IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

NeuroLG-Seg: Attention-Guided Local–Global Feature Learning for Brain Tumor Segmentation in Super-Resolved MRI

Sudha Rani Chikkalwar*
Assistant Professor, Department of Computer Science and Informatics,
University College of Engineering and Technology,
Mahatma Gandhi University, Nalgonda, Telangana, India
https://ror.org/05n7bzj69 chsudharani.mgu@gmail.com
https://orcid.org/0009-0001-4929-9555

Daripelly Raju
Department of Computer Science,
University Post Graduate College, Secunderabad, Osmania University, Telangana, India
https://ror.org/030sjb889 daripellyraju2000@gmail.com
https://orcid.org/0009-0009-7119-5523

Publication History
Manuscript Reference No: IJIRIS/RS/Vol.12/Issue07/OCIS10084
Research Article | Open Access | Double-Blind Peer-Reviewed| Article ID: IJIRAE/RS/Vol.12/Issue07/OCIS10084
Received:09,September 2026,Revised:30,September 2026,Accepted:08,October 2026,Published Online:03,November 2026
https://www.ijiris.com/volumes/Vol12/iss-07/05.OCIS10084.pdf
Article Citation: Dr.Sudha,Daripelly(2026),NeuroLG-Seg: Attention-Guided Local–Global Feature Learning for Brain TumorSegmentation in Super-Resolved MRI. IJIRIS: International Journal of Innovative Research in Information Security, Volume 12, Issue 07 of 2026 pages 367-377 https://doi.org/10.26562/ijiris.2026.v1207.05
BibTeX Key:Dr.Sudha@2026NeuroLG-Seg
IJIRIS papers should be cited as IJIRIS(International Journal of Innovative Research in Information Security, AM Publications, India 2026, ISSN 2349-7017, https://doi.org/10.26562/ijiris.2026.v1207.05 The journal's official abbreviation is IJIRIS.
ORCiD: https://orcid.org/0009-0004-9398-7488
About the License: Copyright©2026 copyright by the authors. This article is an open access and license under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/)

Abstract: Accurate delineation of brain tumor sub-regions on magnetic resonance imaging (MRI) underpins surgical planning, radiotherapy targeting and response assessment, yet manual segmentation is slow and varies between raters. This study presents NeuroLG-Seg, an attention-guided network that jointly learns local and global tumor features from super resolved MRI. Recurrent residual convolutional blocks and UNet++-style dense nested skip pathways capture local boundary detail across scales, a multi-head self-attention bottleneck models global context, and attention gates guide the fusion of the two in the decoder. Input volumes are up-sampled from 1.0 mm to 0.5 mm isotropic resolution with a pretrained enhanced deep super-resolution (EDSR) network, and six hyperparameters are selected by a Hybrid Tuatara Numbat Algorithm (HTNA). On a 187-case held-out test split of the BraTS 2021 cohort (1,251 cases), NeuroLG-Seg achieved Dice similarity coefficients of 0.924, 0.891 and 0.873 for the whole tumor (WT), tumor core (TC) and enhancing tumor (ET), compared with 0.896, 0.855 and 0.827 for TransUNet, the strongest of five baselines trained under the same protocol. It also reduced the ET 95th-percentile Hausdorff distance from 4.83 mm to 3.21 mm (Wilcoxon signed-rank test with Bonferroni correction, $p < 0.01$). In a cumulative ablation, the local components accounted for about half of the 0.068 WT Dice gain over a U-Net baseline, while the global transformer bottleneck gave the largest single gain for the enhancing tumor. On 262 TCGA-GBM cases, evaluated as an additional cross-dataset test, the model retained 97.5–98.3% of its BraTS Dice scores. Comparison with self-configuring 3D pipelines and an explicit super-resolution ablation remain necessary next steps.
Keywords: Brain tumor segmentation; local–global feature learning; super-resolution; transformer; attention gate; metaheuristic hyperparameter optimization

I. INTRODUCTION

Primary brain and central nervous system tumors accounted for roughly 308,000 new cases worldwide in 2020 [1]. Glioblastoma, the most aggressive primary brain tumor, has a median survival of about 15 months even with maximal resection, radiotherapy and temozolomide [2]. Because management is guided by tumor extent and sub-region composition, the accurate and reproducible delineation of the tumor on MRI is central to surgical planning, radiotherapy target definition and longitudinal response assessment [3].

Segmentation is difficult. Gliomas vary widely in shape, size and signal across the four routine sequences (T1, contrast-enhanced T1 (T1Gd), T2 and T2-FLAIR); peritumoral edema blurs the boundary between infiltrated and healthy tissue; and clinical acquisition protocols limit spatial resolution. Manual delineation is time-consuming, and in the multi-institutional BRATS benchmark the agreement between expert raters, measured by Dice score, ranged from 74% to 85% [4]. Automated methods are therefore needed for consistent, scalable delineation.

E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

Deep learning has steadily raised segmentation accuracy. Fully convolutional networks [5] and the U-Net encoder–decoder with skip connections [6] were extended to volumetric data [7], [8], to residual [9] and recurrent residual [10] blocks, to dense nested skip pathways [11] and to attention gating [12]. More recently, hybrid convolution–transformer models [13]–[15] have added global self-attention, and self-configuring pipelines such as nnU-Net [16] have become strong general-purpose baselines.

Three questions remain comparatively under-explored. First, most segmentation networks are developed and evaluated at native acquisition resolution, while deep MRI super-resolution (SR) [17], [18] is increasingly available as a preprocessing step; how segmentation networks behave on SR-reconstructed inputs has received less attention. Second, the many hyperparameters of hybrid encoder–decoder networks are usually set manually or by grid search, although automated alternatives exist [19], [20]. Third, recurrent residual encoding, dense skip pathways, transformer attention and attention gating are usually studied separately, and their combined and incremental contributions have, to our knowledge, rarely been measured within a single controlled pipeline.

This study addresses these questions with NeuroLG-Seg, an attention-guided network for local–global feature learning whose hyperparameters are selected by a Hybrid Tuatara–Numbat Algorithm (HTNA). The contributions are as follows:

- NeuroLG-Seg, an encoder–decoder architecture that learns local features with recurrent residual convolutional blocks and dense nested skip pathways, models global context with a multi-head self-attention bottleneck, and uses attention gates to guide decoding, trained end to end on SR-reconstructed MRI.
- A population-based hybrid search that alternates an exploration operator and an exploitation operator to tune six hyperparameters, using a low-fidelity fitness (validation Dice after five epochs) to keep the search affordable.
- A controlled evaluation on BraTS 2021 against five baselines drawn from the architecture's lineage, with a cumulative ablation, paired non-parametric significance tests and an additional evaluation on TCGA-GBM.
- A transparent account of the search cost and of the limitations that bound the conclusions.

Section II reviews related work, Section III describes the data, architecture, optimizer and evaluation protocol, Section IV reports the results, Section V discusses them, and Section VI concludes.

II. RELATED WORK

A. Encoder–Decoder Networks for Brain Tumor Segmentation

Before deep learning, brain tumor segmentation relied on intensity thresholds, atlas registration and graph-based methods, which were interpretable but sensitive to tumor heterogeneity and scanner variation [21]. Classical pipelines of this kind remain in use; for example, median-filtered fuzzy c-means clustering followed by level-set evolution has been applied to MRI cerebral tissue segmentation [22]. Fully convolutional networks [5] and U-Net [6] enabled end-to-end learning from modest annotated datasets, and 3D U-Net [7] and V-Net [8] exploited inter-slice context. Residual learning [9] eased the optimization of deep networks; R2U-Net [10] placed recurrent residual convolutions in each stage to refine features iteratively without a proportional increase in parameters; and UNet++ [11] replaced plain skip connections with dense nested sub-networks that narrow the semantic gap between encoder and decoder features and support deep supervision. Convolutional neural networks have also been applied to image-level brain tumor detection from MRI, reporting high classification accuracy [23], [24]; such detectors label whole images rather than delineating tumor sub-regions voxel by voxel, which is the task addressed in this study.

