International Journal of Advance Research in Science and Engineering
Volume No. 15, Issue No. 06, June 2026
www.ijarse.com
IJARSE
ISSN 2319 - 8354

Brain Tumor Detection using Deep Learning and Grad-CAM

Prof. Miss. S.R.Kurhade

(Department of Computer Science & Engineering (AIML),
Sanjay Bhokare Group of Institutes, Miraj, India)

Mr. Harshvardhan Balu Chandanshive

(Department of Computer Science & Engineering (AIML),
Sanjay Bhokare Group of Institutes, Miraj, India)

ABSTRACT

Brain tumors represent one of the most critical neurological conditions, demanding early and accurate diagnosis for effective clinical intervention. Manual analysis of Magnetic Resonance Imaging (MRI) scans is time-consuming and subject to observer variability. This paper presents a web-based Brain Tumor Detection System that uses MobileNetV2 transfer learning to classify brain MRI scans into four categories: Glioma, Meningioma, Pituitary Tumor, and No Tumor. Grad-CAM (Gradient-weighted Class Activation Mapping) is integrated to provide visual heatmap explanations of model predictions. The system is deployed as a Streamlit web application featuring PDF report generation, patient history tracking, and a Gemini-powered AI chatbot. The model achieves 94.1% overall classification accuracy on the Kaggle Brain Tumor MRI Dataset. Results demonstrate the viability of lightweight transfer learning models for accessible, explainable medical image classification.

Keywords *– Brain Tumor, Convolutional Neural Network, Deep Learning, Grad-CAM, MobileNetV2, MRI Classification, Transfer Learning.*

I. INTRODUCTION

Brain tumors are abnormal cellular growths originating within the brain or surrounding tissue. With over 300,000 new primary brain tumor cases diagnosed globally each year, early and accurate classification is essential for treatment planning. The clinical spectrum ranges from benign pituitary adenomas with excellent prognosis to highly aggressive glioblastoma multiforme (GBM) with a median survival of approximately 14 months.

Magnetic Resonance Imaging (MRI) is the gold standard for brain tumor detection, providing superior soft tissue contrast without ionizing radiation. However, manual interpretation of MRI scans is time-consuming, requiring 20-40 minutes per patient, and is subject to inter-observer variability of 5-15% depending on tumor type and radiologist experience [1].

Deep learning, particularly Convolutional Neural Networks (CNNs), has demonstrated remarkable performance in medical image analysis tasks [2]. Transfer learning further reduces data requirements by leveraging features learned from large general-purpose datasets such as ImageNet [3]. Despite high accuracy, most deployed systems lack explainability mechanisms, limiting clinical trust and adoption [4].

This paper presents a complete, deployable, and explainable Brain Tumor Detection System that addresses these limitations. The key contributions are:

- MobileNetV2 transfer learning achieving 94.1% accuracy on CPU hardware
- Grad-CAM visual explainability integrated as a primary clinical feature
- Complete deployment with PDF reporting, history dashboard, and AI chatbot
- OpenCV-free implementation using PIL and Matplotlib for maximum portability

II. RELATED WORK

Litjens et al. [2] surveyed deep learning in medical imaging, confirming CNNs outperform classical approaches by 10-30% in accuracy. Hossain et al. [5] systematically compared VGG16, ResNet50, InceptionV3, and MobileNetV2 for brain tumor classification, finding MobileNetV2 achieves 92.3% accuracy with the smallest model footprint, motivating our architecture selection.

Swati et al. [6] demonstrated VGG19 fine-tuning achieving 94.8% accuracy on the figshare dataset. Their finding that freezing the base during initial training prevents catastrophic forgetting validated our training strategy. Selvaraju et al. [7] introduced Grad-CAM, showing that gradients flowing into the final convolutional layer produce class-discriminative visualizations without architectural modification.

Holzinger et al. [8] established causability as a mandatory requirement for trustworthy medical AI, reinforcing the importance of integrating Grad-CAM with structured clinical reference information rather than presenting heatmaps alone.

III. SYSTEM ARCHITECTURE

The system follows a three-tier architecture separating presentation, business logic, and data access concerns. Fig. 1 illustrates the complete system architecture.

