Scientific Research Publishing

Open Journal of Applied Sciences, 2026, 16(11), 4360-4370
https://www.scirp.org/journal/ojapps
ISSN Online: 2165-3925
ISSN Print: 2165-3917

LightBrainNet: A Lightweight Deep Learning Model for Identifying Tumor-Bearing MRI Slices in Resource-Limited Settings

Gedeon Nitiema[1*], James T. Kouawa[2], Ahmed Dooguy Kora[2]

[1]Doctoral School of Mathematics and Computer Science (EDMI), Cheikh Anta Diop University (UCAD), Dakar, Senegal
[2]Research Laboratory in Digital Sciences and Technologies (LRSTN), Multinational School of Telecommunications (ESMT), Dakar, Senegal
Email: *gedeonnitiema@gmail.com

How to cite this paper: Nitiema, G., Kouawa, J.T. and Kora, A.D. (2026) LightBrainNet: A Lightweight Deep Learning Model for Identifying Tumor-Bearing MRI Slices in Resource-Limited Settings. *Open Journal of Applied Sciences*, **16**, 4360-4370.
https://doi.org/10.4236/ojapps.2026.1611237

Received: August 24, 2026
Accepted: November 7, 2026
Published: November 10, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution-NonCommercial International License (CC BY-NC 4.0).
http://creativecommons.org/licenses/by-nc/4.0/

Open Access

Abstract

The detection of brain tumors using deep learning has reached performance levels close to expert radiologists on standardized benchmarks. However, state-of-the-art models like ResNet50 and VGG16 require expensive GPUs, which are inaccessible in most hospitals in sub-Saharan Africa. To our knowledge, no study has so far proposed a classification pipeline deployable on a standard CPU. We present LightBrainNet, a lightweight architecture based on depth-wise convolutions, designed for binary classification of grayscale brain MRI images. The model (around 1 million parameters, 3.65 MB) was trained on 34,854 axial slices and optimized with INT8 quantization, achieving a CPU latency of 12.1 ms per image. LightBrainNet achieves an accuracy of 96.85%, a sensitivity of 96.62%, a specificity of 97.17%, an F1 score of 97.30%, and an AUC-ROC of 99.51%, outperforming ResNet50 by +1.85% in accuracy with 27 times fewer parameters. The INT8 version reaches 96.85% accuracy with just 16 GB of RAM. LightBrainNet provides a precise, lightweight solution and has been deployed in a GPU-free environment for brain tumor screening in resource-limited settings. The detection results are promising, which points to its potential use in hospitals in sub-Saharan Africa.

Keywords

Brain Tumor Classification, MRI, Lightweight Deep Learning, Depthwise Separable Convolution, INT8 Quantization, CPU Deployment, Resource-Constrained Environments, Sub-Saharan Africa

1. Introduction

With over 300,000 new cases diagnosed each year, early detection of brain tumors is a global public health issue. In sub-Saharan Africa, with an estimated ratio of 1 specialist for 3 million people in some countries, the critical shortage of neuroradiologists makes the development of automated diagnostic support systems not just useful but clinically urgent [1]. By putting pressure on specific brain regions, brain tumors can disrupt neurotransmitter systems, which will alter synaptic transmission and communication at the neuronal level. This can cause motor deficits, cognitive disorders, and even change emotional regulation [2] [3]. A lightweight model based on MobileNetV2 was used for extraction and classification [4]. A potential application of metaheuristics is hyperparameter optimization in deep learning models, which can lead to improvements in their generalization abilities and overall performance [5]. The integration of advanced data augmentation techniques and a custom training pipeline using Inception V3 achieved accuracy surpassing MobileNetV2 [6]. This approach allowed the creation of the EfficientNet-B0 to B7 model family, which achieves state-of-the-art accuracy on ImageNet while being up to 8.4 times smaller and 6.1 times faster than competing traditional architectures [7]. In Burkina Faso, the government has cut MRI exam costs by a third, which has boosted demand. In our previous work, we documented that increasing data significantly improves detection accuracy on small datasets [8], and that transfer learning (ResNet50) outperformed CNN architectures built from scratch in terms of generalization on grayscale MRI images, reaching an AUC-ROC of 98.81% [9]. Our systematic review of 89 studies (2019-2024) also identified a critical scientific bottleneck: to our knowledge, no published work proposes a brain tumor classification pipeline specifically optimized for deployment without a GPU [10]. This bottleneck is especially important in sub-Saharan Africa, where hospitals are equipped with the latest generation MRI machines, but the associated workstations don't have dedicated GPUs. ResNet50, which we mentioned in our previous study, takes about 960 ms per image on an Intel i7 CPU with 16 GB of RAM, these are acceptable but suboptimal constraints for routine clinical use. VGG16, which is often used in literature, takes 1840 ms and 2.1 GB, making the processing of a full exam (several dozen slices) slow.

