International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4834-4840
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0636

Enhanced Brain Tumor Detection and Classification Using Transfer Learning and Explainable AI with Clinical Report Generator

Kavya R[1], *Mala M*[2]
[1]*PG Scholar, Department of Computer Science, AMC Engineering College, Bangalore, 560083, India.*
[2]*Asst. Professor, Department of Computer Science, AMC Engineering College, Bangalore, 560083, India.*
Emails: *kavya.r.273.3@gmail.com*[1], *mala.naik@amceducation.in*[2]

Abstract
Correctly sorting brain tumors captured through Magnetic Resonance Imaging (MRI) plays a vital role in timely diagnosis and sound clinical decision-making, since a delayed or wrong call can seriously harm patient outcomes. This work introduces AMC-NeuroDx, an upgraded deep-learning pipeline that pits a Custom Convolutional Neural Network (CNN) built from the ground up against a ResNet50 transfer-learning model on a four-way brain tumor classification task (Glioma, Meningioma, Pituitary Tumor, and No Tumor), drawing on 7,223 MRI scans taken from the Kaggle Brain Tumor MRI dataset. ResNet50 follows a two-stage routine—training with a frozen base before full fine-tuning—and reaches 93.16% validation accuracy after only 20 epochs, whereas the Custom CNN needs 50 epochs to hit 90.70%. The framework also embeds Grad-CAM so that spatial heatmaps can show which brain regions drove each prediction, meeting the clinical demand for transparency. On top of this, a Streamlit interface lets users upload MRI scans in real time and automatically produces patient-specific PDF clinical reports holding the diagnosis, confidence scores, Grad-CAM overlays, and recommended next steps. The results obtained here show that transfer learning yields higher accuracy and quicker convergence than building a CNN from the ground up, and together with the built-in explainability and automated report writing, this makes AMC-NeuroDx a practical candidate for real clinical settings.

Keywords: *Convolutional Neural Network; ResNet50; Transfer Learning; Grad-CAM; Explainable AI; Deep Learning.*

1. Introduction
Among neurological illnesses, brain tumors rank among the most dangerous, driving substantial sickness and death across every age group around the world. The International Agency for Research on Cancer (IARC) reports that over a million fresh brain tumor cases surface annually, and the death toll tied to the disease keeps climbing [1]. MRI remains the preferred non-invasive route for diagnosing brain tumors thanks to its strong soft-tissue contrast and ability to image in multiple planes [2]. Yet radiologists examining MRI scans by hand face a process that is slow, subjective, and inconsistent from one observer to another, which often delays diagnosis, particularly where healthcare resources are limited [3]. Deep Learning (DL), and CNNs specifically, have reshaped medical image analysis by making automated, high-precision feature extraction and classification possible [4]. There is a notable split between CNNs trained entirely from scratch and Transfer Learning (TL) methods, which reuse models such as ResNet50 or VGG16 already trained on ImageNet's pool of over a million images [5]. Bringing in transfer learning cuts down both the data needed and the compute cost, frequently while pushing performance even higher [6]. One drawback common to models judged purely on accuracy is that they behave as black boxes, which chips away at clinicians' trust. Explainable Artificial Intelligence (XAI) methods—Gradient-weighted Class Activation Mapping (Grad-CAM) in particular—address this gap by generating visual heatmaps that point to the exact image regions driving a given prediction [7]. These visual cues let radiologists double-check AI-assisted calls and support regulatory compliance for medical AI systems. This paper introduces AMC–NeuroDx, an integrated brain

tumor classification and reporting pipeline, contributing the following:

- A side-by-side evaluation of a Custom CNN against ResNet50 TL on the Kaggle Brain Tumor MRI dataset.
- A two-stage training routine for ResNet50 covering feature extraction followed by fine-tuning.
- Use of Grad-CAM for explainable AI, highlighting tumor-related regions within MRI scans.
- Building of a Streamlit web app that automatically produces a personalized PDF clinical report for every patient.

