BRAIN TUMOR CLASSIFICATION

A Meta-Learner Ensemble Model for Brain Tumor Classification from MRI Scans

Sercan KÜLCÜ

Computer Engineering Department, Giresun University, Türkiye

7th International Conference on Engineering and Applied Natural Sciences (ICEANS 2025) | October 25-26, 2025, Konya, Turkey

Introduction & Motivation

- Brain tumor classification from MRI is critical for early diagnosis and treatment planning.
- Manual interpretation is time-consuming and prone to inter-observer error.
- Ensemble deep learning has shown strong results (Jader et al. 99.16%; Tandel et al. 98.88%) but overfitting, dataset limits, and computational cost remain open challenges.
- This study proposes a meta-learner ensemble combining EfficientNet-B0, DenseNet201, and ResNet50 on the BRISC dataset.

WHY IT MATTERS

6,000 contrast-enhanced T1-weighted MRI scans

4 tumor classes: glioma, meningioma, pituitary, no-tumor

Dataset: BRISC

- 6,000 contrast-enhanced T1-weighted MRI scans, annotated by certified radiologists.
- 5,000 training images and 1,000 testing images.
- Four classes: glioma, meningioma, pituitary tumor, and non-tumorous.
- Multi-plane views — axial, sagittal, coronal — with balanced annotations.
- Addresses gaps in BraTS / Figshare by offering consistent, diverse data and excluding less relevant modalities.

Glioma

Meningioma

Pituitary Tumor

No Tumor

Image Preprocessing Pipeline

1

RGB → Grayscale

Convert via OpenCV (RGB→BGR then grayscale)

2

Noise Reduction

Median blur with a 5×5 kernel

3

Otsu's Thresholding

Segment the tumor region from background

4

ROI Cropping

Largest contour + bounding box crop, resized to 244×244

Normalization: PyTorch transform pipeline, mean = std = 0.5 across all channels, final input converted back to RGB.

Ensemble Model Architecture

EfficientNet-B0

timm library, pretrained
Dropout(0.4) → FC(512) → ReLU
→ BN → Dropout(0.2) → FC(classes)

DenseNet201

torchvision, pretrained
Dense feature reuse
Same custom classifier head

ResNet50

torchvision, pretrained
Residual connections
Same custom classifier head

MLP Meta-Learner

Concatenated outputs → Linear(→1024) → ReLU → Dropout(0.3) → Linear(num classes)

Learns relationships among base-model outputs, acting as a decision-maker that optimizes the combined prediction.

Training Setup

Hardware	NVIDIA GeForce RTX 3050 (CUDA 11.8)
Framework	PyTorch, with Automatic Mixed Precision (AMP)
Batch Size	8, via PyTorch DataLoader with shuffling
Loss	Cross-entropy with label smoothing (0.1)
Optimizer	AdamW, learning rate 1e-4
Scheduler	Cosine annealing learning-rate schedule
Precision	autocast (forward pass) + GradScaler (gradient scaling)
Epochs	10

Results: Performance Comparison

Model	Precision	Recall	F1-Score
DenseNet201	0.9871	0.9870	0.9870
ResNet50	0.9881	0.9880	0.9880
EfficientNet-B0	0.9910	0.9910	0.9910
Ensemble (proposed)	**0.9980**	**0.9980**	**0.9980**

Table 1 — Precision, Recall, and F1-score (weighted average) on the 1,000-image test set.

WEIGHTED AVG F1

0.998

Best score among all models — outperforming all three individual base networks.

Glioma: 1.000 / 1.000 / 1.000
Precision / Recall / F1

Confusion Matrix & Training Progress

Confusion Matrix (Ensemble)

	Glioma	Meningioma	No Tumor	Pituitary
Glioma	**249**	3	0	2
Meningioma	2	**293**	9	2
No Tumor	0	0	**140**	0
Pituitary	0	4	0	**296**

Diagonal = correct predictions. Misclassifications are minimal (2–9 instances per class).

Training Progress (10 Epochs)

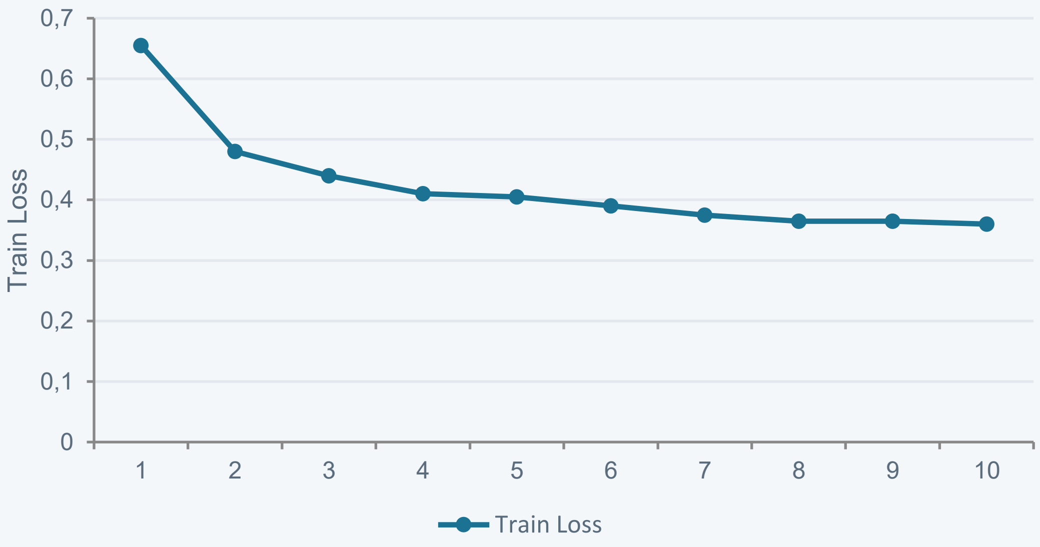

Loss decreases from 0.65 → 0.36; training accuracy rises from 45% to ~100% by epoch 10, confirming effective learning.

Discussion

1 Why the Ensemble Wins

The meta-learner integrates complementary strengths of EfficientNet-B0, DenseNet201, and ResNet50, correcting individual model weaknesses and reducing misclassification.

2 Beats Prior Work

0.998 weighted F1 exceeds comparable ensemble studies (e.g., Jader et al., 99.16% accuracy), with fewer confusion-matrix errors.

3 Efficient by Design

Automatic Mixed Precision (AMP) keeps training computationally efficient, supporting use in resource-limited clinical settings.

4 Preprocessing Matters

Otsu's thresholding and ROI cropping sharpen focus on tumor regions, aided by BRISC's diverse, high-quality, balanced annotations.

Conclusion

- Meta-learner ensemble of EfficientNet-B0, DenseNet201, and ResNet50 achieves 0.998 weighted F1-score on the BRISC dataset.
- Outperforms individual base models and prior published ensembles, with minimal misclassifications.
- AMP-based training keeps the approach computationally practical for resource-limited clinical environments.
- Offers a reliable, non-invasive decision-support tool for radiologists, supporting earlier and more accurate brain tumor diagnosis.

0.998

Weighted F1-Score

6,000
MRI scans (BRISC)

3 CNNs + 1
MLP meta-learner

Contact: sercan.kulcu@giresun.edu.tr | ORCID: 0000-0002-4871-709X