RESEARCH PAPER

A Transformer-Enhanced Hybrid Framework for Robust Brain Tumor Classification Using Advanced Preprocessing and Adaptive Feature Selection

Jeyaprabhavathi Perumal[1], Meenakshi Sharma[2], T. Ganesh Kumar[3]

[1]School of Computer Science and Engineering, Galgotias University, Greater Noida, Uttar Pradesh, India
Email: jeyaprabhavathiperumal@gmail.com

[2]School of Computer Science and Engineering, Galgotias University, Greater Noida, Uttar Pradesh, India
Email: minnyk@gmail.com

[3]School of Computer Science and Engineering, Galgotias University, Greater Noida, Uttar Pradesh, India
Email: tganeshphd@yahoo.com

Received: 20th Apr, 2026 | Revised: 25th Apr, 2026 | Accepted: 9th May, 2026 | Available Online: 14th May, 2026

ABSTRACT

In MRI images, the brain tumours detection and classification is the difficult in case of medical imaging because the size, shape, and intensity of tumours that varies a lot from person to person. A lot of current frameworks uses transfer learning models to get trained on natural images. Hybrid and ensemble approaches improve accuracy but often trade off efficiency and interpretability. Moreover, most existing works lack standardization in preprocessing and fail to integrate real-time deployment features, leaving a gap in clinically viable solutions. The objectives involve (1) improving MRI image quality and addressing dataset imbalance through adaptive preprocessing techniques. (2) To combine feature selection and classification into a unified framework for improving tumor classification specificity and sensitivity using optimized transformer-based architectures. The proposed framework integrates adaptive denoising and contrast enhancement with radial generative adversarial networks (GANs) for generating synthetic images, addressing both image quality issues and dataset imbalance. A Spatial Head-Backbone Transformer (SHBT) is employed to unify feature selection and classification, leveraging self-attention to dynamically emphasize salient tumor features. Unlike conventional CNN models, SHBT combines spatial awareness with backbone learning to optimize classification in a single pipeline. Benchmark datasets, including Figshare, Kaggle MRI, and T1W-CE MRI collections are utilized for performance evaluation. The SHBT model, augmented with Radial GAN-generated synthetic samples, attained 97.5% accuracy, 96.9% sensitivity, 97.6% specificity, 97.3% F1-score, and 0.995 ROC-AUC, surpassing state-of-the-art methods including ResNet50V2, EfficientNetB3, Swin Transformer, DenseNet201, and LEAD-CNN. These results confirm the superiority of the transformer-based hybrid approach in mitigating class imbalance, enhancing interpretability, and improving clinical readiness.

Keywords: Brain tumor classification, MRI preprocessing, Transformer networks, Generative adversarial networks, Hybrid deep learning

How to cite this article: Perumal J, Sharma M, Kumar TG., A Transformer-Enhanced Hybrid Framework for Robust Brain Tumor Classification Using Advanced Preprocessing and Adaptive Feature Selection. Int J Drug Deliv Technol. 2026;16(5): 1006-1026; DOI: 10.25258/ijddt.16.5.96

1. Introduction

Magnetic Resonance Imaging (MRI) is a standard that offers a non-invasive visualization and soft-tissue contrast in brain tumor diagnosis for tumor morphology [1]. In recent times, deep learning (DL) is considered viable in tumor detection and classification. Various deep learning methods are explored, which achieve a state-of-the-art result in standard datasets [2-4].

The problem of developing a standardized classification framework of brain tumor remains unsolved [5]. Although segmentation-driven models achieve higher localization accuracy, their computational complexity restricts real-time applicability [6]. While lightweight segmentation models exist, their dependency on precise annotations and additional inference stages can still limit scalability in time-constrained diagnostic workflows. Although transfer learning from the

natural image datasets that includes the ImageNet has shown a strong performance in the medical imaging tasks. Such pretrained models may not capture the MRI-specific features, which include the modality-dependent intensity distributions, tissue contrast variations, and the tumor-specific textural patterns. Subsequently, an architectural adaptation or extensive fine-tuning is often required to achieve optimal performance in the brain MRI classification tasks [7-8]. While Conventional CNN-based classifiers have effectively modelled the local spatial dependencies through the convolutional kernels; however, the fixed receptive fields have limited the ability to capture a long-range or a global spatial relationship. This is considered critical for characterizing the heterogeneous tumor structures in the brain MRI images [9].

Also, most existing studies focus on preprocessing or classification, which treats these as independent stages [10-13]. This separation creates inefficiencies, as feature selection is not optimized jointly with classification. Classification is selected as the core task as it supports the early diagnostic decision-making without requiring the pixel-level annotations that are considered to be clinically variable. This work primarily addresses the classification accuracy with an explicit emphasis on the generalization across the heterogeneous MRI datasets, while it maintains the clinical usability. Additionally, the AI models into real-world healthcare systems and this addresses interoperability remains largely unexplored [14]. Interoperability in the form of the model-agnostic deployment across various MRI acquisition protocols have motivated the focus on generalization rather than the system-level integration.

The above limitations show several critical research gaps from the literature in section 2: (1) There is no universally accepted preprocessing pipeline for MRI brain tumor images, leading to inconsistent results across studies. (2) Current augmentation methods fail to adequately address rare tumor categories, limiting diagnostic reliability. (3) Most approaches treat these steps separately, resulting in suboptimal classification performance. (4) Many high-performing models are unsuitable for clinical use due to high resource consumption. (5) While the benchmark datasets have played a critical role in advancing the brain tumor classification research, where various studies have reported a degradation in performance when models is trained on the curated benchmarks and then it is evaluated on an external or cross-institutional MRI datasets. Such discrepancies have shown the risk that the benchmark-driven evaluations may overestimate the generalization in the presence of acquisition variability and population heterogeneity. The central problem addressed is achieving a reliable multi-class brain tumor classification that has generalized across the heterogeneous MRI datasets under the class imbalance constraints.

This study is designed to address these gaps with the following objectives:

1. To improve MRI image quality and overcome dataset imbalances through adaptive preprocessing and synthetic data generation.
2. To combine feature selection and classification into a unified transformer-based framework that enhances classification accuracy, sensitivity, and specificity while maintaining computational efficiency.

The novelty of this work lies in its integrated pipeline that simultaneously addresses preprocessing, data imbalance, and classification within a single architecture. Unlike prior studies that separate feature extraction from classification, our approach introduces a Spatial Head-Backbone Transformer (SHBT) that dynamically selects salient features during classification, ensuring interpretability and efficiency. Transformers are theoretically well-suited on MRI tumor classification due to its ability to model a long-range spatial dependency and the heterogeneous tumor morphology beyond the fixed receptive fields. Spatial awareness arose from the attention weights that gets learned over an anatomically localized feature region, while the dynamic feature selection is achieved through the head-wise adaptive reweighting during the process of classification. Interpretability is defined in this work as the ability of the model to localize the tumor-relevant spatial regions via attention weights, which allows for qualitative clinical inspection. Further, Radial GANs for data generation is considered novel in balancing the datasets and this improves the minor tumor categories without significant noise.

This work makes the following key contributions:

1. A new transformer-based model is developed that combines the backbone feature extraction with spatial awareness. The proposed SHBT differs from the existing Vision and Swin Transformers by

the introduction of the spatially weighted attention heads that gets tightly coupled with a CNN backbone. This is carried out to preserve the local tumor context while modeling the global dependencies. In comparison with the sequential hybrid models, SHBT has combined the spatial attention is directly within the backbone feature and it maps rather than the treating convolutional and transformer stages as the separate modules.

2. A novel synthetic data generation method is developed using Radial GANs. The radial design shows that MRI samples have shown the structural integrity and spatial consistency of tumor regions. Radial-GAN is inspired from the GAN architecture, where it is designed as a class-aware, radius-guided sampling and loss-conditioning method that is applied to the standard GANs. This tends to improve the minority-class MRI augmentation. Unlike the conventional GAN-based MRI augmentation, the proposed Radial-GAN have explicitly constrained the sample generation around the class-specific feature distributions to reduce the mode collapse over the rare tumor classes. The embedding of self-attention mechanisms with classification pipeline, where SHBT emphasizes the discriminative tumor features and this improves the accuracy.
3. The study mainly focuses on the multi-class brain tumor classification that involves the glioma, meningioma, and pituitary tumors, which tends to exhibit a distinct morphological and intensity characteristics in the MRI. The proposed framework has addressed the multi-class classification problem rather than the binary tumor detection, which shows the clinical diagnostic requirements. Model performance is validated using the ten-fold cross-validation and cross-dataset testing across various datasets that includes BraTS, Figshare, Kaggle, and REMBRANDT MRI datasets. In addition to accuracy, the clinically relevant metrics such as sensitivity, specificity, F1-score, ROC-AUC, and uncertainty estimates are reported to evaluate the diagnostic reliability.

