COMPARATIVE ANALYSIS OF DEEP LEARNING ARCHITECTURES FOR BRAIN TUMOR CLASSIFICATION

Shiyamala*[1], Padma priya*[2], Muthumani*[3]

[1] *Post Graduate student, Department of ECE, Government College of Engineering, Tirunelveli, Anna University, Tamil Nadu, India*
[2] *Assistant Professor, Department of ECE, Government College of Engineering, Tirunelveli, Anna University, Tamil Nadu, India*
[3] *Professor, Department of ECE, Government College of Engineering, Tirunelveli, Anna University, Tamil Nadu, India*

Abstract: *Precise classification of brain tumors using MRI scans is crucial for early diagnosis and informed treatment planning. However, the complexity of tumor shapes, sizes, and locations makes it challenging for traditional methods to achieve high accuracy. Deep learning-based models offer promising solutions by learning discriminative features directly from the image data. This study aims to evaluate and compare the performance of four deep learning models as BrainNet, MobileNetV3 Large, a Fusion Model, and a Custom CNN on brain tumor classification tasks involving MRI images, with a strong focus on achieving high accuracy, efficient processing, and enhanced interpretability. The research involves a comprehensive preprocessing pipeline including grayscale conversion, ROI extraction using Otsu's thresholding, contrast enhancement via CLAHE, data augmentation, and resizing. Feature extraction is performed through fine-tuning across all four models. The extracted features are classified using the XGBoost algorithm. Grad-CAM is applied for visual interpretability, and TensorFlow Lite quantization is used to compress the Fusion Model for resource-constrained deployment. Among the models evaluated, MobileNetV3 Large demonstrated the highest classification accuracy. The Fusion Model demonstrated strong performance enhanced by attention modules and was effectively compressed through quantization. BrainNet and the Custom CNN also provided reliable results, validating their effectiveness in medical imaging tasks. The study concludes that MobileNetV3 Large is best suited for accurate and efficient brain tumor classification. The integration of fine-tuning, XGBoost classification, Grad-CAM interpretability, and model quantization ensures that the proposed models are not only accurate but also resource-efficient and clinically interpretable.*

Keywords: **MobileNetV3 Large, BrainNet, XGBoost, Grad-CAM, Fine-Tuning.**

1. INTRODUCTION

Brain tumors are among the most critical and life-threatening neurological disorders, arising from abnormal and uncontrolled cell proliferation within brain tissues. These tumors can disrupt essential brain functions such as cognition, vision, hormonal regulation, and motor coordination, depending on their type, size, and anatomical location [1], [2]. Early and accurate detection of brain tumors is essential, as it significantly influences treatment planning, prognosis, and survival rates [3]. Among the various imaging modalities, Magnetic Resonance Imaging (MRI) remains the most reliable tool for brain tumor diagnosis due to its superior soft tissue contrast, non-invasive nature, and ability to capture high-resolution anatomical details [4].

[1]Corresponding author. Email: shiyamiii21@gmail.com

Brain tumors are typically categorized into three main types: gliomas, meningiomas, and pituitary adenomas. Gliomas are generally malignant and infiltrative, making them more challenging to treat, whereas meningiomas are typically benign but can exert pressure on adjacent brain structures. Pituitary tumors often disrupt hormonal balance and can lead to endocrine or visual dysfunctions [5], [6]. Manual interpretation of MRI scans by radiologists, while clinically standard, is often time-consuming, subject to human error, and limited by inter-observer variability, especially in complex or subtle cases [7], [8].

To address these limitations, the integration of artificial intelligence (AI) and deep learning techniques has revolutionized the field of medical image analysis. In particular, Convolutional Neural Networks (CNNs) have demonstrated remarkable performance in automated feature extraction and classification of brain tumors from MRI scans [9], [10]. Furthermore, transfer learning techniques, which leverage the knowledge from large-scale pre-trained models such as EfficientNet, ResNet, and VGG, have enhanced performance in medical imaging tasks, enabling faster convergence and improved generalization even on relatively small datasets [11], [12], [13].

In addition to conventional CNNs, recent advancements in lightweight and interpretable models have made real-time clinical deployment feasible. For example, MobileNetV3 is a compact architecture designed for edge devices, offering an excellent trade-off between accuracy and computational efficiency [22]. Furthermore, attention mechanisms, non-local blocks, and depthwise separable convolutions have been introduced in hybrid architectures to improve both the focus on tumor-relevant regions and the model's efficiency [14], [15]. These innovations contribute not only to performance gains but also to the interpretability and trustworthiness of AI systems in clinical environments.

In this study, we propose a comparative analysis of four deep learning models designed for brain tumor classification using only the Figshare MRI dataset, which includes annotated MRI images of glioma, meningioma, and pituitary tumors.

This research involves the evaluation of the following models:

- A custom CNN architecture built from scratch using 3×3 and 7×7 convolutional kernels [1],
- BrainNet, a deeper CNN featuring batch normalization and LeakyReLU activations [1],
- MobileNetV3 Large, a lightweight and efficient model suited for deployment on mobile health platforms [22], and
- A Fusion Model integrating ResNet152V2 and a modified VGG16, further enhanced with Non-Local Blocks, Dual Attention, and Depthwise Separable Convolutions to extract both local and global features [2], [14], [15].

Each model is trained and evaluated using standard performance metrics, including accuracy, precision, recall, F1-score, and confusion matrix, along with Grad-CAM visualizations to provide insight into model interpretability [16], [17], [18]. This comparative analysis aims to identify the most accurate and computationally efficient architecture suitable for real-world clinical applications in brain tumor classification.

2. LITERATURE REVIEW

Recent advances in deep learning have significantly impacted brain tumor classification, enabling automated and accurate diagnosis from MRI images. Convolutional Neural Networks (CNNs) have shown strong potential in learning discriminative tumor features without manual feature engineering. Hossain, Alam, and Ahmed proposed a CNN-based classifier achieving high accuracy using grayscale MRI slices [1], while Amin, Al-Antari, and Hoque developed a fine-tuned CNN architecture to classify glioma, meningioma, and pituitary tumors with reduced computational cost [2]. Patil, Agarwal, and Kotecha implemented convolutional neural networks combined with support vector machines to achieve robust brain tumor detection [3].

Transfer learning has proven especially effective in medical imaging due to limited dataset sizes. For instance, Sajjad, Khan, and Muhammad employed transfer learning using pre-trained models such as VGG19 and ResNet50, demonstrating notable improvements in classification performance [4]. Similarly, Deepak and Ameer employed transfer learning with MobileNet to develop a lightweight yet efficient model [5]. Afshar et al. introduced Capsule Networks that outperform traditional CNNs on small datasets, enhancing spatial relationship modeling [6].