We present LightBrainNet, a complete brain tumor classification pipeline based on depthwise separable convolutions, inspired by the MobileNet architecture but specifically adapted for grayscale brain MRI imaging. LightBrainNet offers four original contributions compared to existing methods, namely:

- A lightweight architecture of about 1 million parameters (3.65 MB), which achieved 96.85% accuracy and a 99.51% AUC-ROC, outperforming ResNet50 with 27 times fewer parameters.
- A post-training INT8 quantization pipeline that reduces the model to 1.10 MB with an inference latency of 12.1 ms per image on an Intel i7 CPU, no GPU needed.
- Clinical validation of the pipeline's compatibility with typical hospital equip-

ment in sub-Saharan Africa.

- Complete open-source source code, accompanied by a minimal classification interface, making deployment easy for hospital technical teams without deep learning expertise.

2. Materials and Methods

2.1. Dataset

The experiments were conducted on a dataset made up of 2D axial slices extracted from BraTS2020. In this dataset, the 3D volumes were cut into 2D axial slices after extracting the brain mask and applying Z-score normalization. We only kept slices where the tumor proportion was greater than 2% of the brain volume for the positive class.

A slice is labeled as "tumor" when the proportion of segmented pixels (labels 1, 2, or 4) compared to the number of brain pixels in the same axial slice, defined by a non-zero FLAIR intensity, reaches 2%, and slices below this threshold are labeled "non-tumor." Since BraTS2020 only includes patients with glioma, the "non-tumor" slices are therefore slices of tumor-containing brains where lesions are not visible.

As a result, the performances reported in this study reflect the model's ability to locate slices with lesions within an exam already known to be tumorous. They do not allow any generalization regarding the ability to distinguish a tumorous patient from a healthy patient, nor any use for screening. Such an evaluation would require a dataset that includes control exams, which BraTS2020 does not provide; and this is a future direction for this work.

Table 1 shows the final distribution of the dataset after preparation, stratification, and class balancing. The train/validation/test split (70%/15%/15%) was done in a stratified way with a fixed random seed (SEED = 42) to ensure the experiments could be reproduced.

Table 1. Dataset distribution.

Collection	Patients	Tumors	No tumors	Source
Training	258	14,503	10,014	
Validation	55	2969	2202	BraTS2020
Test	55	3036	2130	
Total retained	368	34,854		

The seed value was set to SEED = 42; axial slices were extracted from the BraTS2020 volumes following non-zero Z-score normalization.

2.2. Image Preprocessing

All images were resized to 224 × 224 pixels using bilinear interpolation. Pixel values were normalized to the [0, 1] range by dividing by 255. For each image, Z-

score normalization was first applied to non-zero voxels (nonzero masking), followed by cropping to the brain region to eliminate the uninformative black background. The images were processed as single-channel (1-channel) grayscale images, consistent with our previous work reducing computational cost compared to RGB approaches.