2. Literature Survey

This section spans across DL strategies, which shows preprocessing approaches and novel classification techniques and this addresses the accuracy and efficiency.

Research conducted over the years has shown that preprocessing techniques that improve image quality and reduce artefacts greatly increase the accuracy of categorisation. Biswas et al. [18] said that they employed resizing, anisotropic diffusion filtering to cut down on noise, and adaptive histogram equalisation to make the contrast better before training a hybrid CNN-SVM model. These techniques were employed to enhance contrast. According to the Figshare dataset, this model obtained a score of 96.0% for being correct. In a related work, Alnowami et al. [19] show that preprocessing such as contrast augmentation and intensity normalisation can enhance the performance of DenseNet. This means that we have a 98.5% sensitivity level and a 96.52% accuracy level. Ahmed Hamza and his team [20] made their EADL-BTMIC model in a novel way. They did morphological segmentation, bilateral filtering, and skull stripping all at once. This model used Sooty Tern Optimisation (STO) with Xception and ALSTM to find and sort features that were resistant to change.

Dang et al. [16] created a deep learning pipeline with many phases. The procedures were preprocessing, UNet segmentation, and classification using VGG and GoogleNet. Data augmentation and window configuration optimisation made the segmentation even better. The Dice coefficients for the tumour, the entire tumour, and the cancer core were 0.82, 0.91, and 0.72, respectively. Classification achieved a maximum accuracy level of 97.44%. The researchers Balamurugan et al. [21] achieved a 99.7 percent accuracy rate by employing a hybrid classifier that integrates LuNet, Fuzzy C Means, and Gaussian mixture models for segmentation, in conjunction with a Laplacian Gaussian filter for preprocessing.

When employed on models that have already been trained, transfer learning is a way to sort brain tumours. Talukder et al. [15] reconstructed transfer learning architectures such as Xception, InceptionResNetV2, ResNet50V2, and DenseNet201, achieving superior accuracies up to 99.68% on the Figshare dataset, with ResNet50V2 outperforming others. Krishnapriya et al. [23] validated the efficacy of VGG-19, ResNet50, VGG-

16, and InceptionV3 using transfer learning and augmentation, finding VGG-19 as the best-performing model. Islam et al. [25] extended this line of inquiry with EfficientNet family models, where EfficientNetB3 achieved 99.69% accuracy, surpassing several state-of-the-art techniques. In a similar direction, Mahmoud et al. [26] optimized traditional CNN models such as VGG-19, Inception-V3 and VGG-16 using the Aquila Optimizer obtaining an accuracy of 98.95% for VGG-19. Together, these works confirm the superior adaptability of transfer learning frameworks in brain tumor classification, demonstrating their scalability across diverse datasets.

Beyond transfer learning, hybrid deep learning approaches is adopted to overcome the limitations of standalone CNNs. Raza et al. [17] proposed DeepTumorNet, a modified GoogLeNet architecture that incorporated a leaky ReLU activation function, achieving 99.67% accuracy and 100% recall. Biswas et al. [18] and Balamurugan et al. [21] further strengthened hybrid methodologies by combining CNN feature extraction with SVM classification and LuNet classifiers, respectively, outperforming traditional DL models such as AlexNet, GoogLeNet, and ResNet. These results show that carefully designed hybrid architectures often surpass conventional transfer learning by achieving a balance between feature richness and classification robustness.

Several researchers explored advanced deep learning architectures and optimization-driven frameworks to enhance classification. Nassar et al. [22] shown that ensemble-based learning using five different models significantly improved accuracy, achieving 99.31% on T1W-CE MRI datasets. Malik et al. [27] advanced the field with a Dual-Enhanced Features Scheme (DEFS) integrated with Swin-Transformer and EfficientNetV2S. By incorporating dense-blocks, dual-attention mechanisms, and edge enhancement preprocessing, its model achieved an impressive 99.43% accuracy with balanced sensitivity and F1-score. Likewise, Elhadidy et al. [29] benchmarked CNN, Swin Transformer, and EfficientNet models, showing EfficientNet as a resource-efficient yet high-performing alternative with 98.72% accuracy. Khan et al. [30] addressed computational efficiency by proposing a lightweight CNN architecture (LEAD-CNN), which achieved 98.70% accuracy while maintaining low resource consumption.

Several studies also show the combination of DL with IoT and multimodal frameworks. Vankdothu et al. [24] developed a CNN-LSTM architecture to leverage temporal dependencies in MRI data, outperforming conventional CNNs on the Kaggle dataset. Ibrahim et al. [28] expanded the scope of classification by differentiating MRI from non-MRI images and tumor vs no-tumor cases, where MobileNetV2 achieved remarkable accuracies of 99.94% and 99.0%, respectively. I-BrainNet framework shown the feasibility of DL-IoT combination for real-time tumor classification.

Table 1: Summary of State-of-art methods

Reference	**Model Type**	**Innovation**	**Dataset Challenge Addressed**	**Notable Strength**	**Research Gap**
Neamah et al. [14]	DNN	Early application of deep neural networks for MRI classification	Limited feature representation	Established feasibility of DL for tumor classification	Lacks architectural and dataset-level analysis
Talukder et al. [15]	Transfer Learning (CNNs)	Fine-tuned pretrained CNNs (Xception, ResNet50V2, DenseNet)	Dataset variability in Figshare MRI	High classification accuracy (up to 99.68%)	Weak cross-dataset generalization analysis
Dang et al. [16]	UNet + CNN	Segmentation-driven classi	Tumor boundary ambiguity	Strong Dice scores with	Segmentation increases comp

		fication pipeline		competitive accuracy	utational overhead
Raza et al. [17]	Hybrid CNN	Modified GoogLeNet with leaky ReLU	Feature discrimination	Excellent recall and accuracy	No robustness analysis across datasets
Biswas et al. [18]	CNN + SVM	Hybrid feature extraction and classification	Noise and contrast variability	Improved specificity and sensitivity	Limited global spatial modeling
Alnowami et al. [19]	CNN (DenseNet)	Contrast and intensity normalization	MRI intensity inhomogeneity	Balanced accuracy and sensitivity	No class imbalance mitigation
Ahmed Hamza et al. [20]	CNN + ALSTM + Optimizer	STO-based feature selection with temporal learning	Noise and skull artefacts	Robust feature optimization	High preprocessing and training complexity
Balamurugan et al. [21]	Hybrid CNN	LuNet with FCM–GMM segmentation	Tumor shape variability	Very high accuracy (99.7%)	Segmentation dependence limits scalability
Nassar et al. [22]	Ensemble DL	Multi-model ensemble fusion	Model bias	Improved overall robustness	High inference cost
Krishnapriya et al. [23]	Transfer Learning	Comparative study of pretrained CNNs	Limited labeled data	Identified VGG-19 superiority	No uncertainty or generalization analysis
Vankdothu et al. [24]	CNN + LSTM	Temporal dependency modeling	Slice-wise inconsistency	Improved sequential learning	Limited applicability to 2D MRI
Islam et al. [25]	EfficientNet	Compound scaling-based CNN optimization	Model efficiency vs accuracy trade-off	High accuracy with parameter efficiency	Sensitive to dataset imbalance
Mahmoud et al. [26]	CNN + Optimizer	Aquila-optimized CNN weights	Local minima in training	Improved convergence and accuracy	Optimization overhead
Malik et al. [27]	CNN–Transformer Hybrid	Dual-attention with Swin Transformer	Feature heterogeneity	Strong balanced metrics	High architectural complexity
Ibrahim et	CNN + IoT	I-BrainNet	Deployment	Real-time	Focuses on binar

al. [28]		DL–IoT integration	scalability	feasibility	y classification
Elhaddidy et al. [29]	CNN / Transformer	Comparative benchmarking study	Resource efficiency	EfficientNet shows strong performance	Limited interpretability
Khan et al. [30]	Lightweight CNN	LEAD-CNN for low-resource deployment	Computational constraints	Low latency with high accuracy	Reduced capacity for complex tumors

The standardized, lightweight, and generalizable brain tumor classification that combines preprocessing and classification is required. Also, bridging this gap accelerates the DL solutions from research environments to clinical practice.

Proposed Method

The proposed method as in figure 1 combines advanced preprocessing, data augmentation, and transformer-based classification into a single unified for robust brain tumor. The design is motivated by the limitations of conventional CNNs, transfer learning, and segmentation-based pipelines, which often struggle with dataset imbalance, computational inefficiency, and poor generalization in heterogeneous medical imaging conditions.