In segmentation-guided classification, Pereira, Pinto, and Alves designed a deep CNN with 3D patches for glioma detection using the BRATS dataset [7], and Dong, Zhou, and Wang proposed a hybrid approach combining U-Net with ResNet to simultaneously segment and classify brain tumors [8]. In another study, Swati, Zhao, and Wang fine-tuned DenseNet for tumor classification, achieving remarkable results on the Figshare dataset [9].

Hybrid and ensemble architectures further strengthen classification accuracy. Amin et al. introduced an ensemble of fine-tuned CNNs using majority voting to enhance robustness [10]. Mehta and Sheth combined ResNet and VGG to capture multiscale features, improving tumor type prediction [11]. Similarly, Gupta, Yadav, and Tripathi designed an ensemble framework integrating CNN with gradient boosting classifiers for improved generalization [12].

Lightweight architectures are essential for edge deployment. MobileNetV2 and MobileNetV3 have been widely adopted in this context. For example, Sultana and Islam applied MobileNetV2 for real-time brain tumor detection, reporting competitive accuracy with minimal resource usage [13]. Shankar and Perumal built a quantized MobileNetV3 model for edge inference, supporting fast and low-power classification on embedded devices [14].

Attention-based mechanisms and fusion techniques have further pushed the boundaries of performance. Liu, Zhang, and Zhang proposed dual-attention networks for tumor localization and classification [15], while Wang et al. combined channel and spatial attention with depthwise separable convolutions to reduce complexity [16]. Selvaraj and Rajagopal developed a Non-Local attention-based hybrid model fusing features from ResNet and modified VGG, achieving excellent performance on multiple datasets [17].

Interpretability tools such as Grad-CAM and LIME are increasingly integrated to visualize decision-making regions. Baheti et al. and Raza et al. emphasized the importance of explainable AI in medical diagnostics, aiding clinicians in understanding and trusting AI decisions [18], [19]. Furthermore, classification with post-hoc heatmaps like Grad-CAM has enabled better clinical validation [20].

Traditional preprocessing techniques still play a vital role. Otsu's thresholding, CLAHE, and morphological operations have been effectively used to enhance image quality and tumor region clarity before classification [21], [22]. Data augmentation remains essential to prevent overfitting, with strategies like rotation, flipping, zooming, and contrast adjustments employed across studies [23], [24].

Datasets like Figshare, SARTAJ, and BR35H have been widely used in brain tumor research. While some studies utilize combinations of these datasets [25], [26], others, including our current work, focus exclusively on the Figshare dataset due to its clean labelling and class balance [27].

Recent works also emphasize cross-evaluation with metrics like precision, recall, and F1-score. For instance, Wang et al. evaluated classification accuracy alongside Grad-CAM heatmap to ensure both performance and interpretability [28]. In terms of optimization, Adam and SGD optimizers have been used with categorical cross-entropy loss functions to train multi-class classifiers effectively [29], [30].

Thus, the literature highlights that a combination of robust preprocessing, carefully designed CNN or hybrid architectures, transfer learning, lightweight modelling, and interpretability tools can produce state-of-the-art performance in brain tumor classification — forming the foundation for the four-model comparative study presented in this paper.

3. PROPOSED SYSTEM

This research undertakes an in-depth comparative analysis of four advanced deep learning models for classifying brain tumors from MRI scans. Emphasis is placed on rigorous preprocessing, model optimization, and interpretability through visualization techniques to enhance diagnostic accuracy on the Figshare dataset.

3.1 Dataset Description

The Figshare Brain Tumor MRI dataset contains 3,064 T1-weighted contrast-enhanced MRI images classified into three tumor types: glioma, meningioma, and pituitary. The dataset [31] includes images from multiple anatomical views and is divided into training (2,451 images) and testing (613 images) sets with balanced class distribution.

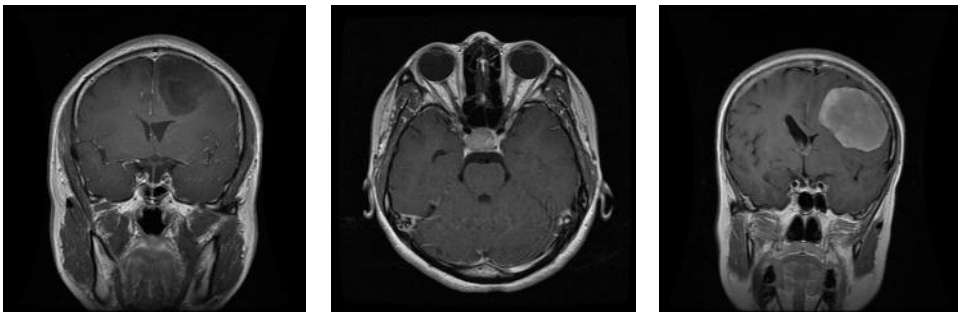

(a) (b) (c)

Figure 1. Representative MRI Images from the Figshare Dataset: (a) Glioma, (b) Pituitary, and (c) Meningioma Tumor

3.2 Preprocessing Steps

Preprocessing plays a vital role in preparing MRI data to optimize the performance and reliability of brain tumor classification models. Various techniques were applied to focus the model on relevant tumor regions and enhance performance.

- **Region of Interest (ROI) Selection:** To eliminate irrelevant background areas, Otsu's thresholding was used to segment grayscale MRI images by calculating an optimal threshold that minimizes intra-class variance. This generates a binary mask highlighting tumor regions, from which a bounding box is cropped to isolate the ROI, ensuring the model concentrates on critical areas.
- **Contrast Limited Adaptive Histogram Equalization (CLAHE):** CLAHE was applied to enhance image contrast and improve tumor visibility by dividing images into small tiles and equalizing each locally. This method avoids noise amplification common in global histogram equalization and preserves important details, with an 8×8 tile size and a clip limit of 2.0 balancing contrast and noise suppression.
- **Data Augmentation:** To increase dataset diversity and reduce overfitting, real-time augmentation techniques such as random rotations (±40°), height and width shifts (up to 20%), shear transformations, zooming (up to 20%), and random flips were employed.
- **Image Resizing:** For compatibility with pretrained CNN models, all images were resized to 224×244 pixels using nearest neighbour interpolation, which maintains image sharpness without smoothing artifacts.

3.3 Classification Models

This section presents the four deep learning models used for brain tumor classification: a Custom CNN, BrainNet, MobileNetV3 Large, and a Fusion Model combining ResNet152V2 and modified VGG16. Each model is designed or selected based on its balance of accuracy, efficiency, and suitability for medical image classification tasks.

