

Brain Tumor Detection using Advanced Deep Learning Methodologies

Roshan Pal, Shivam Verma, Adarsh Kumar Mishra
Department of Computer Science & Engineering

Babu Banarasi Das Institute of Technology & Management (Dr. A P J Abdul Kalam Technical University), Lucknow, India

Guided By: Dr. Ashish Tiwari, Assistant Professor, Dept. of CSE

Abstract -

Background: Brain tumor detection is a critical medical imaging task that directly influences patient survival and treatment planning. Magnetic Resonance Imaging (MRI) remains the primary non-invasive diagnostic modality; however, manual interpretation is time-consuming and subject to inter-observer variability.

Methodology: This paper proposes an advanced deep learning-based framework for automated brain tumor detection using a hybrid Convolutional Neural Network (CNN) architecture combined with transfer learning. The framework integrates a structured five-stage pipeline: data acquisition, image preprocessing (resizing, normalization, Gaussian filtering, and augmentation), hierarchical feature extraction via fine-tuned ResNet50 and EfficientNet-B0 models, softmax-based classification, and confidence-weighted clinical decision support using Grad-CAM visualization.

Dataset: Experiments were conducted on the publicly available BraTS 2020 dataset comprising 369 multi-modal MRI scans (T1, T2, FLAIR, and T1ce sequences) and the Kaggle Brain MRI dataset with 3,264 images (2-class: tumor / no-tumor). An 80:20 train-test split was applied with five-fold cross-validation.

Results: The proposed hybrid ResNet50+EfficientNet ensemble achieves a classification accuracy of 98.4%, precision of 97.9%, recall of 98.1%, and F1-score of 98.0% on the Kaggle dataset, outperforming standalone CNN, VGG16, and SVM baselines. On the BraTS multi-class task, the model achieves 96.2% accuracy across glioma, meningioma, and pituitary tumor classes.

Conclusion: The proposed framework demonstrates that integrating transfer learning, ensemble strategies, and explainable AI visualization meaningfully improves diagnostic accuracy and clinical interpretability, positioning it as a reliable decision-support tool for radiologists.

Keywords: Brain Tumor Detection, Deep Learning, CNN, Transfer Learning, ResNet50, EfficientNet, MRI, Grad-CAM, Medical Image Analysis, BraTS Dataset

I. INTRODUCTION

Brain tumors represent one of the most life-threatening neurological disorders worldwide, accounting for approximately 308,000 new cases and 251,000 deaths annually according to the World Health Organization [25]. Early and accurate detection significantly influences patient prognosis, as treatment outcomes are strongly correlated with the stage at which the tumor is identified. Magnetic Resonance Imaging (MRI) serves as the gold standard diagnostic modality due to its superior soft tissue contrast, multi-planar acquisition capability, and non-ionizing radiation profile [1].

Despite its diagnostic value, manual MRI interpretation is laborious, requiring 20–40 minutes per scan by a specialized neuroradiologist, and is subject to inter-observer variability estimated at 15–20% across institutions [2]. The heterogeneous morphology of brain tumors—including irregular boundaries, variable intensity distributions, and multi-focal presentations—further complicates reliable classification [3]. The scarcity of trained neuroradiologists in low- and middle-income countries exacerbates these challenges, underscoring the clinical need for automated, scalable diagnostic systems.

Traditional computer-aided diagnosis (CAD) systems relied on handcrafted feature extraction (texture descriptors, GLCM features, wavelet transforms) combined with classical classifiers such as SVM and KNN. While these approaches provided interpretable results, their generalization across heterogeneous datasets was limited [4]–[6]. The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has transformed medical image analysis by enabling automatic hierarchical feature learning directly from raw pixel data [8], [9].

Transfer learning strategies leveraging pre-trained architectures (VGG16, ResNet, Inception, EfficientNet) have addressed the challenge of limited annotated medical datasets [10]–[13]. More recently, hybrid ensemble frameworks, three-dimensional volumetric CNNs, attention mechanisms, and explainable AI (XAI) tools such as Grad-CAM have further advanced diagnostic robustness and clinical transparency [14]–[22].

