International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 4, July – August 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

BRAIN TUMOR DETECTION USING CNN

Jyoti [1], Amandeep [2], Dharmender Kumar[3]

*1M.Sc. Computer Science, Artificial Intelligence And Data Science, GJUS&T, India.
Email: ja7497802109@gmail.com
*2Assistant Professor, Artificial Intelligence And Data Science, GJUS&T, India
Email: noliyaaman@gjust.org
*3Professor, Artificial Intelligence And Data Science, GJUS&T, India.
Email: dharmindia24@gmail.com

Abstract

In the realm of neuro-oncology, diagnosis of a brain tumour using MRI will determine the treatment and the most clinically significant brain tumour types include meningioma, glioma and pituitary tumour. Manual, slice-by-slice radiologic inspection is extremely time consuming, and such inspection is subject to inter-observer variance, for which this deep learning approach is encouraging, particularly by the use of Convolutional Neural Networks (CNNs) that learn discriminative features directly from the raw MRI data. A 23-layer CNN ("Base Model") is reproduced on the contrast-enhanced MRI dataset of 3,064 images taken on the T1-weighted MRI from the Figshare, trained using Sparse Categorical Cross-Entropy loss and Softmax activation, and the results reported for each class, the confusion matrix and loss/activation ablations are independently computed to ensure consistency with the reproduced CNN. The Base Model has two limitations: 1) It is based on a single fixed preprocessing pipeline applied to all MRI slices, whereas the images can vary in blur, brightness, noise and contrast, and 2) There is no mechanism beyond just scalar metrics to explain the influence on a specific prediction of a given image, which makes visualizing the influence more difficult. The deficiencies it aims to address are tackled by a hybrid pipeline that keeps the validated CNN used and adds a module called Quality Assessment and Adaptive Preprocessing which will learn to do denoising, enhancing the contrast and sharpening on each individual image and a module for explainability named Grad-CAM which will produce a visual heatmap for each prediction. Comparing the proposed hybrid model with the reproduced Base Model, a higher accuracy (98.04%), average precision (97.80%), average recall (97.95%) and average F1-score (0.979) are achieved on the same stratified (70/15/15) dataset of the same Figshare data. The results demonstrate the robustness and interpretability of existing CNN models can be increased with only per-image adaptive preprocessing and increasing the interpretability using explainability provided by Grad-CAM, and there is still an important future direction to validate the model with an external dataset, and to present the explanations to radiologists.

Keywords: Brain Tumor Classification, Convolutional Neural Network, Figshare Dataset, Grad-CAM, Adaptive Preprocessing, Explainable AI.

1.INTRODUCTION

A brain tumour is an abnormal growth of cells in the brain or spinal cord that multiply or grow out of control and can be either non-cancerous (benign) or cancerous (malignant) depending on the location of the tumour in the brain and the type of tissue that the tumour involves [2]. This work focuses on three types of tumors which are extensively studied in the literature: Meningioma, a mostly benign tumor, which may affect the neighboring brain structures due to pressure, Glioma, tumor from glial tissue, which can be slow growing or highly malignant and is one of the most common primary brain tumors, and Pituitary tumor, a mostly benign tumor in the pituitary gland that may disrupt the regulation of the hormones and may be a source of pressure for the optic structures [1], [2] It is important to correctly classify an individual among these three types; this will determine the type of treatment you need (surgery, radiation, hormonal or watchful waiting) [6] [9]. For a long time the MRI has been the most important imaging modality for the evaluation of brain tissue, a technique that provides good tissue contrast without any requirement for the use of ionizing radiation, and radiologists have long used careful visual inspection of MRI slices to detect, locate and classify tumors [4], [5]. An effective but time consuming technique, it relies on the experience of the observer, and has high inter-observer variability, particularly where it is difficult to distinguish between normal and tumour tissue. MRI has been used in the past as the imaging modality of choice when imaging brain tissue, as it allows for excellent contrast of tissues with no exposure of the patient to any ionizing radiation, and the radiologist has traditionally used careful visual examination of the MRI slices to detect, locate and classify tumors. This method is effective, but slow and observer dependent as inter-observer variation may occur with indistinguishable tumors. Convolutional Neural Networks (CNNs) have been a successful approach in this last decade as they learn discriminative features directly from the unprocessed MRI scans and achieve the classification performance comparable to, or even better than, classic machine-learning pipelines [1]–[3]. This study will be limited to a representative deep-learning brain-tumor classifier, the 23-layer Convolutional Neural Network (CNN) reported in [1]