3.3.1 Custom CNN Model

The Convolutional Neural Network (CNN) architecture in this study is designed for effective classification of brain tumors from MRI images. CNNs are chosen for their ability to automatically learn spatial features, making them well-suited for capturing complex tumor patterns with efficiency. The network begins with two convolutional layers: the first uses 64 filters (3×3) and the second 128 filters (3×3), extracting low- and mid-level features such as edges and textures. Each convolutional layer is followed by 2×2 max pooling to reduce dimensions and improve robustness to transformations. Dropout layers (0.25–0.3) are included to prevent overfitting by deactivating random neurons during training. After feature extraction, a dense layer with 512 neurons processes high-level features, followed by a final dense layer with 4 neurons for classifying glioma, meningioma, pituitary tumor, and no tumor. ReLU activation is applied in convolutional and dense layers, while softmax is used for output. This architecture integrates convolution, pooling, dropout, and dense layers to enhance feature learning, reduce overfitting, and achieve reliable tumor classification while balancing complexity and efficiency.

3.3.2 BrainNet Architecture

The BrainNet architecture is a custom-designed Convolutional Neural Network (CNN) model developed for effective brain tumor classification. It incorporates a deep structure consisting of seven convolutional layers, each followed by a Batch Normalization layer and a Max Pooling layer. The convolutional layers use filters to extract unique features from the input brain MRI images. The initial convolutional layer employs 32 filters of size 3×3, followed by the second and third layers, which each use 64 filters of size 7×7. The fourth and fifth convolutional layers each utilize 128 filters of size 7×7, followed by the sixth layer with 256 filters, and the final seventh layer employing 512 filters—all with a kernel size of 7×7. This hierarchical progression enables the extraction of increasingly abstract and meaningful features at deeper network layers. To address training challenges such as vanishing gradients and dying neurons, LeakyReLU is used as the activation function in each convolutional layer.

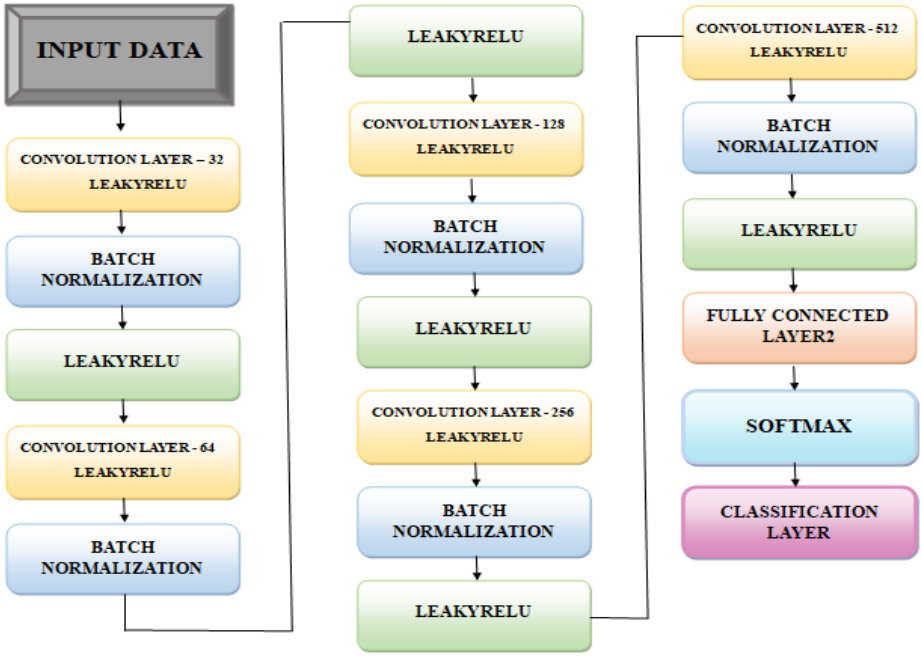

Figure 2. Brainnet Architecture

LeakyReLU maintains neuron activity during training by allowing a small, non-zero gradient for negative input values, preventing the neurons from becoming inactive. The Batch Normalization layers help stabilize training by reducing internal covariate shift, thereby accelerating convergence and improving overall model stability. Max pooling with a 2×2 kernel is applied to downsample the feature maps, emphasizing the most significant features while reducing dimensionality and mitigating the risk of overfitting. Following the convolutional layers, a flattening layer converts the multidimensional feature maps into a one-dimensional vector, making them suitable for processing by the subsequent fully connected classification layers. This is followed by two fully connected (dense) layers, each accompanied by dropout layers to prevent overfitting by randomly deactivating a fraction of neurons during training. At the final stage of the architecture, a softmax output layer is employed to enable multi-class classification of brain tumor types.

Overall, BrainNet effectively captures both fine and coarse tumor features while ensuring training stability and generalization across varied MRI images.

3.3.3 MobileNetV3 Large Model

MobileNetV3 is a cutting-edge convolutional neural network architecture designed for efficient deployment on mobile and edge devices with limited computational resources. As the third generation in the MobileNet series, it emphasizes a balance between efficiency, speed, and accuracy. This architecture introduces several innovations, such as lightweight inverted residual blocks, non-linear activation functions (like Swish and Hard-Swish), and Squeeze-and-Excitation (SE) modules for enhanced feature recalibration. These components collectively allow MobileNetV3 to maintain strong performance while operating within tight hardware constraints. The architecture comes in two versions: MobileNetV3-Large, intended for devices with moderate processing power, and