2. Literature Review

Classifying brain tumors from MRI scans through deep learning has gained wide traction in recent research, and the approach has repeatedly shown how much it can change the outlook for brain tumor diagnosis from imaging data. Sen et al. [8] put forward a lightweight CNN that hit 95.28% accuracy on a sizeable Bangladeshi clinical dataset of 10,234 MRI images, beating out ResNet50 (83.31%), VGG16 (89.23%), and MobileNetV2 (91.90%). Their work relied on Grad-CAM to explain predictions and pointed out why lightweight architectures matter for deployment in low-resource settings. Kumar et al. [9] ran a standard CNN on a Kaggle binary brain tumor dataset (3,264 images), landing 82% test accuracy with 0.85 precision and 0.84 recall for tumor-positive cases, and their findings stressed how much augmentation and preprocessing help generalization. Kowasalya et al. [10] built a hybrid CNN-RCNN pipeline, using Discrete Wavelet Transform (DWT) for denoising and GLCM/LBP for feature extraction ahead of classification, reaching 95.20% accuracy and an 88.12% F1-score—evidence that pairing handcrafted features with deep learning pays off. Bhuvanika and Thamilselvi [11] introduced NeuroFusion, benchmarking ResNet50, VGG16, DenseNet121, and a Custom CNN, where DenseNet121 topped the table at 98.7% and ResNet50 followed at 98.2%, reinforcing that deep, densely connected pretrained networks suit medical image classification well. Sakib et al. [12] combined five transfer-learning models—VGG16, InceptionV3, ResNet50, DenseNet121, and MobileNetV3 Small—into an ensemble on a modest 253-image brain MRI dataset, reaching 94.05% combined accuracy and showing that ensembling lifts results beyond what single models achieve. Gupta and Sandhu [13] applied the Xception transfer-learning architecture to the Kaggle four-class brain tumor dataset (7,023 images) and reached 99.7% testing accuracy, adding further weight to the case for transfer learning in multiclass brain tumor work. Hu et al. [14] looked across several deep learning architectures—VGG, AlexNet, GoogleNet, ResNet—together with YOLO-based object detection for brain tumor diagnosis, finding that plain CNNs without preprocessing enhancements only reached around 76% validation accuracy, which underlines how much proper preprocessing pipelines matter. Tang and Teoh [15] tested ResNet18 on a binary brain tumor task and got 88.33% accuracy, ahead of CapsNet (84.75%) and GoogLeNet (86.67%), pointing to the benefit residual connections bring to gradient flow in medical imaging. Looking across this body of work, three gaps stand out repeatedly: (1) explainability and clinical trust receive limited attention, (2) end-to-end deployment paired with patient-specific report generation is largely absent, and (3) there is not enough head-to-head comparison between custom architectures and transfer learning on standardized datasets. AMC-NeuroDx is built specifically to close these gaps.

3. Dataset Description

This study draws on the publicly available Brain-Tumor MRI Dataset put together by Masoud Nickparvar on Kaggle, which combines and relabels images pulled from three separate datasets for consistency.

3.1. Dataset Overview

The dataset spans 7,223 brain MRI images across four classes—Glioma, Meningioma, No Tumor, and Pituitary Tumor—and comes pre-split into 5,600 images for training and 1,623 for testing, roughly a 77:23 split.

3.2.Class Description

Glioma — A malignant growth that develops in the glial cells of the brain or spinal cord, regarded as one of the most aggressive of all brain tumor variants.

This class holds 1,800 images (1,400 for training, 400 for testing). Meningioma — Usually a non-cancerous growth that forms in the meninges, the protective layers wrapping the brain and spinal cord. This class has 1,823 images (1,400 for training, 423 for testing). No Tumor — Scans from healthy brains showing no tumor at all, totaling 1,800 images (1,400 for training, 400 for testing). Pituitary Tumor — A growth that forms within the pituitary gland, frequently triggering hormone-related imbalances. This class includes 1,800 images (1,400 for training, 400 for testing). As shown in Table 1 Dataset distribution.

