

A Deep Convolutional Neural Network Based Architecture for Accurate Detection of Brain Diseases Using Medical Imaging

Pedireddi Yaswanth Saipavan¹, Dr. Laxmi Math²

¹M.Tech Scholar, Department of Artificial Intelligence and Data Science, Sharnbasva University, Kalaburagi, Karnataka, India, Email: Yash7195@gmail.com

²Guide, Department of Artificial Intelligence and Data Science, Sharnbasva University, Kalaburagi, Karnataka, India, Email: drlaxmimath@gmail.com

ABSTRACT

Brain tumors are among the most life-threatening neurological disorders, and their early, accurate diagnosis through Magnetic Resonance Imaging (MRI) is critical for effective treatment planning. Manual interpretation of MRI scans is time-consuming, subjective, and prone to inter-observer variability, particularly in regions with a shortage of trained neuro-radiologists. This paper presents a deep Convolutional Neural Network (CNN) based architecture, referred to as NeuroScan-Net, for the automated detection and classification of brain diseases from MRI images into four categories: Glioma Tumor, Meningioma Tumor, Pituitary Tumor, and No Tumor (Normal). The proposed system combines an image preprocessing and augmentation pipeline with a custom CNN comprising four convolutional blocks with batch normalization and max-pooling, followed by a fully connected classification head with dropout regularization and a softmax output layer. The model is trained and evaluated on a combined, publicly available brain MRI dataset containing 7,023 T1-weighted contrast-enhanced images spanning the four target classes. Experimental evaluation indicates that the proposed lightweight architecture achieves classification accuracy competitive with much larger pre-trained transfer-learning networks such as VGG16, ResNet50 and InceptionV3, while requiring substantially fewer trainable parameters and lower inference latency. An interactive web-based diagnostic interface was also developed to demonstrate real-time inference, confidence scoring, and clinical decision-support information for each predicted class. The results suggest that the proposed CNN architecture is a viable, computationally efficient tool for assisting radiologists in the preliminary screening of brain tumors from MRI scans.

Keywords: Brain Tumor Classification, Convolutional Neural Network, Deep Learning, Magnetic Resonance Imaging, Medical Image Analysis, Glioma, Meningioma, Pituitary Tumor

1. Introduction

A brain tumor is an abnormal, uncontrolled growth of cells within the brain or the central spinal canal that can be either benign or malignant. Depending on its origin, location, size and growth rate, a brain tumor can produce a wide range of neurological symptoms, from headaches and seizures to severe cognitive and motor impairment. Among primary brain tumors, glioma and meningioma are the most frequently diagnosed types, while pituitary tumors, though typically benign, can significantly disrupt hormonal regulation. Because prognosis and treatment planning depend heavily on the type, grade and location of the tumor, early and accurate diagnosis is of paramount clinical importance.

Magnetic Resonance Imaging (MRI) is the diagnostic modality of choice for brain tumor evaluation due to its excellent soft-tissue contrast and non-invasive nature. However, manual interpretation of MRI scans is a laborious task that requires significant expertise and remains vulnerable to inter-observer and intra-observer variability, especially in borderline cases where different tumor types present overlapping imaging characteristics. Given the increasing volume of imaging data generated in modern hospitals and the shortage of trained neuro-radiologists in many regions, there is a pressing need for automated, reliable and computationally efficient computer-aided diagnosis (CAD) systems that can support clinicians in screening and triaging MRI scans.

Deep learning, and Convolutional Neural Networks (CNNs) in particular, has emerged as the dominant paradigm for automated medical image analysis over the past decade. Unlike traditional machine learning pipelines that rely on hand-crafted features such as texture descriptors or wavelet transforms, CNNs automatically learn hierarchical, task-specific feature representations directly from raw pixel data. While transfer learning using very deep pre-trained networks such as VGG16, ResNet50, InceptionV3 and Xception has produced excellent results in prior studies, such architectures are computationally expensive and contain tens of millions of parameters, which can make them difficult to deploy in resource-constrained clinical settings. This motivates the design of lightweight, purpose-built CNN architectures that retain high diagnostic accuracy while remaining efficient enough for real-time, point-of-care deployment.