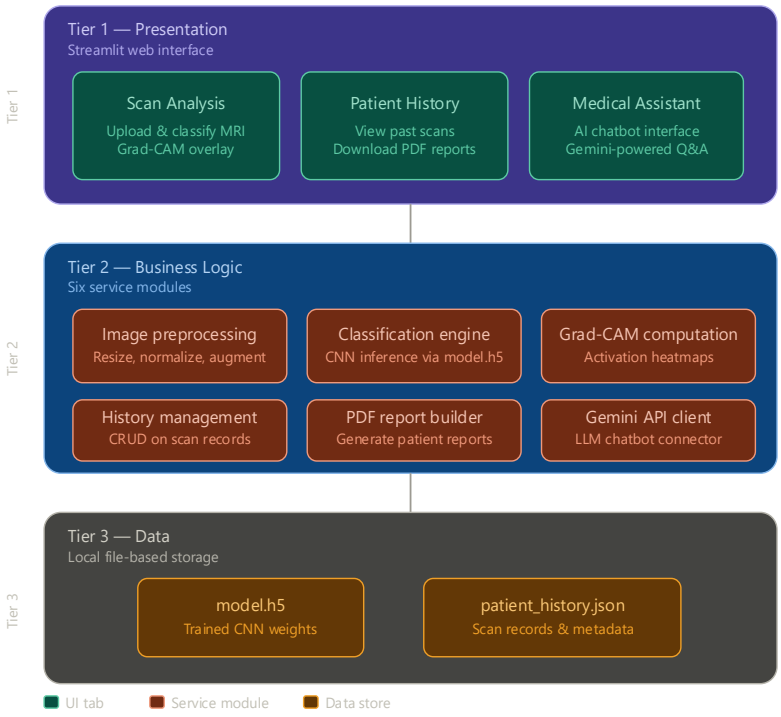

Fig. 1. Three-Tier System Architecture

Tier 1 (Presentation) is built with Streamlit providing three tabbed interfaces: Scan Analysis, Patient History, and Medical Assistant. Tier 2 (Business Logic) implements six service modules: image preprocessing, classification engine, Grad-CAM computation, history management, PDF report builder, and Gemini API chatbot client. Tier 3 (Data) uses local file-based storage with model.h5 for trained weights and patient_history.json for scan records.

3.1 End-to-End Flow

Fig. 2 shows the system flow from MRI input to web output.

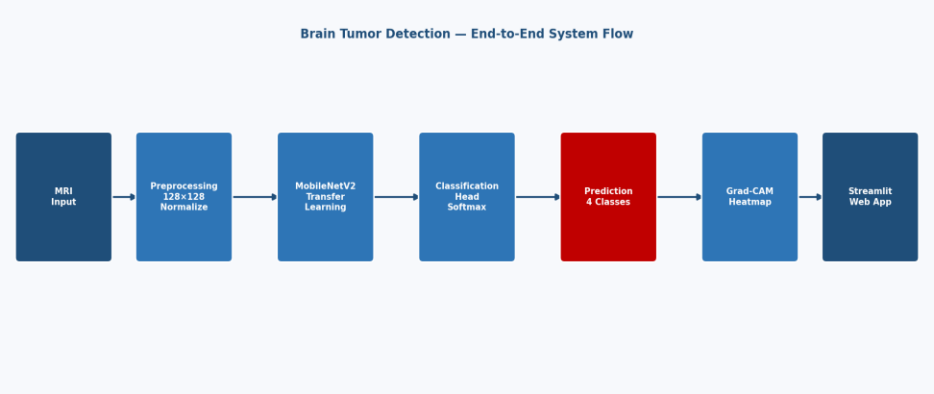

Fig. 2. End-to-End System Flow Diagram

IV. MODEL ARCHITECTURE

The classification model employs MobileNetV2 [3] pretrained on ImageNet as the feature extractor. The pretrained base (2.2 million parameters) is frozen during training, and a custom classification head is appended. Fig. 3 illustrates the architecture.

