lost during downsampling. This design proved exceptionally well suited to dense segmentation tasks and was rapidly adopted and extended for brain tumor delineation. Çiçek and colleagues generalized the architecture to volumetric data with 3D U-Net, replacing two-dimensional operations with their three-dimensional counterparts to learn directly from sparsely annotated volumes (Çiçek et al., 2016). In parallel, Milletari and colleagues proposed V-Net, a fully volumetric encoder-decoder network that introduced residual connections and the Dice-based objective function, which directly optimizes the overlap metric used for evaluation and helps address the severe class imbalance between tumor and healthy tissue (Milletari et al., 2016).

Subsequent work refined the encoder-decoder paradigm in numerous directions. Isensee and colleagues demonstrated that a well-configured U-Net, with carefully tuned preprocessing, training schedule, and post-processing, could match or exceed more elaborate architectures, and their self-configuring nnU-Net framework established a robust, generalizable baseline that automatically adapts to dataset properties and has won multiple medical segmentation challenges (Isensee et al., 2021). Other refinements introduced nested and densely connected skip pathways to reduce the semantic gap between encoder and decoder features, deep supervision to improve gradient flow in deep volumetric networks, and cascaded multi-stage designs that first localize the whole tumor and then progressively refine the internal sub-regions. Myronenko introduced an encoder-decoder network augmented with a variational autoencoder branch that regularized the shared encoder, winning the BraTS 2018 challenge and illustrating how auxiliary reconstruction objectives can improve segmentation when training data are limited (Myronenko, 2019). The enduring dominance of the U-Net family reflects a favorable balance between representational power, training stability, and adaptability, and its core principles continue to underpin even the most recent hybrid and transformer-based systems.

2.3 Attention Mechanisms and Transformer-Based Models

A recurring limitation of purely convolutional architectures is the locality of the convolution operation, which restricts the receptive field and impedes the modeling of long-range spatial dependencies. Attention mechanisms emerged as a means to address this limitation by enabling networks to weight features according to their relevance, focusing representational capacity on

International Journal of Engineering, Science and Humanities
An international peer reviewed, refereed, open-access journal
Impact Factor 8.3 www.ijesh.com ISSN: 2250-3552

informative regions. Oktay and colleagues introduced Attention U-Net, which incorporated attention gates into the skip connections so that the decoder could suppress irrelevant background activations and emphasize salient tumor structures, improving segmentation of small and variable targets (Oktay et al., 2018). Channel and spatial attention modules were subsequently integrated into a variety of encoder-decoder backbones to similar effect, sharpening boundary delineation and improving robustness to anatomical variability.

The introduction of the transformer architecture, originally developed for natural language processing and built entirely on self-attention, prompted a substantial reorientation of the field. Self-attention models global relationships among all elements of an input directly, offering a principled mechanism for capturing the long-range context that convolutions struggle to represent. Hatamizadeh and colleagues proposed UNETR, which employs a pure transformer encoder to process volumetric MRI as a sequence of three-dimensional patches and connects it to a convolutional decoder, demonstrating competitive performance on volumetric medical segmentation including brain tumors (Hatamizadeh et al., 2022). Recognizing that the most effective designs often combine the global modeling capacity of transformers with the local inductive biases of convolutions, hybrid architectures gained prominence. Wang and colleagues introduced TransBTS, which embeds a transformer within an encoder-decoder framework to model global context over CNN-extracted features for multi-modal brain tumor segmentation (Wang et al., 2021). Building on hierarchical transformer design, Hatamizadeh and colleagues further proposed Swin UNETR, which uses a shifted-window transformer encoder to compute self-attention efficiently at multiple resolutions and achieved leading results on the BraTS benchmark (Hatamizadeh et al., 2021). These developments indicate a clear trajectory toward architectures that integrate convolutional and attention-based components to capture both fine local detail and global anatomical structure.

2.4 Data Efficiency, Augmentation, and Generalization

Beyond architectural innovation, a substantial portion of the literature addresses the practical challenges of training robust models with limited, imbalanced, and heterogeneous data. The standardization introduced by the BraTS challenge has been instrumental in enabling fair comparison and cumulative progress; Menze and colleagues established the original benchmark and its multi-region evaluation protocol, while Bakas and colleagues expanded the dataset and provided extensive expert annotations and survival information that became the de facto standard for the field (Menze et al., 2015; Bakas et al., 2017). Even with these resources, the quantity of expertly annotated volumetric data remains modest relative to the demands of deep networks, and severe class imbalance—tumor regions occupy a small fraction of each volume—complicates training.

