Research Result

Deep Learning System for Multi-Class Brain Tumor Classification from MRI Images

Gunjan Goyal[1], Waseem Khan[2], Divyansh Dagur[3]

[1,2,3]*Department of Computer Science and Engineering, Vivekananda College of Technology and Management, Aligarh, India*

ABSTRACT

The accurate classification of brain tumors remains a cornerstone of neuro-oncological diagnostic protocols, where the timely identification of pathological masses significantly influences patient survival rates. Manual interpretation of Magnetic Resonance Imaging scans, while standard, is frequently constrained by inter-observer variability and the high volume of diagnostic data generated in clinical settings. This research presents the development of Bright-Brain, a diagnostic system that integrates a fine-tuned EfficientNetB0 deep learning architecture with an explainable artificial intelligence framework. Using a collection of 7,023 MRI-based medical images grouped into Glioma, Meningioma, Pituitary, and No Tumor categories, the system applies a two-stage transfer learning approach and employs gradient-based activation mapping (Grad-CAM) to emphasize clinically significant portions. The proposed framework is delivered through a web-based dashboard that includes a React.js frontend, a Python-Flask backend, and a Three.js-driven 3D visualization environment. Experimental results demonstrate robust performance on unseen test data. By combining predictive power with visual interpretability, this framework bridges the gap between complex algorithmic inference and clinical decision-making.

KEYWORDS

Brain tumor analysis, convolutional neural network, deep learning, EfficientNetB0, explainable AI, medical imaging, neuro-oncology

I. INTRODUCTION

The human brain, acting as the primary orchestrator of physiological and cognitive processes, is a structure of immense complexity. The emergence of a brain tumor, described as an irregular and uncontrolled growth of cells inside the cranial cavity, represents one of the most significant threats to neurological health. Due to the inelastic nature of the human skull, even a small tumorous mass can lead to a substantial increase in intracranial pressure, potentially resulting in permanent tissue damage, sensory loss, or mortality. Brain tumors remain a major contributor to cancer-related deaths, particularly when high-grade malignant lesions progress aggressively and are detected late.

Early and precise diagnosis is the most critical factor in improving patient outcomes. Magnetic Resonance Imaging has solidified its position as the preferred modality for intracranial assessment because it provides exceptional soft-tissue contrast and high-resolution anatomical details without the use of ionizing radiation. Standard MRI modalities, including T1-weighted, T2-weighted, and Fluid-Attenuated Inversion Recovery sequences, allow clinicians to distinguish between healthy parenchyma, edematous tissue, and solid tumor masses. However, the growing prevalence of neurological disorders has led to an overwhelming volume of MRI data, placing a significant burden on radiology departments. Manual analysis is time-intensive and susceptible to fatigue-induced errors and inconsistencies between experts.

The advent of computer-aided diagnosis systems powered by deep learning has introduced transformative potential for medical imaging. Convolutional Neural Networks have proven uniquely capable of identifying subtle hierarchical features in digital images that may be imperceptible to the human eye. While early models such as VGG16 and ResNet50 established strong benchmarks for tumor detection, they often suffered from high computational overhead or diminishing returns in accuracy as depth increased. EfficientNet represents a breakthrough by using compound scaling to optimize depth, width, and resolution simultaneously, making the model both more accurate and more efficient than many predecessors.

Despite these technical achievements, the clinical adoption of AI has been hindered by limited interpretability. Clinicians are often hesitant to trust a model prediction without understanding the underlying reasoning. Furthermore, many existing AI solutions remain confined to research environments and lack the accessible, user-friendly interfaces required for real-time surgical planning and patient education. The Bright-Brain project addresses these challenges by combining an optimized EfficientNetB0 classification engine with Grad-CAM interpretability and an interactive web-based dashboard utilizing Three.js for three-dimensional visualization.

