

MRI Brain Tumor Detection using an Explainable Ensemble Deep Learning Framework

Vaibhav U. Mane

Department of Data Science, S.I.E.S. College, Mumbai, India

ABSTRACT: Reliable and swift identification of intracranial neoplasms ranks among the foremost determinants of survival in patients with neurological disease. Although Magnetic Resonance Imaging (MRI) yields superior anatomical resolution, manual scan interpretation by radiologists is inherently time-consuming and susceptible to inter-reader disagreement. This study presents a performance-weighted hybrid ensemble deep learning system that fuses VGG16 and ResNet50 convolutional neural networks for fully automated four-category MRI brain tumor discrimination, encompassing glioma, meningioma, pituitary adenoma, and scan-negative cases. Gradient-weighted Class Activation Mapping (Grad-CAM) is woven throughout the inference pipeline to furnish spatially precise, clinically actionable interpretability overlays. A live Streamlit browser interface provides end-user access. Benchmarked against the open-access Kaggle Brain Tumor MRI Dataset (7,023 scans), the fusion ensemble achieves 97.8% classification accuracy and 0.995 AUC-ROC, exceeding both the standalone VGG16 (95.6%) and standalone ResNet50 (96.4%). All four per-category F1 scores surpass 0.95.

KEYWORDS: brain tumor classification; MRI; VGG16; ResNet50; ensemble deep learning; transfer learning; Grad-CAM; Streamlit; medical AI; computer-aided diagnosis.

I. INTRODUCTION

Primary brain tumors constitute one of the gravest categories of neurological pathology, with the World Health Organization recording over 300,000 new diagnoses globally each year [1]. Three histological variants dominate clinical incidence—glioma, meningioma, and pituitary adenoma—each carrying distinct proliferative behavior, preferred anatomical site, and treatment protocol. Unambiguous subtype assignment therefore underpins individualized therapeutic decision-making and long-term prognosis.

MRI has risen to preeminence as the neuroimaging standard, delivering unmatched soft-tissue contrast entirely free from ionizing radiation. However, manual radiological reading is labor-intensive, demands deep subspecialty expertise, and is vulnerable to reader-level inconsistency—most acutely for low-grade or atypically presenting lesions. In high-volume imaging centers these bottlenecks propagate reporting delays with measurable downstream harm to patient outcomes.

Deep learning, and convolutional neural networks (CNNs) in particular, has catalyzed a paradigm shift in computational medical image analysis over the past decade. Transfer learning from large photographic repositories has alleviated the data scarcity endemic to clinical imaging, enabling high-fidelity feature extraction without reinitialization from random weights. Despite impressive accuracy gains, two obstacles continue to limit clinical uptake: achieving uniformly strong performance across all tumor subtypes simultaneously, and generating transparent, human-interpretable rationale for every model output.

This work directly confronts both obstacles through a hybrid ensemble merging VGG16 and ResNet50 with embedded Grad-CAM interpretability, encapsulated in a browser-accessible application. The primary contributions of this paper are as follows:

- A rigorous head-to-head transfer learning assessment of VGG16 and ResNet50 as independent feature extraction engines for four-class tumor identification on a large, standardized public repository.
- A validation-F1-weighted ensemble fusion strategy that capitalizes on the complementary inductive biases of both architectures to exceed individual model accuracy.

- Full Grad-CAM integration yielding spatially localized explanation maps that expose which MRI zones govern each classification verdict—critical for regulatory acceptance and clinician trust.
- A production-ready Streamlit application supporting live scan upload, model toggling, confidence visualization, and Grad-CAM overlay display without any end-user programming requirement.
- Comprehensive per-class quantitative evaluation—precision, recall, F1-score, macro AUC-ROC, confusion matrices—benchmarked against contemporary peer-reviewed literature.

II. RELATED WORK

A. Handcrafted Feature Pipelines

Before deep learning became practical, automated MRI tumor analysis relied on manually engineered feature extractors—intensity histograms, Gabor-filtered responses, gray-level co-occurrence matrices—fed to conventional classifiers such as SVMs and k-NN models [2]. While these pipelines proved the feasibility of machine-assisted neuro-imaging, their expressiveness was fundamentally capped by the semantic poverty of hand-designed descriptors, and cross-institution generalization was hampered by sensitivity to scan protocol variation.