Data augmentation is the most widely used remedy, encompassing spatial transformations such as rotation, scaling, flipping, and elastic deformation, as well as intensity perturbations that simulate acquisition variability. Loss function design constitutes a complementary strategy: the Dice loss and its variants, along with focal and compound losses, were developed specifically to counteract foreground-background imbalance and to align optimization with the overlap-based metrics used for evaluation. Transfer learning and self-supervised pretraining further improve data efficiency by initializing models with representations learned from larger or unlabeled corpora, reducing the annotation burden for the target task. A distinct and clinically critical challenge is generalization across the domain shift induced by different scanners, acquisition protocols, and patient populations, which frequently degrades the performance of models trained at a single institution. Domain adaptation and harmonization techniques, as well as multi-institutional training, have been investigated to improve cross-site robustness. Federated learning has attracted particular attention as a means of training on data distributed across institutions without centralizing sensitive patient information; Sheller and colleagues demonstrated that collaborative federated training could approach the performance of centralized training for brain tumor segmentation while preserving data privacy (Sheller et al., 2020). Complementing these efforts, attention to uncertainty estimation and model calibration has grown, reflecting the recognition that clinically deployable systems must not only be accurate but must also communicate the reliability of their predictions to support safe decision-making.

3. Conclusion

Deep learning has fundamentally reshaped brain tumor segmentation from volumetric MRI, advancing the field from fragile, hand-crafted pipelines to automated systems that approach expert-level performance on standardized benchmarks. This review traced that evolution across four interconnected strands. Convolutional neural networks, and especially their three-dimensional variants, established the capacity to learn discriminative hierarchical features directly from multi-modal MRI, while exposing the computational constraints inherent in volumetric processing. The U-Net family and its encoder-decoder descendants became the dominant design paradigm, with self-configuring frameworks such as nnU-Net demonstrating that disciplined engineering of a well-understood architecture can rival or surpass more complex alternatives. Attention mechanisms and transformer-based models subsequently addressed the locality limitation of convolutions, and the convergence on hybrid architectures that fuse convolutional and self-attention components reflects a maturing understanding of how to capture both local detail and global anatomical context. Finally, a rich body of work on augmentation, loss design, transfer and self-supervised learning, domain adaptation, and federated training has confronted the practical realities of limited, imbalanced, and heterogeneous clinical data.

Despite these substantial achievements, important gaps remain between benchmark performance and reliable clinical deployment. Models that excel on curated datasets often degrade under the domain shift encountered across scanners and institutions, and the computational demands of volumetric processing complicate integration into time-sensitive clinical workflows. Equally significant is the need for trustworthiness: clinical adoption requires not only accuracy but also calibrated uncertainty, interpretability, and demonstrable robustness across diverse patient populations. The field has therefore reached a stage at which methodological sophistication must be matched by attention to generalization, efficiency, transparency, and validation in real clinical settings. The trajectory of progress nonetheless gives strong reason for optimism that automated segmentation will become an increasingly dependable component of neuro-oncological care.

4. Future Work

Several research directions hold particular promise for closing the gap between current capabilities and robust clinical utility. Foremost among these is the development of methods that generalize reliably across institutions and acquisition settings. Federated and collaborative learning frameworks, which enable training across distributed datasets without sharing sensitive patient data, offer a path toward models exposed to the diversity of real-world imaging while respecting privacy and regulatory constraints. Coupled with domain adaptation and harmonization techniques, these approaches could substantially improve cross-site robustness, which remains a central obstacle to deployment.

A second direction concerns data efficiency. Because expert annotation of volumetric MRI is costly and scarce, self-supervised and semi-supervised learning that exploit large quantities of unlabeled imaging to learn transferable representations are likely to play an increasingly important role. Related to this is the emergence of foundation models for medical imaging—large models pretrained on broad datasets and adapted to specific tasks with minimal labeled data—which may provide powerful and reusable backbones for segmentation. A third priority is trustworthiness and clinical integration. Advances in explainable artificial intelligence, principled uncertainty quantification, and model calibration are needed so that automated systems can communicate the reliability of their outputs and earn the confidence of clinicians. Prospective, multi-institutional clinical validation, alongside attention to computational efficiency for practical deployment on standard hardware, will be essential to translate research advances into routine practice.

Finally, the integration of segmentation with broader clinical objectives represents a fertile avenue. Multimodal fusion that combines imaging with genomic, histopathological, and clinical data could enable richer characterization of tumors and support integrated tasks such as grading, molecular subtyping, treatment-response prediction, and survival estimation. Moving from isolated segmentation toward holistic, multimodal decision support, while maintaining transparency and rigorous validation, constitutes a compelling long-term goal for the field. Progress along these intertwined directions—generalization, data efficiency, trustworthiness, and clinical integration—will determine the extent to which deep learning fulfills its considerable promise in the care of patients with brain tumors.