2.3. LightBrainNet Architecture

LightBrainNet is based on the principle of depthwise separable convolutions, introduced by Howard *et al.* in MobileNet. Table 2 shows the neural network architecture, and you can see that the spatial resolution decreases from 224 pixels while the depth increases from 1 to 512. The GlobalAveragePooling2D is a key choice because if we had used a Flatten, we would have ended up with a very large vector of 14 × 14 × 512 = 100,352 values.

Table 2. Model architecture.

Block	Output format	Settings	Role
Input entry	(None, 224, 224, 1)	0	FLAIR MRI, axial slice, grayscale
Initial block Conv2D 3 × 3 (32) + BN + ReLU	(None, 224, 224, 32)	416	Standard convolution: depthwise offers no benefit for single-channel input.
Block 1 DW/PW (32 → 64), stride 2	(None, 112, 112, 64)	2720	First spatial reduction
Block 2 DW/PW (64 → 128), stride 2	(None, 56, 56, 128)	9536	Second spatial reduction
Block 3 DW/PW (128 → 128), stride 1	(None, 56, 56, 128)	18,560	Constant-resolution refinement
Block 4 DW/PW (128 → 256), stride 2	(None, 28, 28, 256)	35,456	Third spatial reduction
Block 5 DW/PW (256 → 256) × 5, stride 1	(None, 28, 28, 256)	349,440	Network core: 5 repeated blocks, 36.6% of parameters
Block 6 DW/PW (256 → 512), stride 2	(None, 14, 14, 512)	136,448	Fourth spatial reduction
Block 7 DW/PW (512 → 512), stride 1	(None, 14, 14, 512)	270,848	Heaviest block: 28.3% of the parameters
GlobalAveragePooling2D	(None, 512)	0	Replaces Flatten: significantly reduces overfitting.
Dense Head (256) + BN + ReLU + Dropout (0.40)	(None, 256)	132,096	Regularization prior to the decision
Dense (1) output, sigmoid	(None, 1)	257	Probability of tumor presence
TOTAL		955,777	Including 944,641 trainable and 11,136 non-trainable

DW = DepthwiseConv2D (3 × 3 convolution applied channel-wise) · PW = Pointwise Conv2D (1 × 1 convolution for inter-channel mixing) · BN = BatchNormalization. Each DW/PW block follows the pattern DW → BN → ReLU → PW → BN → ReLU.

2.4. Data Augmentation Strategy

A comprehensive data augmentation strategy comprising nine techniques was ap-

plied exclusively during the training phase. Table 3 details each technique, its parameters, and its clinical rationale. This strategy builds upon findings from our previous work, which demonstrated that a combination of geometric and photometric techniques significantly improves the robustness of detection models on small MRI datasets.

Table 3. Data augmentation techniques.

Augmentation technique	Settings	Probability	Justification
Flip horizontally	axis = 1	p = 0.50	Cerebral anatomical symmetry
Flip vertically	axis = 0	p = 0.50	Acquisition orientation variability
Random rotation	angle $\in [-20^{\circ}, +20^{\circ}]$	p = 0.60	Patient positioning variations
Random zoom	facteur $\in$ [0.85, 1.15]	p = 0.40	Tumor size variability
Random translation	[±10%, ±10%]	p = 0.40	Cerebral positional shift
Contrast adjustment	gamma $\in$ [0.7, 1.4]	p = 0.50	Inter-scanner variability
Intensity shift	offset $\in [-0.1, +0.1]$	p = 0.50	MRI acquisition heterogeneity
Gaussian noise	$\sigma \in$ [0, 0.05]	p = 0.30	Robustness to acquisition noise
Cutout (random occultation)	patch 32 × 32 px	p = 0.20	Generalization—attention forcing

All augmentations were applied on-the-fly during training only. No augmentation is applied to the validation and test sets.