Cheng et al. [3] pushed this paradigm forward by applying a sparse-coded augmented bag-of-words model, reaching 91.28% accuracy on a multi-class benchmark and establishing that richer intermediate representations offer meaningful gains even without deep learning.

B. CNN and Transfer Learning

Abiwinanda et al. [4] established the feasibility of end-to-end learned representations for neuro-imaging when their compact CNN reached 84.19% accuracy on a three-class MRI problem without any hand-engineered features. Subsequently, transfer from ImageNet-pretrained networks dramatically raised the performance bar. Khan et al. [5] reported 94.82% four-class accuracy by fine-tuning VGG16, confirming that low-level visual statistics acquired from photographic pre-training transfer constructively to clinical MRI despite the apparent domain mismatch.

Deepak and Ameer [6] coupled GoogLeNet feature extraction with an SVM and recorded 98.00% three-class accuracy; however, that result must be weighed against a substantially smaller dataset and a simpler classification task than the four-class Kaggle benchmark adopted here.

C. Ensemble Methods in Medical Imaging

Ensemble approaches rest on the established principle that aggregating outputs from architecturally heterogeneous models with statistically independent error distributions reduces variance and elevates generalization beyond any single constituent. Sultan et al. [7] demonstrated a 3–5 percentage-point error reduction through multi-CNN output averaging versus the best individual model on comparable MRI data, a finding that extends consistently across medical imaging modalities.

D. Explainable AI in Clinical Diagnosis

The intrinsic opacity of deep models presents a critical barrier to regulatory clearance and clinician adoption of AI diagnostic tools. Grad-CAM [8], proposed by Selvaraju et al., derives class-discriminative localization maps directly from gradient signals propagated to the final convolutional layer without requiring architectural surgery. A large-scale survey by Tjoa and Guan [9] verified that such visualizations measurably increase physician readiness to act on AI-generated recommendations. Notwithstanding this recognized clinical importance, the combination of ensemble-level accuracy with deeply integrated Grad-CAM explainability inside a deployable clinical instrument remains an underexplored frontier.

III. DATASET

A. Source and Composition

All experiments draw on the Kaggle Brain Tumor MRI Dataset [10], an open-access compilation of 7,023 T1-weighted contrast-enhanced MRI scans spanning four diagnostic strata: glioma, meningioma, pituitary adenoma, and tumor-negative. Images are sourced from three distinct neuroimaging archives, producing a heterogeneous mixture of scanner makes, acquisition field strengths, and patient demographics that mirrors real-world clinical variability. Native spatial

resolutions range from 120×120 to 512×512 pixels. The prescribed training/testing split is honored in full to allow fair comparison with antecedent publications. Table I details the per-class sample counts.

Class	Train	Test	Total	%
Glioma	1,321	300	1,621	23.1
Meningioma	1,339	306	1,645	25.4
No Tumor	1,595	405	2,000	28.5
Pituitary	1,457	300	1,757	25.0
Total	5,712	1,311	7,023	100

TABLE I: Dataset Class Distribution

B. Preprocessing Pipeline

A fixed-order preprocessing chain is applied uniformly to every image prior to model input. The pipeline consists of five sequential stages:

- RGB three-channel conversion ensuring compatibility with ImageNet-initialized weights.
- Bicubic resampling to a standardized 128×128 spatial grid, balancing resolution fidelity with computational budget.
- Pixel-wise rescaling to [0, 1] through division by 255.0, matching the normalization used during ImageNet pre-training.
- Stochastic on-the-fly augmentation during training—random horizontal flipping ($p=0.5$), rotation within $\pm 15^\circ$, zoom within $\pm 10\%$, and brightness variation within $\pm 20\%$.
- Class-frequency-inverse sample weighting to offset the moderate overrepresentation of scan-negative cases (28.5% of the training corpus).

IV. METHODOLOGY

A. System Architecture Overview

The proposed framework is partitioned into four tightly integrated processing stages: (1) scan ingestion and standardized preprocessing, (2) parallel deep feature extraction through two independently fine-tuned CNN backbones, (3) weighted probabilistic ensemble fusion, and (4) Grad-CAM explanation generation. The complete end-to-end dataflow follows a sequential pipeline where each MRI scan passes through preprocessing before being forwarded simultaneously to both VGG16 and ResNet50 branches, whose softmax outputs are then fused and interpreted via Grad-CAM heatmaps.