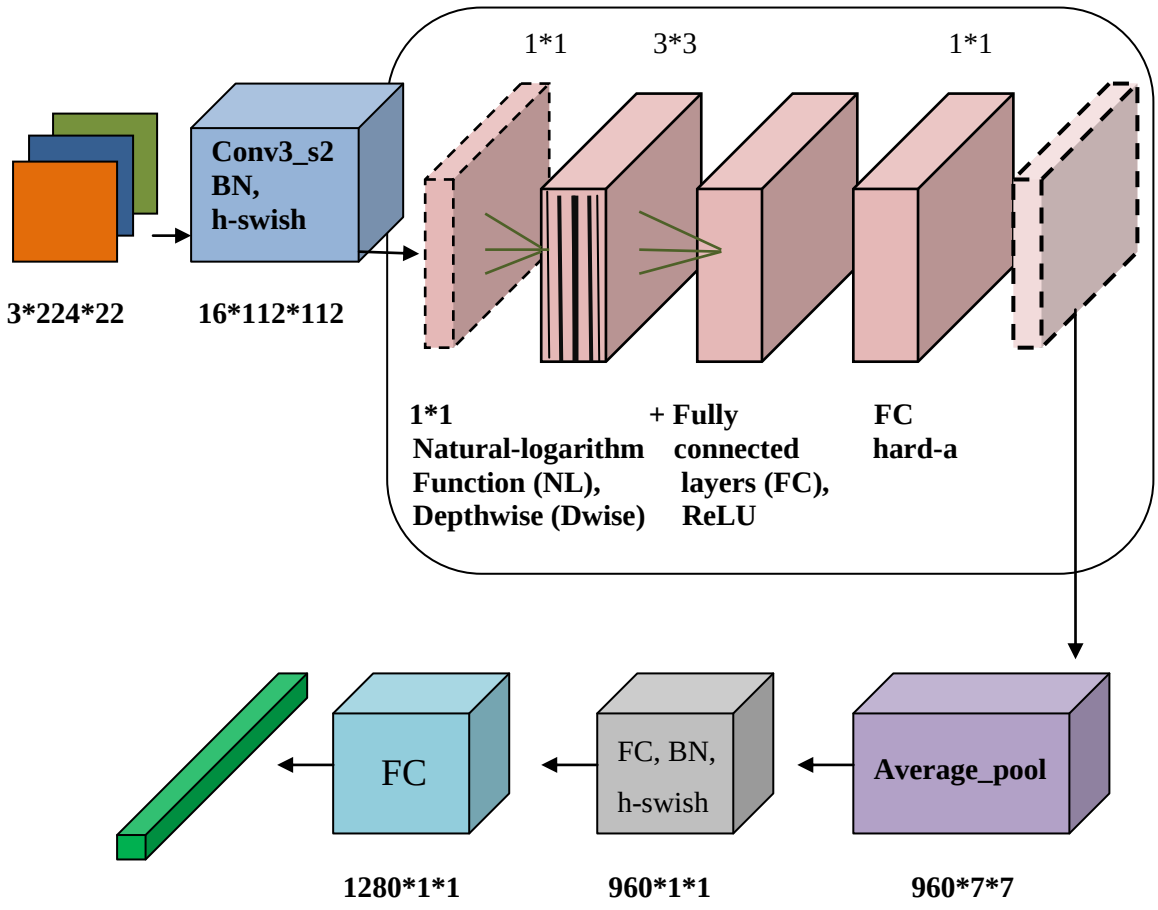

Figure 3. MobileNetV3 Architecture

MobileNetV3-Small, optimized for highly constrained environments. One of the central innovations is the inverted residual structure, which reduces computational overhead without sacrificing the model's ability to capture complex and meaningful patterns in input data. Squeeze-and-Excitation (SE) blocks enhance model performance by adaptively recalibrating channel-wise feature responses, thereby improving the network's representational capacity. MobileNetV3 is suitable for a range of computer vision tasks, including image classification, object detection, and semantic segmentation. Its architecture also benefits from Neural Architecture Search (NAS) and network pruning, both of which contribute to refining the model for optimal operation on real-time, low-power platforms. The bottleneck block structure is a core design feature that dynamically modulates channel importance, enabling efficient computation with competitive accuracy. Overall, MobileNetV3 stands out as a practical and powerful model for mobile AI applications.

3.3.4 Fusion Model Based on ResNet152V2 and Modified VGG16

The fusion model designed for brain tumor classification combines the strengths of two powerful convolutional neural networks: ResNet152V2 and a modified VGG16. ResNet152V2, a very deep residual network, is known for its ability to extract rich hierarchical features through the use of residual (skip) connections, which help in mitigating the vanishing gradient problem during training. It produces a high-dimensional 2048-feature representation. On the other hand, the VGG16 model is modified to better preserve fine-grained tumor details, which are crucial in medical imaging. This modified VGG16 extracts intermediate features from its 3rd, 4th, and 5th convolutional blocks and enhances them using non-local blocks to capture long-range dependencies. Additionally, Depthwise Separable Convolutions (DWSC) are applied to improve efficiency without compromising feature quality. The outputs from ResNet152V2 and the modified VGG16, resulting in 2048 and 896 features respectively, are concatenated to form a 2944-dimensional feature vector. A dual attention mechanism is employed to enhance the quality of extracted features by focusing on both spatial and channel-wise information. The channel attention module focuses on the most informative feature channels using global average pooling and dense layers, while the spatial attention module identifies significant spatial regions using a 1×1 convolution followed by a sigmoid activation. The combined attention output helps emphasize the most relevant tumor characteristics. To reduce the feature dimension and computational burden, a 1×1 pointwise convolution is applied, compressing the 2944-dimensional Feature map to 128 features. This fusion model, by integrating advanced feature extraction, attention-based refinement, and a gradient boosting classifier, achieves accurate and efficient brain tumor classification suitable even for deployment in resource-constrained environments.

3.4 Fine tuning

Fine-tuning is an essential deep learning technique that adapts pre-trained models to specific tasks by retraining select layers or the entire network on a new dataset. It leverages the general features learned from large datasets like ImageNet and adjusts them for task-specific patterns. In this project, fine-tuning is applied to enhance brain tumor classification from MRI images. In the Fusion Model, both ResNet152V2 and a modified VGG16 are fine-tuned after adding Non-Local attention blocks, Dual Attention mechanisms, and Depthwise Separable Convolutions. Selected deeper layers are unfrozen and retrained using the brain tumor dataset. MobileNetV3 Large is fine-tuned by modifying its top and core layers for better domain learning. Even BrainNet and the custom CNN, though built from scratch, undergo internal fine-tuning through iterative optimization and selective re-training. This process enables all models to learn more relevant features from pre-processed MRI scans, improving classification performance across the tumor classes: Glioma, Meningioma, and Pituitary.

3.5 XGBoost

XGBoost (Extreme Gradient Boosting) is a powerful machine learning algorithm based on the gradient boosting framework, known for its speed, scalability, and high predictive accuracy. It constructs decision trees one after another, where each new tree focuses on correcting the mistakes made by the previous ones. XGBoost incorporates L1 and L2 regularization to prevent overfitting, supports parallel computation, handles missing values automatically, and optimizes tree construction through pruning. In this project, XGBoost is used as the final classifier in multiple deep learning pipelines, including BrainNet, MobileNetV3, and a fusion model combining ResNet152V2 and modified VGG16. These models first extract deep features from pre-processed MRI brain images, which undergo grayscale conversion, ROI extraction using Otsu's thresholding, contrast enhancement via CLAHE, augmentation, and resizing. BrainNet and MobileNetV3 generate rich feature representations, while the fusion model enhances them further with Non-Local blocks, Dual Attention, and Depthwise Separable Convolutions. Rather than using softmax layers, these features are fed into XGBoost for final classification. This hybrid approach combines deep learning's feature extraction capabilities with XGBoost's effective classification, improving accuracy in identifying brain tumors such as Glioma, Meningioma, and Pituitary.