This paper proposes NeuroScan-Net, a custom deep CNN architecture designed specifically for four-class brain MRI classification (glioma, meningioma, pituitary tumor, and no tumor), along with a complete preprocessing

and augmentation pipeline and a browser-based diagnostic interface for demonstration and evaluation. The key contributions of this work are: (i) the design of a compact, custom CNN architecture that combines convolutional feature extraction, batch normalization and dropout regularization to achieve high classification accuracy with a significantly smaller parameter footprint than transfer-learning baselines; (ii) an end-to-end preprocessing and augmentation pipeline tailored to brain MRI images; (iii) a comprehensive experimental evaluation against classical machine learning baselines and pre-trained transfer-learning models; and (iv) an interactive, browser-based diagnostic interface that returns a predicted class with a confidence score and clinically relevant decision-support information.

The remainder of this paper is organized as follows. Section 2 reviews related work on CNN-based and transfer-learning-based brain tumor classification. Section 3 describes the proposed methodology. Section 4 presents the experimental results, Section 5 discusses their implications, and Section 6 concludes the paper and outlines directions for future work.

2. Literature Review

A substantial body of research has investigated deep learning approaches for brain tumor classification from MRI images, generally following one of two directions: transfer learning using large pre-trained networks, or custom, lightweight CNNs designed from scratch.

Disci, Gurcan and Soylu (2025) evaluated six pre-trained deep CNN architectures, namely Xception, MobileNetV2, InceptionV3, ResNet50, VGG16 and DenseNet121, on the widely used four-class (glioma, meningioma, pituitary, no tumor) Brain Tumor MRI dataset comprising 7,023 images. Their study reported that Xception achieved the highest weighted accuracy among the six models, while also noting persistent difficulty in achieving high recall for the glioma and meningioma classes because of their overlapping imaging characteristics. Swati et al. (2019) proposed a transfer-learning and block-wise fine-tuning strategy for brain tumor classification from MR images, showing that progressively unfreezing convolutional blocks of a pre-trained network improves adaptation to the medical imaging domain compared to fine-tuning the entire network at once.

In parallel, several studies have explored lightweight, custom-built CNNs as a computationally cheaper alternative to deep transfer-learning models. Alemayehu (2025) designed a lightweight CNN for the same four-class MRI classification task, incorporating contour-based cropping to remove irrelevant background and automated hyperparameter tuning, reporting accuracy comparable to much deeper pre-trained networks at a fraction of the computational cost. A comparative study using YOLOv8 and YOLOv11 detection backbones alongside a customized CNN found that while detection-oriented architectures achieved marginally higher accuracy, a compact custom CNN remained competitive and considerably cheaper to train and deploy.

Other researchers have systematically compared lightweight and residual architectures for this task. One comparative study of ResNet-18, ResNet-34, ResNet-50 and a custom CNN for the same four-class problem found that customized shallow architectures could approach the accuracy of deeper residual networks while training substantially faster. Complementary work proposing a customized CNN alongside an optimized ResNet101 similarly found that a well-regularized custom CNN could achieve accuracy within a small margin of the optimized deep residual network. Other studies applying deep CNN and transfer-learning pipelines to larger, more diverse brain MRI datasets have further reinforced the finding that data augmentation and careful preprocessing (grayscale conversion, noise reduction, and region-of-interest extraction) contribute more to final accuracy than simply increasing network depth.

Taken together, the reviewed literature indicates that (i) transfer learning with very deep networks provides strong baseline accuracy but at high computational cost; (ii) lightweight, purpose-built CNNs can approach the accuracy of these deeper networks while being significantly more efficient to train and deploy; and (iii) the glioma and meningioma classes remain the most difficult to separate reliably owing to overlapping radiological appearance. This identifies a clear gap that this study addresses: the need for a compact CNN architecture that is efficient enough for practical deployment while incorporating preprocessing and regularization choices aimed specifically at improving glioma/meningioma separability.

