

BRAI-XNet: A Deep Learning Framework for Brain Tumor Classification from MRI

Patricia-Ioana Boros
Babeş-Bolyai University

WeADL 2026 Workshop

The workshop is organized under the umbrella of WinDMiL, project funded by CCCDI-UEFISCDI, project number [PN-IV-P7-7.1-PED-2024-0121](#), within PNCDI IV

- 1 Motivation and challenges in brain tumor MRI classification.
- 2 The proposed BRAI-XNet framework and architectural enhancements.
- 3 Dataset preparation, training strategy, and evaluation methodology.
- 4 Comparative analysis of CNN, Transformer, and ensemble-based approaches.
- 5 Explainability analysis using Grad-CAM++ and LIME.
- 6 Deployment of the selected model within the DeepBrain platform.

Brain Tumor Classification from MRI

Clinical motivation

- Brain tumors remain among the most complex neurological pathologies.
- MRI is the preferred imaging modality due to its high soft-tissue contrast.
- Manual interpretation is time-consuming, prone to errors, and highly dependent on clinical expertise.

Machine learning perspective

- Multi-class classification is substantially harder than binary tumor classification.
- Tumors differ in morphology, texture, size, and anatomical location.
- Reliable automated systems should also include explainability methods to offer more transparency regarding the diagnosis.

Main Contributions

- 1 Proposed the BRAI-XNet deep learning framework.
- 2 Systematically evaluated CNN and Transformer architectures.
- 3 Introduced ResPath, SBR, and MA enhancement modules.
- 4 Investigated ensemble learning strategies.
- 5 Integrated explainability through Grad-CAM++ and LIME.
- 6 Developed and deployed the DeepBrain software platform.

Target Classes and Dataset

Classification targets

- Glioma tumor
- Meningioma tumor
- Pituitary tumor
- No tumor

Dataset source


Public 3264 MRI brain tumor dataset collected from Kaggle and organized into predefined training and testing subsets

Subset	Images	Purpose
Training	2870	Model development
Testing	394	Final evaluation
Total	3264	Complete dataset

Preprocessing

- RGB conversion
- Resize to 224×224
- Aspect-ratio-preserving padding
- ImageNet normalization
- Small augmentations

The Proposed BRAI-XNet Framework



BRAI-XNet Architecture

Backbone architectures

- EfficientNetB0
- MobileNetV2
- VGG16
- Vision Transformer

Enhancement modules

- Residual Path (**ResPath**)
- Spatial Boundary Refinement (**SBR**)
- Mixed Attention (**MA**)

Transfer learning

- Pretrained ImageNet weights
- Frozen backbone initialization
- Progressive fine-tuning

Design objective

Improve feature extraction quality while preserving discriminative tumor-related spatial information.

Training and Evaluation Protocol

Training strategy

- 1 Freeze pretrained backbone
- 2 Train classification head
- 3 Fine-tune upper layers

Optimization

- Adam optimizer
- Categorical cross-entropy loss
- Early stopping
- Learning-rate scheduling


Evaluation protocol

- Stratified 10-fold cross-validation
- Independent test set evaluation for the best performing models
- Final full-dataset cross-validation evaluation for the best model (chosen for deployment)


Metrics

- Accuracy, Precision, Recall, F1-score, Specificity, AUC

Training Dynamics



Accuracy evolution during training



Loss evolution during training

Baseline Architecture Comparison

Model	Accuracy	Macro F1	Precision	Recall	Specificity	AUC
EfficientNetB0	0.9624	0.9639	0.9664	0.9629	0.9869	0.9981
MobileNetV2	0.8889	0.8890	0.9045	0.8882	0.9616	0.9875
VGG16	0.4763	0.3875	0.4380	0.4340	0.8166	0.7771
Vision Transformer	0.7401	0.7420	0.7429	0.7491	0.9111	0.9196

Observation

EfficientNetB0 achieved the strongest and most stable overall baseline performance across all evaluation metrics.

EfficientNetB0 Architectural Variants

Architecture Variant	Accuracy	Macro F1-score
EfficientNetB0	0.9624	0.9639
EfficientNetB0 + MA	0.9634	0.9653
EfficientNetB0 + ResPath	0.9589	0.9612
EfficientNetB0 + SBR	0.9641	0.9657
EfficientNetB0 + ResPath + SBR	0.9672	0.9692
EfficientNetB0 + ResPath + MA	0.9526	0.9542
EfficientNetB0 + MA + SBR	0.9561	0.9589
EfficientNetB0 + MA + ResPath + SBR	0.9456	0.9474

Key finding

The combination of **ResPath + SBR** provided the best performance improvement for EfficientNetB0.

Independent Test-Set Evaluation

Metric	Performance
Accuracy	0.8325
Macro F1-score	0.8136
Macro Precision	0.8889
Macro Recall	0.8350
Macro Specificity	0.9415
Macro AUC	0.9534

Generalization analysis

The test set proved significantly more challenging than the stratified cross-validation folds. The top-5 best performing models on the cross-validation were also evaluated on the test set, and the one with the best results was EfficientNetB0 + ResPath + SBR.

Ensemble Learning Analysis


Investigated ensemble strategies

- Soft voting ensembles
- Weighted voting ensembles
- Hybrid ensembles
- Stacked ensembles


Observation

Although ensemble methods produced competitive results, none consistently outperformed the selected EfficientNetB0 + ResPath + SBR architecture on the independent test set.

Explainability Analysis



Grad-CAM++ visualization



LIME explanation

Purpose

Provide visual evidence supporting the classification decision which can be examined by clinical professionals to ensure the model's reliability.


Original MRI Images vs Tumor Crops

Metric	Original MRI	Tumor Crop
Accuracy	0.8325	0.7487
Macro F1-score	0.8136	0.7094
Macro AUC	0.9534	0.9386


Important finding

Tumor-only crops removed contextual anatomical information that contributed to classification performance. This experiment was made to see if only the discriminative regions are enough for correct classifications. The model used was the final one (EfficientNetB0 + ResPath + SBR).


DeepBrain User Workflow



DeepBrain User Workflow



MRI upload interface



Prediction and explainability results

Final outcome

EfficientNetB0 + ResPath + SBR achieved the best balance between predictive performance, robustness, and explainability.

- Strong baseline CNN performance
- Effective architectural enhancement modules
- Competitive test set generalization
- Explainability-aware analysis
- End-to-end deployment through DeepBrain

- Extension toward full 3D MRI processing
- Integration of multi-modal MRI sequences
- Hybrid CNN–Transformer architectures
- Quantitative explainability validation
- Larger-scale clinical evaluation

Thank you!

Questions?

Patricia-Ioana Boros
Babeş-Bolyai University