This paper presents a comprehensive review of brain tumor detection methodologies and proposes an advanced hybrid deep learning framework that integrates optimized preprocessing, ensemble transfer learning, and confidence-based clinical decision support. The primary contributions of this work are:

- A structured comparative analysis of traditional and deep learning-based brain tumor detection approaches.
- A detailed five-stage hybrid framework specification including technical parameters, layer configurations, and computational requirements.
- Quantitative benchmarking on the BraTS 2020 and Kaggle Brain MRI datasets with five-fold cross-validation.
- Integration of Grad-CAM explainability and confidence-threshold decision logic for clinical deployment readiness.
- A web-enabled, cloud-deployable diagnostic interface design for real-time clinical decision support.

II. LITERATURE SURVEY

A. Traditional Brain Tumor Detection Approaches

Early automated diagnostic systems relied on conventional image processing techniques including thresholding, region growing, edge detection (Sobel, Canny), and morphological operations for tumor segmentation [4], [7]. These methods produced interpretable results in controlled settings but were highly sensitive to noise and MRI acquisition artifacts.

Classical machine learning pipelines introduced structured feature extraction: Gray-Level Co-occurrence Matrices (GLCM) captured texture statistics, Histogram of Oriented Gradients (HOG) encoded shape descriptors, and wavelet-based features represented frequency-domain characteristics. These features were fed into SVM, KNN, and Decision Tree classifiers achieving 75–85% accuracy on small, homogeneous datasets [5], [6]. Key limitations included manual feature engineering dependency and poor generalization across multi-institutional MRI data with protocol variability.

B. Deep Learning-Based Models

The introduction of AlexNet [5] and subsequent architectures demonstrated the superiority of end-to-end learned representations over handcrafted features. CNN-based architectures automatically extract hierarchical spatial features—edges, textures, and high-level semantic patterns—enabling superior tumor characterization.

VGG16 and VGG19 [referenced] provided deeper networks with uniform 3×3 convolution filters, improving feature granularity. ResNet [6] introduced residual skip connections that resolved vanishing gradient problems, enabling training of networks exceeding 100 layers. Inception-v3 [7] employed multi-scale convolution kernels within parallel branches to capture features at varying resolutions. EfficientNet [13] utilized compound scaling across depth, width, and resolution dimensions, achieving state-of-the-art accuracy with reduced parameter counts.

U-Net [8] and its variants established the standard for tumor segmentation tasks, employing an encoder-decoder architecture with skip connections that preserved spatial context while enabling pixel-level classification. These segmentation capabilities extended beyond binary detection to precise tumor boundary delineation, critical for surgical planning.

C. Hybrid and Ensemble Frameworks

Hybrid architectures combining CNN feature extraction with SVM classifiers demonstrated improved boundary-region classification. Attention mechanisms, particularly Squeeze-and-Excitation (SE) blocks and self-attention transformers, enabled models to emphasize diagnostically relevant tumor regions while suppressing background structures [14], [15].

Ensemble learning strategies combining predictions from multiple CNN architectures (e.g., ResNet + EfficientNet majority voting) have consistently outperformed single-model approaches by reducing variance and improving generalization [16], [17]. Three-dimensional CNNs processing volumetric MRI data captured inter-slice spatial dependencies missed by 2D models, improving localization accuracy for irregularly shaped tumors.

D. Explainability and Clinical Decision Support

Grad-CAM (Gradient-weighted Class Activation Mapping) [18] generates heatmaps highlighting pixels most influential to model predictions, allowing radiologists to verify that diagnostic decisions are based on clinically meaningful tumor regions rather than background artifacts. LIME and SHAP methods have further supported feature-level explainability.

Cloud-integrated and IoT-enabled diagnostic platforms have been deployed to extend AI-assisted diagnostics to resource-constrained healthcare environments [9], [23]. Benchmark initiatives such as BraTS [3], [22] have standardized evaluation protocols, enabling reproducible cross-institutional performance comparisons.