B. Attention and Transformers

Attention U-Net [12] introduced trainable gates on the skip connections so that the decoder emphasizes relevant regions and suppresses background. Attention-augmented U-Net variants have since been applied beyond the brain, for example to liver segmentation in CT volumes [25]. Vision transformers [26] model global relationships through self-attention over image patches. TransUNet [13] combined a convolutional encoder with a transformer to capture both local detail and global context, Swin-Unet [14] used hierarchical shifted-window attention, and Swin UNETR [15] applied a Swin transformer encoder to 3D brain tumor segmentation. nnU-Net [16] showed that a carefully self-configured U-Net remains highly competitive, which makes it an important reference point for any new architecture.

C. Super-Resolution as Preprocessing

Deep SR progressed from shallow convolutional models [27] to deep residual networks such as EDSR [28] and to adversarial approaches [29]. In MRI, ensemble learning with complementary priors [17] and self-supervised anti-aliasing and super-resolution (SMORE) [18] recover detail from low-resolution acquisitions. Whether downstream segmentation networks can exploit the added spatial detail, and whether SR artifacts harm them, has been examined much less.

D. Hyperparameter Optimization

Bayesian optimization [19] and random search [20] are established alternatives to manual tuning, and population-based metaheuristics such as particle swarm optimization [30] and the grey wolf optimizer [31] are widely used because they need no gradient information. Since no optimizer is best on every problem [32], hybrids that combine a broad exploration operator with a focused exploitation operator are a common design. The HTNA used here draws its exploration and exploitation operators from the Tuatara [33] and Numbat [34] optimization algorithms.

E. Summary

The literature provides each building block separately, but we found few reports that (i) combine recurrent residual encoding, dense nested skips, a transformer bottleneck and attention gating in one network, (ii) tune such a network with an automated population-based search, and (iii) train and evaluate it on SR-reconstructed MRI with a component-wise ablation. This study contributes evidence on all three points within the controlled setting described below.

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

III. MATERIALS AND METHODS

A. Datasets

1) BraTS 2021: The BraTS 2021 segmentation cohort [4], [35], [36] used in this study comprised 1,251 adult glioma cases with four co-registered structural MRI modalities: native T1-weighted (T1), post-contrast T1-weighted (T1Gd), T2-weighted (T2), and T2-weighted fluid-attenuated inversion recovery (T2-FLAIR). The dataset provides manually generated tumor annotations for the relevant tumor sub-regions, from which the whole tumor (WT), tumor core (TC), and enhancing tumor (ET) regions were derived following the BraTS convention. The cases were partitioned at the patient level into 876 training cases (70%), 188 validation cases (15%), and 187 test cases (15%). A fixed random seed was used to ensure reproducibility of the patient-level split.
2) TCGA-GBM: An additional evaluation was performed using 262 glioblastoma cases from the TCGA-GBM collection [37] of The Cancer Imaging Archive (TCIA) [38]. The cohort contains multi-parametric MRI acquired under heterogeneous clinical imaging protocols. Tumor segmentation annotations associated with the TCGA-GBM/BraTS resources [36], [39] were used as reference labels for evaluating the whole tumor (WT), tumor core (TC), and enhancing tumor (ET) regions. Because TCGA-derived cases have contributed to BraTS datasets, complete patient-level independence between this evaluation cohort and the BraTS 2021 development cohort could not be established from the available metadata. Therefore, the TCGA-GBM experiment is treated as an additional cross-dataset evaluation rather than an independent external validation.

B. Preprocessing and Super-Resolution

The BraTS 2021 segmentation data were provided as brain-extracted, co-registered multi-parametric MRI volumes; therefore, skull stripping and spatial registration were not repeated for the BraTS cohort. Each subject contained four aligned MRI modalities: T1, T1Gd, T2 and T2-FLAIR. Intensity normalization was performed independently for each modality using z-score normalization over the non-background brain voxels, and the normalized volumes were then passed to the super-resolution stage (Fig. 1). A pretrained enhanced deep super-resolution (EDSR) network [28] was used as the super-resolution module. It increased the spatial sampling resolution by a factor of two, from 1.0 mm to 0.5 mm isotropic, so that the segmentation network was trained and tested consistently on super-resolved inputs. Training-time augmentation, applied to the training set only, comprised random axial rotations of up to 15°, horizontal and vertical flips, elastic deformation, and additive Gaussian noise with a standard deviation sampled uniformly from [0, 0.1].

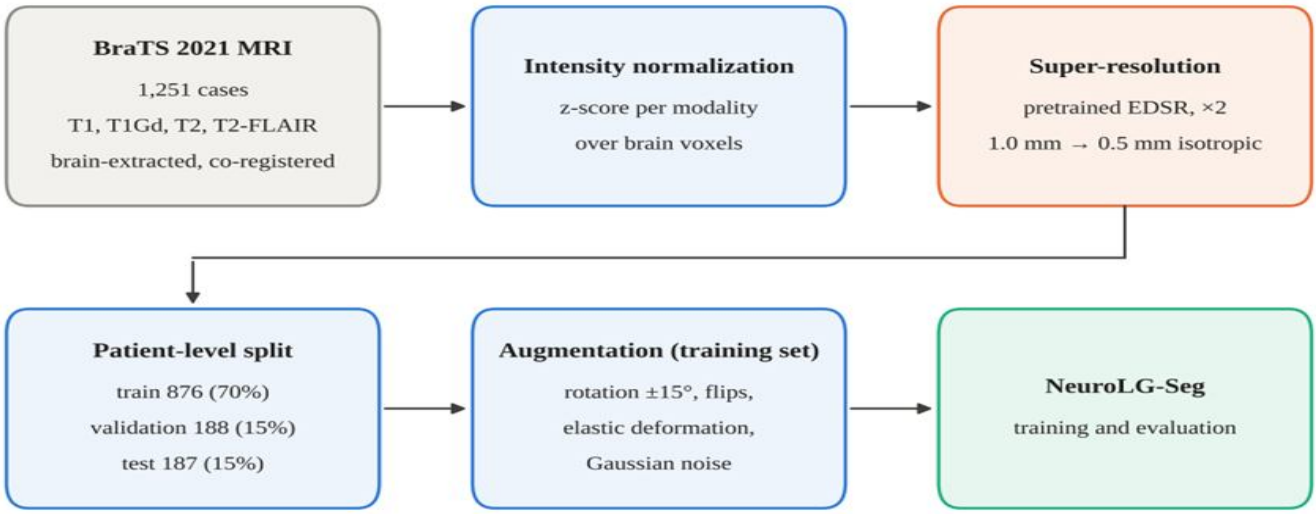


Fig.1 Preprocessing pipeline for BraTS 2021: per-modality z-score normalization of the brain-extracted, co-registered volumes, EDSR super-resolution (×2, 1.0 mm to 0.5 mm isotropic), patient-level split into training (876), validation (188) and test (187) cases, and training-set augmentation.