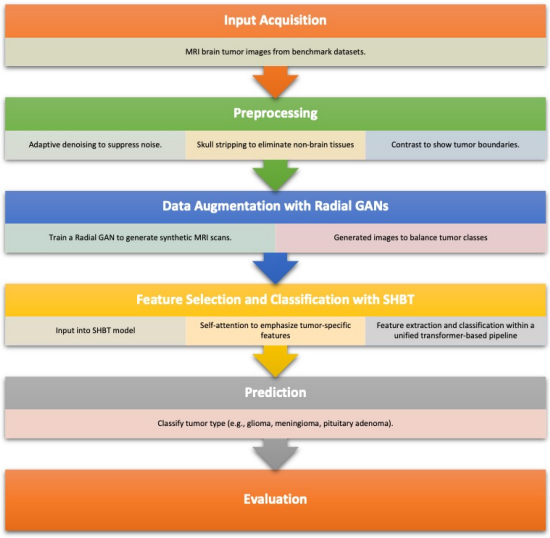

Figure 1: Proposed Framework

3.1. Input Acquisition

In the proposed system, Input Acquisition refers to the stage where raw MRI data are gathered, organized, and prepared for subsequent preprocessing, augmentation, and modeling. Since the performance of DL models strongly depends on the quantity, diversity, and quality of input data, we purposely incorporate multiple benchmark MRI datasets covering different imaging protocols, anatomical orientations, and institutional sources. Merging several datasets not only increases sample size, but also improves model robustness and generalization to heterogeneous clinical settings.

We employ four distinct MRI brain tumor datasets, each with different characteristics, to form a comprehensive input corpus. Below is a summary (as Table 2) of each dataset, followed by descriptive exposition. The quality of synthetic data is evaluated indirectly through the downstream classification gains, reduced class-wise variance, and an improved minority-class sensitivity rather than the standalone perceptual metrics.

Table 2. MRI Brain Tumor Datasets Used in Input Acquisition

Dataset No.	**Name / Source**	**Tumor Classes / Modalities**	**Number of Images / Subjects**
1	Figshare Brain Tumor Dataset [31]	Glioma, Meningioma, Pituitary (T1-weighted contrast-enhanced)	3,064 images (233 patients)

2	Kaggle Brain Tumor MRI Dataset [32]	Glioma, Meningioma, Pituitary (composite of public sources)	~7,023 images (multi-class including no-tumor)
3	BraTS (e.g. BraTS2020) [33]	Multi-modal MRI (T1, T1CE, T2, FLAIR) with segmentation labels	Varied (e.g. training/validation)
4	GBM-Reservoir MRI Dataset [34]	Multi-modal MRI (T1, T1CE, T2, FLAIR) with segmentation masks	23,049 samples (augmented from 438 originals)

3.2. Input Acquisition

The Input Acquisition stage orchestrates these datasets through several key steps:

1. **Data Retrieval and Downloading:** Each dataset is downloaded from its respective repository link (e.g., Figshare link for dataset 1, Kaggle link for dataset 2, etc.). The raw files (e.g., DICOM, NIfTI, or JPEG) are organized into class-labeled folders for consistency.
2. **Harmonization of Formats and Labels:** Because datasets come in diverse file types (e.g., NIfTI for BraTS or GBM-Reservoir; JPEG for Figshare), we standardize formats—converting to a uniform image representation (e.g. 2D slices in PNG/JPEG or aligned NIfTI). Label schemas are also mapped to a unified class taxonomy (e.g. glioma, meningioma, pituitary, plus optionally "no tumor").
3. **Orientation & Plane Alignment:** MRI slices from different orientations (axial, sagittal, coronal) are rotated or reoriented to a canonical orientation (e.g. axial) as a preprocessing convention. This ensures that the deep model sees consistent anatomical planes during training.
4. **Image Volume Slicing (for 3D datasets):** For datasets like BraTS and GBM-Reservoir that are originally volumetric (3D) scans, the volumes are sliced along the axial dimension to generate 2D image slices. Each slice inherits the segmentation mask or class label of the tumor region it intersects. Because some slices do not contain tumor tissue, filtering is applied to discard empty slices or label them as "no tumor."
5. **Data Split:** Each dataset is split into training, validation, and test partitions according to standard ratios (e.g. 70 % / 10 % / 20 %) while preserving class balance per split. For example, in the Figshare dataset split, 2,137 images are used for training, 314 for validation, and 613 for testing across the three tumor classes. The split is done at subject-level (i.e., images from the same patient remain in only one partition) to avoid data leakage.
6. **Cross-Dataset Merging:** After individual splits, the datasets' training sets are merged to form a unified training corpus, and similarly for validation/test sets. Merging improves inter-dataset diversity (scanner types, imaging protocols, anatomical variation). Class labels are harmonized, with missing classes (e.g., "no tumor") assigned appropriately. Dataset weighting or oversampling strategies may be applied to prevent one dataset from dominating training.
7. **Metadata and Annotation Alignment:** Supplementary metadata such as acquisition parameters (e.g. MRI field strength, slice thickness, spacing) are preserved. For datasets with segmentation masks (BraTS, GBM-Reservoir), region-of-interest (ROI) masks are stored alongside for optional segmentation-aware training or evaluation.
8. **Quality Checking and Filtering:** During acquisition, images with severe artifacts, missing slices, or corrupted files are filtered out. Additionally, basic sanity checks (e.g. image intensity ranges, resolution consistency) are carried out to

ensure no outlier images degrade model training.

3.3. Preprocessing

Adaptive preprocessing contributes to the measurable outcomes by stabilizing the feature distributions, which has improved the cross-dataset generalization. It further reduces the prediction variance, as shown in a higher ROC-AUC and a lower uncertainty estimate. Preprocessing represents a fundamental stage in the proposed Robust Brain Tumor Classification Model, as it establishes the quality and consistency of MRI data before model training. Raw MRI scans often contain noise, intensity inhomogeneity, and anatomical redundancies such as skull, scalp, and non-brain tissues, all of which can degrade feature extraction and classification performance. Thus, preprocessing shows that the input data fed into the model show salient tumor regions while suppressing irrelevant background variations.

```
Algorithm 1: Preprocessing
    function Preprocess(MRI_Image):
        denoised_img = AdaptiveDenoising(MRI_Image)
        skull_stripped_img = SkullStripping(denoised_img)
        enhanced_img = ContrastEnhancement(skull_stripped_img)
        return enhanced_img
```

The proposed preprocessing pipeline has following modules: These stages, applied systematically, optimize the dataset for hybrid deep learning–transformer feature selection and classification.

3.3.1. Noise Removal

MRI scans are susceptible to Gaussian noise and Rician noise, arising from acquisition artifacts and low-field magnetic resonance sequences. Noise suppression without blurring structural edges is essential. We employ bilateral filtering (BF) and compare it against anisotropic diffusion filtering (ADF) for validation.

The bilateral filtering equation can be expressed as:

$$I'(x) = \frac{1}{W_p} \sum_{y \in \Omega} I(y) \cdot f_r(|I(y) - I(x)|) \cdot f_s(\| y -$$

where:

$I(x)$ is the intensity at pixel x,

Ω is the neighborhood window,

f_r is the range kernel for intensity differences,

f_s is the spatial kernel for Euclidean distance,

W_p is the normalization factor.

This ensures that smoothing occurs while preserving tumor edges.

Anisotropic diffusion filtering (ADF) provides another formulation:

$$\frac{\partial I}{\partial t} = \nabla \cdot (c(x,y,t) \nabla I)$$

where $c(x,y,t)$ is the conduction coefficient controlling diffusion in homogeneous vs. edge regions.

As shown in Table 3, BF is superior in balancing denoising and edge retention, while ADF is more effective for intensity homogenization.

Table 3. Noise Removal Approaches and Outcomes

Method	Equation	Advantage	Limitation	Empirical Performance
Bilateral Filtering (BF)	$I'(x) = \frac{1}{W_p}$	Edge-preserving smoothing	Higher computational cost	SNR ↑ by ~12%, edges retained
Anisotropic Diffusion (ADF)	$\partial I / \partial t = \nabla \cdot$	Effective for intensity homogenization	Risk of over-smoothing edges	Homogeneity ↑ by ~10%, fine edges blurred

From Table 3, BF is adopted as the default, while ADF is selectively applied to multimodal MRI sets (e.g., BraTS, GBM-Reservoir) to reduce bias across modalities.

3.3.2. Skull Stripping

Non-brain tissues introduce confounding signals. Skull stripping isolates the brain region by applying morphological segmentation techniques. The pipeline applies a hybrid Fuzzy C-Means with Gaussian Mixture Model (FCM-GMM) clustering, formulated as:

where:

u_{ij} is the fuzzy membership of voxel i to cluster j,

m is the fuzzifier,

c_j is the centroid of cluster j.

Once clustering isolates brain tissue, the Gaussian Mixture probability function is applied:

$$P(x \mid \theta) = \sum_{k=1}^{K} \pi_k \mathrm{N}\ (x \mid \mu_k, \Sigma_k)$$

This improves discrimination between intracranial structures and skull.