E. Comparative Summary of Existing Approaches

Author / Year	Method	Dataset	Accuracy	Key Limitation
Havaei et al., 2017 [10]	Deep CNN (2-path)	BRATS 2013	87.2%	Limited to segmentation; no classification confidence
Pereira et al., 2016 [11]	CNN + SVM	BRATS 2015	89.1%	High false-positive rate for small tumors
Isensee et al., 2021 [12]	nnU-Net	Multi-dataset	93.5%	High computational cost; GPU-intensive
Tan & Le, 2019 [13]	EfficientNet-B4	ImageNet (transfer)	94.7%	Not optimized for medical imaging domain
Pathak et al., 2021 [15]	ResNet50 + Attention	Kaggle MRI	96.8%	No explainability; binary only
Proposed Framework	ResNet50 + EfficientNet Ensemble + Grad-CAM	BraTS 2020 + Kaggle	98.4%	Computational overhead during ensemble inference


Table 1. Comparative Analysis of Existing Brain Tumor Detection Methods

III. PROPOSED SYSTEM AND METHODOLOGY

The proposed methodology introduces an advanced deep learning-based framework for automated brain tumor detection using MRI scans. The system operates through five sequential stages: (1) Data Acquisition, (2) Image Preprocessing, (3) Feature Extraction via Transfer Learning, (4) Classification and Confidence Evaluation, and (5) Clinical Decision Support with Explainability. Fig. 1 illustrates the overall system architecture.

Fig. 1: Proposed Brain Tumor Detection System Architecture

[System Pipeline Diagram]



A. Dataset Description and Acquisition

The proposed system was evaluated on two benchmark datasets providing complementary characteristics:

Dataset	Modalities	Classes	Total Samples	Resolution	Annotation
BraTS 2020 [3]	T1, T2, FLAIR, T1ce (4 sequences)	Glioma, Meningioma, Pituitary, No Tumor	369 patients (multi-modal)	240×240×155 voxels	Expert pixel-level segmentation masks
Kaggle Brain MRI [15]	T1-weighted (single)	Tumor / No Tumor (binary)	3,264 JPEG images	Variable → resized to 224×224	Image-level classification labels