C. NeuroLG-Seg Architecture

NeuroLG-Seg combines a local pathway and a global pathway. In the local pathway, recurrent residual convolutions and dense nested skip connections capture fine boundary detail at several scales; in the global pathway, a transformer bottleneck models long-range context across the whole feature map; and attention gates guide how the two are fused in the decoder. The network (Fig. 2) has four encoder–decoder levels with C, 2C, 4C and 8C feature channels, where the base width C is one of the tuned hyperparameters (C = 48 in the final model, giving 48, 96, 192 and 384 channels). Each encoder level applies a recurrent residual convolutional block (RRCB) followed by max pooling. Within an RRCB at level l, a convolution with batch normalization and ReLU is unrolled over T = 2 recurrent steps [10]:

$$O_t = ReLU(W_l * (x_l + O_{t-1}) + b_l),\ t = 1,\dots,T,\ O_0 = 0 \quad (1)$$

and a residual shortcut gives the block output:

$$y_l = x_l + O_T \quad (2)$$

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

At the bottleneck, the encoded feature map is flattened into a sequence of feature tokens and processed by a transformer module with eight-head self-attention, learnable positional embeddings that retain spatial information, and a feed-forward sub-layer with GELU activation [13], [26]. This module captures long-range dependencies among spatially distributed tumor features; its output is reshaped into the spatial domain and passed to the nested decoder pathway. Skip connections follow the dense nested design of UNet++ [11]: intermediate nodes aggregate the features of all shallower nodes at the same level and the up-sampled features from the level below before passing them to the decoder. At each decoder stage an attention gate [12] computes a spatial coefficient α_i for each location i from the skip feature x_i and the up-sampled decoder (gating) feature g_i:

$$\alpha_i = \sigma(W_\alpha \cdot ReLU(W_g g_i + W_x x_i + b_g) + b_\alpha) \quad (3)$$

$$z_i = \alpha_i \cdot x_i \quad (4)$$

where σ is the sigmoid function, z_i is the gated skip feature passed to the decoder, and W_α, W_g, W_x, b_g and b_α are learned 1×1 convolution weights and biases. A final 1×1 convolution with softmax produces a four-class map (background, necrotic core, edema, enhancing tumor), and auxiliary deep-supervision outputs are attached to the intermediate decoder nodes during training. Relative to a conventional U-Net, the recurrent residual blocks, dense nested decoder pathways and transformer bottleneck add computational cost, which should be weighed against the accuracy gains reported in Section IV.

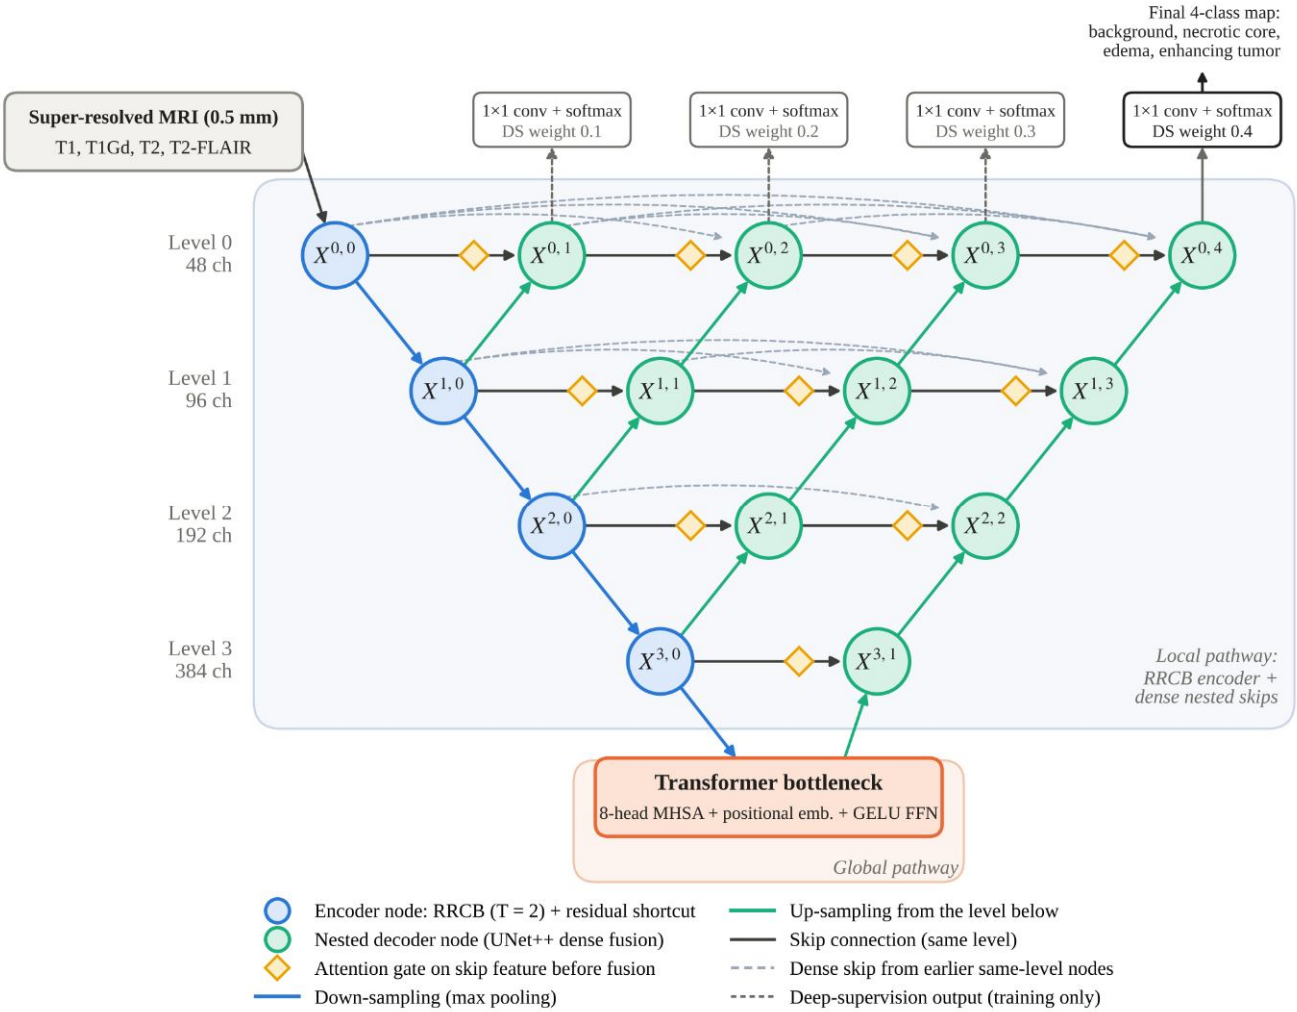


Fig.2 NeuroLG-Seg architecture. The local pathway (recurrent residual encoder and UNet++-style dense nested skips with attention gates) and the global pathway (transformer bottleneck) are fused in the nested decoder; deep-supervision outputs are used during training only.

D. Loss Function

Class imbalance was handled with a weighted combination of a soft Dice loss [8] and a focal cross-entropy loss [40]. With $p_{i,k}$ the predicted probability and $g_{i,k}$ the one-hot label of class k at voxel i, N voxels, K classes and a smoothing constant ε:

$$L_{Dice} = 1 - \frac{1}{K}\sum_k \frac{2\sum_i p_{i,k} g_{i,k} + \varepsilon}{\sum_i p_{i,k} + \sum_i g_{i,k} + \varepsilon} \quad (5)$$

$$L_{FCE} = -\frac{1}{N}\sum_i \sum_k (1 - p_{i,k})^\gamma \, g_{i,k} \log p_{i,k} \quad (6)$$

$$L = \lambda_{Dice} L_{Dice} + \lambda_{Focal} L_{FCE} \quad (7)$$

E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

In the experiments, the soft Dice and focal cross-entropy terms were weighted equally, with $\lambda_{Dice} = 0.5$ and $\lambda_{Focal} = 0.5$. The focal focusing parameter was set to $\gamma = 2.0$, and a smoothing constant $\varepsilon = 10^{-6}$ was used for numerical stability. Deep supervision was applied to the intermediate decoder outputs with normalized weights of 0.1, 0.2, 0.3 and 0.4 from the shallowest to the final decoder output.

