

TumorSense AI: An Explainable Deep Learning Framework for Brain Tumor Detection and Classification using MRI Images

Mrs. Murugalakshmi A , M.E., (co-author)

Assistant Professor, Department of Electronics and
Communication Engineering, Renganayagi Varatharaj College of Engineering, Sivakasi,

S. Sudha, V. Karpagam, M. Mahalakshmi

Department of Electronics and Communication Engineering
Renganayagi Varatharaj College of Engineering, Sivakasi – 626 128, Tamil Nadu, India

Abstract - Brain tumors represent one of the most life-threatening neurological disorders, where early and accurate detection is critical for improving patient outcomes. Manual interpretation of Magnetic Resonance Imaging (MRI) scans is time-consuming, subjective, and susceptible to inter-observer variability. This paper presents TumorSense AI, an intelligent, cloud-deployed framework for automated brain tumor detection and multi-class classification using deep learning. The proposed system employs a transfer-learning-based ResNet50 architecture fine-tuned on a curated dataset of approximately 15,000 brain MRI images categorized into four classes: glioma, meningioma, pituitary tumor, and no tumor. Advanced preprocessing strategies including image normalization, resizing to 224×224 pixels, and data augmentation are applied to improve model generalization. To address the critical limitation of black-box decision-making in deep learning models, the system integrates Explainable AI (XAI) via Gradient-weighted Class Activation Mapping (Grad-CAM), which generates visual heatmaps highlighting tumor-relevant regions in MRI scans. The trained model is deployed as a real-time web application using Streamlit and hosted on cloud platforms, enabling accessible, near-instantaneous predictions. The system achieves an overall test accuracy of approximately 88%, with high precision and recall for pituitary and no-tumor categories. The results validate that combining transfer learning, explainability, and cloud deployment yields a practical and reliable diagnostic support system for clinical and research applications.

Keywords - Brain Tumor Detection; Convolutional Neural Network; Transfer Learning; ResNet50; Explainable AI; Grad-CAM; Medical Image Classification; Deep Learning; MRI Analysis.

I. INTRODUCTION

Brain tumors arise from the abnormal and uncontrolled proliferation of cells within the brain or central nervous system, leading to severe neurological complications including seizures, cognitive impairment, and, in many cases, death. Globally, more than 300,000 new cases of brain and central nervous system tumors are diagnosed annually, with an estimated incidence rate of approximately 10.82 cases per 100,000 person-years [1]. The most clinically significant tumor categories include glioma, meningioma, and pituitary tumors, each differing in origin, aggressiveness, and treatment requirements.

Magnetic Resonance Imaging (MRI) remains the gold standard for non-invasive brain tumor diagnosis due to its superior soft-tissue contrast and high spatial resolution. However, the interpretation of MRI scans demands considerable radiological expertise and is inherently time-consuming, particularly given the rapidly increasing volume of medical imaging data in modern healthcare systems [2]. Traditional computer-aided diagnostic (CAD) systems, relying on handcrafted features and conventional machine learning algorithms, have demonstrated limited generalizability and classification performance in complex clinical scenarios.

The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has fundamentally transformed medical image analysis. These models autonomously learn hierarchical feature representations directly from imaging data, often matching or surpassing specialist-level diagnostic performance [3]. Nevertheless, a major barrier to clinical adoption remains the lack of interpretability: most deep learning models function as "black boxes," providing predictions without discernible reasoning, which undermines trust among medical professionals.

This paper introduces TumorSense AI, a comprehensive brain tumor detection and classification framework that addresses the dual challenges of accuracy and interpretability. The system integrates a fine-tuned ResNet50 architecture for multi-class MRI

classification with Explainable AI (XAI) via Gradient-weighted Class Activation Mapping (Grad-CAM), and deploys the solution as a real-time cloud-based web application. The key contributions of this work are: (i) a transfer learning pipeline achieving ~88% multi-class accuracy; (ii) integration of Grad-CAM for clinically meaningful visual explanations; and (iii) a fully deployable Streamlit-based interface for real-world accessibility.

II. LITERATURE REVIEW

Research in AI-based brain tumor detection has evolved substantially over the past decade. Early approaches employed traditional machine learning algorithms such as Support Vector Machines (SVM) and Random Forests on manually engineered features extracted from MRI images. While these methods provided modest automation, they lacked the representational capacity to capture the complex spatial patterns inherent to tumor pathology.

The advent of deep learning shifted the paradigm significantly. Pereira et al. [3] demonstrated that CNN-based architectures trained end-to-end on MRI data outperformed conventional methods in both segmentation and classification tasks. Subsequent studies established CNN variants—including VGGNet, Inception, and ResNet—as benchmark architectures for medical image recognition, with reported classification accuracies frequently exceeding 90% [4].