Table 1 Dataset distribution

Class	Training	Testing	Total
Glioma	1,400	400	1,800
Meningioma	1,400	423	1,823
No Tumor	1,400	400	1,800
Pituitary	1,400	400	1,800
Total	5,600	1,623	7,223

4. Proposed Methodology

4.1. System Architecture Overview

AMC-NeuroDx is organized as a modular pipeline made up of six stages: (1) data acquisition and preprocessing, (2) data augmentation, (3) model architecture design, (4) the two-phase training strategy, (5) Grad-CAM explainability, and (6) Streamlit-based deployment paired with PDF clinical report generation.

4.2. Preprocessing Pipeline

Ahead of training, each MRI image was resized to 224 × 224 pixels so it would match what ResNet50 expects as input. For the Custom CNN, pixel intensities were normalized by dividing every value by 255, which maps intensity values onto a [0,1] scale. ResNet50, on the other hand, used its own dedicated preprocess_input function from Keras, which carries out channel-wise mean subtraction matched to ImageNet's statistics. This difference matters a great deal—applying the generic 1/255 rescaling to ResNet50 instead causes a sharp accuracy drop because of the mismatch in input distribution.

4.3. Data Augmentation

To make the model more robust and cut down on the chance of overfitting, the training images were augmented on fly using Keras's ImageDataGenerator.

- **Rotation range:** ±15 degrees
- **Width and height shift:** 10% of image dimension
- **Zoom range:** 20%
- **Horizontal flip:** enabled
- **Brightness range:** [0.8, 1.2]

No augmentation was applied to the testing set to preserve evaluation integrity.

4.4. Custom CNN Architecture

The Custom CNN was built around four convolutional blocks chained in sequence, ending in fully connected layers. Each block runs a Conv2D layer with a growing filter count (32→64→128→256), same padding, a ReLU activation, Batch Normalization to keep training stable, and 2×2 MaxPooling2D to shrink the spatial dimensions. Once flattened, the resulting feature maps feed into a 512-unit Dense layer with ReLU, then a Dropout layer (rate=0.3) for regularization, and finally a softmax layer producing the four-class probabilities. In total, the network carries roughly 19.27 million trainable parameters.

4.5. ResNet50 Transfer Learning Architecture

The backbone chosen was ResNet50, pretrained on ImageNet's 1.28-million-image, 1,000-class dataset. Its residual connections help get around the vanishing-gradient problem, making it possible to train a network 50 layers deep without trouble. In place of the original classification head, a custom head was attached, made up of GlobalAveragePooling2D (to collapse the spatial dimensions), a 256-unit Dense layer with ReLU, a Dropout layer (rate=0.5), and a Softmax layer for the four output classes. Training proceeded in two phases:

4.5.1. Phase 1 — Feature Extraction:

In Phase 1, the ResNet50 base was kept frozen (trainable = False), training only the newly attached classification head for 20 epochs with the Adam optimizer at a learning rate of 0.001. This step quickly tunes the output layers toward the brain tumor domain.

4.5.2. Phase 2 — Fine-Tuning:

In Phase 2, the full network—base and head

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4834-4840
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0636

together—was unfrozen and retrained for 45 epochs at a much smaller learning rate (0.00001), gently adjusting the pretrained weights toward the specific features found in MRI scans.

4.6.Training Configuration

The Custom CNN and ResNet50 were both trained with the Adam optimizer and categorical cross-entropy loss. Three callbacks supported training: EarlyStopping (patience between 5 and 10, watching val_loss), ModelCheckpoint (keeping only the best-performing weights), and ReduceLROnPlateau (cutting the learning rate by half if val_loss stalls for 3–5 epochs). Both models used a batch size of 32.

4.7.Grad-CAM Explainability

To make model predictions easier to interpret, Grad-CAM was set up using TensorFlow's GradientTape. Gradients for the target class were taken with respect to the activations of the last convolutional layer (conv5_block3_out for ResNet50), then averaged globally to gauge how much each feature mattered, and merged with the activation maps to build a localization heatmap. This heatmap was scaled back up to the original 224 × 224 resolution and overlaid on the MRI scan with a JET colormap, highlighting the regions that drove the tumor classification. The end result is that clinicians can verify whether a prediction is grounded in genuinely pathological regions of the scan rather than incidental artifacts.