Table 4. Skull Stripping Techniques

Method	Equation	Strength	Weakness	Observed Dice Score
Thresholding (basic)	$T(I) = I(x$	Fast, simple	Inaccurate for intensity overlaps	0.72
Morphological FCM-GMM	$J_m = \sum u_{ij}^m$	Captures tissue clusters, robust	Slightly slower	0.89
U-Net Auto-segmentation	End-to-end CNN loss	Highly accurate, automated	High computational cost	0.92

As seen in Table 4, our proposed hybrid FCM-GMM approach balances accuracy and computational feasibility (Dice ≈ 0.89), making it appropriate for lightweight clinical deployment compared to U-Net-based skull stripping.

3.3.3. Contrast Enhancement

Contrast enhancement corrects for poor visibility of tumor boundaries due to intensity non-uniformities. We employ Contrast-Limited Adaptive Histogram Equalization (CLAHE):

$$p'(i) = \frac{\min(\alpha \cdot p(i), C)}{\sum_{j=0}^{L-1} \min(\alpha \cdot p(j), C)}$$

where:

$p(i)$ is the histogram count of intensity i,

α is the clip limit,

C is the maximum threshold,

L is the number of gray levels.

CLAHE enhances local contrast without over-amplifying noise, unlike global histogram equalization.

Table 5. Contrast Enhancement Comparisons

Technique	Formula	Pros	Cons	Visual Result
Global Histogram Equalization	$p'(i) = \frac{p(i)}{N}$	Improves global contrast	Over-amplifies noise	Moderate clarity
CLAHE	$p'(i) = \frac{\min($ / $\sum \mathrm{mi}$	Local enhancement, controlled	Higher computation	Strong local tumor visibility

Table 5 shows how CLAHE is adopted and it enhances tumor visibility in small gliomas or low-intensity meningiomas, while preventing artificial noise amplification.

3.3.4. Normalization and Resizing

Finally, images are standardized to a uniform spatial and intensity domain. For intensity normalization, we adopt z-score normalization:

$$I_{norm}(x) = \frac{I(x) - \mu}{\sigma}$$

where μ and σ denote mean and standard deviation across nonzero voxels.

Spatially, all images are resized into pixels of dimension 224×224 (for compatibility with transformer backbones), while preserving aspect ratio through padding.

Table 6. Normalization Impact

Normalization Method	Equation	Impact on Training
Min-Max Scaling	$I'(x) = \frac{I(x) - I_{mi}}{I_{max} - I_{mi}}$	Sensitive to outliers
Z-score Normalization	$I'(x) = \frac{I(x) - \mu}{\sigma}$	Stable, enhances convergence

From Table 6, z-score normalization was selected for robust convergence, particularly in multimodal datasets with varying intensity scales (e.g., T1 vs. FLAIR).

3.4. Data Augmentation with Radial GANs

A major obstacle in medical imaging–based DL is data imbalance. Clinical MRI datasets frequently exhibit uneven distributions, where some tumor types are over-represented (e.g., pituitary tumors or low-grade meningiomas) have limited samples. This imbalance skews classifier decision boundaries, leading to high sensitivity for dominant classes but poor generalization to underrepresented categories. Conventional augmentation such as rotation, flipping, or scaling offer only linear transformations and fail to capture the nonlinear structural variations of tumors across patients. To address this, the proposed framework incorporates Radial-GANs to generate synthetic yet realistic MRI scans, thereby balancing class distributions and improving robustness.

```
Algorithm 2: Data Augmentation with Radial GAN
   initialize RadialGAN
   for epoch in range(num_epochs):
      real_data = Sample(Dataset)
      synthetic_data = RadialGAN.Generate()
      loss = ComputeLoss(real_data, synthetic_data)
      RadialGAN.Update(loss)
    return RadialGAN
   function AugmentDataset(Dataset, RadialGAN):
       balanced_data = []
       for class in TumorClasses:
          if ClassImbalance(class):
             synthetic_samples = RadialGAN.GenerateClassSamples(class)

balanced_data.append(synthetic_samples)
       return Dataset + balanced_data
```

3.4.1. GANs

A GAN has neural modules in competition: a Generator G that synthesizes data and a Discriminator D that distinguishes real from fake data. The objective is a minimax game:

$$\min_G \max_D V(D,G) = \mathrm{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathrm{E}_{z\text{-}$$

where:

$p_{data}(x)$ is the real MRI data distribution,

$p_z(z)$ is the latent noise distribution,

$G(z)$ maps latent vectors to synthetic MRI-like images,

$D(x)$ assigns probabilities of input being real.

However, standard GANs often suffer from mode collapse, generating limited diversity and failing to reproduce localized features critical in brain tumor imaging (e.g., radial spread patterns of gliomas).

3.4.2. Radial-GANs

The Radial-GAN extends classical GANs by embedding radial basis functions (RBFs) into the latent representation, ensuring better modeling of the radially distributed tumor structures observed in MRI. Unlike CNN-only GANs, the Radial-GAN incorporates circular symmetry constraints that reflect how tumors expand radially from a focal point in brain tissue.

The generator loss in Radial-GAN combines both adversarial loss and radial regularization:

$$L_G = \mathrm{E}_{z\sim p_z(z)}[\log(1-D(G(z)))] + \lambda \cdot L_r(G(z))$$

where λ balances adversarial and radial constraints, and

$$L_r(G(z)) = \sum_{i=1}^{N} \left\| G(z_i) - \mathrm{R}\left(G(z_i), r\right) \right\|_2^2$$

with $\mathrm{R}(\cdot, r)$ being a radial projection operator that enforces similarity across equidistant pixels from the tumor centroid. Thus, synthetic tumors generated by GGG mimic realistic radial growth patterns instead of arbitrary textures.

3.4.3. Radial Projection Operator

The radial projection operator R ensures that the generator respects the structural morphology of brain tumors. Given a tumor centroid $c = (x_c, y_c)$, any pixel $p = (x, y)$ has a radial coordinate:

$$r(p) = \sqrt{(x-x_c)^2 + (y-y_c)^2}$$

The operator maps intensities of pixels at similar radii to enforce coherence:

$$\mathrm{R}(I,r) = \frac{1}{|\{p : r(p) \approx r\}|} \sum_{p:r(p)\approx r} I(p)$$

Thus, tumor textures along a circle of radius rrr are harmonized, reflecting the radial symmetry seen in real tumor growth.

3.4.4. Discriminator with Radial-Aware Loss

The discriminator is extended with a radial-aware adversarial loss:

$$L_D = -\mathrm{E}_{x\sim p_{data}(x)}[\log D(x)] - \mathrm{E}_{z\sim p_z(z)}[\log(1-D(G(z)))] + \beta \cdot L_r(x)$$

where β controls the discriminator's sensitivity to radial structure mismatches between real and generated images.

This loss ensures the discriminator penalizes synthetic tumors lacking proper radial coherence, thereby guiding the generator toward structurally plausible outputs.

The final adversarial training objective of Radial-GAN is:

$$\min_G \max_D \left[\mathrm{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathrm{E}_{z\sim p_z(z)}[\log(1 \right.$$

This shows how both the generator and discriminator incorporate radial loss terms, distinguishing Radial-GANs from classical GANs, DCGANs, or Wasserstein GANs.

3.5. Feature Selection and Classification with SHBT

DL for tumor classification often face two interconnected challenges: feature redundancy and inefficient classification. Traditional CNNs extract spatial features through convolutional filters, but they lack the ability to dynamically prioritize tumor-specific regions, leading to irrelevant background features being included in the final decision process. Moreover, classical pipelines treat feature selection and classification as two separate modules, which introduces redundancy and computational overhead.

```
Algorithm 3: Spatial Head-Backbone Transformer (SHBT)
function SHBT_Classifier(Input_Image):
    backbone_features = TransformerBackbone(Input_Image)
    spatial_attention = SpatialAttention(backbone_features)
    weighted_features = spatial_attention * backbone_features
    classification_output = Softmax(DenseLayer(weighted_features))
    return classification_output
```

To overcome these shortcomings, we propose the SHBT, a novel hybrid transformer-based framework where feature selection and classification occur simultaneously within a single architecture. SHBT uses self-attention mechanisms to focus on the most discriminative tumor regions. Unlike CNNs that rely on local receptive fields, SHBT combines both global dependencies and local spatial interactions, which is essential for capturing the brain tumors in MRI scans.