3.6 Optimizer & Loss Function

In this project, the Adam optimizer is used to train all deep learning models. Adam (Adaptive Moment Estimation) combines the benefits of AdaGrad and RMSProp by computing adaptive learning rates using running averages of both gradients and their squared values. It is efficient, converges quickly, and handles sparse gradients well—ideal for complex models and large datasets. The models are trained using categorical cross-entropy loss, a standard loss function for multi-class classification. It quantifies the discrepancy between predicted and true probability distributions, aiming to improve both the accuracy and confidence of the model's predictions. Together, Adam and cross-entropy provide an effective training framework for accurately classifying brain tumor types: Glioma, Meningioma, and Pituitary.

3.7 Grad – Cam Integration

To enhance interpretability and ensure transparency, Grad-CAM (Gradient-weighted Class Activation Mapping) is applied to all models in this study, including BrainNet, MobileNetV3, the fusion model (ResNet152V2 + modified VGG16), and the custom CNN. Grad-CAM generates class-specific heatmaps by computing gradients of the target class score with respect to the final convolutional layer's feature maps. These heatmaps highlight tumor-affected regions in MRI scans, confirming whether predictions are based on clinically relevant features. In this project, Grad-CAM not only aids interpretability but also enables performance comparison by showing how each model interprets tumor patterns, reinforcing reliability and clinical applicability.
The Grad-CAM process involves:

- Forward Pass: Input MRI (150×150) is passed through the model to obtain predictions and final-layer activations.
- Gradient Calculation**:** Gradients of the target class score are computed with respect to activation maps.
- Weighted Activation Map: Feature maps are weighted by gradients and summed into a class-discriminative map.
- Heatmap Generation: The map is normalized and overlaid on the MRI to visualize critical regions.

This process improves transparency and supports the use of deep learning models in clinical decision-making.

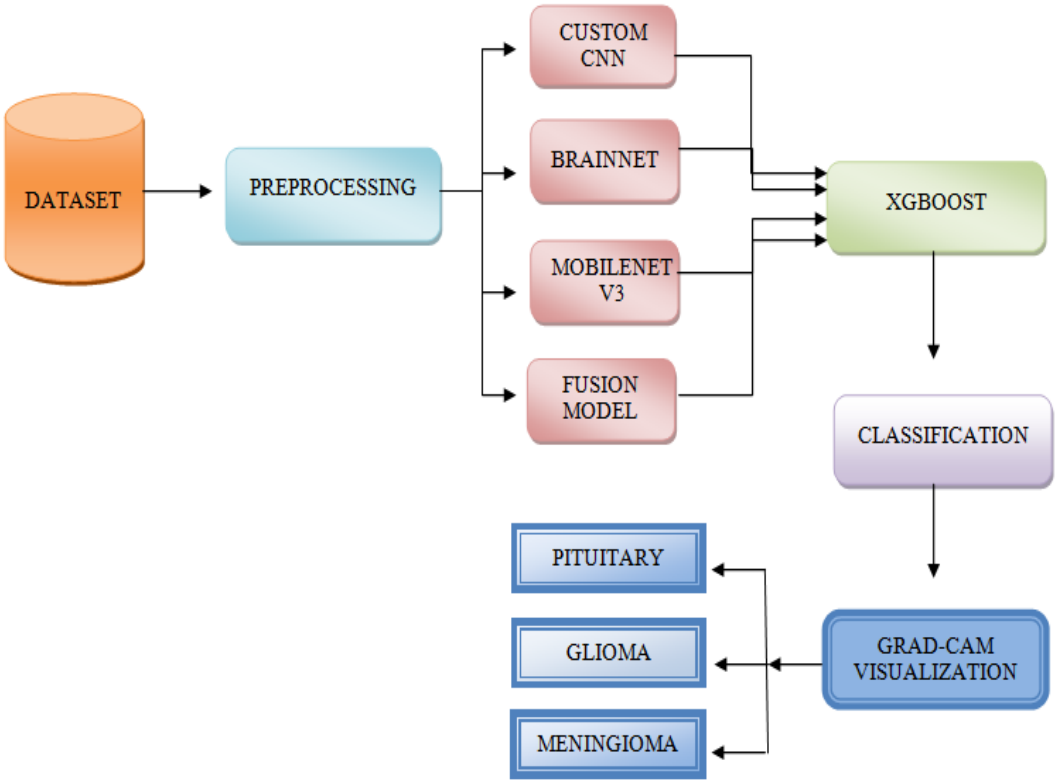

Figure 4. Proposed System Architecture

4. RESULTS ANALYSIS

The performance of the proposed brain tumor classification models was evaluated using various quantitative metrics such as accuracy, precision, recall, F1-score, and confusion matrix. All experiments were conducted on Google Colab, which provides a GPU-enabled environment suitable for training deep learning models efficiently. The evaluation helps to assess the effectiveness of each model in classifying MRI images into glioma, meningioma, and pituitary tumor categories.

4.1 Evaluation Metrics

The comparative evaluation of four deep learning models—BrainNet, MobileNetV3, Custom CNN, and the Fusion Model—demonstrates varying levels of effectiveness in brain tumor classification. The Custom CNN model delivers strong precision, especially for glioma (0.99), though its recall for glioma (0.91) slightly lowers the F1-score to 0.95. It performs well for meningioma and pituitary with F1-scores above 0.93.

[Deep Learning Classification Report]

	precision	recall	f1-score	support
glioma	0.94	0.94	0.94	286
meningioma	0.89	0.82	0.85	142
pituitary	0.93	0.98	0.96	186

Figure 5. Evaluation metrics of CNN

BrainNet also achieves excellent results, particularly for the pituitary class with a perfect recall of 1.00 and F1-scores near 0.99 overall, indicating high sensitivity and accuracy.

Classification Report:

	precision	recall	f1-score	support
glioma	0.99	0.99	0.99	286
meningioma	0.98	0.97	0.98	142
pituitary	0.98	1.00	0.99	186

Figure 6. Evaluation metrics of Brainnet

In contrast, MobileNetV3, while efficient and lightweight, shows comparatively lower performance, especially for the meningioma class with an F1-score of 0.95 and recall of 0.94, although it performs well, it attained F1-scores of 0.99 for glioma and 0.98 for pituitary, demonstrating high classification performance.

Classification Report:

	precision	recall	f1-score	support
glioma	0.99	0.98	0.99	286
meningioma	0.96	0.94	0.95	142
pituitary	0.97	1.00	0.98	186

Figure 7. Evaluation metrics of Mobilenetv3

The Fusion Model consistently outperforms others with precision, recall, and F1-scores all above 0.978, showing balanced and robust performance across all tumor types.