4.8.Streamlit Web Application and Clinical Report Generation

The web front-end was built in Streamlit and served via an ngrok tunnel from Google Colab. Users enter patient details (name, age, gender) and upload an MRI scan; once submitted, the app runs inference through both the CNN and ResNet50, builds the Grad-CAM visualizations, shows confidence scores, and offers medical recommendations tied to the predicted diagnosis. It then auto-generates a personalized PDF clinical report via the fpdf2 library, covering patient demographics, a Report ID, the original MRI scan, and the Grad-CAM overlay, each model's diagnosis, the final call, a confidence percentage, and a clinical recommendation—closing the loop between AI-driven detection and something a clinician can actually use.

5. Experimental Results

5.1.Model Performance

Table 2 Model performance comparison

Metric	Custom CNN	ResNet50
Val. Accuracy	90.70%	93.16%
Train Accuracy	95.61%	97.16%
Epochs	50	20
F1-Score	91%	93%

ResNet50 beat the Custom CNN by 2.46 percentage points in validation accuracy while needing 56% fewer training epochs, underscoring how efficient transfer learning can be as shown in Table 2 Model performance comparison.

5.2.Accuracy Curves

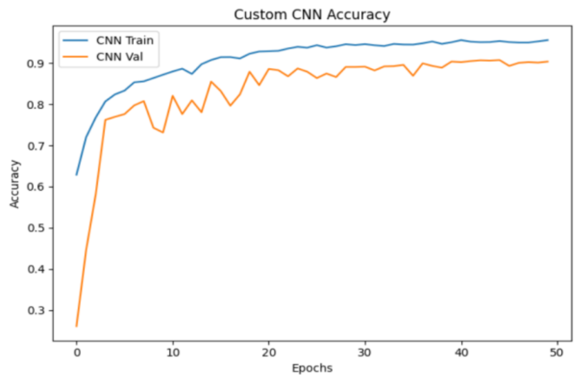

Figure 1 Training and Validation Accuracy: CNN

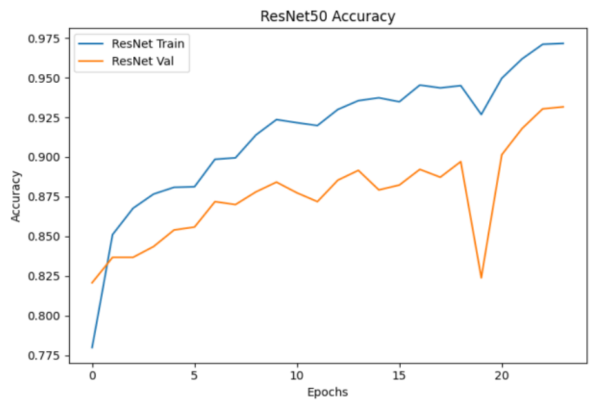

Figure 2 Training and Validation Accuracy Across Phase 1 And Phase 2: Resnet50

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4834-4840
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0636

In its first epoch, the Custom CNN started at a 62.88% training accuracy and climbed steadily to 95.61% by Epoch 50, with validation accuracy settling around 90.70%. ResNet50's Phase 1 ended at 89.71% validation accuracy after 20 epochs, and Phase 2 fine-tuning lifted that to 93.16% in just 5 more epochs—a clear sign of how efficient transfer learning is as shown in Figure 1 Training and Validation Accuracy: Cnn, Figure 2 Training and Validation Accuracy Across Phase 1 And Phase 2: Resnet50.

5.3.Grad-CAM visualization

Grad-CAM output shows the model attending to anatomically sensible regions in each case: peripheral margins for Glioma, the central parenchyma for Meningioma, the inferior region for Pituitary tumors, and a broadly distributed pattern for No Tumor scans—all consistent with what would be clinically expected.As shown in Figure 5 ResNet50 Confusion Matrix (89% on test set; 93.16% val_accuracy), Figure 4 Custom CNN Confusion Matrix (91% overall accuracy), Table 3 CNN class-wise report, Figure 3 Grad-CAM Heatmaps For Glioma, Meningioma, No Tumor, And Pituitary (Left: Original MRI; Center: Heatmap; Right: Overlay), Figure 6 CNN vs ResNet50 Validation Accuracy Comparison.