B. VGG16 Feature Extractor

VGG16 [11] is a 16-layer CNN organized into five convolutional blocks, each terminated by 2×2 max-pooling, succeeded by three fully connected layers. Its exclusive use of 3×3 convolutional kernels at every depth generates dense multi-scale textural and structural encodings well matched to the fine-grained morphological distinctions required for tumor subtype discrimination.

The pretrained convolutional trunk is retained while the original classification head is replaced with: global average pooling reducing spatial maps to a 1D descriptor; a 512-unit dense layer with ReLU activation and 0.5 dropout; and a four-neuron softmax output. Training unfolds in two phases: an initial alignment stage (10 epochs, all backbone weights frozen, learning rate 1×10^{-4}) followed by selective fine-tuning of the two deepest convolutional blocks (40 epochs, learning rate 1×10^{-5}).

C. ResNet50 Feature Extractor

ResNet50 [12] resolves the vanishing-gradient obstacle in deep plain networks through identity shortcut connections that route gradients around one or more intermediate layers, preserving clean backward-pass signal regardless of

network depth. The architecture comprises a 7×7 convolutional entry block with 2×2 max-pooling, four successive residual stages containing 3, 4, 6, and 3 bottleneck units respectively, and global average pooling at the apex.

The shortcut pathways enable ResNet50 to simultaneously preserve low-level spatial cues—edge geometry, textural detail—alongside high-level semantic encodings of tumor mass and peritumoral tissue displacement across the entire network depth. This dual-scale representational richness is especially valuable for subtype discrimination where both fine-grained local texture and broader spatial organization carry independent diagnostic weight. The two-phase fine-tuning schedule applied to VGG16 is replicated identically for ResNet50.

D. Hybrid Ensemble Fusion

VGG16 and ResNet50 embody contrasting architectural philosophies: the former achieves superior fine-grained texture encoding through uniform shallow-kernel homogeneity, while the latter constructs richer hierarchical abstractions through skip-facilitated depth. A validation-score-weighted ensemble fusion strategy is adopted to exploit this architectural duality and systematically offset the characteristic failure modes of each constituent.

Let $P_1 \in \mathbb{R}^4$ and $P_2 \in \mathbb{R}^4$ denote the softmax probability vectors emitted by VGG16 and ResNet50, respectively. The ensemble output P_e is computed as the convex combination:

$$P_e = w_1 P_1 + w_2 P_2 \quad \text{subject to: } w_1 + w_2 = 1, \quad w_1, w_2 \geq 0$$

Fusion weights are proportional to each backbone's macro validation F1-score, yielding $w_1 = 0.47$ (VGG16) and $w_2 = 0.53$ (ResNet50). The predicted class is determined by $\text{argmax}(P_e)$.

E. Grad-CAM Interpretability Module

Transparent AI reasoning is indispensable for regulatory compliance and physician trust in clinical decision-support systems. Grad-CAM [8] delivers this without modifying the deployed architecture. For a designated target class c , the gradient of the scalar class activation y^c with respect to feature map A^k of the terminal convolutional layer is computed, then spatially pooled to derive neuron-level importance weights α^{kc} :

$$\alpha^{kc} = (1/Z) \cdot \sum_i \sum_j (\partial y^c / \partial A^{k_{ij}})$$

A ReLU-rectified weighted sum over all activation channels yields the final class-discriminative localization heatmap L^c , retaining only activations that positively influence the prediction:

$$L^c = \text{ReLU}(\sum^k \alpha^{kc} \cdot A^k)$$

Bilinear upsampling restores L^c to the original input resolution. A jet colormap overlay superimposed on the source MRI scan generates a per-inference localization visualization that highlights anatomically salient regions for direct clinical inspection.