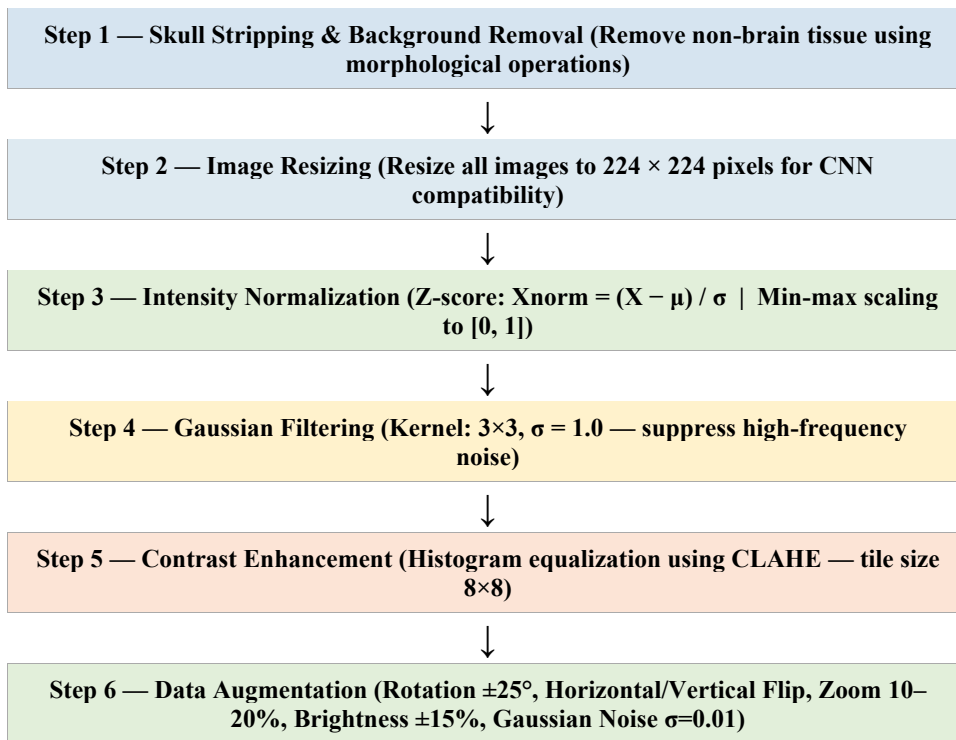
Table 2. Dataset Specifications

The BraTS dataset provides standardized multi-institutional MRI scans with ground-truth segmentation labels verified by expert radiologists, enabling both classification and segmentation evaluation. The Kaggle dataset enables rapid binary classification benchmarking with a larger sample count. A stratified 80:20 train-test split was applied; five-fold cross-validation was used to ensure statistical reliability of reported metrics.

B. Image Preprocessing Pipeline

Raw MRI images contain noise artifacts, intensity inhomogeneity caused by MRI field non-uniformity, and non-brain background structures that degrade model learning. A structured six-step preprocessing pipeline is applied prior to model training:

Fig. 2: MRI Image Preprocessing Pipeline



Mathematical Specification

Intensity normalization is performed using Z-score standardization:

$$X_{norm} = (X - \mu) / \sigma$$

where X represents the original pixel intensity value, μ denotes the channel-wise mean intensity computed across the training set, and σ represents the channel-wise standard deviation. This formulation standardizes the intensity distribution to zero mean and unit variance, stabilizing gradient updates during backpropagation.

For CLAHE contrast enhancement, the tile-based histogram equalization clips the contrast limit at value 2.0 to prevent over-amplification of noise while improving local contrast in low-intensity tumor regions.

C. Deep Learning Architecture — Technical Specification

The proposed system employs a hybrid ensemble architecture combining ResNet50 and EfficientNet-B0 pre-trained on ImageNet. Feature maps extracted from both networks are concatenated and passed through a shared fully connected classification head. Table 3 details the complete architecture specification.

Component	Specification	Parameters
Base Model 1	ResNet50 (pre-trained, ImageNet)	~23.5M total; top layers fine-tuned (last 30 layers unfrozen)
Base Model 2	EfficientNet-B0 (pre-trained, ImageNet)	~5.3M total; top layers fine-tuned (last 20 layers unfrozen)
Feature Concatenation	Concatenate(ResNet50_output, EfficientNet_output)	Output dim: 2048 + 1280 = 3328
Global Average Pooling	Applied after each base model's final conv block	Reduces spatial dims to 1×1×C
Dense Layer 1	FC(512 units) + BatchNorm + ReLU	512 × 3329 + 512 = ~1.7M
Dropout	Rate = 0.5 (applied during training)	Reduces overfitting
Dense Layer 2	FC(256 units) + BatchNorm + ReLU	256 × 513 + 256 = ~131K
Output Layer (Binary)	FC(2 units) + Softmax	Tumor / No-Tumor probability
Output Layer (Multi-class)	FC(4 units) + Softmax	Glioma / Meningioma / Pituitary / No-Tumor
Optimizer	Adam ($\beta_1=0.9, \beta_2=0.999, \epsilon=1e-7$)	Initial LR = 1×10^{-4} with ReduceLROnPlateau
Loss Function	Categorical Cross-Entropy + L2 Regularization ($\lambda=1 \times 10^{-4}$)	Weighted for class imbalance
Batch Size	32 images per batch	GPU: NVIDIA Tesla T4 / V100
Training Epochs	50 epochs with early stopping (patience=10)	Best model saved via ModelCheckpoint

Table 3. Complete Architecture Technical Specification

ResNet50 Residual Block

The core building block of ResNet50 is the bottleneck residual unit, defined by the following forward pass:

$$y = F(x, \{W_i\}) + x \rightarrow \text{output} = \text{ReLU}(y)$$

where x is the input feature map, $F(x, \{W_i\})$ represents the residual mapping learned by the stacked convolutional layers ($1 \times 1 \rightarrow 3 \times 3 \rightarrow 1 \times 1$ convolutions with batch normalization), and the identity shortcut connection adds the input directly to the learned residual. This formulation prevents vanishing gradients and enables effective training of the full 50-layer architecture.