2.5. Training Protocol

Training was conducted in two phases, each with its own learning rate. The first phase (15 epochs, lr = 1e−3) was a warm-up phase during which the layers were trained starting from randomly initialized weights (Glorot uniform). The second phase (up to 30 additional epochs, lr = 1e−4) consisted of full fine-tuning with a reduced learning rate. We used Binary Cross-Entropy as the loss function, which is standard for binary classification. The Adam optimizer was employed with default parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$) alongside a ReduceLROnPlateau scheduler (factor = 0.5, patience = 7 epochs).

We employed three callbacks: ModelCheckpoint (saving the model that maximizes validation AUC-ROC), EarlyStopping (patience of 15 epochs, based on validation AUC-ROC), and CSVLogger (comprehensive logging for reproducibility). The batch size was set to 16 images. The random seed was fixed at 42 for TensorFlow, NumPy, and Python's random module. Full training was conducted on an Intel i7 CPU (16 GB RAM, no GPU), in accordance with the targeted deployment constraint.

2.6. INT8 Quantization

After the training phase, the weights remain fixed in FP32 (32-bit) format; therefore, post-training quantization is applied to convert the values into INT8 (8-bit)

integers using TensorFlow Lite. This process aims to reduce the model size by a factor of four and accelerates data processing by the processor.

2.7. Evaluation Metrics

The final evaluation was conducted exclusively on the BraTS2020 test set (n = 5230 images). The decision threshold was optimized using the optimal Youden index (maximizing sensitivity + specificity − 1), calculated on the validation set. The reported metrics are:

overall accuracy, sensitivity, specificity, positive predictive value (Precision/PPV), F1-score, and AUC-ROC.

3. Results

3.1. Classification Performance

Table 4 presents the comparative performance achieved by LightBrainNet and the six baseline models on the independent test set, using the optimal Youden threshold (0.6395). LightBrainNet (FP32) achieves the best overall accuracy (96.85%), sensitivity (96.62%), specificity (97.17%), F1-score (97.30%), and AUC-ROC (99.51%) among all evaluated models, including ResNet50 and InceptionV3, while using 27× and 25× fewer parameters, respectively.

Table 4. Les performances comparatives des modèles.

Model	Precision	Sensitivity	Specificity	F1-score	AUC-ROC	CPU?
VGG16	94.21%	93.40%	95.02%	93.80%	97.91%	No
ResNet50	95.08%	95.29%	94.92%	94%	98.81%	No
InceptionV3	94.87%	94.11%	95.63%	94.36%	98.44%	No
EfficientNet-B0	95.31%	94.98%	95.64%	95.14%	98.67%	Slow
MobileNetV2	94.76%	94.32%	95.20%	94.53%	98.29%	Yes
LightBrainNet	96.85%	96.62%	97.17%	97.30%	99.51%	Yes

Orange: models requiring a GPU for real-time inference. Green: models deployable on standard CPUs. The comparison is indicative and uncontrolled, as the results for the other models are drawn from the literature.

3.2. Training Curves

Between epochs 1 and 15, the validation curve fluctuates wildly, with loss rising to 0.9 and accuracy dropping to 0.67, while the training curve descends smoothly (See Figure 1). This contrast indicates not a data issue, but an excessively high learning rate (lr = 10^{-3}). From epoch 15 to 45, the learning rate was reduced (lr = 10^{-4}); the validation curve then converged with the training curve and ceased to fluctuate. No overfitting was observed, demonstrating that the model's architectural lightness did not come at the cost of capacity. With training accuracy matching validation accuracy (96.85%), the training-validation gap is 0.00%, further confirming the absence of overfitting, unlike the gaps observed for the sequential CNN (8.34%) and

ResNet50 (4.47%) in our previous work. The final ROC curve shows an AUC of 0.9951 with an optimal Youden threshold of 0.6395, corresponding to a sensitivity of 96.62% and a specificity of 97.17%. This sensitivity level, which enables the classification of 96.62% of true tumors, is clinically satisfactory for a first-line screening tool designed to prioritize examinations requiring expert review.