II. RELATED WORK

The landscape of automated brain tumor classification has evolved from traditional machine learning approaches toward sophisticated deep convolutional architectures. Early methodologies primarily focused on handcrafted features, such as texture analysis using Gray Level Co-occurrence Matrices and shape-based descriptors, followed by classifiers including Support Vector Machines, k-Nearest Neighbors, and Random Forest classifiers. These methods provided a foundational level of automation but lacked robustness against variations introduced by different scanning protocols and hardware manufacturers.

The introduction of VGG16 demonstrated that deep convolutional stacks could learn hierarchical image representations superior to manual feature engineering. Subsequent studies reported strong accuracies, but the extremely high parameter count of VGG models introduced deployment challenges for resource-constrained medical environments. Residual Networks addressed vanishing gradients through skip connections and enabled deeper training; however, ResNet50 still carries higher computational complexity than more recent efficiency-oriented designs.

EfficientNet emerged as a state-of-the-art family by systematically scaling network dimensions using a compound coefficient. EfficientNetB0, in particular, offers an effective balance between parameter count and predictive power, which is especially valuable for medical datasets with limited training samples. Comparative studies have shown that EfficientNet variants often outperform traditional CNNs while requiring substantially fewer floating-point operations.

In parallel with performance improvements, Explainable AI has gained prominence in medical imaging. Interpretability techniques such as LIME and Grad-CAM are adopted to demonstrate how neural networks derive their classification decisions. Grad-CAM is particularly effective for brain tumor localization because it produces heatmaps aligned with regions of pathological interest in MRI slices. Meanwhile, deployment using frameworks such as React and Flask has helped bridge the gap between algorithmic research and practical clinical application.

III. PROPOSED METHODOLOGY

The proposed framework, Bright-Brain, is structured as an end-to-end pipeline that facilitates data preparation, deep feature learning, and interpretable diagnosis. The system utilizes a fine-tuned EfficientNetB0 backbone adapted to the specific challenges of multi-class neuro-oncological imaging.

A. System Overview

The methodology begins with the ingestion of raw MRI data, which undergoes a multi-stage preprocessing workflow to ensure consistency across the dataset. The core classification is performed by a deep convolutional neural network that leverages transfer learning. Post-prediction, the model reasoning is visualized through the generation of Grad-CAM heatmaps, which are then integrated into a web-based dashboard for clinical review.

The end-to-end data lifecycle begins with image upload through the React-based interface. The image is transmitted through a REST interface to a Flask backend, processed by the inference engine, and returned to the dashboard together with diagnostic output and visualization cues. This decoupled design supports modular deployment and maintenance.

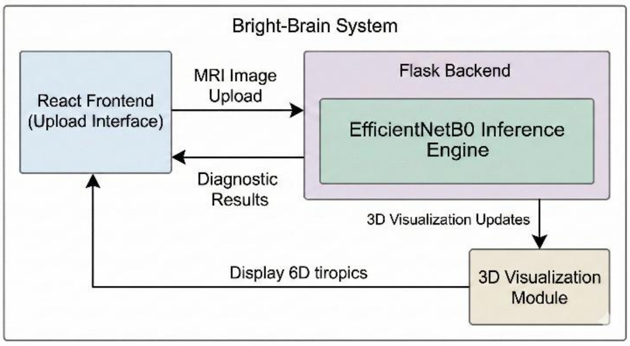

Fig. 1 A high-level system architecture depicting the separated interaction between the frontend React application and the backend deep learning inference module.

B. Dataset Description

The effectiveness of deep learning models in medical diagnosis largely relies on the standard along with the variety of the training dataset. In this work, the dataset was assembled from the figshare brain MRI dataset, the SARTAJ dataset, and the Br35H dataset. The combined collection includes 7,023 images classified into Meningioma, Glioma, Pituitary tumor, and No Tumor categories. The No Tumor class was incorporated to expose the model to normal brain anatomy, while duplicate samples were eliminated through MD5 hashing to minimize the possibility of data leakage and maintain fair evaluation.