E. Hybrid Tuatara–Numbat Hyperparameter Optimization

HTNA searched the six-dimensional space in Table I. A population of M = 30 candidates was initialized by Latin hypercube sampling [41] for even coverage. In each iteration t, every candidate X_i was moved by one of two operators (Fig. 3). The exploration (Tuatara) operator moved the candidate with a step size s(t) that decreased as the search progressed:

$$X_i(t+1) = X_i(t) + s(t) \cdot r \cdot (X_{best} - X_i(t)) \quad (8)$$

and the exploitation (Numbat) operator moved it toward the population best with a distance-normalized cooperative step:

$$X_i(t+1) = X_i(t) + \varphi \cdot \frac{X_{best} - X_i(t)}{d(X_i(t), X_{best}) + \varepsilon} \quad (9)$$

where r is a random vector sampled independently from U(0, 1) for each candidate and search dimension, $\varphi = 0.5$ is the communication factor and $d(\cdot,\cdot)$ is the Euclidean distance in the normalized search space. The step size decreased linearly as $s(t) = 1 - t/t_{max}$, where t is the current iteration and $t_{max} = 25$ is the maximum number of iterations. For each candidate, one of the two operators was selected according to the switching probability $p(t) = 1 - t/t_{max}$, so exploration was more likely in early iterations and exploitation dominated as the search approached termination (Fig. 3). After each move, candidates were clipped to the search bounds and integer hyperparameters were rounded. The fitness of a candidate was its validation Dice score after five training epochs, which traded evaluation fidelity for computational cost.

Table I - HTNA Search Space and the Configuration Selected for the Final Model

Hyperparameter	Range	Sampling	Selected
Learning rate	1×10^{-5} – 1×10^{-2}	Log-uniform	3.2×10^{-4}
Dropout rate	0.1 – 0.5	Uniform	0.28
Encoder base filters, C	16 – 64	Integer	48
Transformer heads	4 – 16	Integer	8
Attention-gate intermediate channels	8 – 64	Integer	32
Batch size	4 – 16	Integer	8

Population M = 30 (Latin hypercube initialization); fitness = validation Dice after five epochs.

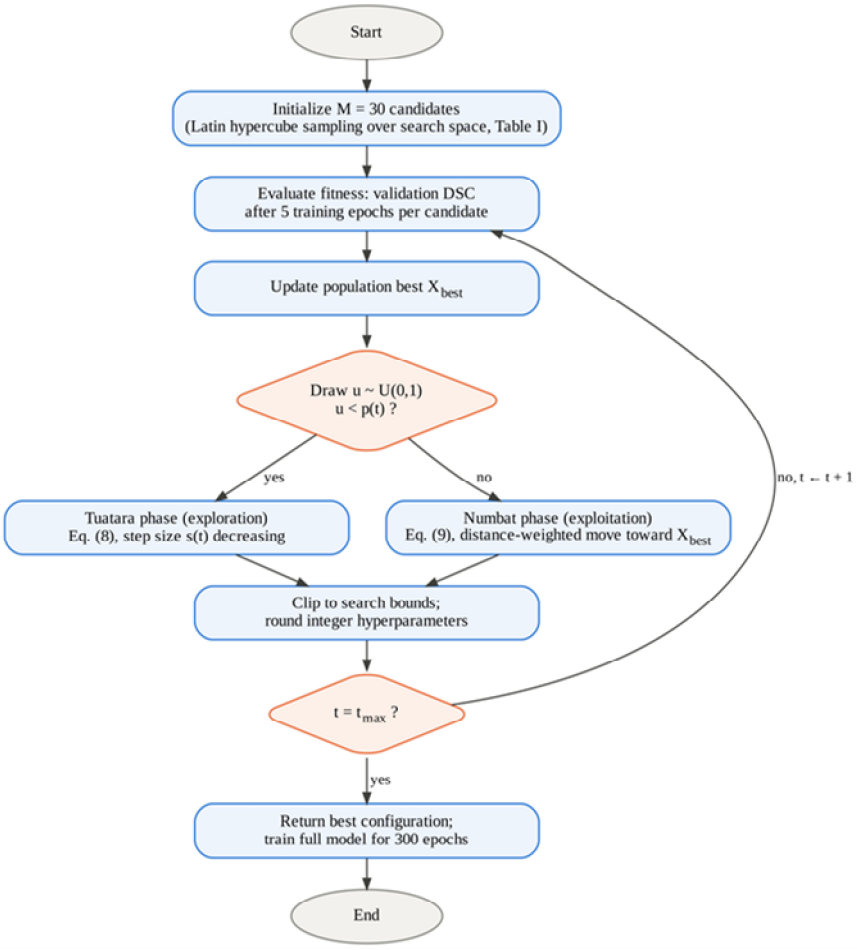


Fig.3 Workflow of the Hybrid Tuatara–Numbat Algorithm (HTNA) used to select the hyperparameters in Table I.

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

F. Training and Evaluation

The configuration selected by HTNA was used to train the full network from scratch for 300 epochs on an NVIDIA A100 (80 GB) GPU with PyTorch 2.1, using the Adam optimizer [42] and a cosine-annealing learning-rate schedule [43] with a ten-epoch linear warm-up. Segmentation quality was measured with the Dice similarity coefficient (DSC), the intersection over union (IoU), sensitivity and the 95th-percentile Hausdorff distance (HD95) [44]. With P and G the predicted and ground-truth voxel sets and ∂P, ∂G their boundaries:

$$DSC = \frac{2|P \cap G|}{|P| + |G|}, \quad IoU = \frac{|P \cap G|}{|P \cup G|}, \quad Sensitivity = \frac{TP}{TP + FN} \tag{10}$$

$$HD95 = \max\{P_{95}\, d(p, \partial G) : p \in \partial P,\ P_{95}\, d(g, \partial P) : g \in \partial G\} \tag{11}$$

where P□□ denotes the 95th percentile of the point-to-surface distances. Metrics were computed per case and averaged over the test set. Differences between the proposed model and each baseline were tested with the two-sided Wilcoxon signed-rank test on per-case scores [45], with Bonferroni correction for multiple comparisons at $\alpha = 0.05$. For a controlled comparison, all baseline models were re-implemented within the same PyTorch framework and trained using the identical patient-level train/validation/test partition. The same super-resolution preprocessing, intensity normalization, data augmentation, loss formulation, and 300-epoch training schedule were applied to all models. Model-specific architectural settings followed their original formulations, while the remaining training conditions were kept unchanged wherever applicable. This protocol ensures that performance differences reflect the segmentation architectures rather than differences in preprocessing or evaluation.

IV. RESULTS

A. Comparison with Baselines on BraTS 2021

Table II and Figs. 4–5 compare the proposed model with five baselines on the BraTS 2021 test split. The proposed model achieved DSC values of 0.924 (WT), 0.891 (TC) and 0.873 (ET). Relative to TransUNet, the strongest baseline, this is an absolute gain of 2.8, 3.6 and 4.6 percentage points (relative gains of 3.1%, 4.2% and 5.6%). The gain therefore grows from the whole tumor to the enhancing tumor, the smallest and most irregular sub-region. The ET HD95 fell from 4.83 mm to 3.21 mm, a reduction of 1.62 mm (34%), and the ET IoU rose from 0.711 to 0.776. All differences were statistically significant (Wilcoxon signed-rank test, Bonferroni-corrected $p < 0.01$).