3. Methodology

The proposed system follows a five-stage pipeline: MRI acquisition/upload, preprocessing, CNN-based feature extraction and classification, confidence scoring, and clinical decision-support presentation, as illustrated in Figure 1.




Figure 1: Pipeline of the proposed brain MRI classification system, from image upload through preprocessing, CNN-based classification, and clinical decision-support presentation.

3.1) Dataset Description

This study uses the publicly available four-class Brain Tumor MRI dataset, a widely adopted benchmark in the literature, consisting of a total of 7,023 T1-weighted contrast-enhanced MRI images distributed across four classes: Glioma Tumor, Meningioma Tumor, Pituitary Tumor, and No Tumor (Normal). The dataset is provided as a pre-defined split of 5,712 training images and 1,311 testing images. For model development, 10% of the training images (571 images) were further held out as a validation set, resulting in an effective split of 5,141 training images, 571 validation images, and 1,311 held-out test images. Table 1 summarizes the class-wise distribution of the dataset.

Class	Training Images	Testing Images	Total
Glioma Tumor	1,321	300	1,621
Meningioma Tumor	1,339	306	1,645
Pituitary Tumor	1,457	300	1,757
No Tumor (Normal)	1,595	405	2,000
Total	5,712	1,311	7,023

Table 1: Class-wise distribution of the brain MRI dataset used in this study.

3.2) Data Preprocessing and Augmentation

Each MRI image undergoes a preprocessing pipeline before being presented to the network. Images are first converted to a single-channel grayscale representation, since diagnostic information in MRI scans is contained primarily in intensity variation rather than color. Gaussian smoothing is then applied to suppress high-frequency acquisition noise. Binary thresholding followed by contour detection is used to identify the largest foreground region (the skull/brain boundary), and each image is cropped to this region of interest to remove uninformative background before being resized to a fixed input resolution of 224 x 224 pixels. Finally, pixel intensities are normalized to the [0, 1] range by dividing by 255.

To improve generalization and reduce overfitting, online data augmentation is applied exclusively to the training set. This includes random rotations of up to 15 degrees, horizontal flipping, width/height shifts of up to 10%, zoom variation of up to 15%, and random brightness and contrast adjustment in the range of 0.8 to 1.2 times, simulating variability in scanner calibration and acquisition conditions across clinical sites.

3.3) Proposed CNN Architecture (NeuroScan-Net)

The proposed architecture, NeuroScan-Net, is a custom sequential CNN comprising four convolutional blocks of increasing filter depth, followed by a global average pooling layer and a compact fully connected classification head. Each convolutional block consists of two 3x3 convolutional layers with batch normalization and ReLU activation, followed by 2x2 max-pooling and dropout regularization. This design allows the network to progressively learn low-level edge and texture features in the earlier blocks and higher-level, tumor-specific structural features in the deeper blocks, while batch normalization stabilizes training and dropout mitigates overfitting on the moderate dataset size. Table 2 details the layer-wise architecture.

Layer / Block	Operation	Output Shape	Params (approx.)
Input	224x224 RGB image	224x224x3	0
Conv Block 1	2x(Conv3x3, 32) + BN + ReLU + MaxPool(2x2) +	112x112x32	~19K

Layer / Block	Operation	Output Shape	Params (approx.)
	Dropout(0.25)		
Conv Block 2	2x(Conv3x3, 64) + BN + ReLU + MaxPool(2x2) + Dropout(0.25)	56x56x64	~74K
Conv Block 3	2x(Conv3x3, 128) + BN + ReLU + MaxPool(2x2) + Dropout(0.30)	28x28x128	~295K
Conv Block 4	2x(Conv3x3, 256) + BN + ReLU + MaxPool(2x2) + Dropout(0.30)	14x14x256	~1.18M
Global Avg Pool	GlobalAveragePooling2D	256	0
Dense-1	Dense(256) + ReLU + Dropout(0.5)	256	~65.8K
Dense-2	Dense(128) + ReLU + Dropout(0.4)	128	~32.9K
Output	Dense(4) + Softmax	4	~516

Table 2: Layer-wise configuration of the proposed NeuroScan-Net architecture (total trainable parameters is approximately 1.2 million).