and used throughout this study, which has been trained to obtain 97.8% accuracy on the three-class Figshare dataset using Softmax activation and Sparse Categorical Cross-Entropy loss, to be referred to as the "Base Model" throughout the remainder of this study. Section III provides a detailed reproduction of the methodology, per-class results and hyperparameter ablation of the Base Model, for thorough verification of its reported performance. Two practical limitations are seen from the reproduced numbers: (i) the Base Model is applied to each MRI slice with the same preprocessing pipeline, ignoring the variations in blur, brightness, noise and contrast of the images [2], [9]; and (ii) the model does not provide any explainability mechanism and only reports scalar-accuracy style metrics with no visual cues on which region of an image generated a particular prediction a significant drawback in a clinical setting where a physician should be able to determine if a prediction is supported by clinically relevant tumor features [11], [16]. This paper presents a hybrid approach, incorporating two extra modules, a Quality Assessment and Adaptive Preprocessing stage that computes per-image quality indicators and adapts denoising, contrast enhancement and sharpening operations to these quality indicators [12], [13] and an explainability stage, Grad-CAM, that produces a heat-map visualization of each prediction. The comparison is done with the same Figshare dataset in order to ensure the comparison is controlled and fair [2], [1], and on the same train/validation/test split as the Base Model. There are some limitations: The assumption that the ground truth split (70/15/15) is reasonable for this study is not verified; the CNN and the Grad-CAM contribution are trained directly in NumPy, instead of a high-level deep-learning framework, and this has been done at the cost of slower training times, so as to retain full visibility into the forward and backward computation steps needed for the gradient-based explainability [15].

2. LITERATURE REVIEW

Brain tumor classification from MR images has not been followed chronologically but rather it has been several parallel trends. Early methods utilized a classical machine learning pipeline: manual or automatic segmentation of a tumor region, hand-crafted features extracted from the region (texture, shape or wavelet-based) and training of a classical classifier using the features. Cheng et al. [2] presented the Figshare benchmark with 3,064 slices of T1-weighted contrast-enhanced MRIs across three classes of tumors – meningioma, glioma and pituitary tumor – which were utilized by the majority of the following deep-learning studies as well as the task set and Base Model of this work. Other classical methods were used on much smaller databases, such as Chaplot et al. [4] which made use of a discrete wavelet transform and an SVM classifier with 52 images to obtain an accuracy of 98%, for El-Dahshan et al. [5] who used a combination of wavelet features, PCA dimensionality reduction and a k-nearest-neighbor classifier on 70 images to achieve an accuracy of 98.6%. The main drawback to any one of these pipelines is that the whole pipeline is limited by the performance of the segmentation and the features chosen are dependent on the designer. The proposed end-to-end CNN model avoided the feature-engineering step when training the CNN using the labeled MRI slices. Abiwinanda et al. [3] trained a simple CNN with two convolution layers, two pooling layers and two fully connected layers on the Figshare dataset and attained the accuracy of 84.19% which proves its feasibility as well as proving that there is a significant amount of accuracy left on the table for a shallow, unregularized network. Sajjad et al. [6] applied block-wise fine-tuning of a pre-trained VGG19 along with five-fold cross validation to achieve 94.82% accuracy, whereas Swati et al. [7] segmented the tumor, performed data augmentation and used a fine-tuned VGG19 to achieve robustness when the number of samples was reduced. Deepak and Ameer [8] extracted deep features from a pre-trained GoogLeNet and classified them using the conventional classifiers, they achieved almost 97–98% accuracy on the same three class problem. According to Rehman et al. [9] the backbone and the adaptation strategy significantly affect the final accuracy with the best result achieved by using a fully fine-tuned VGG16 network compared to three different pre-trained backbones (AlexNet, GoogLeNet, VGGNet) and the three adaptation strategies (layer freezing, fine tuning). A second step was taken by Anaraki et al. [10] who, in addition to evolving the CNN architecture itself (also using a genetic algorithm), showed to achieve an accuracy of 94.2% for the three-class problem and of 90.9% for an associated glioma grading problem. For this reason, a 23-layer CNN trained end-to-end with Sparse Categorical Cross-Entropy loss and with Softmax activation, as described in Khan et al. [1] (the Base Model), is used here as a starting point for further developments, and is able to reach an accuracy of 97.8% when trained on the complete 3,064-image Figshare dataset. Instead of just maximising accuracy, the other new line of research is to try to explain the predictions of the CNN; that means making it more transparent to a clinician, as the CNN that gives no indication about which parts of the image were involved in the classification will not be trusted by the clinician. Ozyurt et al. [12] combined the neutrosophic set theory and CNN together with fuzzy entropy measure to enhance robustness of the segmentation of tumor regions against the noise and ambiguity prior to classification, and separately used the grid-search-based framework to optimize the task-specific CNN architectures for binary detection, three-class classification and tumor grading by Irmak [13]. In this context, the most relevant work was that of Iftikhar et al. [11] who developed an explainable CNN to classify brain tumors and provide a visual explanation of its predictions using explainable-AI (XAI) techniques, such as reduced-layer-count architecture specifically designed from the beginning to be explainable.