Table II - Segmentation Performance on the BraTS 2021 Test Split

Model	WT DSC	TC DSC	ET DSC	ET IoU	ET HD95 (mm)	ET Sensitivity
U-Net [6]	0.856	0.801	0.762	0.618	7.21	0.831
R2U-Net [10]	0.871	0.819	0.783	0.643	6.54	0.849
Attention U-Net [12]	0.882	0.836	0.801	0.672	5.88	0.861
UNet++ [11]	0.894	0.852	0.822	0.703	5.12	0.874
TransUNet [13]	0.896	0.855	0.827	0.711	4.83	0.879
Proposed	0.924	0.891	0.873	0.776	3.21	0.912

WT, whole tumor; TC, tumor core; ET, enhancing tumor. IoU and sensitivity are reported for the enhancing tumor. Best values are in bold.

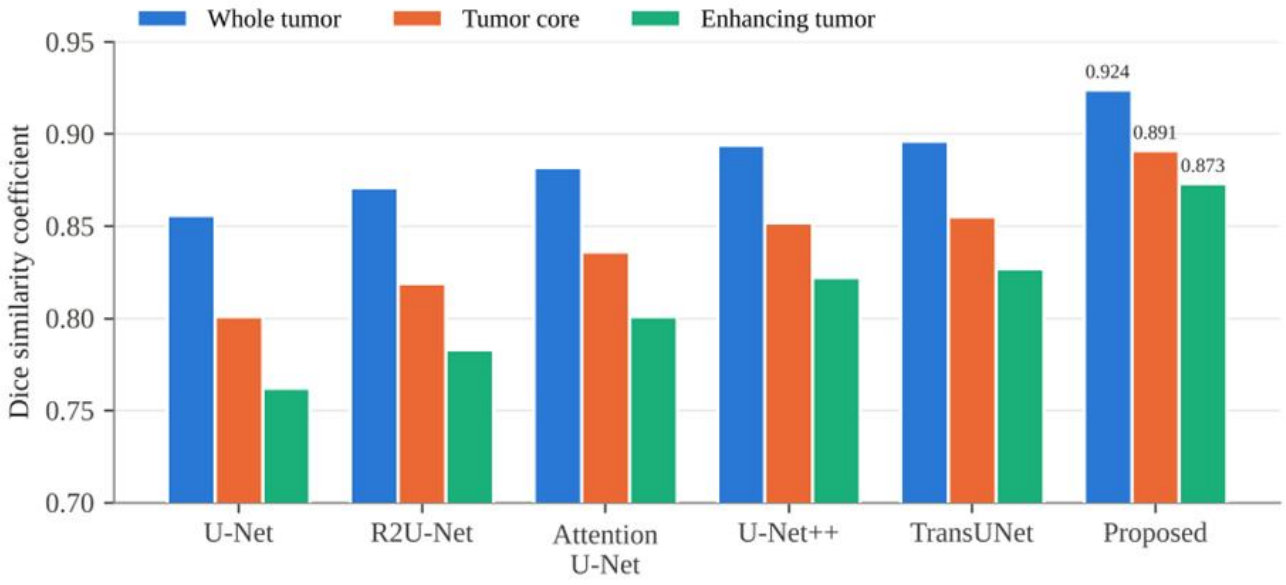


Fig 4 Dice similarity coefficient of each model for the whole tumor, tumor core and enhancing tumor on the BraTS 2021 test split; values for the proposed model are labeled (all values in Table II).

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

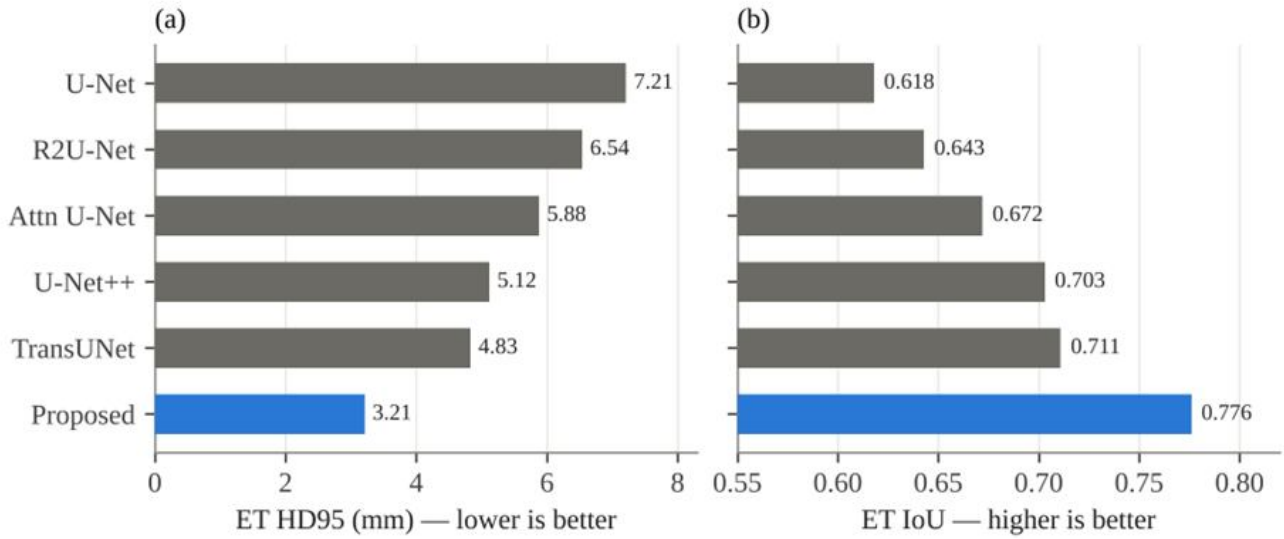


Fig.5 Enhancing-tumor boundary and overlap metrics on the BraTS 2021 test split: (a) HD95 in mm (lower is better); (b) IoU (higher is better).

B. Ablation Study

Table III and Fig. 6 report a cumulative ablation in which components are added one at a time to a U-Net baseline. In WT DSC, the recurrent residual blocks added 1.9 percentage points, the dense skip pathways 1.6, the transformer bottleneck 1.2, the attention gates 0.8 and HTNA tuning 0.9. For the enhancing tumor the pattern differs: the transformer bottleneck gave the largest single gain (2.7 points, from 0.804 to 0.831) and the largest single HD95 reduction (0.75 mm). Each addition was statistically significant ($p < 0.01$). Taken together, the components raised WT DSC by 6.8 points and ET DSC by 11.1 points over the baseline, and reduced ET HD95 from 7.21 mm to 3.21 mm. Because the components were added cumulatively in a fixed order, each increment is conditional on the components already present, and the ablation does not show what would be lost by removing any single component from the full model.