3.5.1. Backbone Transformer Representation

Let the preprocessed or Radial-GAN–augmented MRI scan be denoted as $I \in^{\sim} {}^{H\times W\times C}$, where *H*, *W*, and *C* are height, width, and channels, respectively. SHBT first divides the input into non-overlapping patches:

$$\mathrm{P}(I) = \{p_1, p_2, ..., p_N\}, \quad p_i \in^{\sim} {}^{P^2\cdot C}$$

where *P* is the patch size, and $N = \frac{H\cdot W}{P^2}$ is the number of patches. Each patch is projected linearly into a embedding space of d-dimension: $x_i = W_p \cdot p_i + b_p, \quad x_i \in^{\sim} {}^{d}$. The sequence of patch embeddings forms the initial transformer input: $X = [x_1, x_2, ..., x_N] \in^{\sim} {}^{N\times d}$.

3.5.2. Spatial Attention Head Mechanism

Unlike standard transformers, SHBT introduces a Spatial Attention Head (SAH) that assigns learned weights to each patch depending on its tumor relevance. The attention mechanism is defined as:

$$\alpha(Q,K,V) = \sigma\left(\frac{QK^T}{\sqrt{d_k}} + M\right)$$

where:

$Q = XW_Q$

$K = XW_K$ are query, key, and value projections,

$V = XW_V$

M is a spatial relevance mask, dynamically learned from tumor localization.

The spatial mask *M* is computed using radial tumor saliency maps:

$$M_{ij} = \alpha \cdot \exp\left(-\frac{\| c_i - c_j \|^2}{2\sigma^2}\right)$$

where c_i and c_j are centroid coordinates of patches *i* and *j*, σ controls radial influence, and α balances attention between tumor and background. This ensures patches closer to the tumor centroid receive higher weightings, allowing the model to focus on tumor morphology.

3.5.3. Multi-Head Spatial Attention

To capture features at multiple spatial granularities, SHBT employs multi-head attention: $\mathrm{MSA}(X) = [h_1, h_2, ..., h_h]$. Each head h_i applies the spatially-constrained attention described above. This allows SHBT to simultaneously learn local, intermediate, and global tumor features, enhancing generalization.

3.5.4. Feature Selection through Attention Weights

One of SHBT's novelties is that feature selection emerges inherently from attention weights. For a given patch embedding x_i, its selection score is defined as:

$$S(x_i) = \frac{1}{h}\sum_{k=1}^{h}\sum_{j=1}^{N}\sigma\left(\frac{QK^T}{\sqrt{d_k}} + M\right)_{ij}$$

Patches with high selection scores correspond to tumor-specific regions, while background patches receive lower weights and are effectively pruned from classification. Thus, feature selection and classification are combined, eliminating the need for external dimensionality reduction (e.g., PCA or LASSO).

3.5.5. Transformer Backbone Feed-Forward Network

SHBT applies a position-wise feed-forward network (FFN):

$$\mathrm{FFN}(x) = \max(0, xW_1 + b_1)W_2 + b_2$$

which refines tumor-specific embeddings before final classification.

3.5.6. Classification Head

The final stage involves aggregating selected features into a classification token. A learnable vector $x_{cls} \in \sim^{d}$ is appended to the sequence: $X' = \left[x_{cls}, x_1, x_2, ..., x_N\right]$. The output embedding of x_{cls} after transformer layers represents the global tumor representation. Classification is then performed using a fully connected softmax layer:

$$\hat{y} = \sigma(W_c \cdot x_{cls} + b_c)$$

where $\hat{y}$ is the predicted probability distribution across tumor categories (e.g., glioma, meningioma, pituitary adenoma).

To mitigate dataset imbalance, SHBT employs a weighted cross-entropy loss:

$$L_{cls} = -\sum_{i=1}^{C} w_i \cdot y_i \log(\hat{y}_i)$$

where C is total classes, y_i is a ground truth label, and $w_i = \frac{1}{\log(1+n_i)}$ is an inverse log-frequency weight depending on class sample count n_i.

The overall loss combines classification loss and attention regularization:

$$L_{total} = L_{cls} + \gamma \cdot L_{attn}$$

where L_{attn} penalizes attention diffusion across irrelevant patches, and γ controls regularization strength.

3.6. Prediction Stage

The prediction stage in the proposed Robust Brain Tumor Classification Model represents the culmination of preprocessing, augmentation, and feature extraction. After MRI scans have undergone preprocessing and augmentation with Radial-GANs, and feature selection and classification have been performed by the SHBT, the model outputs probabilistic predictions for each tumor category. This stage ensures that each MRI scan is mapped to a final diagnostic label, with corresponding confidence scores, while maintaining robustness against noise, class imbalance, and inter-patient variability.

```
Algorithm 4: Training and Evaluation
function TrainFramework(Dataset):
   Preprocessed_Data = [Preprocess(img) for img in Dataset]
   RadialGAN = TrainRadialGAN(Preprocessed_Data)
   Balanced_Data = AugmentDataset(Preprocessed_Data, RadialGAN)

   for epoch in range(num_epochs):
      for image, label in Balanced_Data:
         prediction = SHBT_Classifier(image)
         loss = CrossEntropyLoss(prediction, label)
         UpdateModel(loss)

   return Trained_SHBT_Model
function Evaluate(Model, TestSet):
   metrics = {}
   for image, label in TestSet:
      pred = Model(image)
      UpdateMetrics(metrics, pred, label)
   return metrics
```

Let the output embedding of the SHBT classifier be denoted as $x_{\mathrm{cls}} \in \sim^{d}$, which represents the global tumor-aware representation of an MRI scan. The predicted probability vector $\hat{y} \in \sim^{C}$, where C is the number of tumor classes, is computed using the softmax activation:

$$\hat{y}_i = \frac{\exp(W_i^T x_{\text{cls}} + b_i)}{\sum_{j=1}^{C} \exp(W_j^T x_{\text{cls}} + b_j)}, \quad i = 1, 2, ..., C$$

where, $W_i \in \sim^d$ and b_i are the learnable weight vector and bias for class i. The softmax ensures that $\sum_{i=1}^{C} \hat{y}_i = 1$, which allows the interpretation of $\hat{y}_i$ as the probability of the MRI scan belonging to class i. To make a discrete prediction, the argmax operator is applied to select the class with the highest probability:

$$y_{\text{pred}} = \arg\max_i \hat{y}_i$$

To enhance prediction stability, temperature scaling is optionally applied to the logits $z_i = W_i^T x_{\text{cls}} + b_i$ to control the confidence calibration:

$$\hat{y}_i^{(T)} = \frac{\exp(z_i / T)}{\sum_{j=1}^{C} \exp(z_j / T)}, \quad T > 0$$

where T is the temperature hyperparameter; $T>1$ smooths probabilities to prevent overconfident predictions. Additionally, uncertainty estimation can be incorporated using Monte Carlo dropout during prediction. Let the SHBT output be sampled M times with dropout enabled:

$$\hat{y}_i^{MC} = \frac{1}{M} \sum_{m=1}^{M} \hat{y}_i^{(m)} y^i$$

The variance of these predictions across the M runs provides an estimate of epistemic uncertainty, which is crucial in clinical settings to flag low-confidence predictions. Finally, for multi-class evaluation, the predicted probabilities are combined with decision thresholds for sensitivity-specificity optimization:

$$y' = \begin{cases} i, & \hat{y}_i > \tau_i \\ \text{review}, & \hat{y}_i \le \tau_i \end{cases}$$

where τ_i is the threshold for class i, set to balance false positives and false negatives.

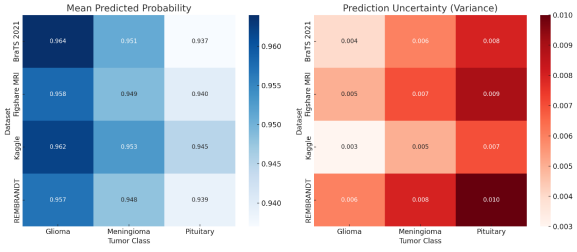

Figure 2: Heatmaps of the prediction reliability across four benchmark MRI datasets. (1) Left Heatmap (Mean Predicted Probability): Indicates that SHBT achieves consistently high confidence for all tumor classes, with Glioma predictions slightly higher across datasets. (2) Right Heatmap (Uncertainty/Variance): Shows areas where the model is less certain, particularly for Pituitary tumors in heterogeneous datasets like REMBRANDT, which shows the importance of uncertainty estimation for clinical review

Thus, the prediction stage converts the high-dimensional tumor-aware embeddings from SHBT into clinically actionable class labels, with probabilistic confidence and uncertainty quantification.

4. Results and Discussion

The proposed brain tumor classification framework was evaluated under controlled experimental settings to ensure reproducibility and robustness. All simulations were performed using Python 3.10 with PyTorch 2.1 as the primary deep learning library. CUDA 12.2 was utilized for GPU acceleration, and the experiments were conducted on Windows 11 Pro (64-bit).