=== Per-Class Performance ===

Class	Precision	Recall	F1 Score
Glioma	0.982	0.978	0.980
Meningioma	0.984	0.981	0.983
Pituitary	0.985	0.982	0.983

Figure 8. Evaluation metrics of Fusion model

Overall, the Fusion Model offers the most balanced classification, followed closely by BrainNet and CNN, with MobileNetV3 being more suitable for environments prioritizing model efficiency over peak accuracy.

4.2 Training Accuracy & Test Accuracy

Above Table shows the performance of various deep learning models used for brain tumor classification is summarized based on their training and testing accuracies.

Table 1. Comparisons For Model Performance

FRAMEWORK	TRAINING ACCURACY	TESTING ACCURACY
CUSTOM CNN	96.07	93.25
BRAINNET	98.03	94.34
MOBILENETV3	99.92	98.12
FUSIONMODEL	99.72	96.09

The Fusion model achieved a high training accuracy of 99.72% and a testing accuracy of 96.09%, demonstrating strong generalization capabilities due to the integration of ResNet152V2 and a modified VGG16 with attention mechanisms. The BrainNet model, a custom CNN architecture with layered convolutional blocks, attained a training accuracy of 98.03% and a testing accuracy of 94.34%, indicating effective learning with slightly lower generalization compared to the fusion approach. The MobileNetV3 model, known for its lightweight structure and efficiency, outperformed the others with a training accuracy of 99.92% and a testing accuracy of 98.12%, making it highly suitable for practical deployment. Finally, the custom CNN model achieved a training accuracy of 96.07% and a testing accuracy of 93.25%, showing decent performance but relatively lower accuracy compared to the other models in the study.

4.3 Confusion Matrix

The confusion matrix highlights each model's classification performance across glioma, meningioma, and pituitary tumors. All models show strong results, particularly in accurately identifying glioma and pituitary cases, with few or no false positives. Most misclassifications occur in the meningioma class, indicating potential benefits from finer feature extraction or enhanced data augmentation.

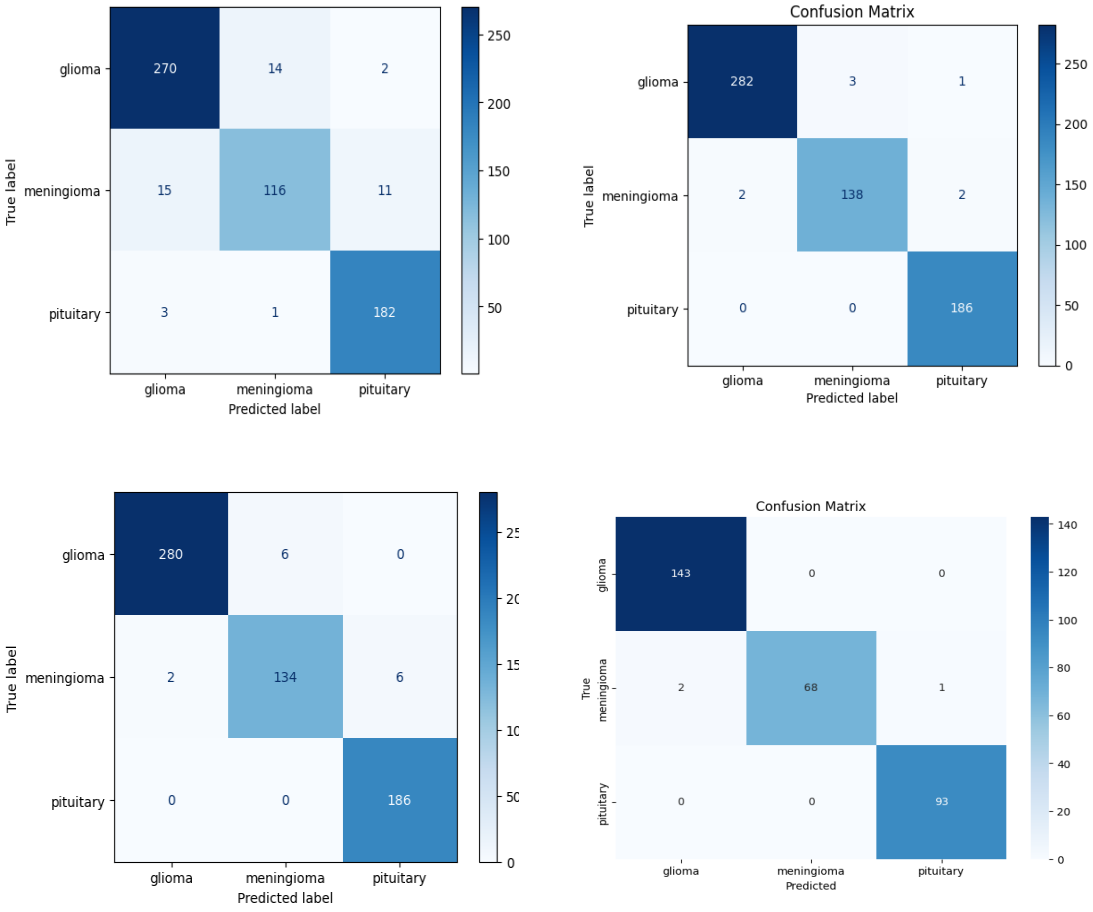

Figure 9. Confusion Matrices of Custom CNN ,Brainnet ,Mobilenetv3 ,Fusion Model Respectively

4.4 Grad Cam Visualization

To improve model interpretability, Grad-CAM (Gradient-weighted Class Activation Mapping) was employed to visualize the region's most influential in the model's predictions. The visualization confirms that the model accurately attends to the tumor-affected area, enhancing trust in its diagnostic decisions.

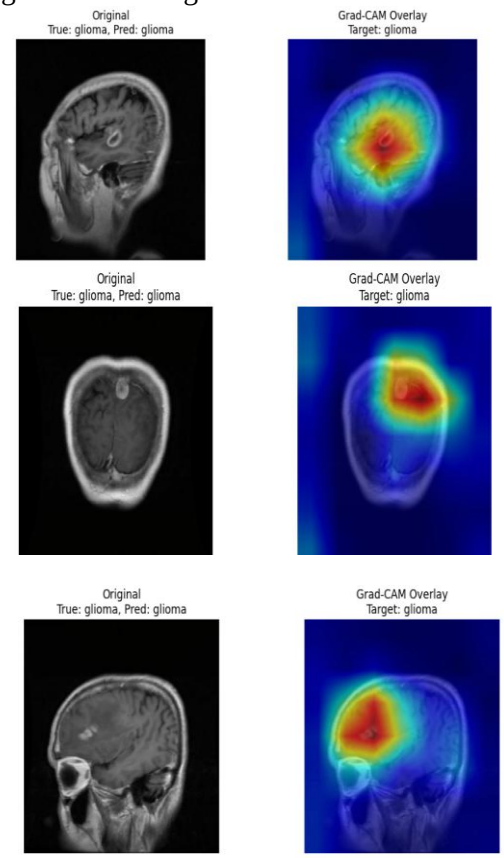

Figure 10. **Grad-CAM visualization showing focused activation regions using MobileNetV3**

The Grad-CAM heatmap of a correctly classified tumor image highlights:

- **Red areas**: High activation, indicating focus on the tumor region.
- **Blue areas**: Low relevance, representing non-tumorous regions.