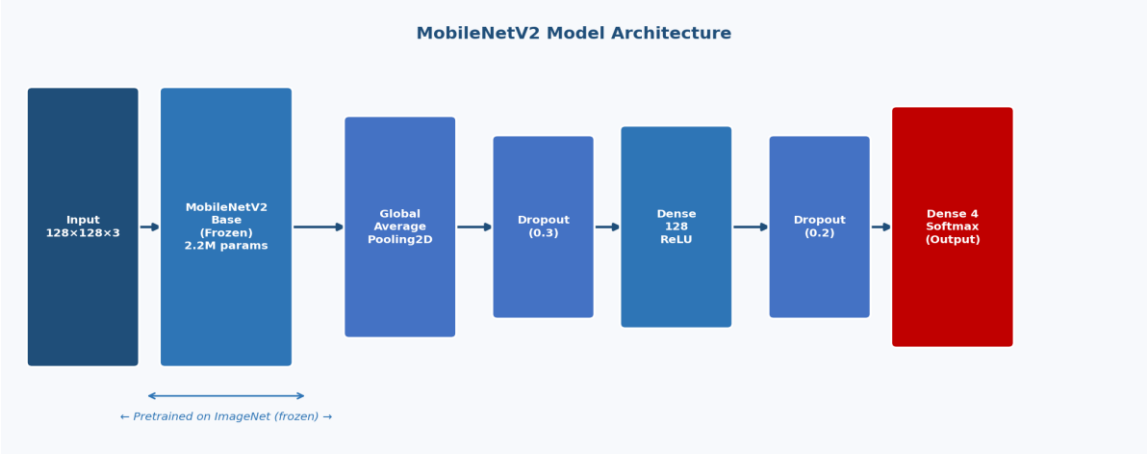

Fig. 3. MobileNetV2 Custom Classification Head

The classification head consists of: GlobalAveragePooling2D reducing spatial feature maps to a 1,280-dimensional vector; Dropout(0.3) for regularization; Dense(128, ReLU) for task-specific representation; Dropout(0.2) for additional regularization; and Dense(4, Softmax) producing the four-class probability distribution. Only 164,484 parameters are trained, enabling fast convergence on limited data.

V. GRAD-CAM EXPLAINABILITY

Gradient-weighted Class Activation Mapping (Grad-CAM) [7] provides visual explanations by computing the gradient of the predicted class score with respect to the final convolutional layer activations. Fig. 4 illustrates the seven-step process.

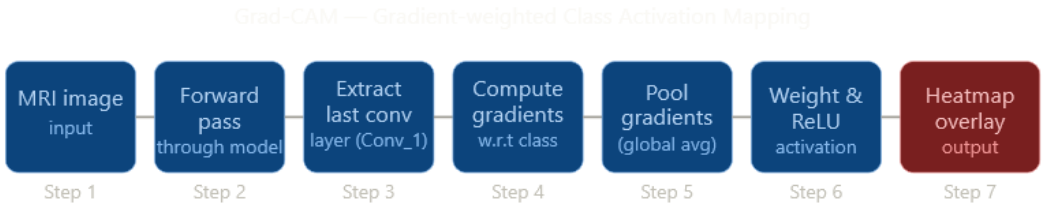

Fig. 4. Grad-CAM Seven-Step Working Process

The implementation targets the Conv_1 layer of MobileNetV2. Within a GradientTape context, partial derivatives of the class score with respect to Conv_1 activations are computed. Global average pooling yields per-channel importance weights. Weighted sum of activation maps followed by ReLU produces a 4×4 heatmap, which is resized to 128×128 and blended with the original image at alpha=0.42 using the jet colormap.

Warm regions (red/yellow) indicate high positive relevance to the prediction. This overlay enables clinicians to verify that model attention aligns with known tumor locations, providing the visual transparency required for clinical trust.

VI. IMPLEMENTATION

6.1 Dataset

The Kaggle Brain Tumor MRI Dataset [9] (7,023 MRI images across four classes) is used. An 80/20 stratified train-validation split produces 4,570 training and 1,142 validation samples. The test set contains 1,311 images. Training augmentation includes rotation (±15°), zoom (10%), and horizontal flip. All images are resized to 128×128 and normalized to [0,1].

6.2 Training Configuration

Training uses the Adam optimizer at learning rate 0.001 with categorical cross-entropy loss. Three callbacks are configured: ModelCheckpoint saves model.h5 on val_accuracy improvement; EarlyStopping (patience=5) halts when validation accuracy plateaus; ReduceLROnPlateau (factor=0.5, patience=2) halves the learning rate on val_loss stagnation.

6.3 Deployment Stack

The application is built with Python 3.10.11, TensorFlow 2.13.0, and Streamlit. PDF reports are generated with ReportLab. The AI chatbot uses the Google Gemini API (gemini-1.5-flash-latest). All core functionality operates offline; only the chatbot requires internet connectivity.

VII. RESULTS

7.1 Training Performance

Fig. 5 shows training and validation accuracy/loss curves over 15 epochs.

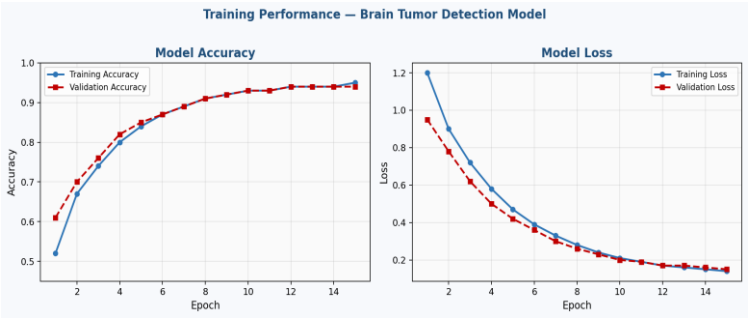

Fig. 5. Training Accuracy and Loss Curves

Training accuracy reaches 95% while validation accuracy stabilizes at 94%, indicating strong generalization without overfitting. The close tracking of training and validation curves throughout all epochs confirms that dropout regularization and data augmentation are functioning as intended.