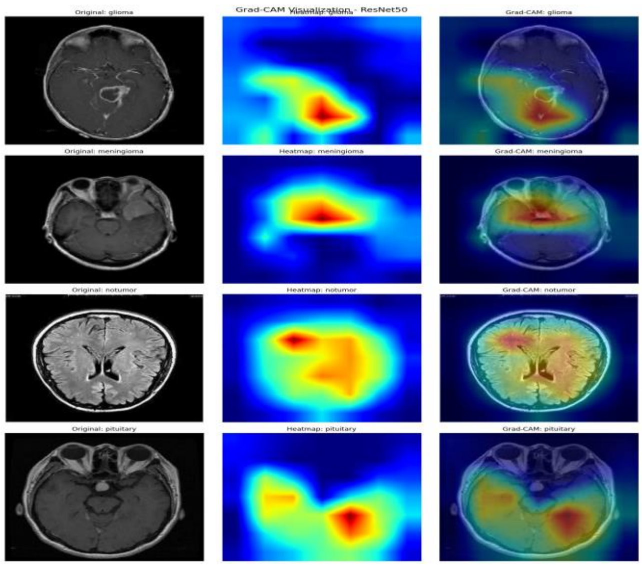

Figure 3 Grad-CAM Heatmaps for Glioma, Meningioma, No Tumor, And Pituitary (Left: Original MRI; Center: Heatmap; Right: Overlay)

5.4.Class-wise Performance — Custom CNN

Table 3 CNN class-wise report

Class	Precision	Recall	F1-Score	Support
Glioma	0.97	0.75	0.85	400
Meningioma	0.90	0.89	0.89	423
No Tumor	0.88	1.00	0.93	400
Pituitary	0.90	0.99	0.95	400
Overall	0.91	0.91	0.91	1623

5.5. Confusion Matrices

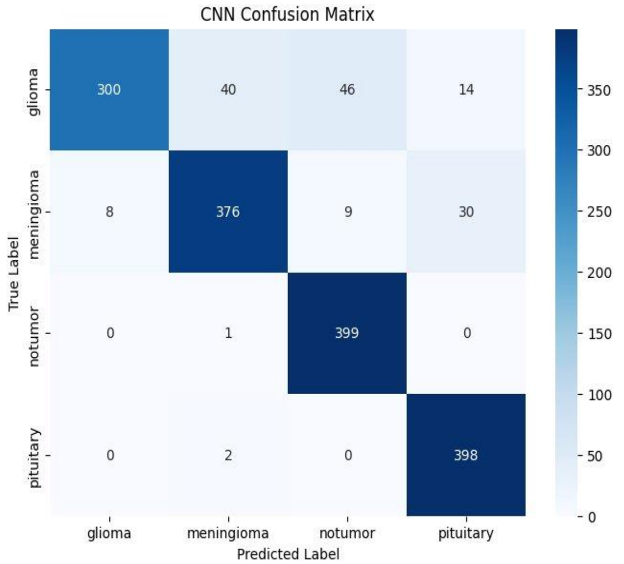

Figure 4 Custom CNN Confusion Matrix (91% overall accuracy)

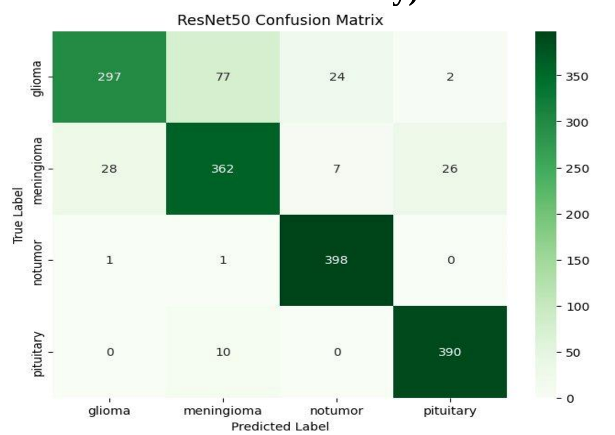

Figure 5 ResNet50 Confusion Matrix (89% on test set; 93.16% val_accuracy)