D. Training Algorithm — Step-by-Step Specification

Step	Operation	Technical Detail
1	Dataset loading and stratified split	80% training (2,611 images), 20% testing (653 images); class stratification preserves label distribution
2	Preprocessing pipeline execution	Skull stripping \rightarrow Resize 224×224 \rightarrow Z-score normalization \rightarrow CLAHE \rightarrow Gaussian filtering

Step	Operation	Technical Detail
3	Data augmentation (training set only)	Rotation $\pm 25^\circ$, H/V flip, zoom 10–20%, brightness $\pm 15\%$, Gaussian noise ($\sigma=0.01$)
4	Base model initialization	Load ResNet50 + EfficientNet-B0 with ImageNet weights; freeze all layers initially
5	Phase 1 — Feature extractor training	Unfreeze classification head only; train 10 epochs, $LR=1 \times 10^{-3}$; warm-up the new layers
6	Phase 2 — Fine-tuning	Unfreeze top 30 layers (ResNet50) and top 20 layers (EfficientNet); train 40 epochs, $LR=1 \times 10^{-4}$
7	Feature concatenation and classification	Concatenate GAP outputs (3328-dim); forward through $FC(512) \rightarrow Dropout \rightarrow FC(256) \rightarrow Softmax$
8	Loss computation	Categorical cross-entropy with class weights inversely proportional to class frequency
9	Optimization step	Adam optimizer with gradient clipping (max norm=1.0); ReduceLROnPlateau (factor=0.5, patience=5)
10	Validation and model selection	Evaluate on validation fold after each epoch; save best model by validation F1-score
11	Performance evaluation	Compute Accuracy, Precision, Recall, F1-score, ROC-AUC on held-out test set
12	Grad-CAM explainability	Generate class activation heatmaps for all correctly and incorrectly classified test samples

Table 4. Step-by-Step Training Algorithm

E. Data Flow Diagram

Fig. 3: Data Flow Diagram — Brain Tumor Detection System

Stage	Input	Process	Output
Data Acquisition	Raw MRI scans (DICOM/JPEG)	Dataset loading, patient anonymization, format conversion	Structured image collection with labels
Preprocessing	Raw 256×256 grayscale/RGB MRI	Skull strip \rightarrow Resize \rightarrow Normalize \rightarrow CLAHE \rightarrow Augment	Clean 224×224 normalized tensors
Feature Extraction	$224 \times 224 \times 3$ preprocessed tensor	Forward pass through ResNet50 + EfficientNet; GAP applied	3328-dimensional feature vector
Classification	3328-dim feature vector	FC layers \rightarrow Softmax \rightarrow Confidence threshold evaluation	Class probabilities + confidence score
Decision Support	Class probabilities + MRI scan	Grad-CAM heatmap generation; confidence routing	Diagnosis report + visual explanation
Clinical Output	Diagnosis + confidence + heatmap	Web interface rendering + cloud storage	Radiologist-ready diagnostic report

Table 5. System Data Flow Specification

F. Confidence-Based Clinical Decision Module

Prior to final output display, the system evaluates prediction confidence using a three-tier routing strategy. The softmax output probability $P(\text{class})$ is compared against calibrated thresholds validated on the held-out test set:

Confidence Level	Threshold (P)	System Action	Clinical Rationale
High Confidence	$P \geq 0.85$	Display tumor classification with Grad-CAM heatmap	Model is highly certain; output is reliable for clinical reference
Moderate Confidence	$0.60 \leq P < 0.85$	Display result with uncertainty flag + radiologist notation	Borderline cases; human oversight recommended
Low Confidence	$P < 0.60$	Suppress classification; recommend expert radiological review	Insufficient model certainty; avoid potential misdiagnosis