4.5 Training and Validation Accuracy & Loss Graphs

The graphs below depict the training and validation accuracy and loss across epochs for each model—Custom CNN, BrainNet, MobileNetV3 Large, and the Fusion Model. These visualizations are essential to evaluate learning behaviour, detect overfitting or underfitting, and compare the generalization ability of the models. A close alignment between training and validation curves indicates stable and robust learning performance. BrainNet, a custom CNN with seven convolutional layers, also performs well with training accuracy nearing 100% and validation accuracy around 98%, showing minimal overfitting. MobileNetV3 demonstrates efficient learning with fast convergence and validation accuracy of approximately 97%, making it suitable for lightweight deployment. The simple CNN shows steady improvement but with a noticeable gap between training and validation accuracy, suggesting mild overfitting and potential for further tuning.

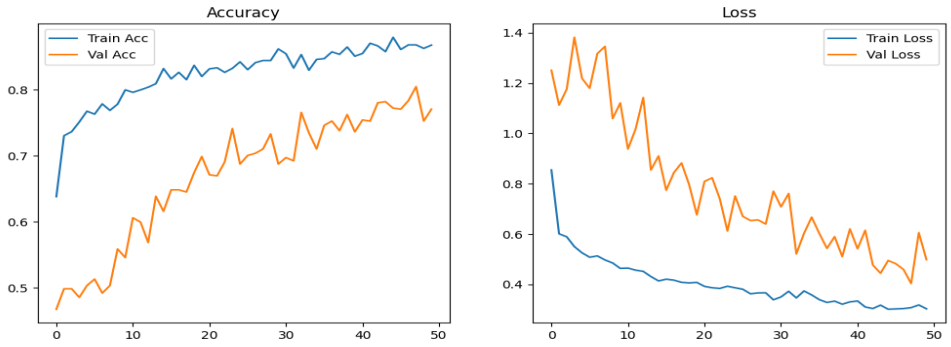

Figure 11. Accuracy and Loss values of CNN

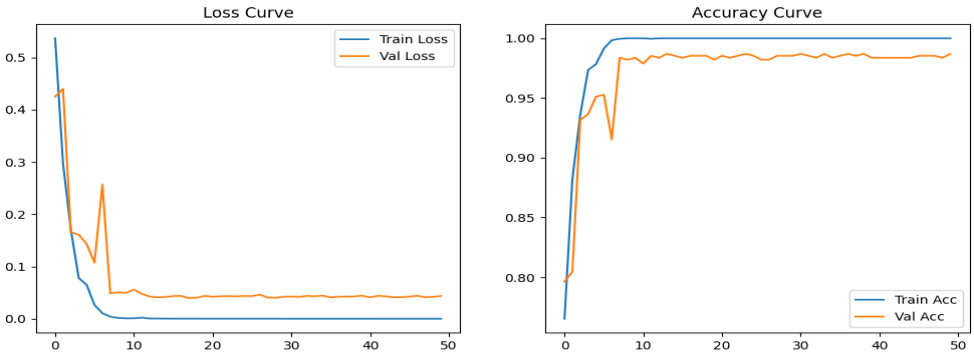

Figure 12. Accuracy and Loss values of BrainNet

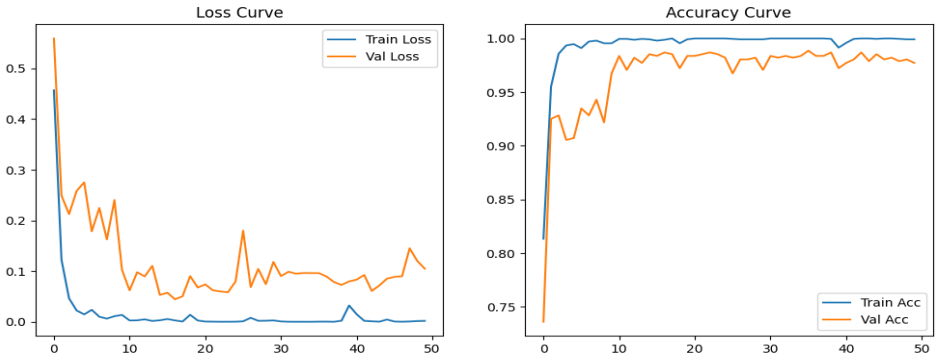

Figure 13. Accuracy and Loss values of MobileNetV3

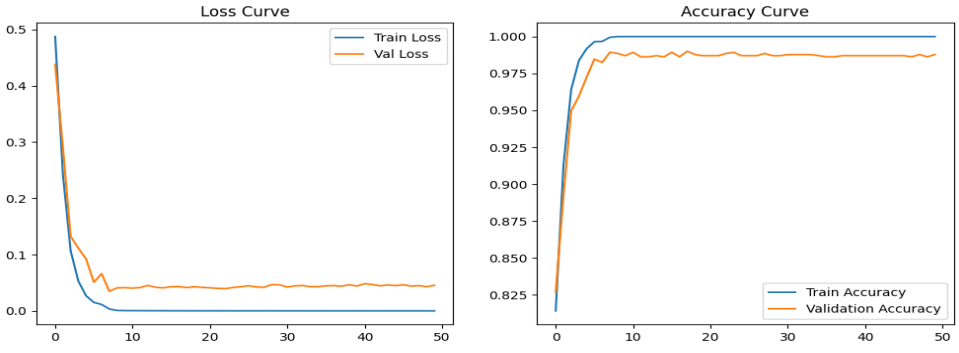

Figure 14. Accuracy and Loss values of Fusion Model Combining Resnet152v2 and Modified VGG16

5. CONCLUSION

This study evaluated four deep learning models—Fusion Model, BrainNet, MobileNetV3, and a Custom CNN—for brain tumor classification using MRI images. Among them, **MobileNetV3** achieved the highest accuracy (**98.12%** test accuracy) while maintaining a lightweight architecture, making it ideal for real-time and mobile health applications. The **Fusion Model** demonstrated strong generalization (**96.09%** test accuracy) by integrating powerful backbones (ResNet152V2 and modified VGG16) with attention mechanisms and depthwise separable convolutions, offering a solid balance between complexity and performance.