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 04 Issue: 07 July 2026
Page No: 4834-4840
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2026.0636

5.6.Model Comparison Chart

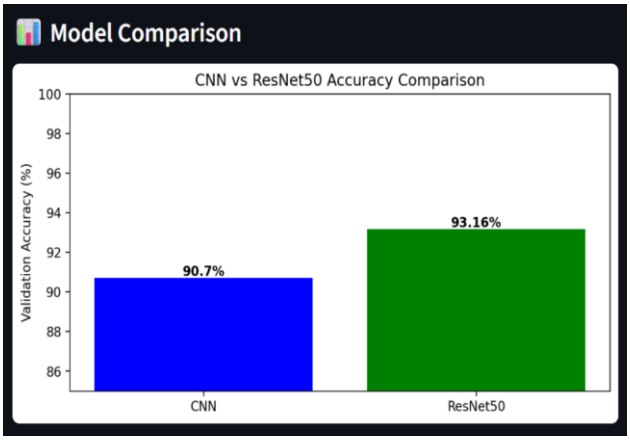

Figure 6 CNN vs ResNet50 Validation Accuracy Comparison

6. Discussion

These results back up the idea that ResNet50-based transfer learning beats the Custom CNN on validation accuracy (93.16% against 90.70%) while needing far fewer training epochs (20 versus 50). The pretrained ImageNet weights give ResNet50 a head start with rich low-level features—edges, textures, shapes—that carry over well to MRI analysis, speeding up convergence and improving generalization. One important technical takeaway was just how much model-specific preprocessing matters for transfer learning: swapping in generic 1/255 rescaling for ResNet50 instead of the required preprocess_input function dropped accuracy from above 93% down to roughly 60%, matching what prior work has reported [8]. The Grad-CAM analysis confirms that both models are looking at tumor-affected regions when making their classification, with the Glioma prediction anchored on peripheral tumor margins, Meningioma on the central brain parenchyma, and Pituitary tumors on the inferior region near the sella turcica. No Tumor cases show activation spread more broadly, reflecting the model scanning multiple areas to rule out pathology. This match between the model's attention and known tumor locations backs up the diagnostic reliability of both models [7]. What makes AMC-NeuroDx clinically meaningful is its end-to-end deployment—unlike most published work, which reports accuracy figures without any deployment path, AMC-NeuroDx delivers a live prediction through a web interface together with automated clinical report generation, putting it within reach of radiologists in day-to-day practice. The PDF report covers everything needed for clinical documentation: patient demographics, predictions from both models with confidence scores, Grad-CAM evidence, and standardized recommendations. The one weak spot was Glioma's recall (0.75 for the CNN), which traces back to how visually similar it is to Meningioma—a difficulty that even seasoned clinicians in neuro-oncology run into [9]. Larger datasets and more advanced architectures going forward should help close this gap.

Conclusion and Future Scope

This paper has laid out AMC-NeuroDx, a complete brain tumor MRI classification system that brings together a Custom CNN, ResNet50 transfer learning, Grad-CAM explainability, and automated clinical report generation. ResNet50 reached 93.16% validation accuracy in just 20 training epochs, ahead of the Custom CNN's 90.70% after 50 epochs, making clear how much transfer learning helps when annotated medical images are limited. The two-phase training approach makes good use of pretrained ImageNet representations while still adapting to the specifics of MRI data. Grad-CAM visualization backs up that the model's predictions are anchored in clinically meaningful tumor regions, which builds confidence in the diagnosis. Pairing the Streamlit interface with automated PDF clinical reports takes the project beyond a typical model-development exercise and toward something usable in real clinical settings. Looking ahead, future work could pursue:

- Integration of more advanced architectures such as EfficientNet, DenseNet121, and Vision Transformers (ViT) for further accuracy improvement.
- Extension to multi-modal imaging by incorporating CT scan data alongside MRI for improved diagnostic sensitivity
- Application of federated learning techniques to enable training across
- Clinical validation with radiologist feedback in real hospital environments to assess real-world utility and regulatory compliance.