V. EXPERIMENTAL SETUP

A. Implementation Environment

All experiments were conducted under the following hardware and software configuration:

- Platform: TensorFlow 2.12 / Keras, CUDA 11.8, Google Colab Pro.
- Hardware: NVIDIA Tesla T4 GPU (16 GB VRAM), Intel Xeon 2.20 GHz, 32 GB RAM.
- Optimizer: Adam ($\beta_1=0.9$, $\beta_2=0.999$, $\epsilon=1 \times 10^{-7}$); learning rate 1×10^{-4} per phase one and 1×10^{-5} per phase two.
- Learning Rate Schedule: ReduceLROnPlateau (factor=0.5, patience=5, min_lr= 1×10^{-7}).
- Training: Batch size 32; up to 50 epochs per phase; early stopping (patience=10) on validation loss.
- Regularization: L2 weight decay $\lambda=1 \times 10^{-4}$ on dense layers; dropout rate 0.50.
- Loss Function: Categorical cross-entropy with inverse-frequency class-balanced sample weights.
- Deployment: Streamlit 1.28 with @st.cache_resource model caching; Cloudflare tunnel for public access.

B. Evaluation Protocol

All performance figures are derived exclusively from the held-out test partition (1,311 scans), which remained entirely unseen during training and hyperparameter selection. Reported metrics encompass macro-averaged accuracy, precision, recall, and F1-score, together with one-versus-rest AUC-ROC providing a threshold-free discriminability index for each diagnostic class. This strict separation between training, validation, and test data ensures that reported metrics constitute genuine generalisation estimates free from any form of data leakage.

VI. RESULTS AND DISCUSSION

A. Overall Model Comparison

Table II reports the five-metric comparison of VGG16, ResNet50, and the Hybrid Ensemble on the 1,311-scan test set. The ensemble yields consistent gains across every criterion, substantiating the complementarity premise that motivates the fusion approach. The 1.4 percentage-point accuracy increment over the stronger individual backbone (ResNet50) is practically significant at this evaluation scale.

Model	Acc.	Prec.	Rec.	F1	AUC
VGG16	95.6%	95.1%	94.8%	0.950	0.987
ResNet50	96.4%	96.1%	95.9%	0.960	0.991
Ensemble	97.8%	97.5%	97.3%	0.974	0.995

TABLE II: Performance Comparison — Held-Out Test Set (1,311 scans)

B. Per-Class Performance

Table III provides the class-level breakdown for the Hybrid Ensemble. Glioma and No Tumor categories achieve the highest scores, reflecting their comparatively distinctive morphological signatures in T1-contrast MRI. Meningioma recall is marginally lower at 95.4%, in keeping with its well-documented WHO-grade-dependent and anatomical-location-dependent imaging heterogeneity: convexity, sphenoid-wing, and skull-base presentations differ substantially in enhancement pattern—variation that a uniform 128×128 input may only partially resolve.

Class	Precision	Recall	F1-Score	N
Glioma	98.3%	97.7%	0.980	300
Meningioma	96.2%	95.4%	0.958	306
No Tumor	98.5%	98.8%	0.987	405
Pituitary	97.3%	97.7%	0.975	300
Average	97.5%	97.3%	0.974	1,311

TABLE III: Per-Class Hybrid Ensemble Metrics — Test Set

C. Literature Comparison

Table IV positions the proposed framework against landmark peer-reviewed studies on comparable brain tumor MRI tasks. The ensemble attains competitive accuracy under the most demanding experimental configuration—four diagnostic categories and the largest, most heterogeneous dataset in the comparison set. Numerical comparisons with three-class studies on smaller cohorts (notably Deepak and Ameer [6] at 98.00%) must account for the substantially reduced task difficulty in those settings.

Study	Architecture	Classes	Accuracy
Abiwinanda et al. [4]	Custom CNN	3	84.19%
Deepak & Ameer [6]	GoogLeNet + SVM	3	98.00%
Khan et al. [5]	VGG16 Fine-tune	4	94.82%
Sultan et al. [7]	CNN Ensemble	4	96.13%
Proposed (This Work)	VGG16 + ResNet50 Ensemble	4	97.80%

TABLE IV: State-of-the-Art Comparison

D. Discussion

The 1.4-point ensemble premium over ResNet50 and 2.2-point premium over VGG16 confirm that fusing architecturally heterogeneous models systematically cancels prediction errors arising from each backbone's distinctive inductive biases. The 0.995 AUC-ROC indicates near-complete probabilistic class separability across all four categories, suggesting well-calibrated softmax outputs suitable for clinical confidence thresholding.

Grad-CAM attribution maps consistently focus the ensemble's attention on tumor core and peritumoral edema zones in a manner anatomically concordant with radiological convention, providing qualitative evidence that the learned representations encode diagnostically meaningful information rather than scanner-specific artefacts or background correlates. This spatial fidelity is of paramount importance for clinical translation, as it allows radiologists to rapidly verify whether the model's decision is grounded in the correct anatomical region.