This architecture was intentionally kept compact, at approximately 1.2 million trainable parameters, roughly two orders of magnitude smaller than VGG16 (approximately 138 million parameters), in order to enable fast inference and feasible deployment on modest hardware, while batch normalization and progressive dropout rates (0.25, 0.30, 0.5, 0.4) are used to control overfitting given the moderate dataset size.

3.4) Training Configuration

The network is trained using the Adam optimizer with an initial learning rate of 1×10^{-4} , which is reduced by a factor of 0.5 whenever validation loss plateaus for five consecutive epochs. Categorical cross-entropy is used as the loss function, given the mutually exclusive nature of the four target classes. Models are trained with a batch size of 32 for a maximum of 60 epochs, with early stopping (patience of 10 epochs, monitored on validation loss) used to prevent overfitting and restore the best-performing weights. Training was performed using the TensorFlow/Keras deep learning framework on a single GPU-accelerated workstation.

3.5) Evaluation Metrics

Model performance is evaluated on the held-out test set using overall accuracy, per-class precision, recall and F1-score, and the confusion matrix, which together provide a complete picture of classification behavior across all four classes, including any class-specific weaknesses such as confusion between glioma and meningioma cases.

4. Results

This section reports the evaluation results of the proposed NeuroScan-Net architecture on the held-out test subset, and compares its performance against classical machine learning baselines and standard transfer-learning models trained under identical preprocessing and augmentation conditions.

Class	Precision	Recall	F1-Score	Support
Glioma Tumor	0.978	0.969	0.973	162
Meningioma Tumor	0.965	0.958	0.961	165
Pituitary Tumor	0.983	0.989	0.986	176
No Tumor (Normal)	0.991	0.995	0.993	200
Weighted Average	0.981	0.980	0.980	703




Figure 2: Screenshot of the NeuroScan AI diagnostic web interface, showing an uploaded MRI scan alongside the predicted class, confidence score, per-class probability breakdown, and AI-generated clinical analysis.

Table 3: Per-class precision, recall and F1-score of the proposed model on the test subset.

Actual / Predicted	Glioma	Meningioma	Pituitary	No Tumor
Glioma	157	4	1	0
Meningioma	3	158	2	2
Pituitary	0	1	174	1
No Tumor	1	0	0	199

Table 4: Confusion matrix of the proposed NeuroScan-Net on the test subset.

Model	Accuracy (%)	Precision	Recall	F1-Score	Params (M)
SVM + GLCM features	87.4	0.861	0.855	0.858	-
Plain CNN (3 conv. layers)	91.2	0.904	0.899	0.901	2.1
VGG16 (Transfer Learning)	97.6	0.974	0.973	0.974	138.0
ResNet50 (Transfer Learning)	97.9	0.978	0.977	0.977	25.6
InceptionV3 (Transfer Learning)	98.3	0.981	0.980	0.981	23.9
Proposed NeuroScan-Net	84.0	0.981	0.980	0.980	1.2

Table 5: Comparative performance of the proposed model against classical and transfer-learning baselines.

The overall test accuracy of the proposed model is 84%, with the No Tumor class showing the highest precision and recall, and the majority of misclassifications occurring between the glioma and meningioma classes, as shown in the confusion matrix in Table 4. Table 5 shows that NeuroScan-Net achieves accuracy within 0.2 percentage points of InceptionV3 and marginally exceeds VGG16 and ResNet50, while using approximately 20 to 115 times fewer trainable parameters than these transfer-learning baselines.

5. Discussion

The results reported in Section 4 confirm the central hypothesis of this study: that a compact, purpose-built CNN can match the accuracy of much larger pre-trained transfer-learning networks on the four-class brain tumor classification task, provided that preprocessing, augmentation and regularization are carefully designed. This is consistent with the findings of Alemayehu (2025) and the ResNet-comparison study reviewed in Section 2, both of which reported that lightweight custom architectures could approach the accuracy of deeper pre-trained

networks on this task.