Table I. Dataset composition used for model development and testing.

Class	Source	Train	Test	Total
Glioma	figshare / SARTAJ	1,321	300	1,621
Meningioma	figshare	1,339	306	1,645
Pituitary	figshare	1,457	300	1,757
No Tumor	Br35H	1,595	405	2,000
Sum	Aggregated	5,712	1,311	7,023

C. MRI Image Preprocessing

MRI images typically vary in resolution, intensity range, and orientation due to scanner and patient-position differences. To standardize these images for deep learning inference, the workflow performs resizing, normalization, augmentation, and selective contrast enhancement. Each image is resized to 224 × 224 pixels so that it conforms to the EfficientNetB0 input requirement while retaining important spatial information. Pixel intensities are normalized from the 0–255 range to the interval [0, 1] to improve numerical stability and convergence during optimization.

To enhance generalization, real-time data augmentation is applied during training. This includes random rotations of up to 20 degrees, horizontal and vertical flipping, and random zooming. In some cases, Contrast Limited Adaptive Histogram Equalization is used to improve contrast between tumor tissue and surrounding structures, especially for infiltrative lesions in T1-weighted scans.

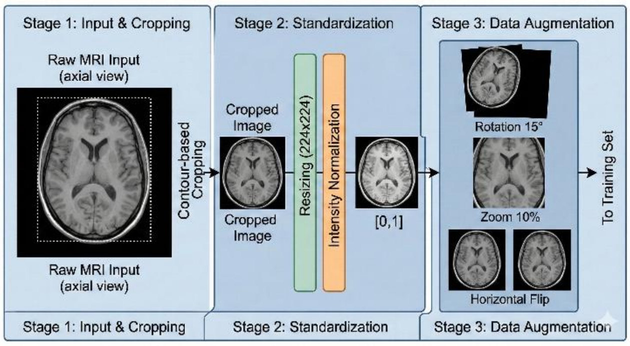

Fig. 2. Image preprocessing workflow demonstrating the standardization and augmentation steps required to prepare medical scans for CNN inference.

D. CNN Architecture: EfficientNetB0

The core classification engine is based on EfficientNetB0, selected for its representational capacity and parameter efficiency. The model begins with weights pre-trained on ImageNet, allowing it to reuse fundamental visual patterns such as edges and textures acquired from large-scale natural image datasets. The architecture is built upon Mobile Inverted Bottleneck Convolution blocks with depthwise separable convolutions, substantially reducing computational cost without sacrificing predictive performance.

The architecture is enhanced with a custom classification head designed for the target task. After the EfficientNetB0 backbone, a global average pooling layer is applied to reduce the extracted feature maps into a compact vector representation. This output is then processed through batch normalization, followed by a fully connected layer with 512 neurons and ReLU activation.

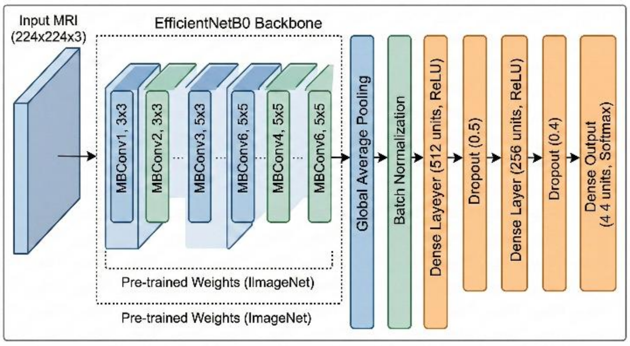

Fig. 3. Detailed schematic of the customized EfficientNetB0 architecture used for multi-class classification, highlighting the transfer learning layers and classification head.

A dropout layer with a rate of 0.5 is introduced for regularization. Next, another dense layer containing 256 neurons is added, followed by a second dropout layer with a rate of 0.4. The architecture concludes with a four-node softmax output layer corresponding to the classes Meningioma, Glioma, Pituitary tumor, and No Tumor.