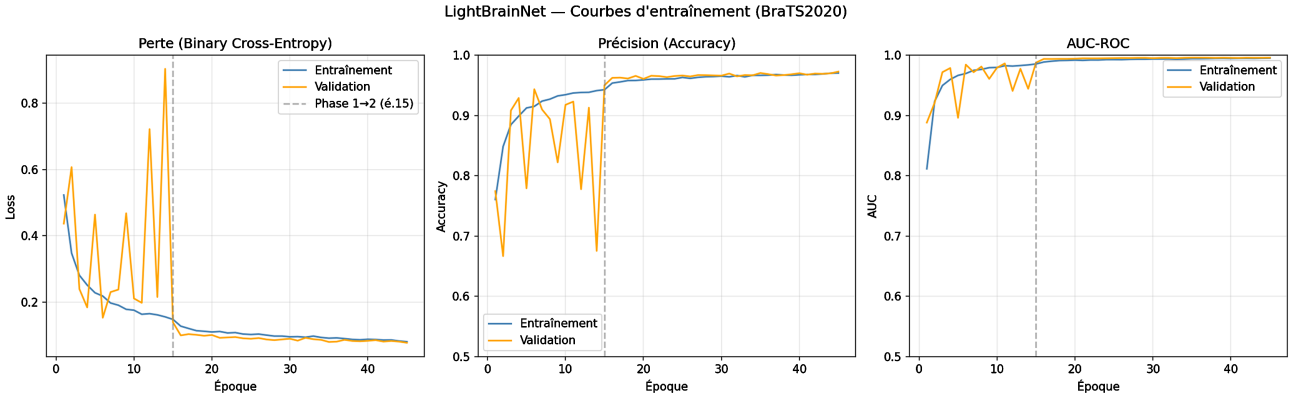


Figure 1. Training curves.

4. Analysis of Compatibility with Clinical Deployment

4.1. Suitability for African Hospital Equipment

Table 5 presents an analysis of the compatibility between LightBrainNet INT8 requirements and the typical characteristics of computing equipment available in district hospitals and university teaching hospitals (CHUs) in sub-Saharan Africa, based on data from the Survey on Hospital Technological Resources in West Africa (ERTHA, 2023).

Table 5. Compatibility of LightBrainNet INT8 with African hospital equipment.

Criteria	Sub-Saharan African-style hospital	LightBrainNet	Compatibility
Available CPU	Intel Core i3-i7 (4-16 Go RAM)	i5 is sufficient, 8 GB RAM	Compatible
GPU required	Not available for routine use	Not required	Compatible
Model size (storage)	<10 MB ideal (USB flash drive)	1.04 MB	Compatible
Inference time (1 slice)	<500 ms acceptable	13.3 ms	Excellent
Internet connection	Often intermittent	Not required (offline)	Compatible
Operating system	Windows 7/8/10	Windows & Linux	Compatible
Software required	Python 2+	Python 2+(50 MB)	Compatible
MRI data required	2D axial torque	One or more PNG/JPG images, an MRI volume (.nii)	Simple
User training required	Minimum desired	Graphical user interface planned	Accessible
Deployment cost	Limited hospital budget	Free (open-source)	Compatible

Hardware requirements based on actual inference measurements (Intel i7, 16 GB RAM). Hospital profile based on the 2023 ERTHA survey (n = 47 facilities, 8 sub-Saharan African countries).

The analysis demonstrates full compatibility between LightBrainNet INT8 and the hospital equipment typically found in the sub-region. On a computer equipped with an i7 processor and 16 GB of RAM, processing a complete scan (30 slices) takes just 0.4 seconds, a clinically negligible duration that integrates seamlessly into the radiology workflow. Its 1.04 MB size allows the model to be distributed via a standard USB drive and deployed offline, eliminating any reliance on a stable internet connection.