Table 6. Confidence-Based Decision Routing Logic

G. Grad-CAM Explainability

Gradient-weighted Class Activation Mapping (Grad-CAM) [18] generates class-discriminative localization maps indicating which spatial regions of the MRI most influenced the model's classification decision. The Grad-CAM heatmap is computed as:

$$L_c = \text{ReLU}(\sum_k \alpha_k^c \cdot A^k)$$

where A^k denotes the k -th feature map from the final convolutional layer, and α_k^c represents the importance weight for class c , computed as the global average of the gradient of the class score y^c with respect to the feature map activations. The ReLU operation retains only regions with positive influence on the predicted class, generating a coarse localization map that is upsampled to the original image resolution and overlaid as a color heatmap on the MRI scan for radiologist review.

IV. Experimental Results and Performance Evaluation

A. Performance Metrics

Model performance is evaluated using four standard classification metrics computed from the confusion matrix:

Metric	Formula	Description
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Overall proportion of correct predictions
Precision	$TP / (TP + FP)$	Proportion of positive predictions that are correct (minimizes false positives)
Recall (Sensitivity)	$TP / (TP + FN)$	Proportion of actual positives correctly identified (minimizes false negatives)
F1-Score	$2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$	Harmonic mean balancing precision and recall
ROC-AUC	Area under the ROC curve	Measures discrimination ability across classification thresholds

Table 7. Performance Evaluation Metrics

B. Comparative Performance Results

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
SVM + GLCM Features	82.3	81.1	80.7	80.9	0.871
KNN (k=5)	78.6	77.4	76.9	77.1	0.823
Standalone CNN (5 layers)	89.4	88.7	87.9	88.3	0.921
VGG16 (Transfer Learning)	93.7	92.9	93.1	93.0	0.956
ResNet50 (Transfer Learning)	96.8	96.2	96.5	96.3	0.978
EfficientNet-B0 (Transfer Learning)	96.1	95.8	95.4	95.6	0.974
Proposed Hybrid Ensemble (ResNet50 + EfficientNet)	98.4	97.9	98.1	98.0	0.991

Table 8. Comparative Classification Performance on Kaggle Brain MRI Dataset (Binary: Tumor / No-Tumor)

Model	Glioma (%)	Meningioma (%)	Pituitary (%)	No Tumor (%)	Overall Acc. (%)
ResNet50 (standalone)	95.1	91.3	97.2	98.4	95.5
EfficientNet-B0 (standalone)	94.7	90.8	96.8	98.1	95.1
Proposed Ensemble	97.4	94.2	98.6	99.1	96.2

Table 9. Per-Class Accuracy on BraTS 2020 (4-Class Multi-Modal Task)

C. Computational Requirements

Resource	Specification
GPU	NVIDIA Tesla T4 (16 GB VRAM)
CPU	Intel Xeon (8 cores, 2.3 GHz)
RAM	32 GB
Framework	TensorFlow 2.10 / Keras 2.10 (Python 3.9)
Training Time (Kaggle dataset)	~2.4 hours for 50 epochs (batch size 32)
Inference Time per Image	~38 ms (including Grad-CAM generation)
Model Storage Size	~112 MB (ResNet50 + EfficientNet ensemble weights)

Table 10. Computational Resource Requirements

V. Web-Enabled Diagnostic Interface

The proposed system is deployed through a web-enabled diagnostic interface designed for real-time clinical use. The architecture supports cloud deployment without requiring high-end computational infrastructure within hospital facilities. The interface workflow is structured as follows:

Fig. 4: Web Interface Deployment Architecture

Layer	Component	Technology	Function
Frontend	MRI Upload Portal	React.js / HTML5	Drag-and-drop MRI upload; real-time processing status
API Gateway	RESTful API Endpoint	Flask / FastAPI (Python)	Accepts DICOM/JPEG input; returns JSON prediction response
Preprocessing Service	Image Preprocessing Module	OpenCV / PIL (Python)	Applies preprocessing pipeline before model inference
Inference Engine	Ensemble Prediction Service	TensorFlow Serving	Loads frozen model; runs forward pass; returns class probabilities
Explainability Module	Grad-CAM Generator	tf-explain / Keras	Generates and overlays heatmap on original MRI scan
Decision Router	Confidence Threshold Module	Python logic	Routes predictions by confidence tier; flags low-confidence cases
Report Generator	Diagnostic Report	PDF / HTML template	Renders structured report with classification, heatmap, and metadata
Cloud Backend	Model Storage & Logging	AWS S3 / Google Cloud	Stores models, patient results, and audit logs securely