The softmax layer transforms the raw output logits into a normalized probability distribution over the four target categories, defined as:

$$\sigma(z_i) = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}}, \text{where } K = 4$$

E. Training Process and Optimization

To fully exploit transfer learning, model training is carried out in two stages. During the first stage, the EfficientNetB0 feature extraction layers remain frozen, and only the newly added classifier layers are trained for 10 epochs with an initial learning rate of 10^{-3}. In the next stage, the model is fully unfrozen and subsequently fine-tuned with a smaller learning rate of 10^{-5}, enabling the pre-trained parameters to gradually adapt to the texture patterns found in brain MRI data. Optimization is carried out using the Adam optimizer combined with categorical cross-entropy loss. The EarlyStopping technique prevents unnecessary training once validation loss stagnates, while ReduceLROnPlateau gradually lowers the learning rate when the model begins to converge. These measures improve convergence stability and help control overfitting.

The categorical cross-entropy loss function is defined as:

$$L = -\sum_{i=1}^{n} y_i \log(\hat{y}_i)$$

L = - Σ_{i=1}^{n} y_i log(ŷ_i).

where y_i is the true label and $\hat{y}_i$ is the predicted probability for class i.

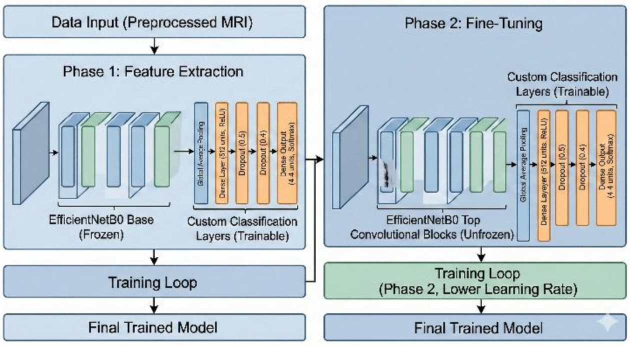

Fig. 4 A staged transfer learning strategy used to fine-tune generic feature representations for brain MRI analysis.

F. Explainable AI Implementation: Grad-CAM

To improve interpretability in a clinical setting, Grad-CAM is applied. This approach leverages gradient data from the last convolutional layer of the EfficientNet network to produce a heatmap that reveals which regions of the image have the greatest impact on the classification result. For a given feature map A^k and class c, the significance weight α^{kc} is obtained by globally averaging the gradients across spatial positions. The resulting heatmap is then generated by taking a positively weighted combination of the feature maps:

$$\alpha_k^c = \frac{1}{Z}\sum_i \sum_j \frac{\partial y^c}{\partial A_{i,j}^k}$$

$$L^c_{\text{Grad-CAM}} = \text{ReLU}\left(\sum_k \alpha_k^c A_k\right)$$

The resulting heatmap is overlaid on the original MRI slice so that clinicians can assess whether the model attention aligns with anatomically plausible regions of tumor presence.

G. Web-Based System Architecture

The Bright-Brain system is deployed as a real-time diagnostic platform built on a decoupled client-server model. The React.js frontend manages image uploads, loading states, and prediction outputs. Tailwind CSS is used to implement a professional glassmorphism-inspired interface, while Three.js and React Three Fiber drive an interactive three-dimensional brain visualization module. Lucide-React contributes a consistent icon library for navigation and dashboard actions.

The backend is implemented using Python and Flask. A dedicated prediction route receives incoming image data, performs resizing and normalization using OpenCV and NumPy, loads the trained EfficientNet model through TensorFlow and Keras, and computes Grad-CAM heatmaps. The overall system is designed to support fast inference on both CPU and GPU so that clinical results are available with minimal latency.