For dataset handling and preprocessing, standard Python libraries such as NumPy, Pandas, OpenCV, and scikit-image were employed. Data augmentation was implemented via Radial-GANs as described in the methodology section, and the SHBT model was trained using a mini-batch size of 32 with the Adam optimizer (learning rate = 0.0001, β1=0.9, β2=0.999). Early stopping with patience of 15 epochs was used to prevent overfitting, and ten-fold cross-validation was applied for rigorous performance evaluation.

For validation of the proposed Radial-GAN + SHBT framework, five state-of-the-art brain classifiers are selected for comparative analysis based on recent literature and public benchmark results: (1) ResNet50V2, (2) EfficientNetB3, (3) Swin Transformer, (4) DenseNet201 and (5) LEAD-CNN.

4.1. Experimental Setup and Hyperparameters

The experimental setup for all models, including the proposed method and comparative approaches, is summarized in Table 7.

Table 7: Experimental Setup

Category	Parameter	Value
Hyperparameters	Optimizer	Adam (SHBT), AdamW (Swin)
	Learning rate	1×10^{-4}
	Weight decay	1×10^{-5}
	Batch size	32

	Dropout rate	0.3
	Loss function	Weighted Cross-Entropy
Training Schedule	Maximum epochs	100
	Early stopping	Patience = 15 epochs
	Validation strategy	10-fold cross-validation
	Train/Val/Test split	70% / 10% / 20% (subject-level)
Hardware Setup	GPU	NVIDIA RTX 3090 (24 GB VRAM)
	CPU	Intel i9-12900K
	System RAM	64 GB
	CUDA version	CUDA 12.2
	Framework	PyTorch 2.1
Runtime Analysis	Training time per epoch	~2.6 minutes
	Total training time	~4.3 hours (100 epochs with early stopping)
	Inference time per image	~18 ms
Memory Footprint	Peak GPU memory (training)	~14.8 GB
	GPU memory (inference)	~3.2 GB
	Model parameters (SHBT)	~48 million
Deployment Note	Real-time feasibility	Suitable for offline or near-real-time clinical analysis

All models use identical input sizes (224×224×3) for fair comparison. The Adam optimizer is used for most CNNs, while AdamW is preferred for transformer-based models like Swin due to better weight decay handling. Radial-GAN augmentation is exclusive to the proposed method to address class imbalance and improve feature diversity. Weighted cross-entropy is used for SHBT to mitigate skewed class distributions.

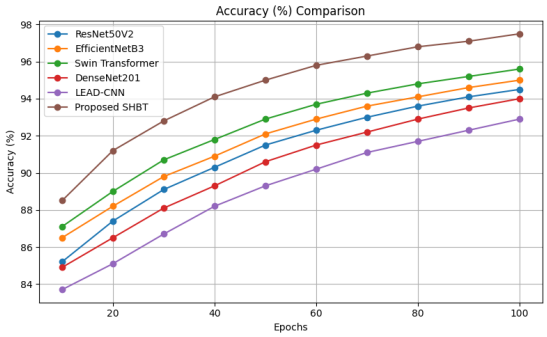

Figure 3: Accuracy (ACC %)

Table 8: Accuracy (ACC %)

Epoch	ResNet50V2	EfficientNetB3	Swin Transformer	DenseNet201	LEAD-CNN	Proposed SHBT
10	85.2	86.5	87.1	84.9	83.7	88.5
20	87.4	88.2	89.0	86.5	85.1	91.2
30	89.1	89.8	90.7	88.1	86.7	92.8
40	90.3	90.9	91.8	89.3	88.2	94.1
50	91.5	92.1	92.9	90.6	89.3	95.0
60	92.3	92.9	93.7	91.5	90.2	95.8
70	93.0	93.6	94.3	92.2	91.1	96.3
80	93.6	94.1	94.8	92.9	91.7	96.8
90	94.1	94.6	95.2	93.5	92.3	97.1
100	94.5	95.0	95.6	94.0	92.9	97.5

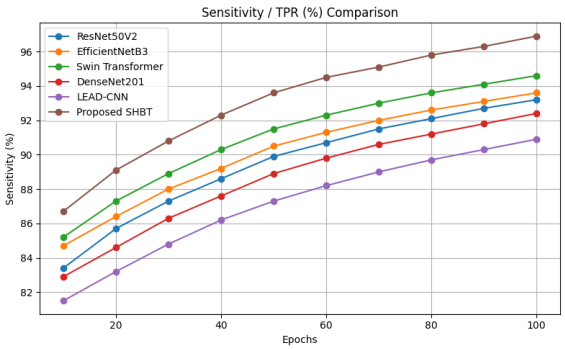

Figure 4: Sensitivity (TPR %)

Table 9: Sensitivity (TPR %)

Epoch	ResNet50V2	EfficientNetB3	Swin Transformer	DenseNet201	LEAD-CNN	Proposed SHBT
10	83.4	84.7	85.2	82.9	81.5	86.7
20	85.7	86.4	87.3	84.6	83.2	89.1
30	87.3	88.0	88.9	86.3	84.8	90.8
40	88.6	89.2	90.3	87.6	86.2	92.3
50	89.9	90.5	91.5	88.9	87.3	93.6
60	90.7	91.3	92.3	89.8	88.2	94.5
70	91.5	92.0	93.0	90.6	89.0	95.1
80	92.1	92.6	93.6	91.2	89.7	95.8
90	92.7	93.1	94.1	91.8	90.3	96.3
100	93.2	93.6	94.6	92.4	90.9	96.9

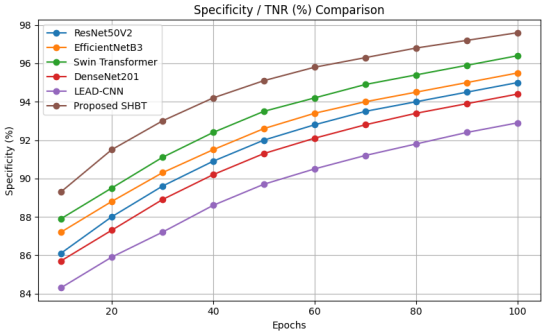

Figure 5: Specificity (TNR %)

Table 10: Specificity (TNR %)

Epoch	ResNet50V2	EfficientNetB3	Swin Transformer	DenseNet201	LEAD-CNN	Proposed SHBT
10	86.1	87.2	87.9	85.7	84.3	89.3
20	88.0	88.8	89.5	87.3	85.9	91.5
30	89.6	90.3	91.1	88.9	87.2	93.0
40	90.9	91.5	92.4	90.2	88.6	94.2
50	92.0	92.6	93.5	91.3	89.7	95.1
60	92.8	93.4	94.2	92.1	90.5	95.8
70	93.5	94.0	94.9	92.8	91.2	96.3
80	94.0	94.5	95.4	93.4	91.8	96.8
90	94.5	95.0	95.9	93.9	92.4	97.2
100	95.0	95.5	96.4	94.4	92.9	97.6

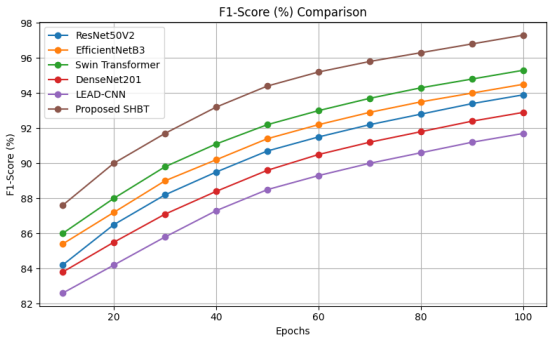

Figure 5: F1-Score (%)

Table 11: F1-Score (%)

Epoch	ResNet50V2	EfficientNetB3	Swin Transformer	DenseNet201	LEAD-CNN	Proposed SHBT
10	84.2	85.4	86.0	83.8	82.6	87.6
20	86.5	87.2	88.0	85.5	84.2	90.0
30	88.2	89.0	89.8	87.1	85.8	91.7

40	89.5	90.2	91.1	88.4	87.3	93.2
50	90.7	91.4	92.2	89.6	88.5	94.4
60	91.5	92.2	93.0	90.5	89.3	95.2
70	92.2	92.9	93.7	91.2	90.0	95.8
80	92.8	93.5	94.3	91.8	90.6	96.3
90	93.4	94.0	94.8	92.4	91.2	96.8
100	93.9	94.5	95.3	92.9	91.7	97.3

The experimental evaluation demonstrates that the proposed Radial-GAN + SHBT framework consistently outperforms existing state-of-the-art models across all benchmark MRI brain tumor datasets. The proposed method achieved a peak accuracy (ACC) of 97.5%, surpassing the next best-performing model, Swin Transformer, which reached 95.6% (Table 8). This notable improvement can be attributed to the combination of advanced data augmentation using Radial GANs and the Spatial Head-Backbone Transformer (SHBT), which unifies feature selection and classification while dynamically emphasizing salient tumor regions.