Transfer learning has emerged as a particularly effective strategy for medical imaging, where labeled data is scarce and domain-specific. Cheng et al. [5] showed that fine-tuning pre-trained networks on MRI datasets substantially improves classification performance compared to training from scratch. More recently, lightweight CNN architectures have been proposed to address computational constraints in resource-limited settings, maintaining competitive accuracy while reducing inference overhead [6].

Explainability has become a critical research focus in medical AI. Selvaraju et al. [7] introduced Grad-CAM, which computes gradient-weighted activations in the final convolutional layer to produce class-discriminative heatmaps. Studies applying Grad-CAM to brain tumor models demonstrate that visual attribution maps correlate well with clinically relevant tumor regions, improving radiologist trust [8]. Despite these advancements, the majority of published systems remain confined to controlled laboratory environments, lacking real-time deployment and user-facing interfaces—gaps that the proposed system directly addresses.

III. METHODOLOGY

A. Dataset Description

The proposed system is trained and evaluated on a publicly available Brain Tumor MRI Dataset comprising approximately 15,000 grayscale MRI images sourced from curated repositories including Kaggle. The dataset is organized into four clinically distinct categories: glioma (aggressive, originating from glial cells), meningioma (typically benign, arising from meningeal layers), pituitary tumor (affecting hormonal balance), and no tumor (healthy control). The dataset is partitioned into training (80%) and testing (20%) subsets with balanced class representation. Standard formats (JPEG/PNG) are employed, and images exhibit variability in tumor size, location, and intensity—conditions reflective of real clinical data.

B. Data Preprocessing and Augmentation

All MRI images are resized to a fixed spatial dimension of 224×224 pixels to conform to the input requirements of the ResNet50 architecture. Pixel intensities are normalized to the range [0, 1] to facilitate stable gradient computation during training. To mitigate overfitting and improve model generalization across unseen MRI conditions, dynamic data augmentation is applied during training, including random rotation (up to 20°), horizontal flipping, zoom (up to 15%), and brightness variation. These transformations simulate the natural variability present in clinical MRI acquisitions.

C. Model Architecture

The classification backbone of TumorSense AI is ResNet50, a 50-layer deep residual network originally pre-trained on the ImageNet large-scale visual recognition dataset. Residual connections in ResNet50 effectively mitigate the vanishing gradient problem that impairs training of very deep networks, enabling robust feature learning. The pre-trained convolutional layers are retained and leveraged for feature extraction, while the final fully connected layers are replaced and fine-tuned for the four-class brain tumor classification task. This transfer learning strategy reduces training data requirements and accelerates convergence.

The modified classification head consists of global average pooling, followed by dropout regularization (rate = 0.5) to prevent overfitting, and a dense softmax output layer with four neurons corresponding to the target classes. The model is compiled using the Adam optimizer with categorical cross-entropy as the loss function. Training employs early stopping with a patience of ten epochs, monitoring validation loss to prevent unnecessary training. Model checkpointing preserves the best-performing weight configuration based on validation accuracy.

D. Explainable AI via Grad-CAM

To overcome the opacity inherent in deep neural network predictions, TumorSense AI incorporates Gradient-weighted Class Activation Mapping (Grad-CAM) as its explainability mechanism. Grad-CAM computes the gradient of the score for the predicted class with respect to the feature maps of the final convolutional layer. These gradients are globally average-pooled to obtain weights indicating the importance of each feature map. A weighted combination of these maps is then passed through a ReLU activation to yield a coarse localization heatmap highlighting tumor-discriminative regions. This heatmap is resized and superimposed as a color overlay on the original MRI scan, providing clinicians with an intuitive visual explanation of the model's decision. The Grad-CAM output is generated in real time alongside each classification result within the deployed application.

E. System Deployment

The complete TumorSense AI pipeline is deployed as a cloud-based web application using the Streamlit framework, enabling users to upload brain MRI images and receive instantaneous classification predictions with confidence scores and Grad-CAM visualizations. The trained model, serialized in .keras format, is hosted externally on Hugging Face to overcome storage constraints and ensure dynamic loading. Streamlit's caching mechanisms ensure that the model is loaded once per session, minimizing inference latency. The system requires no specialized hardware or local installation, and is accessible via standard web browsers on any internet-connected device.

IV. RESULTS AND DISCUSSION

A. Classification Performance

The TumorSense AI model was evaluated on a held-out test set of 2,414 MRI images. The system achieved an overall accuracy of approximately 88%, demonstrating strong multi-class classification capability. Class-wise analysis revealed superior performance for the pituitary tumor and no-tumor categories, with precision and recall values consistently exceeding 0.90. These classes exhibit comparatively distinct morphological signatures in MRI, facilitating reliable discrimination. Moderate confusion was observed between glioma and meningioma categories, which is consistent with the inherent structural similarity of these tumor types—a challenge widely reported in clinical radiology literature. The macro-averaged F1-score across all four classes reflected a balanced trade-off between precision and recall, confirming the robustness of the trained model.