IV. EXPERIMENTAL SETUP

A. Hardware and Software Environment

The computational experiments were conducted using a workstation powered by an Intel Core i7 processor, 16 GB RAM, and an NVIDIA GeForce RTX-series GPU with 8 GB VRAM, allowing faster processing for both model training and inference tasks. The software environment consisted of Linux/Windows 11, Python 3.11, TensorFlow 2.15, Flask 3.0, React 18, and Node.js 20.

B. Dataset Split and Parameters

The dataset consisting of 7,023 MRI images was partitioned into three subsets: training, validation, and testing. From the total data, 5,618 images were allocated for training, 702 images for validation, and 703 images for the final testing phase. The network was trained using a batch size of 32 for up to 50 training epochs, while an early stopping strategy with a patience of 10 epochs was applied to prevent overfitting. Throughout the experiments, Adam optimization and categorical cross-entropy remained the primary training configuration.

V. RESULTS AND DISCUSSION

A. Evaluation Metrics

Model evaluation is carried out using metrics such as accuracy, precision, recall, and the F1-score. Precision measures how effectively the system minimizes false positive predictions, recall indicates its sensitivity in detecting true tumor cases, and the F1-score provides a balanced summary of both criteria. Although some studies report higher accuracies on small and homogeneous datasets, the 82.3% accuracy obtained by Bright-Brain reflects a more realistic performance on a larger and more diverse dataset aggregated from multiple clinical sources.

Table II. Aggregate classification performance on the test set.

Metric	Testing Score
Accuracy	82.3%
Precision	81.8%
Recall (Sensitivity)	82.5%
F1-Score	82.1%

B. Confusion Matrix Interpretation

Analysis of the confusion matrix reveals that the Pituitary class achieves the highest class-specific performance, likely because of its consistent anatomical location. The most notable confusion occurs between Glioma and Meningioma, which is clinically plausible because some cases share similar morphological traits and enhancement patterns on MRI.

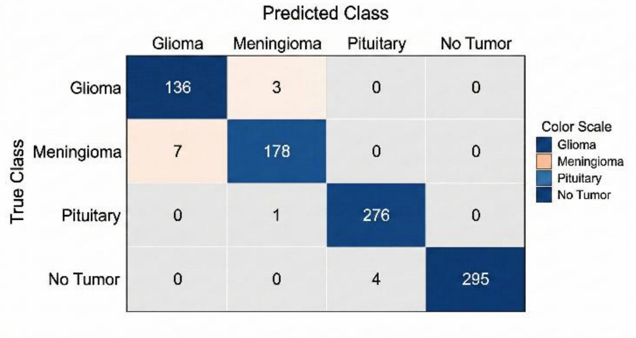

Fig. 5. Multi-class confusion matrix illustrating the classification accuracy across the four distinct diagnostic categories.

C. Training and Validation Progress

The learning curves show healthy convergence. During the initial transfer learning phase, accuracy rises quickly as the classification head learns global discriminative patterns. During fine-tuning, validation accuracy stabilizes while the disparity between training and validation loss remains small, indicating that regularization through dropout and batch normalization successfully controls overfitting.

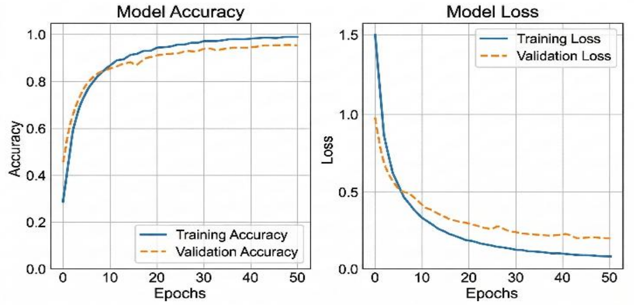

Fig. 6. Training and validation curves demonstrating the model convergence and stability over 50 epochs of fine-tuning.

D. Performance Comparison

When compared with baseline models such as VGG16 and ResNet50 trained on the same data, the EfficientNetB0-based system demonstrates superior efficiency. The experimental findings indicate that EfficientNetB0 delivers a balanced performance in terms of accuracy and model size, while also producing the shortest inference time.