The performance-weighted fusion scheme proved more effective than simple averaging in preliminary experiments, since ResNet50's marginally stronger validation score warrants a slightly higher contribution ($w_2 = 0.53$) to the ensemble output. This data-driven approach to weight assignment is straightforward to implement and avoids the hyperparameter search associated with learned stacking approaches, while still capturing architecture-level complementarity.

VII. DEPLOYED WEB APPLICATION

The full inference stack is packaged within a production-grade Streamlit application that enables real-time MRI analysis by clinical end users without programming expertise. The application provides a unified, browser-accessible interface that abstracts all model complexity, exposing only the clinically meaningful inputs and outputs.

Implemented capabilities include:

- Drag-and-drop scan upload supporting JPEG and PNG formats with pre-inference input validation that rejects corrupted or sub-resolution images before they reach the model.
- Backbone selector toggling between VGG16, ResNet50, and the Hybrid Ensemble, allowing clinicians to compare individual model outputs against the fused prediction.
- Per-class softmax probability bar charts for transparent uncertainty communication, enabling risk-stratified interpretation rather than binary classification outputs.
- @st.cache_resource model caching to minimize per-request latency, ensuring interactive response times even in constrained cloud environments.
- Zero-configuration public access via Cloudflare tunneling, enabling secure remote access without requiring institutional VPN infrastructure.

Future production migration targets HIPAA-compliant cloud infrastructure with persistent TensorFlow Serving endpoints and role-based multi-user access control. Integration with existing hospital PACS systems and DICOM image support are also planned as near-term enhancements.

VIII. LIMITATIONS AND FUTURE WORK

A. Current Limitations

Despite the strong empirical results reported on the Kaggle benchmark, several limitations must be acknowledged before clinical deployment of this system can be considered:

- Single-benchmark evaluation: All metrics originate from one public dataset; prospective external validation across varied scanner platforms and clinical sites is essential before real-world deployment.
- Slice-level independence: Each 2D MRI slice is classified in isolation, discarding the inter-slice volumetric context available in full 3D acquisitions.
- Heuristic fusion weights: Weights derived from a single validation fold lack the statistical robustness of cross-validated or learned alternatives.
- Single-sequence input: Reliance on T1-contrast MRI alone leaves FLAIR, T2, and DWI complementary data unused, potentially limiting sensitivity to certain tumor subtypes.
- No prospective clinical evaluation: Controlled comparison of AI-assisted versus unassisted radiologist reporting has not yet been conducted, making it impossible to quantify clinical diagnostic gain attributable to AI assistance.

B. Proposed Enhancements

The following extensions are proposed to address the identified limitations and advance the framework toward clinical-grade robustness:

- Extend the ensemble to include EfficientNetB3 [13], DenseNet121 [14], and InceptionV3 [15] under a learned stacking meta-classifier to further diversify the inductive bias landscape.
- Add a ViT-B/16 Vision Transformer branch to capture long-range spatial dependencies via multi-head self-attention, complementing the local receptive field of convolutional backbones.
- Transition to 3D volumetric CNNs processing complete MRI volumes to exploit inter-slice spatial context and eliminate the slice-independence limitation.
- Incorporate FLAIR, T2, and DWI sequences through attention-guided late multi-parametric fusion to maximize diagnostic information extraction.
- Deploy federated learning to enable distributed cross-hospital training without centralizing sensitive patient data, ensuring compliance with GDPR and HIPAA mandates.
- Commission a prospective reader study to quantify clinical diagnostic gain attributable to AI assistance, with appropriate institutional review board oversight.

IX. CONCLUSION

This paper introduced a clinically oriented, explainable hybrid ensemble deep learning framework for automated multi-class MRI brain tumor classification. Strategically merging VGG16 and ResNet50 through a validation-proportional weighted fusion mechanism, the system achieves 97.8% classification accuracy and 0.995 AUC-ROC across four clinically relevant tumor categories on the 7,023-image Kaggle benchmark, outperforming both constituent model baselines and comparing favorably with contemporary literature under equivalent experimental conditions.