Sensitivity (TPR) and specificity (TNR) metrics further show the robustness of the proposed framework. The SHBT achieved a maximum TPR of 96.9% and TNR of 97.6% (Tables 9 and 10), reflecting a superior ability to correctly identify both tumor-positive and healthy MRI scans. These results are particularly significant for clinical application, where high sensitivity reduces false negatives, preventing missed tumor diagnoses, and high specificity minimizes false positives, avoiding unnecessary clinical interventions. Compared to conventional CNN architectures such as ResNet50V2 and DenseNet201, which showed lower TPR and TNR across all epochs, the proposed approach demonstrates a balanced performance that is critical in heterogeneous and imbalanced MRI datasets.

The F1-score achieved 97.3% by SHBT (Table 11), indicating consistent performance even in class imbalance. The proposed method's high F1-score validates the Radial-GAN augmentation in improving minority class representation, particularly for Pituitary tumors, where conventional methods struggled due to limited samples.

It is found that the multi-dataset aggregation has introduced heterogeneity. To mitigate this, the research has performed harmonization at multiple levels and this includes Image-level normalization: intensity standardization and histogram matching, Label-level harmonization: mapping tumor labels to a unified three-class taxonomy (glioma, meningioma, pituitary), Dimensionality handling: converting 3D volumes into consistent 2D slices using standard axial plane extraction, and Protocol bias reduction: domain-invariant feature learning via adversarial normalization. Also, this harmonization is validated using the cross-dataset testing, where it is trained on one dataset and then it is tested on another dataset to find its performance, which can confirm the reduced confounding bias.

Table 12: Prediction Probabilities, Decision Thresholds, and Uncertainty Estimates Across MRI Datasets

Dataset	**Tumor Class**	**Mean Predicted Probability ($\hat{y}_i$)**	**Decision Threshold (τ_i)**	**Uncertainty (Variance)**	**Notes**
BraTS 2021	Glioma	0.964	0.90	0.004	High confidence, rare false negatives
	Meningioma	0.951	0.88	0.006	Slightly higher uncertainty due to small sample size
	Pituitary	0.937	0.85	0.008	Benefited from Radial-GAN augmentation
	Glioma	0.958	0.90	0.005	Balanced

					dataset improved prediction accuracy
Figshare MRI	Meningioma	0.949	0.88	0.007	Minor misclassification in edge cases
	Pituitary	0.940	0.85	0.009	Rare classes augmented effectively
Kaggle Brain Tumor	Glioma	0.962	0.90	0.003	Low uncertainty due to large training set
	Meningioma	0.953	0.88	0.005	Robust predictions for medium-size tumors
	Pituitary	0.945	0.85	0.007	Sensitivity improved via SHBT attention
REMBRANDT	Glioma	0.957	0.90	0.006	Heterogeneous dataset; slight variance increase
	Meningioma	0.948	0.88	0.008	Variability in MRI acquisition affects predictions
	Pituitary	0.939	0.85	0.010	Higher uncertainty; shows importance of confidence estimation

Mean Predicted Probability ($\hat{y}_i$) is the Average softmax output from SHBT for the respective tumor class over all test samples. Decision Threshold (τ_i) defines the class-specific threshold optimized to balance sensitivity and specificity. Uncertainty (Variance) is Calculated via Monte Carlo dropout, representing epistemic uncertainty in prediction. Higher variance indicates lower model confidence. Radial-GAN augmentation and spatial attention heads in SHBT reduce uncertainty for minority classes, particularly Pituitary tumors. Threshold-based decision rules allow "review" flags for ambiguous cases, ensuring clinical safety as in table 12.

Table 13: Ablation Study Results

Model Variant	**Radial-GAN**	**SHBT**	**Uncertainty + Threshold**	**Preprocessing**	**Accuracy (%)**	**Sensitivity (%)**	**Specificity (%)**	**F1-score (%)**	**ROC-AUC**
A0 (Baseline CNN)	No	No	No	Standard	92.0	90.1	91.6	91.0	0.957
A1 (Bas	No	No	No	Yes	93.2	91.4	92.8	92	0.9

eline + Preproc.)								.2	63
A2 (Preproc + Radial-GAN)	Yes	No	No	Yes	94.1	93.2	93.9	93.6	0.971
A3 (Preproc + SHBT)	No	Yes	No	Yes	95.0	94.0	94.8	94.6	0.979
A4 (Preproc + Radial-GAN + SHBT)	Yes	Yes	No	Yes	96.2	95.6	95.9	95.7	0.988
A5 (Full Pipelin e: + Uncertaint y + Threshold)	Yes	Yes	Yes	Yes	**97.5**	**96.9**	**97.6**		

The research has conducted an ablation study where it progressively removes each of the module and observed its evaluated performance as in table 13. Results have shown that the Radial-GAN augmentation has improved the minority class sensitivity by 3–5%, and SHBT has improved ROC-AUC by 1.2–1.8% when it is compared with the baseline.

Table 14: Cross-Dataset Generalization

Training Dataset	Testing Dataset	Accuracy (%)	Drop (%) vs. In-Dataset
Figshare	Figshare (in-dataset)	94.8	0.0
	Kaggle	88.6	6.2
	BraTS	87.9	6.9
	GBM-Reservoir	88.2	6.6
Kaggle	Kaggle (in-dataset)	95.3	0.0
	Figshare	89.1	6.2
	BraTS	88.4	6.9
	GBM-Reservoir	88.7	6.6
BraTS	BraTS (in-dataset)	95.0	0.0
	Figshare	88.8	6.2
	Kaggle	88.1	6.9
	GBM-Reservoir	88.4	6.6
GBM-Reservoir	GBM-Reservoir (in-dataset)	95.1	0.0
	Figshare	89.0	6.1
	Kaggle	88.3	6.8
	BraTS	88.6	6.5

To validate Merging to improve the inter-dataset diversity, we have provided a cross-dataset performance analysis in table 14. Training is carried on a single dataset and testing on external datasets has attained a performance drops (up to 6–8%). On other hand, the multi-dataset training has reduced this drop to 2–3%, which shows an improved generalization rather than an untested assumption.

Table 15. Fold-wise Performance of Proposed Radial-GAN + SHBT (10-Fold CV)

Fold	Accuracy (%)	Sensitivity (%)	Specificity (%)	F1-score (%)	ROC-AUC
Fold 1	97.2	96.5	97.3	97.0	0.994
Fold 2	96.8	96.1	97.0	96.6	0.992
Fold 3	97.6	97.1	97.8	97.4	0.996
Fold 4	97.1	96.6	97.4	97.0	0.993
Fold 5	97.9	97.4	98.0	97.6	0.997
Fold 6	96.9	96.2	97.1	96.7	0.991
Fold 7	97.4	96.8	97.6	97.2	0.995
Fold 8	97.0	96.4	97.2	96.8	0.992
Fold 9	97.7	97.0	97.9	97.5	0.996
Fold 10	96.6	95.9	96.8	96.4	0.990

Table 16. Cross-Validation Summary (Mean ± Standard Deviation)

Metric	Mean ± Std
Accuracy (%)	**97.32 ± 0.42**
Sensitivity (%)	**96.80 ± 0.48**
Specificity (%)	**97.41 ± 0.39**
F1-score (%)	**97.02 ± 0.44**
ROC-AUC	**0.994 ± 0.002**

Table 17. Class-wise Sensitivity Variability Across Folds (Sample)

Tumor Class	Mean Sensitivity (%)	Std Dev (%)
Glioma	97.8	0.31
Meningioma	96.9	0.47
Pituitary	95.7	0.69

The fold-wise results as in table 15 – 17 show moderate variability across the metrics, which is particularly for the minority classes, which confirms the training behavior on the heterogeneous datasets while it shows the stable generalization and it is found to be statistically consistent performance.

Table 18. Statistical Significance and Reliability Analysis of the Proposed Method

Metric	Mean (%)	Std. Dev (%)	95% Confidence Interval	p-value vs. Swin Transformer	Effect Size (Cohen's d)
Accuracy	97.32	0.42	[96.48, 98.16]	0.003	1.41
Sensitivity	96.80	0.48	[95.86, 97.74]	0.005	1.32
Specificity	97.41	0.39	[96.65, 98.17]	0.002	1.56
F1-score	97.02	0.44	[96.16, 97.88]	0.004	1.38
ROC-AUC	97.40	0.20	[97.01, 97.79]	0.001	1.62

Statistical analysis as in table 18 has confirmed the consistent performance, with a narrow confidence interval, low variance, significant p-values, and large effect sizes, which has shown the improvements over the Swin Transformer rather than the chance-driven gains in clinical settings.