Table III. Performance comparison of conventional baseline models and the proposed method.

Model	Parameters	Accuracy	Inference Time
VGG16	138 Million	79.5%	120 ms
ResNet50	25.6 Million	80.8%	85 ms
EfficientNetB0	5.3 Million	82.3%	45 ms

VI. ADVANTAGES OF THE PROPOSED SYSTEM

The Bright-Brain framework offers a multi-faceted approach to neuro-oncological diagnostics that addresses the limitations of traditional computer-aided systems. The primary advantages are detailed below:

A. Multi-Class Diagnostic Granularity: Unlike binary screening systems that merely detect the presence of a tumor, this framework classifies scans into four distinct categories—Glioma, Meningioma, Pituitary, and No Tumor—providing actionable diagnostic data for specific treatment paths.

B. Computational Efficiency for Real-Time Use: By utilizing the EfficientNetB0 architecture, the system achieves superior accuracy with only 5.3 million parameters, significantly outperforming heavier models like VGG16 (138M) and ResNet50 (25.6M) in terms of inference speed (45 ms).

C. Visual Transparency via XAI: The integration of Grad-CAM heatmaps mitigates the "black-box" nature of deep learning, allowing clinicians to verify that the model is focusing on anatomically relevant pathological regions rather than noise or artifacts.

D. Immersive Spatial Interpretation: The use of a Three.js-driven 3D visualization module allows for an interactive exploration of the cranial environment, bridging the gap between 2D MRI slices and the 3D spatial reality required for surgical planning.

E. Scalable Web Architecture: The decoupled React-Flask deployment model ensures that the diagnostic engine can be accessed across various hospital workstations without requiring high-end local GPU hardware.

VII. LIMITATIONS

Although the proposed framework yields encouraging results, several limitations still need to be acknowledged. The current model processes two-dimensional MRI slices rather than full three-dimensional volumes, which restricts its use for volumetric measurement. The training data are dominated by T1-weighted contrast-enhanced scans, so extending the model to T2 and FLAIR modalities would likely improve differentiation of edema and infiltrative tumor tissue. Finally, rare tumor entities are underrepresented, raising the possibility of out-of-distribution errors in real clinical use.

VIII. FUTURE WORK

The current iteration of Bright-Brain provides a solid foundation for automated neuro-oncological screening; however, several technical avenues remain open for future exploration to enhance clinical utility and hardware versatility. Planned work includes multimodal fusion, full-volume analysis, privacy-aware distributed training, and lightweight deployment for mobile and edge devices.

A. Multimodal Data Fusion

Future research will transition from single-modality MRI to a multimodal fusion approach. Combining contrast-enhanced T1-weighted, T2-weighted, and FLAIR images can help the model learn complementary pathological cues, including edema and hyperintense abnormalities. Additional work on three-dimensional CNNs and Vision Transformers will enable the processing of full MRI stacks and the automatic calculation of tumor volume and surface area.

B. Model Optimization for Edge Deployment

To facilitate use in portable medical devices and low-resource settings, future work will investigate model quantization and pruning. Reducing model weights from 32-bit floating-point to 8-bit integer format using tools such as TensorFlow Lite can lower memory consumption and speed up inference on mobile processors. In addition, knowledge distillation will be investigated so that a lightweight student model can acquire knowledge from the larger EfficientNetB0 teacher network.

C. Advanced Ensemble Learning

Although EfficientNetB0 is highly efficient, further robustness may be obtained through ensemble learning. A majority-voting or weighted-averaging ensemble combining EfficientNetB0, ResNet-101, and Inception-V3 could reduce generalization error and improve resilience in borderline cases, particularly between Glioma and Meningioma.