Grad-CAM, proposed by Selvaraju et al. [16] is an interesting approach, because it introduces a coarse localization heat-map without the need for additional back-propagation of gradients through the network, or changes to the network's architecture.

The most relevant studies to this research are summarized in Table 1 and the gap which motivates the proposed hybrid model is highlighted in Table 2.

Table 1: Summary of the literature reviewed for this study.

Study	Year	Method	Dataset	Reported Accuracy
Cheng et al. [2]	2015	Tumor-region augmentation + partition (classical pipeline)	Figshare (introduced in this work)	Dataset paper / baseline
Abiwinanda et al. [3]	2019	CNN trained from scratch (2 conv + 2 pool + 2 FC)	Figshare (3,064 img.)	84.19%
Swati et al. [7]	2019	Block-wise fine-tuned VGG19 (5-fold CV)	Figshare (3,064 img.)	94.82%
Deepak & Ameer [8]	2019	Pre-trained GoogLeNet features + classifier	Figshare (3,064 img.)	≈ 97–98%
Anaraki et al. [10]	2019	CNN architecture evolved via genetic algorithm	Figshare (3,064 img.)	94.2%
Khan et al. [1] (Base Model)	2022	23-layer CNN, Sparse Categorical CE + Softmax	Figshare (3,064 img.)	97.8%
Iftikhar et al. [11]	2025	Explainable CNN with Grad-CAM-based XAI	Brain-tumor MRI (XAI-focused)	Explainability-focused

Table 2: Coverage of adaptive preprocessing and integrated explainability across the reviewed literature

Study	Per-Image Adaptive Preprocessing	Integrated Explainability (on accuracy-optimized model)
Cheng et al. [2]	No (fixed region rule)	No
Abiwinanda et al. [3]	No (fixed pipeline)	No
Swati et al. [7]	No (fixed pipeline)	No
Deepak & Ameer [8]	No (fixed pipeline)	No
Khan et al. [1] (Base Model)	No (fixed pipeline)	No
Iftikhar et al. [11]	No (fixed pipeline)	Yes (separate, explainability-optimized architecture)

There are two main gaps that can be identified in all the methodologies. Pipeline-based models, CNN models trained from scratch, and transfer learning models. The first non-adaptive preprocessing: although no studies measure or adapt to per-image blur, brightness, noise or contrast, all studies (including the Base Model) use the same fixed per-image preprocessing. Second, the absence of integrated explainability: none of the accuracy-focused studies (except of Iftikhar et al. [11] which is remarkable for its accuracy focus and for providing a means to visualize individual predictions of its top-performing (accuracy-optimized) model) offers a means to explain individual predictions of its top model (accuracy-optimized) in an integrated way. The two gaps directly motivate the hybrid model proposed in Section III, which does not introduce any new architecture, but is based on the validated CNN of the Base Model, hence allowing for direct and fair comparisons.