Embedding Grad-CAM attribution throughout the inference chain represents a substantive step toward clinically trustworthy AI-assisted neuro-oncological diagnosis, delivering spatially grounded interpretive evidence that aligns with established neuroimaging convention. The Streamlit web deployment eliminates the technical access barrier separating research-grade deep learning from practical clinical utility.

The delineated limitations and proposed extensions—spanning multi-backbone ensembles, Vision Transformer integration, 3D volumetric processing, multi-parametric MRI fusion, federated learning, and prospective clinical validation—chart a rigorous and ambitious trajectory toward robust, generalizable, and regulatory-compliant computer-aided diagnostic infrastructure for neuro-oncological imaging. The work demonstrates that high classification accuracy and human-interpretable AI reasoning are not mutually exclusive objectives; rather, they can be jointly optimized within a unified, deployable clinical framework.

X. ACKNOWLEDGMENTS

The author gratefully acknowledges the contributors to the Kaggle Brain Tumor MRI Dataset for making their curated neuroimaging data openly accessible to the research community. Computational resources were supplied via Google Colaboratory. No external funding supported this work, and no conflicts of interest exist.

REFERENCES

- [1] Q. T. Ostrom et al., "CBTRUS Statistical Report: Primary Brain and CNS Tumors in the United States, 2015–2019," *Neuro-Oncology*, vol. 24, no. S2, pp. v1–v95, 2022.
- [2] Y. Zhang et al., "Review of Brain Tumor MRI Detection and Segmentation Techniques," *Artif. Intell. Med.*, vol. 52, no. 3, pp. 118–130, 2011.
- [3] J. Cheng et al., "Augmented Region Representation for Brain Tumor Classification," *PLoS ONE*, vol. 10, no. 10, e0140381, 2015.
- [4] N. Abiwinanda et al., "Brain Tumor Classification Using Convolutional Neural Network," *IFMBE Proc.*, vol. 68, pp. 183–189, Springer, 2019.
- [5] M. A. Khan et al., "Multi-Class Brain Tumor Classification via Fine-Tuned VGG-16," *Comput. Mater. Continua*, vol. 67, no. 3, pp. 3353–3368, 2020.
- [6] S. Deepak and P. M. Ameer, "Transfer Learning CNN Features for Brain Tumor Identification," *Comput. Biol. Med.*, vol. 111, 103345, 2019.
- [7] H. H. Sultan et al., "Multi-Classification of Brain Tumor MRI Using Deep Neural Networks," *IEEE Access*, vol. 7, pp. 69215–69225, 2019.
- [8] R. R. Selvaraju et al., "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 336–359, 2020.
- [9] E. Tjoa and C. Guan, "A Survey on Explainable AI Toward Medical XAI," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 11, pp. 4793–4813, 2021.
- [10] M. Nickparvar, "Brain Tumor MRI Dataset," Kaggle, 2021. [Online]. Available <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>
- [11] K. Simonyan and A. Zisserman, "Very Deep Convolutional Networks for Large-Scale Image Recognition," *ICLR 2015*, arXiv:1409.1556.
- [12] K. He et al., "Deep Residual Learning for Image Recognition," *Proc. CVPR*, pp. 770–778, IEEE, 2016.
- [13] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for CNNs," *Proc. ICML*, vol. 97, pp. 6105–6114, 2019.
- [14] G. Huang et al., "Densely Connected Convolutional Networks," *Proc. CVPR*, pp. 4700–4708, IEEE, 2017.
- [15] C. Szegedy et al., "Rethinking the Inception Architecture for Computer Vision," *Proc. CVPR*, pp. 2818–2826, IEEE, 2016.

ABOUT THE AUTHOR

Vaibhav U. Mane is a postgraduate student in the Department of Data Science at SIES College of Arts, Science & Commerce (Empowered Autonomous), Mumbai, India. His academic interests include machine learning, deep learning, computer vision, medical image analysis, and explainable artificial intelligence. This paper presents his original research contribution in the domain of healthcare analytics, focusing on an explainable ensemble deep learning framework for accurate detection of brain tumors using MRI images. His work emphasizes the integration of multiple deep learning models with interpretability techniques to enhance diagnostic accuracy and transparency in medical decision-making. He can be reached at vaibhavsmc2526131@asc.sies.edu.in.



INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  ijirset@gmail.com



www.ijirset.com

Scan to save the contact details