- Mobile deployment for on-device inference in resource-limited healthcare settings.

References

[1]. P. McKinney, "Brain tumours: incidence, survival, and aetiology," Journal of Neurology, Neurosurgery & Psychiatry, vol. 75, no. suppl 2, pp. ii12–ii17, 2004.

[2]. G. Katti, S. A. Ara, and A. Shireen, "Magnetic resonance imaging (MRI) – a review," International Journal of Dental Clinics, vol. 3, no. 1, pp. 65–70, 2011.

[3]. T. A. Soomro et al., "Image processing techniques for MRI brain tumor identification using machine learning: A review," IEEE Rev. Biomed. Eng., vol. 15, pp. 156–170, 2022.

[4]. L. O. Chua, "CNN: A vision of complexity," International Journal of Bifurcation and Chaos, vol. 7, no. 10, pp. 2219–2425, 1997.

[5]. B. Koonce, "ResNet 50," in Convolutional Neural Networks with Swift for TensorFlow: Image Recognition and Dataset Categorization, Springer, 2021, pp. 63–72.

[6]. M. Romero, Y. Interian, T. Solberg, and G. Valdes, "Targeted transfer learning to improve performance in small medical physics datasets," arXiv preprint arXiv:1912.05498, 2019.

[7]. 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 Proc. IEEE ICCV, 2017, pp. 618–626.

[8]. T. Sen, G. M. A. Rahaman, and S. Saha, "A lightweight CNN for brain tumor classification in resource-constrained settings," in Proc. IEEE DICTA, 2025, DOI: 10.1109/DICTA68720.2025.11302430.

[9]. K. S. Kumar, Bhoomika, and G. Verma, "Brain tumor detection using CNN: A deep learning approach for MRI classification," in Proc. 6th International Conference for Emerging Technology (INCET), 2025, DOI: 10.1109/INCET64471.2025.11140022.

[10]. R. Kowasalya, A. Rajeshwari, S. Arthi, S. Gethsiyal, and S. Vinothini, "Brain tumor detection using CNN algorithm," in Proc. 4th International Conference on Innovative Mechanisms for Industry Applications (ICIMIA), 2025, DOI: 10.1109/ICIMIA67127.2025.11200696.

[11]. G. Bhuvanika and J. Thamilselvi, "NeuroFusion: Advancing brain tumor diagnosis using CNN architectures," in Proc. 2nd Asia Pacific Conference on Innovation in Technology (APCIT), 2025, DOI: 10.1109/APCIT65661.2025.11411512.

[12]. K. M. S. Sakib, D. Sarker, M. I. Masroor, M. Hasan, and M. A. Al Arman, "Brain MRI image classification using transfer learning-based ensemble approach," in Proc. IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI), 2024, DOI: 10.1109/IATMSI60426.2024.10503556.

[13]. P. Gupta and R. Sandhu, "Implementation of brain tumor detection and classification model using CNN approach," in Proc. 4th International Conference on Sensors and Related Networks (SENNET), 2025, DOI: 10.1109/SENNET64220.2025.11136032.

[14]. H. Hu, W. Yao, X. Li, and Z. Yao, "Brain tumor diagnose applying CNN through MRI," in Proc. 2nd International Conference on Artificial Intelligence and Computer Engineering (ICAICE), 2021, DOI: 10.1109/ICAICE54393.2021.00090.

[15]. M. C. S. Tang and S. S. Teoh, "Brain tumor detection from MRI images based on ResNet18," in Proc. 6th International Conference on Information Systems and Computer Networks (ISCON), 2023, DOI: 10.1109/ISCON57294.2023.10112025.