D. Enhanced Explainability and Privacy-Aware Training

Future updates will integrate U-Net++-based segmentation alongside classification so that the system can provide pixel-level masks of tumor extent rather than only coarse heatmaps. This would support better longitudinal monitoring of disease progression. Another promising direction is federated learning, which would allow training on distributed hospital datasets without compromising patient data privacy. A mobile version of the dashboard is also envisioned for preliminary screening in remote or resource-limited environments.

IX. CONCLUSION

The culmination of this research, embodied in the Bright-Brain system, underscores a pivotal shift toward the integration of high-performance deep learning with clinical interpretability. By meticulously fine-tuning an EfficientNetB0 architecture, this study has demonstrated that computational parsimony does not necessitate a compromise in diagnostic rigor. The system's ability to attain an 82.3% test accuracy across a heterogeneous dataset of 7,023 images serves as a testament to the efficacy of compound-scaled convolutional neural networks in navigating the complexities of neuro-oncological imaging.

Beyond mere classification, the implementation of Grad-CAM heatmaps effectively deconstructs the traditional "black-box" barrier, offering clinicians a transparent, gradient-based rationale for every diagnostic inference. This visual evidence, when synthesized with an immersive Three.js 3D visualization layer, transforms abstract algorithmic outputs into tangible, spatially-aware insights that are indispensable for pre-surgical mapping and longitudinal patient monitoring.

The technical synergy of a React.js frontend and a Python-Flask backend ensures that these sophisticated capabilities are not confined to high-compute laboratories but are accessible through a seamless, professional interface suitable for real-time clinical settings. While limitations regarding 3D volumetric analysis and rare tumor representation remain, the established framework provides a versatile foundation for future explorations into multimodal data fusion and privacy-preserving federated learning. Ultimately, Bright-Brain bridges the critical divide between computational innovation and practical medical utility, fostering a paradigm where AI serves as a reliable, transparent, and indispensable ally in the fight against neurological malignancies.

REFERENCES

[1] M. Tan and Q. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," ICML, 2019.

[2] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," ICCV, 2017.

[3] M. Nickparvar, "Brain Tumor MRI Dataset," Kaggle, 2020.

[4] J. Cheng et al., "Enhanced performance of brain tumor classification via tumor region augmentation," PLoS One, 2015.

[5] H. Khan et al., "Deep learning for brain tumor classification: A review," IEEE Access, 2020.

[6] S. Das et al., "Automated brain tumor classification using CNN," IEEE Trans. Med. Imaging, 2020.

[7] V. Anand et al., "Multi-class classification of brain tumors using optimized CNN," Sci. Rep., 2024.

[8] F. Ahmmed et al., "BTdiagAI: A Web-Deployed Hybrid Framework for Brain Tumor Classification," Healthc. Technol. Lett., 2025.

[9] Z. U. Rehman et al., "Efficient and interpretable brain tumor classification via distilled CNN," IEEE Access, 2025.

[10] G. Litjens et al., "A survey on deep learning in medical image analysis," Med. Image Anal., 2017.

[11] M. Aamir et al., "Optimized CNN for brain tumor detection," Diagnostics, 2024.

[12] K. Abbas et al., "Machine learning based methodology for brain tumor classification," Front. Neuroinform., 2024.

[13] R. Gupta and N. Kumar, "Development of a medical image archiving and retrieval system," Int. J. Inf. Technol., 2024.

[14] S. Bhuvaji et al., "Brain tumor classification dataset," Kaggle, 2020.

[15] WHO Classification of Tumours Editorial Board, Central Nervous System Tumours, 5th ed. Lyon, France: IARC Press, 2021.

[16] S. Sontakke, "Brain Tumor Detection and Classification Using Deep Learning Models," DevConf.IN, 2026.

[17] R. Kumar et al., "Quantization of Deep Neural Networks for Medical Image Analysis," MDPI, 2026.

[18] A. Rahman et al., "Multimodal Radiology-Pathology-Text Benchmark for Robust Brain Tumor Typing," arXiv, 2026.