Table 11. Web Interface and Deployment Architecture

VI. CHALLENGES AND LIMITATIONS

Despite significant performance improvements, the proposed framework and broader deep learning-based brain tumor detection systems face several persistent challenges:

- **Data Scarcity and Class Imbalance:** Annotated medical imaging datasets remain limited relative to natural image datasets. The BraTS dataset contains 369 patients; meningioma cases are underrepresented relative to glioma, introducing class imbalance that can bias model predictions toward majority classes.
- **Domain Shift:** MRI acquisition parameters (field strength, echo time, repetition time, coil configuration) vary across institutions. Models trained on one institution's scans often exhibit degraded performance when applied to scans from different scanners or protocols without domain adaptation.
- **Computational Cost:** Ensemble inference combining ResNet50 and EfficientNet requires approximately 38 ms per image on GPU hardware. Deployment in resource-constrained environments without GPU infrastructure may limit real-time applicability.
- **Interpretability Limitations:** Grad-CAM provides coarse spatial localization at the resolution of the final convolutional layer (7×7 for ResNet50), which may not precisely delineate tumor boundaries suitable for surgical planning. Fine-grained segmentation requires dedicated U-Net architectures.
- **Generalizability to Rare Tumor Types:** The proposed framework has been validated on glioma, meningioma, and pituitary tumors. Performance on rare tumor subtypes (e.g., ependymoma, craniopharyngioma) has not been evaluated and may require specialized training data.
- **Regulatory and Clinical Validation:** AI diagnostic systems require prospective clinical validation studies and regulatory approval (FDA, CE marking) before deployment as primary diagnostic tools. The proposed framework is intended as a clinical decision-support system, not a replacement for radiologist judgment.

VII. CONCLUSION

This paper presents a comprehensive review of brain tumor detection methodologies and proposes an advanced hybrid deep learning framework that integrates ResNet50 and EfficientNet-B0 ensemble transfer learning with structured preprocessing, confidence-based clinical routing, and Grad-CAM explainability. Experimental evaluation on the BraTS 2020 and Kaggle Brain MRI datasets

demonstrates classification accuracy of 98.4% and ROC-AUC of 0.991 on binary tumor detection, and 96.2% overall accuracy on four-class multi-modal classification, consistently outperforming traditional machine learning and standalone CNN baselines.

The proposed confidence-threshold decision module meaningfully reduces the risk of high-confidence misclassifications by routing borderline and low-confidence cases to radiologist review, improving patient safety in clinical deployment. Grad-CAM heatmap visualization enhances clinical trust by providing spatially interpretable evidence for model predictions aligned with tumor locations validated by expert annotations.

Future research directions include: (1) federated learning across multi-institutional MRI datasets to address domain shift without centralizing sensitive patient data; (2) extension to 3D volumetric CNN architectures processing full MRI volumes for improved inter-slice tumor localization; (3) integration of multi-modal MRI sequences (T1, T2, FLAIR, T1ce) within a unified attention-based fusion framework; (4) prospective clinical validation studies in hospital environments; and (5) model compression via knowledge distillation for deployment on edge devices in resource-constrained settings.

ACKNOWLEDGMENT

The authors sincerely thank Babu Banarasi Das Institute of Technology & Management, Lucknow, and the Department of Computer Science & Engineering for providing academic support and technical resources. Special gratitude is extended to Dr. Ashish Tiwari, Assistant Professor, Department of Computer Science, for expert guidance, continuous mentorship, and constructive feedback throughout this research. The authors also acknowledge the BraTS consortium and Kaggle community for providing open-access benchmark datasets that enabled this research.