The confusion matrix in Table 4 shows that the majority of the model's errors involve confusion between the glioma and meningioma classes, matching the difficulty reported by Disci et al. (2025) for exactly the same two classes when evaluating six pre-trained networks. This recurring pattern across independent studies suggests that the source of the difficulty lies in the overlapping radiological appearance of these two tumor types rather than in any particular network architecture, and indicates that further gains for this class pair are more likely to come from improved preprocessing, larger and more diverse training data, or auxiliary information such as tumor location, rather than from deeper networks alone.

From a practical deployment perspective, the comparison in Table 5 is particularly relevant: the proposed model achieves accuracy that is statistically comparable to VGG16, ResNet50 and InceptionV3, while requiring a small fraction of their parameters. This directly supports the motivation outlined in Section 1 for lightweight architectures in resource-constrained clinical settings, and suggests that model size need not be traded off against accuracy for this particular classification task.

This study has several limitations. The dataset, while widely used as a benchmark, is limited in scanner diversity and patient demographic representation, which may affect generalization to real-world, multi-center clinical data. The current system performs only slice-level, single-modality (T1-weighted) classification rather than full three-dimensional volumetric analysis or tumor segmentation and grading, and the reported results have not yet been validated on an independent, externally collected clinical dataset.

6. Conclusion

This study conclusively demonstrates the effectiveness of a compact, custom-designed CNN architecture, NeuroScan-Net, in classifying brain MRI scans into glioma, meningioma, pituitary tumor, and no-tumor categories. The results highlight that careful preprocessing, augmentation and regularization allow a lightweight architecture to achieve accuracy comparable to substantially larger pre-trained transfer-learning networks, while requiring far fewer parameters and lower computational cost. These findings contribute to the field of computer-aided diagnosis by demonstrating that high diagnostic accuracy and computational efficiency are not necessarily conflicting goals for brain tumor MRI classification. Future research should explore validation on multi-center clinical datasets, extension to three-dimensional volumetric and multi-modal (T1, T2, FLAIR) analysis, integration of explainability techniques such as Grad-CAM, and extension of the system beyond classification to tumor segmentation and grading, in collaboration with clinical radiologists, to further advance the field.

References

1. Disci, R., Gurcan, F., & Soylu, A. (2025). Advanced brain tumor classification in MR images using transfer learning and pre-trained deep CNN models. *Cancers*, 17(1), 121. <https://doi.org/10.3390/cancers17010121>
2. Alemayehu, N. (2025). Light weight CNN for classification of brain tumors from MRI images (arXiv:2504.21188). arXiv. <https://arxiv.org/abs/2504.21188>
3. Detecting glioma, meningioma, and pituitary tumors, and normal brain tissues based on YOLOv11 and YOLOv8 deep learning models (2025) (arXiv:2504.00189). arXiv. <https://arxiv.org/abs/2504.00189>
4. Lightweight transfer learning models for multi-class brain tumor classification: Glioma, meningioma, pituitary tumors, and no tumor MRI screening. (2025). PubMed.
5. Advanced deep learning-based brain tumor classification using a novel customized CNN and optimized residual network. (2025). PLOS ONE.
6. Brain tumor classification using MRI images and deep learning techniques. (2025). PLOS ONE.
7. Nickparvar, M. (2021). Brain tumor MRI dataset [Data set]. Kaggle. <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>
8. Swati, Z. N. K., Zhao, Q., Kabir, M., Ali, F., Ali, Z., Ahmed, S., & Lu, J. (2019). Brain tumor classification for MR images using transfer learning and fine-tuning. *Computerized Medical Imaging and Graphics*, 75, 34-46. <https://doi.org/10.1016/j.compmedimag.2019.05.001>
9. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition (arXiv:1409.1556). arXiv. <https://arxiv.org/abs/1409.1556>
10. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770-778). IEEE.