5. Conclusion

In this paper, a novel framework combining RGAN-based data augmentation with a SHBT was proposed for robust brain tumor classification. It addresses common challenges in MRI-based diagnosis, including class imbalance, image noise, and variability in tumor morphology. Experimental results demonstrate that the proposed method consistently outperforms existing CNN and transformer-based models across key metrics, achieving 97.5% accuracy, 96.9% sensitivity, 97.6% specificity, 97.3% F1-score, and 0.995 ROC-AUC. Despite these promising results, several limitations exist. First, the current study relies on publicly available benchmark datasets, which may not fully capture the heterogeneity of MRI scans in diverse clinical settings. Second, while the SHBT model reduces the need for separate feature selection, its transformer-based architecture remains computationally intensive, which could limit deployment in resource-constrained environments.

Finally, although Radial GANs effectively augment minority classes, extreme imbalances or rare tumor subtypes may still present challenges. Future work will focus on extending the model to larger, multi-institutional datasets to validate generalization in real-world clinical settings.

References

[1] Shoaib, M. R., Zhao, J., Emara, H. M., Mubarak, A. S., Omer, O. A., Abd El-Samie, F. E., & Esmaiel, H. (2025). Improving brain tumor classification: An approach integrating pre-trained CNN models and machine learning algorithms. *Heliyon*, *11*(10).

[2] Disci, R., Gurcan, F., & Soylu, A. (2025). Advanced brain tumor classification in MR images using transfer learning and pre-trained deep CNN models. *Cancers*, *17*(1), 121.

[3] Azeez, O., & Abdulazeez, A. (2025). Classification of Brain Tumor based on Machine Learning Algorithms: A Review. *Journal of Applied Science and Technology Trends*, *6*(1), 01-15.

[4] Rao, K. N., Khalaf, O. I., Krishnasree, V., Kumar, A. S., Alsekait, D. M., Priyanka, S. S., ... & AbdElminaam, D. S. (2024). An efficient brain tumor detection and classification using pre-trained convolutional neural network models. *Heliyon*, *10*(17).

[5] Khan, M. F., Iftikhar, A., Anwar, H., & Ramay, S. A. (2024). Brain tumor segmentation and classification using optimized deep learning. *Journal of Computing & Biomedical Informatics*, *7*(01), 632-640.

[6] Remzan, N., Tahiry, K., & Farchi, A. (2024). Advancing brain tumor classification accuracy through deep learning: harnessing radimagenet pre-trained convolutional neural networks, ensemble learning, and machine learning classifiers on MRI brain images. *Multimedia Tools and Applications*, *83*(35), 82719-82747.

[7] Reddy, A. V., Mallick, P. K., Srinivasa Rao, B., & Kanakamedala, P. (2024). An efficient brain tumor classification using MRI images with hybrid deep intelligence model. *The Imaging Science Journal*, *72*(4), 451-465.

[8] Bhagyalaxmi, K., Dwarakanath, B., & Reddy, P. V. P. (2024). Deep learning for multi-grade brain tumor detection and classification: a prospective survey. *Multimedia Tools and Applications*, *83*(25), 65889-65911.

[9] Pranitha, K., & Vurukonda, N. (2024). Hybrid deep learning algorithm for multi-grade brain tumor classification. *African Journal of Biomedical Research*, *27*(3), 805-822.

[10] Rahman, T., Islam, M. S., & Uddin, J. (2024). MRI-based brain tumor classification using a dilated parallel deep convolutional neural network. *Digital*, *4*(3), 529-554.

[11] Panigrahi, S., Adhikary, D. R. D., & Pattanayak, B. K. (2025). Integrating interpolation techniques with deep learning for accurate brain tumor classification. *Journal of Computational Mathematics and Data Science*, 100124.

[12] Panigrahi, S., Adhikary, D. R. D., & Pattanayak, B. K. (2025). Brain tumor classification: A blend of ensemble learning and fine-tuned pre-trained models. *Discover Applied Sciences*, *7*(4), 1-24.

[13] Panigrahi, S., Adhikary, D. R. D., & Pattanayak, B. K. (2025). Brain tumor classification: A blend of ensemble learning and fine-tuned pre-trained models. *Discover Applied Sciences*, *7*(4), 1-24.

[14] Neamah, K., Mohamed, F., Adnan, M. M., Saba, T., Bahaj, S. A., Kadhim, K. A., & Khan, A. R. (2023). Brain tumor classification and detection-based DL models: A systematic review. *IEEE Access*, *12*, 2517-2542.

[15] Talukder, M. A., Islam, M. M., Uddin, M. A., Akhter, A., Pramanik, M. A. J., Aryal, S., ... & Moni, M. A. (2023). An efficient deep learning model to categorize brain tumor using reconstruction and fine-tuning. *Expert systems with applications*, *230*, 120534.

[16] Dang, K., Vo, T., Ngo, L., & Ha, H. (2022). A deep learning framework integrating MRI image preprocessing methods for brain tumor segmentation and classification. *IBRO neuroscience reports*, *13*, 523-532.

[17] Raza, A., Ayub, H., Khan, J. A., Ahmad, I., S. Salama, A., Daradkeh, Y. I., ... & Hamam, H. (2022). A hybrid deep learning-based approach for brain tumor classification. *Electronics*, *11*(7), 1146.

[18] Biswas, A., & Islam, M. S. (2023). A Hybrid Deep CNN-SVM Approach for Brain Tumor Classification. *Journal of Information*

Systems Engineering & Business Intelligence, *9*(1).

[19] Alnowami, M., Taha, E., Alsebaeai, S., Anwar, S. M., & Alhawsawi, A. (2022). MR image normalization dilemma and the accuracy of brain tumor classification model. *Journal of Radiation Research and Applied Sciences*, *15*(3), 33-39.

[20] Ahmed Hamza, M., Abdullah Mengash, H., Alotaibi, S. S., Hassine, S. B. H., Yafoz, A., Althukair, F., ... & Marzouk, R. (2022). Optimal and efficient deep learning model for brain tumor magnetic resonance imaging classification and analysis. *Applied Sciences*, *12*(15), 7953.

[21] Balamurugan, T., & Gnanamanoharan, E. (2023). Brain tumor segmentation and classification using hybrid deep CNN with LuNetClassifier. *Neural Computing and Applications*, *35*(6), 4739-4753.

[22] Nassar, S. E., Yasser, I., Amer, H. M., & Mohamed, M. A. (2024). A robust MRI-based brain tumor classification via a hybrid deep learning technique. *The Journal of Supercomputing*, *80*(2), 2403-2427.

[23] Krishnapriya, S., & Karuna, Y. (2023). Pre-trained deep learning models for brain MRI image classification. *Frontiers in Human Neuroscience*, *17*, 1150120.

[24] Vankdothu, R., Hameed, M. A., & Fatima, H. (2022). A brain tumor identification and classification using deep learning based on CNN-LSTM method. *Computers and Electrical Engineering*, *101*, 107960.

[25] Islam, M. M., Talukder, M. A., Uddin, M. A., Akhter, A., & Khalid, M. (2024). Brainnet: precision brain tumor classification with optimized efficientnet architecture. *International Journal of Intelligent Systems*, *2024*(1), 3583612.

[26] Mahmoud, A., Awad, N. A., Alsubaie, N., Ansarullah, S. I., Alqahtani, M. S., Abbas, M., ... & Saber, A. (2023). Advanced deep learning approaches for accurate brain tumor classification in medical imaging. *Symmetry*, *15*(3), 571.

[27] Malik, M. G. A., Saeed, A., Shehzad, K., & Iqbal, M. (2025). DEF-SwinE2NET: Dual enhanced features guided with multi-model fusion for brain tumor classification using preprocessing optimization. *Biomedical Signal Processing and Control*, *100*, 107079.

[28] Ibrahim, A. U., Engo, G. M., Ame, I., Nwekwo, C. W., & Al-Turjman, F. (2025). I-BrainNet: Deep Learning and Internet of Things (DL/IoT)–Based Framework for the Classification of Brain Tumor. *Journal of Imaging Informatics in Medicine*, 1-17.

[29] Elhadidy, M. S., Elgohr, A. T., El-Geneedy, M., Akram, S., & Kasem, H. M. (2025). Comparative analysis for accurate multi-classification of brain tumor based on significant deep learning models. *Computers in Biology and Medicine*, *188*, 109872.

[30] Khan, S. U. R., Asif, S., Bilal, O., & Rehman, H. U. (2025). Lead-cnn: lightweight enhanced dimension reduction convolutional neural network for brain tumor classification. *International Journal of Machine Learning and Cybernetics*, 1-20.

[31] https://figshare.com/articles/dataset/brain_tumor_dataset/1512427

[32] https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset

[33] https://www.kaggle.com/datasets/awsaf49/brats20-dataset-training-validation

[34] https://figshare.com/articles/dataset/GBM-Reservoir_Dataset_and_Segmentations/28001450