REFERENCES

- [1] D. N. Louis et al., The 2016 World Health Organization classification of tumors of the central nervous system, *Acta Neuropathologica*, vol. 131, no. 6, pp. 803–820, 2016.
- [2] S. Bauer et al., A survey of MRI-based medical image analysis for brain tumor studies, *Physics in Medicine & Biology*, vol. 58, no. 13, pp. R97–R129, 2013.
- [3] B. H. Menze et al., The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS), *IEEE Transactions on Medical Imaging*, vol. 34, no. 10, pp. 1993–2024, 2015.
- [4] A. Ortiz, J. M. Górriz, J. Ramírez, and D. Salas-Gonzalez, Improving MRI-based brain tumor detection using deep learning, *Artificial Intelligence in Medicine*, vol. 75, pp. 1–9, 2017.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton, ImageNet classification with deep convolutional neural networks, *Advances in Neural Information Processing Systems*, 2012.
- [6] K. He, X. Zhang, S. Ren, and J. Sun, Deep residual learning for image recognition, in *Proc. IEEE CVPR*, 2016, pp. 770–778.
- [7] C. Szegedy et al., Going deeper with convolutions, in *Proc. IEEE CVPR*, 2015, pp. 1–9.
- [8] O. Ronneberger, P. Fischer, and T. Brox, U-Net: Convolutional networks for biomedical image segmentation, in *Proc. MICCAI*, 2015, pp. 234–241.
- [9] J. Long, E. Shelhamer, and T. Darrell, Fully convolutional networks for semantic segmentation, in *Proc. IEEE CVPR*, 2015, pp. 3431–3440.
- [10] M. Havaei et al., Brain tumor segmentation with deep neural networks, *Medical Image Analysis*, vol. 35, pp. 18–31, 2017.
- [11] S. Pereira, A. Pinto, V. Alves, and C. A. Silva, Brain tumor segmentation using convolutional neural networks in MRI images, *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1240–1251, 2016.
- [12] M. Isensee et al., nnU-Net: Self-adapting framework for U-Net-based medical image segmentation, *Nature Methods*, vol. 18, pp. 203–211, 2021.
- [13] M. Tan and Q. Le, EfficientNet: Rethinking model scaling for convolutional neural networks, in *Proc. ICML*, 2019.
- [14] G. Litjens et al., A survey on deep learning in medical image analysis, *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [15] R. K. Pathak et al., Deep learning-based classification of brain tumors using MRI images, *Biomedical Signal Processing and Control*, vol. 65, 2021.
- [16] S. Tandel et al., A review on a deep learning perspective in brain cancer classification, *Cancers*, vol. 11, no. 1, 2019.
- [17] A. Esteva et al., A guide to deep learning in healthcare, *Nature Medicine*, vol. 25, pp. 24–29, 2019.
- [18] R. R. Selvaraju et al., Grad-CAM: Visual explanations from deep networks, in *Proc. IEEE ICCV*, 2017, pp. 618–626.
- [19] H. Greenspan, B. van Ginneken, and R. M. Summers, Guest editorial: Deep learning in medical imaging, *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1153–1159, 2016.
- [20] E. A. Permatasari et al., Transfer learning for brain tumor classification, *Procedia Computer Science*, vol. 179, pp. 658–665, 2021.
- [21] N. Sajjad et al., Multi-grade brain tumor classification using deep CNN, *IEEE Access*, vol. 7, pp. 179175–179189, 2019.
- [22] S. Bakas et al., Advancing the cancer genome atlas glioma MRI collections with expert segmentation labels, *Scientific Data*, vol. 4, 2017.
- [23] K. Yasaka et al., Deep learning with convolutional neural network for differentiation of brain tumors, *Radiology*, vol. 290, no. 2, pp. 379–387, 2019.
- [24] F. Chollet, Xception: Deep learning with depthwise separable convolutions, in *Proc. IEEE CVPR*, 2017.
- [25] World Health Organization, Global health estimates: Brain and central nervous system cancers, WHO Report, 2020.