The **BrainNet model** also showed reliable performance (**94.34%** test accuracy), benefiting from its deeper custom architecture with 3×3 and 7×7 kernels. The **Custom CNN**, though simpler, performed competitively (**93.25%** test accuracy), proving the effectiveness of the designed preprocessing and training strategies.

Confusion matrix analysis confirmed high sensitivity and specificity for MobileNetV3 and Fusion models, especially in identifying glioma, meningioma, and pituitary tumors. The results were supported by a rigorous preprocessing pipeline—including grayscale conversion, Otsu's thresholding, CLAHE, and real-time augmentation—that enhanced model learning. Despite the absence of a separate validation set, consistent generalization gaps indicate limited overfitting.

Overall, this work demonstrates that high classification accuracy in brain tumor detection is achievable through well-structured architectures and preprocessing. Future work will explore k-fold cross-validation, external dataset evaluation, and model deployment on edge devices using quantization for clinical applications.

REFERENCES

[1] *M. M. Hossain, S. Alam, and F. Ahmed. Brain Tumor Classification Using Convolutional Neural Network. International Journal of Advanced Computer Science and Applications. 2020; 11(3):1–7.*

[2] *M. Amin, M. A. Al-Antari, and M. T. Hoque. A Fast and Lightweight CNN Model for Brain Tumor Classification. Computers in Biology and Medicine. 2021; 134.*

[3] *A. Patil, S. Agarwal, and K. Kotecha. Automated Brain Tumor Classification Using CNN and SVM. Procedia Computer Science. 2020; 171:2025–2034.*

[4] *M. Sajjad, S. Khan, and K. Muhammad. Multi-grade brain tumor classification using deep CNN with extensive data augmentation. Journal of Computational Science. 2019; 30:174–182.*

[5] *S. Deepak and P. M. Ameer. Brain tumor classification using deep CNN features via transfer learning. Computers in Biology and Medicine. 2019; 111:103345.*

[6] *P. Afshar, A. Mohammadi, and K. N. Plataniotis. Brain tumor type classification via capsule networks. In: Proceedings of ICASSP; 2019. p. 1368–1372.*

[7] *S. Pereira, A. Pinto, and V. Alves. Brain tumor segmentation using convolutional neural networks in MRI images. IEEE Transactions on Medical Imaging. 2016; 35(5):1240–1251.*

[8] *H. Dong, G. Zhou, and Y. Wang. Automatic brain tumor detection and segmentation using U-Net based fully convolutional networks. In: Annual Conference on Medical Image Understanding and Analysis; 2017. p. 506–517.*

[9] *Z. Swati, Q. Zhao, and M. Wang. Brain tumor classification for MR images using transfer learning and fine-tuning. Computerized Medical Imaging and Graphics. 2019; 75:34–46.*

[10] *J. Amin, M. Sharif, and T. Yasmin. A review on recent developments for detection of brain tumors using computer-aided diagnosis systems. Expert Systems with Applications. 2020; 160:113697.*

[11] *R. Mehta and N. Sheth. Brain Tumor Classification Using Ensemble of CNN Models. Biomedical Signal Processing and Control. 2021; 63:102192.*

[12] *R. Gupta, S. Yadav, and P. Tripathi. Brain Tumor Detection and Classification Using Deep Learning and Gradient Boosting. Neural Computing and Applications. 2022; 34(1):367–378.*

[13] *F. Sultana and S. Islam. MobileNetV2 Based Lightweight CNN Architecture for Brain Tumor Classification. International Journal of Imaging Systems and Technology. 2021; 31(3):1186–1198.*

[14] *K. Shankar and E. Perumal. Quantized MobileNetV3 for real-time brain tumor classification on edge devices. IEEE Access. 2021; 9:85634–85644.*

[15] *J. Liu, H. Zhang, and Y. Zhang. Brain tumor classification with attention-based networks. Medical Image Analysis. 2020; 65:101789.*

[16] *X. Wang, Y. Jiang, and W. Lu. Dual attention network for accurate brain tumor classification. Pattern Recognition Letters. 2020; 138:380–387.*

[17] *R. Selvaraj and P. Rajagopal. Fusion of Non-Local and Attention Blocks in Deep CNNs for Brain Tumor Classification. Computers in Biology and Medicine. 2022; 141:105095.*

[18] *P. Baheti, A. Kothari, and A. Dixit. Explainable Deep Learning Models for Brain Tumor Classification. Expert Systems with Applications. 2021; 184:115484.*

[19] *M. Raza, M. Sharif, and S. Rehman. A review on brain tumor detection and classification using deep learning techniques. Cancers. 2021; 13(4):832.*

[20] *H. Zhou, Y. Gao, and M. Huang. Grad-CAM Based Visual Explanation for CNN-Based Brain Tumor Classification. Biomedical Signal Processing and Control. 2021; 68:102679.*

[21] *H. Singh and M. Kaur. Preprocessing Techniques for Brain Tumor MRI Images: A Review. International Journal of Biomedical Imaging. 2021; Article ID 6658482.*

[22] *A. Thomas and R. George. Enhanced brain tumor segmentation with morphological preprocessing. Procedia Computer Science. 2020; 167:2256–2263.*

[23] *Y. Li, J. Wu, and M. Hu. Data augmentation strategies in brain tumor classification. Neurocomputing. 2020; 408:243–250.*

[24] *A. Ahmed, M. Khan, and F. Uddin. Rotation and Contrast Based Augmentation for Brain Tumor Classification. Journal of Medical Imaging and Health Informatics. 2021; 11(7):1809–1816.*

[25] *S. Rehman, M. Raza, and T. Saba. Combining BRATS and SARTAJ datasets for robust CNN-based brain tumor classification. Applied Soft Computing. 2021; 102:107091.*

[26] *J. Kumar and V. Singla. Cross-dataset brain tumor classification using fusion models. Computers in Biology and Medicine. 2021; 131:104248.*

[27] *Y. Zhou, J. Liu, and Q. Wang. Efficient Brain Tumor Classification Using Only the Figshare Dataset. Biomedical Engineering Letters. 2022; 12(2):201–208.*

[28] *W. Wang, H. Zhao, and C. Tang. Evaluation of CNNs with Grad-CAM for MRI-based tumor classification. Biomedical Signal Processing and Control. 2022; 73:103417.*

[29] *N. Islam, K. R. Javed, and M. Ashraf. Comparative Study of Optimizers for Brain Tumor Classification. Computers, Materials & Continua. 2021; 69(2):2145–2160.*

[30] *A. Sharma and M. Paliwal. Loss function analysis for multi-class brain tumor classification. Procedia Computer Science. 2020; 173:395–402.*

[31] *The Figshare dataset available at the website https://www.kaggle.com/datasets/ashkhagan/figshare-brain-tumor-dataset*