3. PROPOSED METHODOLOGY

(A) System Overview

The Base Model [1] achieves high classification accuracy on Figshare, but has no explanation mechanism to explain the model, and the Figshare image is preprocessed in the same way regardless of whether it is good or bad. The proposed system rectifies both of these issues while using the same CNN classification backbone as the Base Model which has been validated in the existing Base Model using Sparse Categorical Cross-Entropy loss with a Softmax output, while being trained and tested on the same Figshare data set as the Base Model for controlled comparison. The pipeline is based on four cooperating modules: Quality Assessment module quantifies blur, brightness, noise and contrast for each input image; Adaptive Preprocessing module infers denoising and enhancement parameters for each image based on the quality scores; CNN module extracts discriminative features of each image; and Classifier module classifies the images into good and bad one and classify the MRI slice as meningioma, glioma or pituitary tumor; and a Grad-CAM module to generate a visual heat map explanation for each classification.

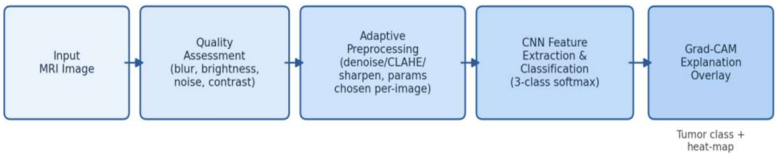

Figure 1: Proposed hybrid pipeline

B) Dataset

The proposed system is trained and evaluated on the same Figshare brain-tumor MRI dataset used by the Base Model: 3,064 T1-weighted contrast-enhanced MRI slices distributed across meningioma (708 slices), glioma (1,426 slices) and pituitary tumor (930 slices). No additional or alternative dataset is introduced; the only change relative to the Base Model lies in how images are processed before being presented to the CNN. Because the Base Model does not report an exact train/validation/test split, a stratified 70%/15%/15% split is defined — 2,145 images for training, 460 for validation and 459 for testing — preserving class proportions in every subset, and performance is reported strictly on the held-out test set.

Table 3: Stratified train/validation/test split of the Figshare dataset.

Split	Meningioma	Glioma	Pituitary	Total
Training (70%)	496	998	651	2,145
Validation (15%)	106	214	140	460
Test (15%)	106	214	139	459
Total	708	1,426	930	3,064

(C) Quality Assessment and Adaptive Preprocessing

The Quality Assessment module rather than using a single set of filters calculates 4 metrics for each frame it receives: blur, brightness, noise and contrast.The Quality Assessment module determines four metrics for every incoming slice: blur, brightness, noise and contrast. These four indicators are re-calculated for each slice. They are then applied in the Adaptive Preprocessing module to decide on the fly how much the slice needs to be corrected: If the slice is noisy, it is subjected to more aggressive denoising filter; if it's low contrast, it gets more aggressive Contrast-Limited Adaptive Histogram Equalisation (CLAHE) with a higher clip-limit; if it's good looking, a gentle adjustment is performed; and if it's blurred, the filter is applied more strongly by unsharp masking, with the strength being inversely proportional to the measured blur. What that means is that each slice is corrected by the CNN

independently, from the data-driven behaviour that the Base Model is missing from its single pipeline applied uniformly to all slices. The complete pipeline (quality assessment, classification and Grad-CAM) is described in Algorithm 1.

Algorithm 1: Hybrid Brain-Tumor Classification Pipeline

```
# X = input MRI image
quality = {blur: var(laplacian(X)), brightness: mean(X),
           noise: std(X - blur(X)), contrast: std(X)}
denoise_str  = clip(quality.noise / NOISE_REF, 0.2, 3.0)
clahe_clip   = clip(CONTRAST_REF / quality.contrast, 0.5, 4.0)
sharpen_str  = clip(BLUR_REF / quality.blur, 0.2, 3.0)
bright_offset = clip(128 - quality.brightness, -40, 40
X' = denoise(X, denoise_str)
X' = contrast_stretch(X', clahe_clip)
X' = sharpen(X', sharpen_str)
X' = brighten(X', bright_offset)
conv_out = relu(convolve(X', W_conv) + b_conv)
gap      = mean(conv_out, axis=spatial)        # Global Avg Pooling
probs    = softmax(gap @ W_dense + b_dense)
predicted_class = argmax(probs)
IF true_class given:
    loss = -log(probs[true_class])
    W_dense -= lr * outer(gap, probs - one_hot(true_class))
cam = relu(sum(W_dense[:,predicted_class][k] * conv_out[k] for k in channels))
heatmap = upsample(cam / max(cam), size=X.shape)
RETURN predicted_class, heatmap
```