Table III - Cumulative Ablation on the BraTS 2021 Test Split

Configuration	WT DSC	ET DSC	ET HD95 (mm)	Δ WT DSC (points)
Baseline U-Net	0.856	0.762	7.21	—
+ Recurrent residual blocks	0.875	0.783	6.44	+1.9
+ Dense skip pathways	0.891	0.804	5.67	+1.6
+ Transformer bottleneck	0.903	0.831	4.92	+1.2
+ Attention gates	0.911	0.848	4.31	+0.8
+ HTNA hyperparameters	0.920	0.861	3.74	+0.9
+ Deep supervision (Full model)	0.924	0.873	3.21	+0.4

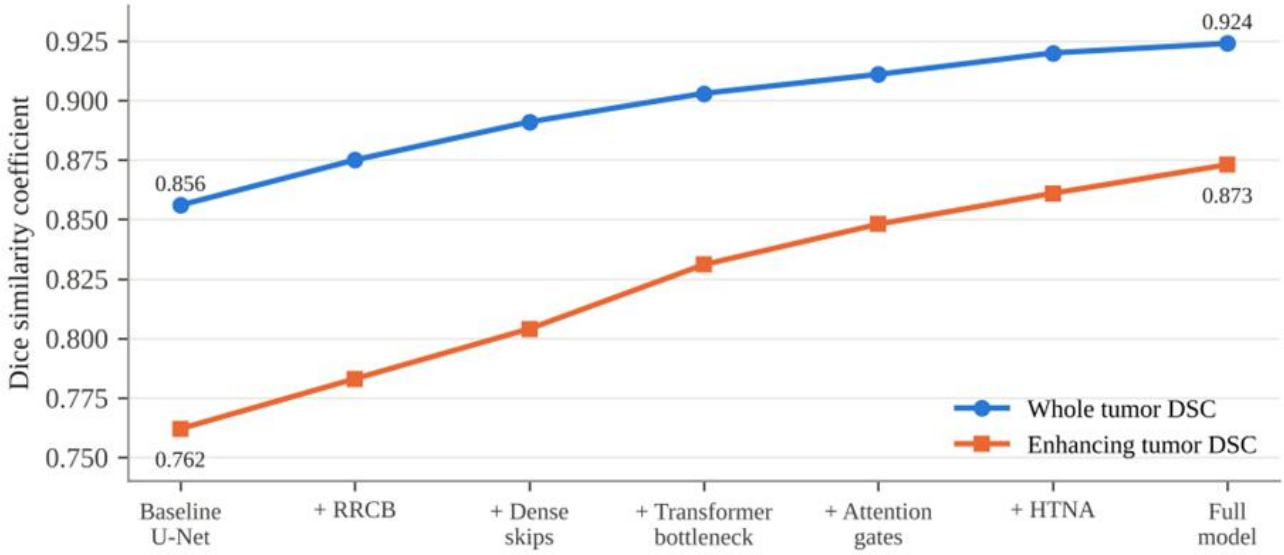


Fig.6 Cumulative ablation: whole-tumor and enhancing-tumor DSC as each component is added (values in Table III).

IJIRIS: International Journal of Innovative Research in Information Security
Volume 12, Issue 07, October 2026
E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

C. Hyperparameter Optimization

HTNA selected a learning rate of 3.2×10^{-4}, a dropout rate of 0.28, a base width of C = 48 filters and eight transformer heads (Table I). With this configuration, the model reached 90% of its peak validation DSC by epoch 85, compared with epoch 124 for the manually tuned configuration, about 31% fewer epochs (Fig. 7). The training and validation curves tracked each other closely, with a loss difference below 0.02 throughout training. The search itself took about 340 GPU-hours.

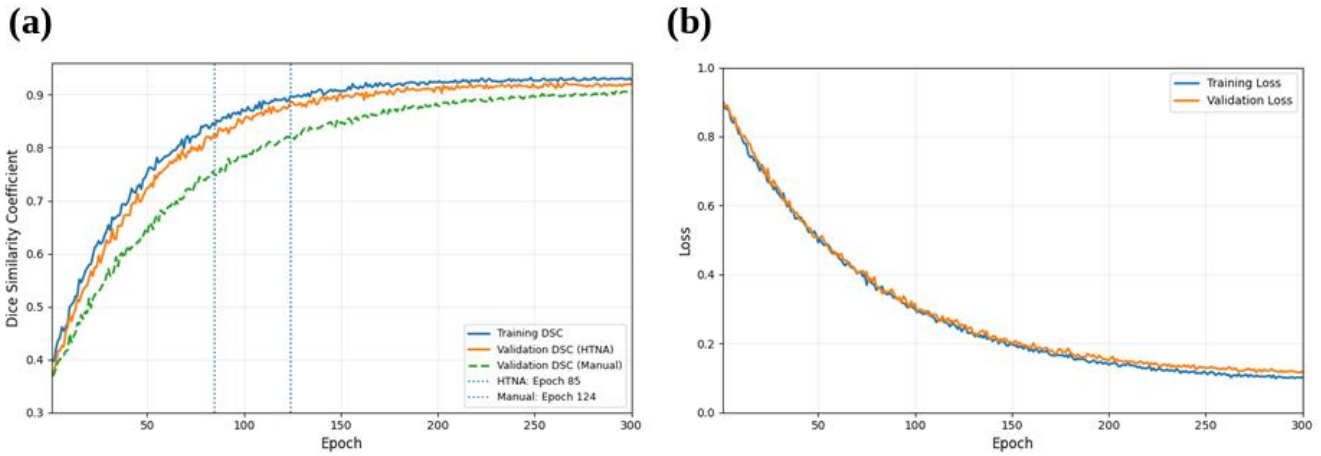


Fig 7 (a) Training and validation DSC of the HTNA-selected configuration and validation DSC of the manually tuned configuration over 300 epochs; dotted lines mark the epochs at which each reaches 90% of its peak validation DSC (85 and 124). **(b)** Training and validation loss of the HTNA-selected configuration.

D. Additional Evaluation on TCGA-GBM

On the 262 TCGA-GBM cases, the model trained on BraTS achieved DSC values of 0.907 (WT), 0.876 (TC) and 0.851 (ET), that is 98.2%, 98.3% and 97.5% of its BraTS test scores (absolute decreases of 1.7, 1.5 and 2.2 points). The largest decrease was for the enhancing tumor, which is consistent with that sub-region being the most sensitive to differences in contrast-enhanced acquisition.

V. DISCUSSION

A. Interpretation of the Gains

The ablation suggests that the local and global components of NeuroLG-Seg address different weaknesses. The recurrent residual blocks and dense skip pathways together contributed about half of the WT improvement (3.5 of 6.8 points), consistent with their role in preserving and refining local, multi-scale detail that successive pooling would otherwise discard [10], [11]. The transformer bottleneck contributed the largest single gain for the enhancing tumor in both DSC and HD95, consistent with the value of global context for a sub-region whose appearance depends on its relation to the surrounding necrosis and edema [13]. The attention gates gave a further, smaller gain, in line with their role in suppressing irrelevant activations [12]. HTNA added 0.9 points of WT DSC and faster convergence; given its cost of about 340 GPU-hours, its value relative to cheaper search strategies [19], [20] remains to be established.

B. Relation to Prior Work

The five baselines were chosen from the architecture's own lineage (U-Net, R2U-Net, Attention U-Net, UNet++ and TransUNet), which makes them well suited to isolating the effect of each component but does not establish state-of-the-art performance. Strong 3D pipelines such as nnU-Net [16] and Swin UNETR [15] were not included, and because the test set is an internal split of the BraTS 2021 training data, the scores are not directly comparable with challenge leaderboard results [35]. The claims of this study are therefore limited to improvement within this controlled comparison.

C. Role of Super-Resolution

All models in this study were trained and evaluated on SR-reconstructed inputs, so the experiments show how the architecture performs on SR data but do not isolate the effect of SR itself. SR can also introduce plausible but spurious fine structure at lesion boundaries, particularly where the SR network was trained on data unlike the target domain [18], [29]. A direct comparison of the same model trained on native-resolution and SR inputs is needed before any conclusion about the benefit of SR can be drawn.

D. Limitations

- The comparison does not include self-configuring or 3D transformer baselines such as nnU-Net and Swin UNETR.
- Results come from a single internal split; variability across random seeds or folds is not yet reported.
- The effect of SR preprocessing is not isolated, and SR may introduce artifacts at lesion boundaries.
- The ablation is cumulative and order-dependent, so it does not measure the contribution of each component to the full model.
- HTNA required about 340 GPU-hours and has not been compared with random search or standard optimizers under an equal budget.
- BraTS 2021 and TCGA-GBM are dominated by high-grade gliomas scanned at 1.5 T or 3 T in academic centers; generalization to low-grade gliomas, meningiomas, pediatric tumors and community-hospital scanners is untested.