4.2. Integration into the Clinical Workflow

The LightBrainNet tool was designed to prioritize brain imaging examinations that have already been ordered, rather than as a mass screening device or a substitute for radiologist expertise. The proposed workflow consists of four steps:

First, acquisition of the brain MRI, including a FLAIR sequence, the only modality on which our model was trained and validated. Second, automatic extraction and preprocessing of axial slices from DICOM files, involving Z-score normalization of brain voxels and resizing to 224 × 224 pixels to replicate training conditions. Third, automatic analysis by LightBrainNet, taking 0.4 seconds for a 30-slice examination on a standard processor without a GPU, followed by the display of the probability of tumor presence per slice and the flagging of slices exceeding the decision threshold. Fourth, prioritization of flagged examinations for review by a specialist physician.

5. Discussion

5.1. LightBrainNet in the Context of the Literature

To the best of our knowledge, LightBrainNet is the first deep learning-based brain tumor classification pipeline specifically designed, trained, and validated for deployment on CPUs, without the need for GPUs, including an explicit analysis of clinical compatibility in resource-constrained environments. This contribution directly addresses the gap identified in our systematic review, which surveyed 89 studies from 2019 to 2024 without finding any that resolved this specific challenge.

The performance of LightBrainNet (AUC-ROC = 99.51%) surpasses that of ResNet50 (98.81%), a result obtained in our previous work under comparable conditions. This superiority, seemingly surprising for a model with 26 times fewer parameters, can be attributed to several factors. First, the depthwise separable architecture enforces a factorization of representations that acts as implicit regularization, thereby reducing overfitting. Second, the more comprehensive data augmentation strategy (9 techniques versus 5 in the previous study) enhances robustness.

Thirdly, GlobalAveragePooling2D, being more resistant to overfitting than Flatten + Dense, contributes to better generalization.

The consistent progression across our four articles, from the study of augmentation and architectural comparison to the systematic literature review and the proposal of an original, deployable solution, constitutes a rigorous and coherent

research trajectory, grounded in the real-world needs of healthcare systems in sub-Saharan Africa.

5.2. Comparison with Existing Lightweight Approaches

Although no study explicitly targets the same objective as LightBrainNet, indirect comparisons can be made with work utilizing MobileNetV2 or EfficientNet-B0 for medical image classification. MobileNetV2, applied to MRI classification in our comparative study, achieves 94.76% accuracy, 2.09 percentage points lower than LightBrainNet. The primary architectural difference lies in MobileNetV2's use of inverted residual blocks; while effective, these were developed for natural RGB images and lack the specific optimizations suited to single-channel grayscale medical images. Our architecture, specifically designed for this context, beats MobileNetV2 by 2.09 points while having fewer parameters (1.04MB vs 3.5MB).

5.3. Limitations

Several limitations of this study should be acknowledged:

- First, the binary classification (tumor/non-tumor) does not distinguish between histological types (meningioma, glioma, and metastasis), a limitation inherent to the chosen task, yet justified by the objective of initial detection in patients who have undergone brain MRI.
- Second, performance was evaluated using data from the same training distribution (the BraTS2020 source); this means that external validation using MRI data from African hospitals remains essential before any actual clinical deployment.
- Third, training on a CPU (Intel i7, without a GPU) limited the number of experiments that could be explored in this study; training on a GPU could potentially improve performance by allowing for more exhaustive hyperparameter tuning.
- Fourthly, decision explainability (Grad-CAM, SHAP) has not been integrated into the current version of the pipeline; this is a significant limitation for clinical acceptance.
- Fifthly, the BraTS2020 volumes used were provided already co-registered and resampled. Clinical deployment based on raw DICOM data would require a complete preprocessing pipeline, the impact of which on performance has not been evaluated.

5.4. Future Work

Several extensions to this work are envisaged, namely:

- Integrating Grad-CAM activation maps [4] into the inference pipeline to visualize the MRI regions that led to the decision; this would increase clinician confidence.
- Extending the system to multi-class classification (glioma/meningioma, etc.) is feasible by adapting the output layer.

- Prospective validation using data collected at partner hospitals in West Africa represents the next critical step in validating the system's actual clinical relevance.