TABLE I
Classification Performance by Tumor Class

Class	Precision	Recall	F1-Score	Accuracy
Glioma	0.84	0.81	0.82	~85%
Meningioma	0.82	0.80	0.81	~83%
Pituitary	0.93	0.94	0.93	~94%
No Tumor	0.91	0.92	0.91	~92%
Overall (Weighted Avg.)	0.88	0.88	0.88	~88%

B. Comparison with Prior Work

Table II presents a comparative analysis of TumorSense AI against representative studies in the brain tumor classification literature. The proposed system achieves competitive accuracy using a simple yet effective transfer-learning architecture, without reliance on ensemble methods or custom segmentation modules.

TABLE II
Comparison with Related Works on Brain MRI Classification

Study / Model	Architecture	Accuracy (%)	XAI / Deployment
Pereira et al. [3]	Custom CNN	~88	No / No
Cheng et al. [5]	CNN + Augmentation	~91	No / No
Litjens et al. [4]	Survey (avg.)	>90	Limited / No
Proposed: TumorSense AI	ResNet50 + Transfer Learning	~88	Grad-CAM / Yes

C. Explainability Analysis

Grad-CAM visualizations generated by TumorSense AI consistently highlighted anatomically plausible regions corresponding to tumor sites in correctly classified images. For pituitary tumor cases, activation maps concentrated in the sella turcica region; for glioma cases, activation patterns aligned with perilesional zones. This spatial correspondence validates that the model has learned diagnostically meaningful features rather than spurious correlations. In borderline cases—particularly glioma vs. meningioma—diffuse or overlapping activation patterns were observed, reflecting the genuine ambiguity of these tumor types in MRI imaging. The side-by-side display of original MRI and Grad-CAM overlay within the Streamlit interface enables clinicians to rapidly assess the model's focus area and make informed decisions regarding further investigation.

D. System Performance

The deployed web application processes each uploaded MRI image and generates predictions within approximately 2–5 seconds under standard cloud computing conditions. Model caching ensures that repeated queries within a session incur negligible additional latency. Confidence scores exceeding 90% were observed in the majority of correctly classified images, while lower confidence values appropriately flagged ambiguous cases. The probability distribution panel, displaying likelihood scores across all four tumor classes, provides users with a quantitative basis for assessing prediction reliability beyond the top-1 output.

V. CONCLUSION

This paper has presented TumorSense AI, a deep learning framework for automated multi-class brain tumor classification from MRI images. The system combines transfer learning via a fine-tuned ResNet50 architecture with Gradient-weighted Class Activation Mapping (Grad-CAM) for explainability, and deploys the complete pipeline as a real-time, cloud-accessible web application. The system achieves approximately 88% overall accuracy on a four-class classification task, with strong performance particularly for pituitary tumor and no-tumor categories. The integration of Grad-CAM as an XAI mechanism directly addresses the clinical trust barrier associated with black-box deep learning models, making the system more suitable for practical diagnostic support.

Future work will focus on expanding dataset diversity across multiple clinical sites and imaging protocols to improve generalization, incorporating advanced architectures such as EfficientNet and Vision Transformers, extending the system to tumor segmentation for precise boundary delineation, and integrating multi-modal data streams including CT imaging and electronic health records. Clinical validation studies and integration with hospital Picture Archiving and Communication Systems (PACS) are also envisioned as pathways toward regulatory approval and real-world deployment.

REFERENCES

- [1] World Health Organization, "Global Cancer Observatory: Brain, Central Nervous System," International Agency for Research on Cancer, 2024.
- [2] G. Litjens et al., "A Survey on Deep Learning in Medical Image Analysis," *Medical Image Analysis*, vol. 42, pp. 60–88, 2017.
- [3] S. Pereira et al., "Brain Tumor Segmentation Using Convolutional Neural Networks in MRI Images," *IEEE Transactions on Medical Imaging*, vol. 35, no. 5, pp. 1240–1251, 2016.
- [4] G. Litjens et al., "Computer-Aided Detection of Prostate Cancer in MRI," *IEEE Transactions on Medical Imaging*, vol. 33, no. 5, pp. 1083–1092, 2014.
- [5] J. Cheng et al., "Enhanced Performance of Brain Tumor Classification via Tumor Region Augmentation and Partition," *PLOS ONE*, vol. 10, no. 10, e0140381, 2015.
- [6] A. Krizhevsky, I. Sutskever, and G. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," *Advances in Neural Information Processing Systems (NeurIPS)*, vol. 25, 2012.
- [7] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," *International Journal of Computer Vision*, vol. 128, pp. 336–359, 2020.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, USA, pp. 770–778, 2016.
- [9] T. Shorten and T. M. Khoshgoftaar, "A Survey on Image Data Augmentation for Deep Learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [10] "Streamlit: The Fastest Way to Build Data Apps in Python," *Streamlit Documentation*, 2024. [Online]. Available: <https://docs.streamlit.io>