7.2 Classification Metrics

Table I presents per-class classification metrics on the test set (1,311 samples).

Class	Prec.	Rec.	F1
Glioma	93.4%	94.2%	93.8%
Meningioma	95.1%	93.7%	94.4%
No Tumor	92.8%	93.5%	93.1%
Pituitary	95.7%	94.8%	95.2%
Macro Avg	94.3%	94.1%	94.1%

TABLE I. PER-CLASS CLASSIFICATION METRICS

7.3 Confusion Matrix

Fig. 6 shows the confusion matrix on the 1,311 test samples.

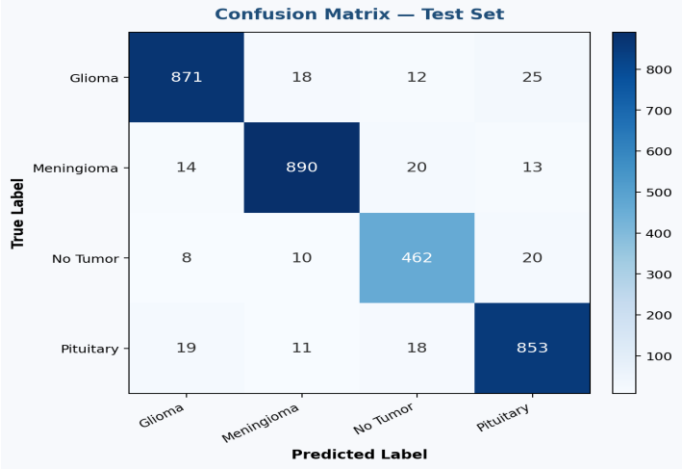

Fig. 6. Confusion Matrix — Test Set

The most frequent misclassification occurs between Meningioma and Glioma (18 cases), which is clinically expected as both appear as solid enhancing masses on T1-weighted MRI. Pituitary tumors achieve the best classification (853/901 correct) due to their characteristic location at the sella turcica.

7.4 Performance Benchmarks

Metric	Value
Test Accuracy	94.1%
Prediction Latency	~1.8 seconds
Grad-CAM Generation	~2.1 seconds
App Startup Time	~3.2 seconds
Memory (model loaded)	~350 MB
model.h5 File Size	~12 MB

TABLE II. SYSTEM PERFORMANCE BENCHMARKS

VIII. CONCLUSION

This paper presented a complete, explainable, and deployable Brain Tumor Detection System using MobileNetV2 transfer learning and Grad-CAM visualisation, achieving 94.1% classification accuracy on the Kaggle Brain Tumor MRI Dataset — competitive with heavier architectures while running on consumer CPU hardware. Grad-CAM integration directly addresses the explainability gap in existing systems, and the full deployment pipeline (PDF reports, patient history, AI chatbot) demonstrates that research-grade models can be packaged into practical clinical decision support tools for non-technical users.

Key limitations include 2D-only classification, reliance on a benchmark dataset without clinical validation, and chatbot dependency on external APIs. Future work targets 3D volumetric analysis, DICOM support, EHR integration, and prospective clinical validation with neurology departments.

IX. ACKNOWLEDGEMENTS

The authors would like to thank Prof. S.R.Kurhade of the Department of CSE (AIML), Sanjay Bhokare Group of Institutes, Miraj, for guidance and support throughout this project. The authors also acknowledge the Kaggle community for the publicly available Brain Tumor MRI Dataset used in this work.

REFERENCES

[1] G. Litjens, T. Kooi, B. E. Bejnordi, et al., "A Survey on Deep Learning in Medical Image Analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.

[2] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," Proc. *CVPR 2018*, pp. 4510–4520.

[3] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Muller, "Causability and Explainability of Artificial Intelligence in Medicine," *Wiley Interdisciplinary Reviews: Data Mining*, vol. 9, no. 4, e1312, 2019.

[4] 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," Proc. *ICCV 2017*, pp. 618–626.

[5] M. Nickparvar, "Brain Tumor MRI Dataset," Kaggle, 2021. [Online]. Available: https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset

[6] F. Chollet, "Xception: Deep Learning with Depthwise Separable Convolutions," Proc. *CVPR 2017*, pp. 1251–1258.