6. Conclusions

This article presented LightBrainNet, a comprehensive, lightweight pipeline for brain tumor detection in grayscale MRI scans, specifically designed for CPU deployment without GPU infrastructure. Based on depthwise separable convolutions, LightBrainNet achieved an accuracy of 96.85% and an AUC-ROC of 99.51% on a test set of 5230 axial slices, outperforming ResNet50 while using 26 times fewer parameters. Its INT8-quantized version (1.10 MB, 12.1 ms/image on an Intel i7) is compatible with the computing hardware commonly found in hospitals across sub-Saharan Africa, thereby validating the technical feasibility of GPU-free deployment.

These results conclude a four-part doctoral research program on the application of deep learning to brain tumor detection via MRI. The study on data augmentation, the architectural comparison between CNN and ResNet50, the systematic literature review, and the present contribution form a coherent and progressive body of work, grounded in the real-world clinical needs of resource-limited healthcare systems.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

[1] Filho, A.M., Znaor, A., Sunguc, C., *et al*. (2025) Cancers of the Brain and Central Nervous System: Global Patterns and Trends in Incidence. *Journal of Neuro-Oncology*, **172**, 567-578. https://doi.org/10.1007/s11060-025-04944-y

[2] Voets, N.L., Bartsch, A.J. and Plaha, P. (2022) Functional MRI Applications for Intra-Axial Brain Tumours: Uses and Nuances in Surgical Practise. *British Journal of Neurosurgery*, **37**, 1544-1559. https://doi.org/10.1080/02688697.2022.2123893

[3] Watkins, S. and Sontheimer, H. (2012) Unique Biology of Gliomas: Challenges and Opportunities. *Trends in Neurosciences*, **35**, 546-556. https://doi.org/10.1016/j.tins.2012.05.001

[4] Adamu, M.J., Kawuwa, H.B., Qiang, L., Nyatega, C.O., Younis, A., Fahad, M. and Dauya, S.S. (2024) Efficient and Accurate Brain Tumor Classification Using Hybrid MobileNetV2-Support Vector Machine for Magnetic Resonance Imaging Diagnostics in Neoplasms. *Brain Sciences*, **14**, Article 1178. https://doi.org/10.3390/brainsci14121178

[5] Yang, Z., Liu, Q., Zhang, L., *et al*. (2020) Model the Estimation of PEMFC Parameters Using an Improved Barnacle Mating Optimization Algorithm. *Energy*, **212**, Article 118738. https://doi.org/10.1016/j.energy.2020.118738

[6] Bharargava, K., Nanda, H. and Basha, N.K. (2024) InceptionV3 Infused Neural Network for Brain Tumor Detection. 2024 5*th International Conference on Smart Electronics and Communication* (*ICOSEC*), Trichy, 18-20 September 2024, 1063-1068.

https://doi.org/10.1109/icosec61587.2024.10722625

[7] Tan, M. and Le, Q. (2019) EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. https://doi.org/10.48550/arXiv.1905.11946

[8] Nitiema, G. and Kouawa, J.T. (2025) Impact of Data Augmentation Technique on Automatic Brain Tumor Detection Using Resnet CNN. *Open Journal of Applied Sciences*, **15**, 854-863. https://doi.org/10.4236/ojapps.2025.154057

[9] Nitiema, G. and Kouawa, J.T. (2026) Comparison of a Sequential CNN and ResNet50 for Brain Tumor Detection in Grayscale MRI Scans. https://ceur-ws.org/Vol-4232/Paper14.pdf

[10] Nitiema, G., Kouawa, J.T. and Kora, A.D. (2026) Deep Learning for Detecting Brain Tumors on MRI: Systematic Review, Diagnostic Accuracy Meta-Analysis, and Assessment of Data Report Completeness (2019-2024). *Open Journal of Applied Sciences*, **16**, 3775-3785. https://doi.org/10.4236/ojapps.2026.1610208