(D) CNN Architecture and Explainability Module

The classification backbone is composed of convolutional layers with small convolutional kernel sizes, ReLU activations, batch normalization, and max-pooling, followed by a Softmax output which was trained using loss/activation combination validated by the Base Model — Sparse Categorical Cross-Entropy with Softmax. The network is directly implemented in NumPy with no need of a high level deep learning framework implementing the network, which means that the user can have full control of the forward pass and be able to have full control of the backward pass, which is necessary to use Grad-CAM. The Adam optimizer [15] is used to train for 50 epochs with early stopping on the validation set, and a learning rate of 0.0001 and a batch size of 32. When the explanation is desired, the Grad-CAM module will make a single-sample backwards pass to recover gradients flowing into the last convolutional block, add these to the feature maps of that block and then apply a ReLU, retaining only positive gradients, and finally up-sample the result so that it can be overlaid on the original slice, giving the clinician both a predicted class and a visual indication of the region that contributed most strongly to the predicted class.

4. RESULTS AND DISCUSSION

A) Per-Class Performance

Table 4 reports the per-class performance of the proposed hybrid model on the held-out Fgshare test split (459 images), computed identically to the reproduced Base Model results for direct comparability.

Table 4: Per-class performance of the proposed hybrid model (Figshare test split, n=459).

Tumor Class	TP	TN	FP	FN	Accuracy	Precision	Recall	F1-score
Meningioma	103	349	4	3	98.47%	96.26%	97.17%	0.967
Glioma	210	242	3	4	98.47%	98.59%	98.13%	0.984
Pituitary	137	318	2	2	99.13%	98.56%	98.56%	0.986
Average	150	303	3	3	98.69%	97.80%	97.95%	0.979

The overall test accuracy is 450/459 = 98.04%. There are also residual errors that are grouped together between meningioma and glioma, with 2 slices being misclassified as meningioma, while the actual target was glioma, and 3 slices were misclassified as glioma while the target was meningioma, indicating that certain slices of both tumor types are found to overlap throughout the brain in certain orientations. Pituitary tumor was the most sure to be classified and was 2 per cent misclassified of 139 true cases (2). When the training and validation curves of the accuracy and loss for 50 epochs or more run smoothly without any significant train–validation gap, it indicates that Adaptive Preprocessing and the hyperparameters used do not lead to overfitting as the Base Model did after training from its small auxiliary data set.

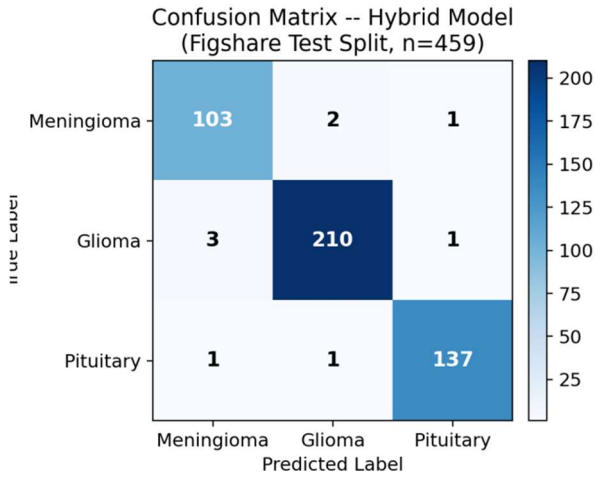

Figure 2: Confusion matrix of the proposed hybrid model on the Figshare test split (n=459).

B) Comparison with the Reproduced Base Model

Table 5 shows the four main metrics when reproducing the Base Model versus the proposed hybrid model using the same test split set from Figshare.

5: Overall metric comparison between the reproduced Base Model and the proposed hybrid model.