E-ISSN: 2349-7017
P-ISSN: 2349-7009
https://www.ijiris.com/archives
https://doi.org/10.26562/ijiris.2026.v1207.04

- No prospective clinical evaluation has been performed, so the effect on surgical or radiotherapy outcomes is unknown.

E. Future Work

Future work should add nnU-Net and Swin UNETR baselines, report results across seeds or folds, isolate the effect of SR with a native-resolution control, and compare HTNA with random search and Bayesian optimization at equal cost. Jointly optimizing the SR and segmentation networks, so that segmentation errors guide reconstruction, is a natural extension, as are federated training across institutions and prospective evaluation against radiologist consensus.

VI. CONCLUSION

This paper presented NeuroLG-Seg, an attention-guided local–global network that combines recurrent residual encoding, dense nested skip pathways, a transformer bottleneck and attention-gated decoding, tuned with a hybrid metaheuristic and trained on super-resolved MRI. On the BraTS 2021 test split it reached DSC values of 0.924, 0.891 and 0.873 for the whole tumor, tumor core and enhancing tumor, and an enhancing-tumor HD95 of 3.21 mm, improving significantly on five baselines from its lineage. The cumulative ablation indicates that each component adds accuracy, with the transformer bottleneck mattering most for the enhancing tumor. Performance was largely retained on TCGA-GBM. Comparison with nnU-Net-class pipelines, a super-resolution ablation, multi-seed reporting and prospective clinical evaluation are the next steps toward establishing the method's practical value.

VII. DECLARATIONS

1. Data Availability: BraTS 2021 data are available from the RSNA-ASNR-MICCAI BraTS 2021 challenge, and TCGA-GBM data from The Cancer Imaging Archive [38].
2. Code Availability: Code will be made available from the corresponding author on reasonable request.
3. Ethics Approval: Only publicly available, de-identified data were used; no additional ethical approval was required.
4. Funding: This research received no external funding.
5. Competing Interests: The authors declare no conflict of interest.

VIII. AUTHOR CONTRIBUTION STATEMENT

Conceptualization: Dr.Sudha Rani Chikkalwar*, Daripelly Raju
Literature Review and Methodology design: Dr.Sudha Rani Chikkalwar*, Daripelly Raju
Software: Daripelly Raju
Validation: Daripelly Raju
Formal Analysis: Daripelly Raju
Investigation: Daripelly Raju
Resources: Dr.Sudha Rani Chikkalwar*
Data Curation: Daripelly Raju
Writing original draft preparation: Daripelly Raju
Writing review and Editing: Dr.Sudha Rani Chikkalwar*
Visualization: Daripelly Raju
Supervision: Dr.Sudha Rani Chikkalwar*
Conflict of interest: The authors declare no conflicts of interest.
Data availability statement: Data supporting these findings are available within the article, at https://doi.org/10.26562/ijiris.2026.v1207.05, or upon request.
Supplementary materials: The supplementary materials are available at https://doi.org/10.26562/ijiris.2026.v1207.05
Publisher's note: AM Publications, India. IJIRIS (International Journal of Innovative Research in Information Security) Journals stays neutral with regard to jurisdictional claims in published maps and institutional affiliations. All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.
https://ampublications-india.com/

REFERENCES

1. H.Sung, J.Ferlay, R.L.Siegel, M.Laversanne, I.Soerjomataram, A.Jemal, and F. Bray, "Global cancer statistics 2020: GLOBOCAN estimates of incidence and mortality worldwide for 36 cancers in 185 countries," CA: A Cancer Journal for Clinicians, vol. 71, no. 3, pp. 209–249, 2021. https://doi.org/10.3322/caac.21660
2. R.Stupp, W.P.Mason, M.J.van den Bent, et al., "Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma," New England Journal of Medicine, vol. 352, no. 10, pp. 987–996, 2005. https://doi.org/10.1056/NEJMoa043330
3. D.N.Louis, A.Perry, P.Wesseling, D.J.Brat, I.A.Cree, D.Figarella-Branger, et al., "The 2021 WHO classification of tumors of the central nervous system: a summary," Neuro-Oncology, vol. 23, no. 8, pp. 1231–1251, 2021. https://doi.org/10.1093/neuonc/noab106
4. B.H.Menze, A.Jakab, S.Bauer, J.Kalpathy-Cramer, K.Farahani, J. Kirby, et al., "The multimodal brain tumor image segmentation benchmark (BRATS)," IEEE Trans. Med. Imaging, vol. 34, no. 10, pp. 1993–2024, 2015. https://doi.org/10.1109/TMI.2014.2377694