References

1. Bakas, S., Akbari, H., Sotiras, A., Bilello, M., Rozycki, M., Kirby, J. S., Freymann, J. B., Farahani, K., & Davatzikos, C. (2017). Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels and radiomic features. Scientific Data, 4, 170117. https://doi.org/10.1038/sdata.2017.117
2. Çiçek, Ö., Abdulkadir, A., Lienkamp, S. S., Brox, T., & Ronneberger, O. (2016). 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2016 (pp. 424–432). Springer. https://doi.org/10.1007/978-3-319-46723-8_49
3. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H. R., & Xu, D. (2021). Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images. In Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (pp. 272–284). Springer. https://doi.org/10.1007/978-3-031-08999-2_22
4. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H. R., & Xu, D. (2022). UNETR: Transformers for 3D medical image segmentation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (pp. 574–584). IEEE
5. Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.-M., & Larochelle, H. (2017). Brain tumor segmentation with deep neural networks. Medical Image Analysis, 35, 18–31.
6. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: A self-configuring method for deep learning-based biomedical image segmentation. Nature Methods, 18(2), 203–211.
7. Kamnitsas, K., Ledig, C., Newcombe, V. F. J., Simpson, J. P., Kane, A. D., Menon, D. K., Rueckert, D., & Glocker, B. (2017). Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation. Medical Image Analysis, 36, 61–78. https://doi.org/10.1016/j.media.2016.10.004
8. Menze, B. H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., Burren, Y., Porz, N., Slotboom, J., Wiest, R., Lanczi, L., Gerstner, E., Weber, M.-A., Arbel, T., Avants, B. B., Ayache, N., Buendia, P., Collins, D. L., Cordier, N., … Van Leemput, K. (2015). The multimodal brain tumor image segmentation benchmark (BRATS). IEEE Transactions on Medical Imaging, 34(10), 1993–2024.
9. Milletari, F., Navab, N., & Ahmadi, S.-A. (2016). V-Net: Fully convolutional neural networks for volumetric medical image segmentation. In Proceedings of the Fourth International Conference on 3D Vision (3DV) (pp. 565–571). IEEE.

10. Myronenko, A. (2019). 3D MRI brain tumor segmentation using autoencoder regularization. In Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (pp. 311–320). Springer.
11. Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., Glocker, B., & Rueckert, D. (2018). Attention U-Net: Learning where to look for the pancreas. In Proceedings of Medical Imaging with Deep Learning (MIDL).
12. Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumor segmentation using convolutional neural networks in MRI images. IEEE Transactions on Medical Imaging, 35(5), 1240–1251.
13. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 (pp. 234–241). Springer. https://doi.org/10.1007/978-3-319-24574-4_28
14. Sheller, M. J., Edwards, B., Reina, G. A., Martin, J., Pati, S., Kotrotsou, A., Milchenko, M., Xu, W., Marcus, D., Colen, R. R., & Bakas, S. (2020). Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data. Scientific Reports, 10, 12598.
15. Wang, W., Chen, C., Ding, M., Yu, H., Zha, S., & Li, J. (2021). TransBTS: Multimodal brain tumor segmentation using transformer. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2021 (pp. 109–119). Springer.
16. Zhou, Z., Rahman Siddiquee, M. M., Tajbakhsh, N., & Liang, J. (2018). UNet++: A nested U-Net architecture for medical image segmentation. In Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support (pp. 3–11). Springer.
17. Zhang, J., Xie, Y., Wang, Y., & Xia, Y. (2021). Inter-slice context residual learning for 3D medical image segmentation. IEEE Transactions on Medical Imaging, 40(2), 661–672.
18. Tang, Y., Yang, D., Li, W., Roth, H. R., Landman, B., Xu, D., Nath, V., & Hatamizadeh, A. (2022). Self-supervised pre-training of Swin transformers for 3D medical image analysis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 20730–20740). IEEE. https://doi.org/10.1109/CVPR52688.2022.02007
19. Baid, U., Ghodasara, S., Mohan, S., Bilello, M., Calabrese, E., Colak, E., Farahani, K., Kalpathy-Cramer, J., Kitamura, F. C., Pati, S., Prevedello, L. M., Rudie, J. D., Sako, C., Shinohara, R. T., Bergquist, T., Chai, R., Eddy, J., Elliott, J., Reade, W., … Bakas, S. (2021). The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification. arXiv.
20. Liu, Z., Tong, L., Chen, L., Jiang, Z., Zhou, F., Zhang, Q., Zhang, X., Jin, Y., & Zhou, H. (2023). Deep learning based brain tumor segmentation: A survey. Complex & Intelligent Systems, 9(1), 1001–1026.