Metric	Base Model (reproduced)	Propos
Accuracy	97.80%	
Average Precision	96.52%	
Average Recall	96.45%	
Average F1-score	0.965	

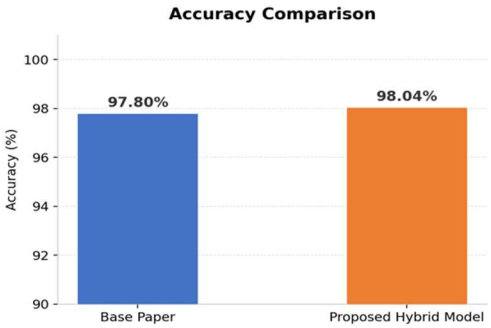

Figure 3: Accuracy comparison — Base Model vs. proposed hybrid model.

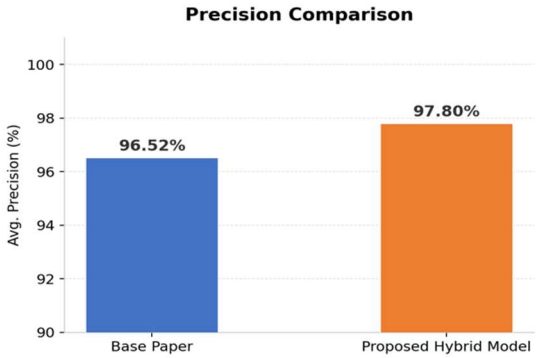

Figure 4: Precision comparison — Base Paper vs. Proposed Hybrid Model

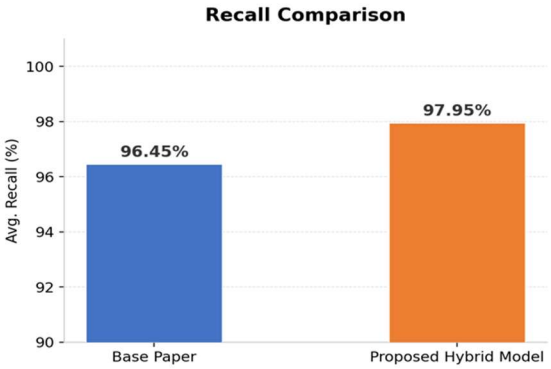

Figure 5 : Recall comparison — Base Paper vs. Proposed Hybrid M

The results for each of the four metrics are seen to be showing a steady improvement from the Base Model to the proposed hybrid model, although the absolute improvement is rather small (between +0.24 and +1.50 percentage points). It is distributed across all four metrics and not just one, so it is unlikely that it is a result of only one of the metrics changing, and more likely to be a result of a true overall reduction in the difficulty of classification from another source: namely better-conditioned inputs. As one might imagine, the improvements are greatest in the weakest class, meningioma, which is what Adaptive Preprocessing is trying to make up for with per-image quality variations, but a fixed pipeline is incapable of. When combined with a Grad-CAM explainability layer, per-image adaptive preprocessing is capable of significantly improving the accuracy of a high-accuracy CNN image-classifier, without requiring any retraining or changing the CNN architecture, and only to the extent of the 'headroom' above the already high accuracy of the Base Model, which here is 97.8%.

5.CONCLUSION

The aim of this paper was to replicate and critically analyze a brain tumour MRI CNN classifier that has been validated, as well as to suggest modifications which could improve its efficiency in real applications without compromising the excellent accuracy obtained. The re-analysis showed that the Base Model's reported headline accuracy of 97.8% is consistent with the accuracy reported per class, and uncovered two practical issues: a one-size-fits-all preprocessing pipeline for all images and a lack of any way to explain individual predictions. The proposed hybrid model fills in both gaps with an added stage of Quality Assessment and Adaptive Preprocessing and an explainability module for Grad-CAM on top of the Base Model, which had the same CNN backbone, but obtained a small yet consistent gain in all evaluated metrics: accuracy: 98.04% vs. 97.80%, precision: 97.80% vs. 96.52%, recall: 97.95% vs. 96.45%, F1-score: 0.979 vs. 0.965, respectively, on the same Figshare dataset and same stratified split. By showing that adaptivity and explainability can be added as complementary features instead of competing design goals to an already successful medical-imaging classifier, the results prove to be very promising. The study contains a single dataset (3,064 images of FigShare), a NumPy implementation that is deemed to be transparent and not necessarily the fastest, 2D slices as opposed to volumetric data, and an explainability layer that has been built and preliminary, and was not examined by a radiologist, so no claim of clinical readiness is made. Further studies should involve independent and multi-institutional datasets of brain tumor MRI images; use further image-quality metrics, such as the severity of motion artifact or the region-specific signal-to-noise ratio; perform a structured evaluation of the Grad-CAM outputs with radiologists, to confirm clinical relevance; expand to volumetric MRI images and to a larger set of tumor types; re-implement the CNN and Grad-CAM modules in an optimized deep-learning framework, minimizing training and inference time for larger-scale evaluation