5. 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. https://doi.org/10.1109/CVPR.2015.7298965
6. O.Ronneberger, P. Fischer, and T. Brox, "U-Net: convolutional networks for biomedical image segmentation," in Proc. MICCAI, LNCS vol. 9351, Springer, 2015, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28
7. Ö.Çiçek, A.Abdulkadir, S.S.Lienkamp, T.Brox, and O.Ronneberger, "3D U-Net: learning dense volumetric segmentation from sparse annotation," in Proc. MICCAI, Springer, 2016, pp. 424–432. https://doi.org/10.1007/978-3-319-46723-8_49
8. F.Milletari, N.Navab, and S.-A.Ahmadi, "V-Net: fully convolutional neural networks for volumetric medical image segmentation," in Proc. 4th Int. Conf. 3D Vision (3DV), IEEE, 2016, pp. 565–571. https://doi.org/10.1109/3DV.2016.79
9. 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. https://doi.org/10.1109/CVPR.2016.90
10. M.Z.Alom, C.Yakopcic, M.Hasan, T.M.Taha, and V.K.Asari, "Recurrent residual U-Net for medical image segmentation," J. Medical Imaging, vol. 6, no. 1, Art. 014006, 2019. https://doi.org/10.1117/1.JMI.6.1.014006
11. Z. Zhou, M. M. Rahman Siddiquee, N. Tajbakhsh, and J. Liang, "UNet++: a nested U-Net architecture for medical image segmentation," in Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, LNCS vol. 11045, Springer, 2018, pp. 3–11. https://doi.org/10.1007/978-3-030-00889-5_1
12. O.Oktay, J.Schlemper, L.Le Folgoc, M. Lee, M. Heinrich, K. Misawa, et al., "Attention U-Net: learning where to look for the pancreas," arXiv:1804.03999, 2018.
13. J.Chen, Y.Lu, Q.Yu, X.Luo, E.Adeli, Y. Wang, L. Lu, A. L. Yuille, and Y. Zhou, "TransUNet: transformers make strong encoders for medical image segmentation," arXiv:2102.04306, 2021.
14. H.Cao, Y.Wang, J.Chen, D.Jiang, X.Zhang, Q.Tian, and M.Wang, "Swin-Unet: Unet-like pure transformer for medical image segmentation," arXiv:2105.05537, 2021.
15. A.Hatamizadeh, V.Nath, Y.Tang, D.Yang, H.R.Roth, and D.Xu, "Swin UNETR: Swin transformers for semantic segmentation of brain tumors in MRI images," in Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries (BrainLes 2021), LNCS vol. 12962, Springer, 2022, pp. 272–284. https://doi.org/10.1007/978-3-031-08999-2_22
16. F.Isensee, P.F.Jaeger, S.A.A.Kohl, J.Petersen, and K.H.Maier-Hein,"nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation," Nature Methods, vol. 18, no. 2, pp. 203–211, 2021. https://doi.org/10.1038/s41592-020-01008-z
17. Q.Lyu, H.Shan, and G.Wang, "MRI super-resolution with ensemble learning and complementary priors," IEEE Trans. Comput. Imaging, vol. 6, pp. 615–624, 2020. https://doi.org/10.1109/TCI.2020.2964201
18. C.Zhao, B.E.Dewey, D.L.Pham, P.A.Calabresi, D.S.Reich, and J.L.Prince, "SMORE: a self-supervised anti-aliasing and super-resolution algorithm for MRI using deep learning," IEEE Trans. Med. Imaging, vol. 40, no. 3, pp. 805–817, 2021. https://doi.org/10.1109/TMI.2020.3037187
19. J.Snoek, H.Larochelle, and R.P.Adams, "Practical Bayesian optimization of machine learning algorithms," in Advances in Neural Information Processing Systems (NeurIPS), vol. 25, 2012, pp. 2951–2959.
20. J.Bergstra and Y.Bengio, "Random search for hyper-parameter optimization," J. Machine Learning Research, vol. 13, pp. 281–305, 2012.
21. N.Gordillo, E.Montseny, and P.Sobrevilla, "State of the art survey on MRI brain tumor segmentation," Magnetic Resonance Imaging, vol. 31, no. 8, pp. 1426–1438, 2013. https://doi.org/10.1016/j.mri.2013.05.002
22. Y.Maolood, S.Lu, Y.E.A.Al-Salhi, S.Al Resheedi, and M.Ince, "Fuzzy segmentation of MRI cerebral tissue using level set algorithm," International Journal of Innovative Research in Information Security (IJIRIS), vol. 5, no. 5, pp. 25–35, 2018, https://doi.org/10.26562/IJIRIS.2018.JYIS10080
23. Nithya Kalyani, M.N.Kavya, A.Nalavadi, and M.Keerthana, "Brain tumor detection using convolutional neural network," International Journal of Innovative Research in Information Security (IJIRIS), vol. 9, no. 3, pp. 97–103, 2023, https://doi.org/10.26562/ijiris.2023.v0903.11
24. S. Indumalar, M.Sathik Raja, and R.Reshma, "Brain tumor detection from MRI image using deep learning," International Journal of Innovative Research in Information Security (IJIRIS), vol. 10, no. 2, pp. 61–64, 2024, https://doi.org/10.26562/ijiris.2024.v1002.04
25. V.Singhal, V.Raj, V.Bhati, and S.Kumar, "Smart healthcare liver segmentation image analysis using AI," International Journal of Innovative Research in Information Security (IJIRIS), vol. 10, no. 6, pp. 677–682, 2024, https://doi.org/10.26562/ijiris.2024.v1006.03
26. Dosovitskiy, L.Beyer, A.Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, et al., "An image is worth 16×16 words: transformers for image recognition at scale," in Proc. Int. Conf. Learning Representations (ICLR), 2021.
27. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," IEEE Trans. Pattern Anal. Mach. Intell., vol. 38, no. 2, pp. 295–307, 2016. https://doi.org/10.1109/TPAMI.2015.2439281
28. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in Proc. IEEE Conf. Computer Vision and Pattern Recognition Workshops (CVPRW), 2017, pp. 136–144. https://doi.org/10.1109/CVPRW.2017.151
29. Ledig, L.Theis, F.Huszár, J. Caballero, A. Cunningham, A. Acosta, et al., "Photo-realistic single image super-resolution using a generative adversarial network," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2017, pp. 4681–4690. https://doi.org/10.1109/CVPR.2017.19

30. Kennedy and R. Eberhart, "Particle swarm optimization," in Proc. Int. Conf. Neural Networks (ICNN'95), vol. 4, 1995, pp. 1942–1948. https://doi.org/10.1109/ICNN.1995.488968
31. S.Mirjalili, S.M.Mirjalili, and A.Lewis, "Grey wolf optimizer," Advances in Engineering Software, vol. 69, pp. 46–61, 2014. https://doi.org/10.1016/j.advengsoft.2013.12.007
32. D.H.Wolpert and W.G.Macready, "No free lunch theorems for optimization," IEEE Trans. Evol. Comput., vol. 1, no. 1, pp. 67–82, 1997. https://doi.org/10.1109/4235.585893
33. H.Migdady, R.Zitar, A.Smerat, Z.Montazeri, M.Dehghani, O.P.Malik, and K.Eguchi, "Tuatara optimization algorithm (TOA): a novel bio-inspired metaheuristic for constrained engineering design optimization," Int. J. Intelligent Engineering and Systems, vol. 19, pp. 421–435, 2026. https://doi.org/10.22266/ijies2026.0228.26
34. Al-Oun, S.Alomari, R.Zitar, M.Ismoilov, A.Smerat, W.Airbow, Z.Montazeri, M.Dehghani, O.P.Malik, and K.Eguchi, "Numbat optimization algorithm (NOA): a bio-inspired metaheuristic for solving optimization problems," Int. J. Intelligent Engineering and Systems, vol. 19, pp. 996–1015, 2026. https://doi.org/10.22266/ijies2026.0228.60
35. U.Baid, S.Ghodasara, S.Mohan, M.Bilello, E.Calabrese, E.Colak, et al., "The RSNA-ASNR-MICCAI BraTS 2021 benchmark on brain tumor segmentation and radiogenomic classification," arXiv:2107.02314, 2021.
36. S.Bakas, H.Akbari, A.Sotiras, M.Bilello, M.Rozycki, J.S.Kirby, et al., "Advancing The Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features," Scientific Data, vol. 4, Art. 170117, 2017. https://doi.org/10.1038/sdata.2017.117
37. Scarpace, T. Mikkelsen, S.Cha, S.Rao, S.Tekchandani, D.Gutman, et al., "Radiology data from The Cancer Genome Atlas Glioblastoma Multiforme (TCGA-GBM) collection," The Cancer Imaging Archive, 2016. [Online]. Available: https://doi.org/10.7937/K9/TCIA.2016.RNYFUYE9
38. Clark, B. Vendt, K. Smith, J. Freymann, J. Kirby, P. Koppel, et al., "The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository," J. Digital Imaging, vol. 26, no. 6, pp. 1045–1057, 2013. https://doi.org/10.1007/s10278-013-9622-7
39. S.Bakas, H.Akbari, A.Sotiras, M. Bilello, M. Rozycki, J. Kirby, et al., "Segmentation labels and radiomic features for the pre-operative scans of the TCGA-GBM collection," The Cancer Imaging Archive, 2017. [Online]. Available: https://doi.org/10.7937/K9/TCIA.2017.KLXWJJ1Q
40. T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in Proc. IEEE Int. Conf. Computer Vision (ICCV), 2017, pp. 2980–2988.
41. D. McKay, R. J. Beckman, and W. J. Conover, "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," Technometrics, vol. 21, no. 2, pp. 239–245, 1979. https://doi.org/10.1080/00401706.1979.10489755
42. D. P. Kingma and J. Ba, "Adam: a method for stochastic optimization," in Proc. Int. Conf. Learning Representations (ICLR), 2015.
43. Loshchilov and F. Hutter, "SGDR: stochastic gradient descent with warm restarts," in Proc. Int. Conf. Learning Representations (ICLR), 2017.
44. A.Taha and A. Hanbury, "Metrics for evaluating 3D medical image segmentation: analysis, selection, and tool," BMC Medical Imaging, vol. 15, Art. 29, 2015. https://doi.org/10.1186/s12880-015-0068-x
45. F. Wilcoxon, "Individual comparisons by ranking methods," Biometrics Bulletin, vol. 1, no. 6, pp. 80–83, 1945. https://doi.org/10.2307/3001968