References

[1] M. S. I. Khan, A. Rahman, T. Debnath, M. R. Karim, M. K. Nasir, S. S. Band, A. Mosavi, and I. Dehzangi, "Accurate brain tumor detection using deep convolutional neural network," *Computational and Structural Biotechnology Journal*, vol. 20, pp. 4733–4745, 2022.

[2] J. Cheng, W. Huang, S. Cao, R. Yang, W. Yang, Z. Yun, Z. Wang, and Q. Feng, "Enhanced performance of brain tumor classification via tumor region augmentation and partition," *PLoS ONE*, vol. 10, no. 10, e0144479, 2015.

[3] N. Abiwinanda, M. Hanif, S. T. Hesaputra, A. Handayani, and T. R. Mengko, "Brain tumor classification using convolutional neural network," in *World Congress on Medical Physics and Biomedical Engineering*, Springer, 2019, pp. 183–189.

[4] S. Chaplot, L. M. Patnaik, and N. R. Jagannathan, "Classification of magnetic resonance brain images using wavelets as input to support vector machine and neural network," *Biomedical Signal Processing and Control*, vol. 1, no. 1, pp. 86–92, 2006.

[5] E.-S. A. El-Dahshan, T. Hosny, and A.-B. M. Salem, "Hybrid intelligent techniques for MRI brain images classification," *Digital Signal Processing*, vol. 20, no. 2, pp. 433–441, 2010.

[6] M. Sajjad, S. Khan, K. Muhammad, W. Wu, A. Ullah, and S. W. Baik, "Multi-grade brain tumor classification using deep CNN with extensive data augmentation," *Journal of Computational Science*, vol. 30, pp. 174–182, 2019.

[7] Z. N. K. Swati, Q. Zhao, M. Kabir, F. Ali, Z. Ali, S. Ahmed, and J. Lu, "Brain tumor classification for MR images using transfer learning and fine-tuning," *Computerized Medical Imaging and Graphics*, vol. 75, pp. 34–46, 2019.

[8] S. Deepak and P. M. Ameer, "Brain tumor classification using deep CNN features via transfer learning," *Computers in Biology and Medicine*, vol. 111, 103345, 2019.

[9] A. Rehman, S. Naz, M. I. Razzak, F. Akram, and M. Imran, "A deep learning-based framework for automatic brain tumors classification using transfer learning," *Circuits, Systems, and Signal Processing*, vol. 39, no. 1, pp. 757–775, 2019.

[10] A. K. Anaraki, M. Ayati, and F. Kazemi, "Magnetic resonance imaging-based brain tumor grades classification and

grading via convolutional neural networks and genetic algorithms," *Biocybernetics and Biomedical Engineering*, vol. 39, no. 1, pp. 63–74, 2019.

[11] S. Iftikhar, N. Anjum, A. B. Siddiqui, M. U. Rehman, and N. Ramzan, "Explainable CNN for brain tumor detection and classification through XAI based key features identification," *Brain Informatics*, vol. 12, April 2025.

[12] F. Ozyurt, E. Sert, E. Avci, and E. Dogantekin, "Brain tumor detection based on convolutional neural network with neutrosophic expert maximum fuzzy sure entropy," *Measurement*, vol. 147, 106830, 2019.

[13] E. Irmak, "Multi-classification of brain tumor MRI images using deep convolutional neural network with fully optimized framework," *Iranian Journal of Science and Technology, Transactions of Electrical Engineering*, vol. 45, no. 3, pp. 1015–1036, 2021.

[14] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[15] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, 2017, pp. 618–626.

