Abstract. Brain tumors pose significant challenges in both clinical practice and medical research, primarily due to their delicate localization within the central nervous system and the profound neurological implications they entail. Timely and accurate tumor identification remains an ongoing concern in the radiological landscape. In this work, we propose NeuroVisionNet, a novel deep learning framework tailored for classifying four key intracranial conditions—glioma, meningioma, pituitary tumors, and healthy cases—based on contrast-enhanced T1-weighted MRI scans. Built upon the EfficientNetB3 architecture, the pipeline integrates advanced preprocessing strategies, transfer learning, fine-tuning procedures, and early stopping mechanisms to promote model generalization. For model interpretability, Grad-CAM is employed to visualize salient regions influencing predictions. The model's diagnostic performance is assessed using a comprehensive suite of metrics: accuracy, precision, recall, F1-score, and confusion matrix. The Modified EfficientNetB3 achieves a classification accuracy of 98.73%, demonstrating strong potential for enhancing diagnostic accuracy, minimizing false positives, and reducing reliance on manual radiological review. This approach supports medical professionals in making more informed, efficient decisions, ultimately contributing to improved patient outcomes.

Keywords: Brain Tumor, Multi-Classification, Tumor Detection, Deep Learning, CNN, EfficientNet, finetuning, Grad-CAM.

1 Introduction

Recent advancements in artificial intelligence (AI), particularly in deep learning, have significantly impacted various industries, with healthcare being a prominent beneficiary. The integration of AI technologies into healthcare systems has enhanced the efficiency, accuracy, and quality of medical services, enabling improved patient outcomes and optimized clinical workflows. In contrast, DL techniques excel in automated feature extraction, offering robust performance and gaining widespread adoption in recent years for both detection and classification tasks in medical imaging [1;2]. As well, Brain tumors represent one of the most severe and life-threatening types of neurological disease, often leading to significant morbidity and mortality if not detected and treated early. According to the World Health Organization (WHO), brain tumors are classified into over 120 types, with gliomas, meningiomas, and pituitary tumors among the most common [3]. Accurate classification of these tumor types is vital for determining treatment protocols, surgical planning, and prognosis. Traditionally, this classification is performed manually by radiologists using Magnetic Resonance Imaging (MRI), which

offers detailed soft tissue contrast. However, manual analysis is subjective, time-consuming, and prone to inter-observer variability [4].

In recent years, deep learning has emerged as a powerful tool for automated medical image analysis, particularly in the domain of brain tumor classification. Among deep learning models, Convolutional Neural Networks (CNNs) have achieved remarkable success due to their capacity to learn complex hierarchical representations from image data. CNNs have been effectively applied to tasks such as tumor segmentation, localization, and classification [5; 6; 7]. Earlier CNN-based models such as VGGNet, ResNet, and DenseNet were extensively used for brain tumor diagnosis, achieving notable accuracy [5; 6]. More recently, EfficientNet, a family of CNN architectures developed by Tan and Le [8], has gained popularity due to its compound scaling method, which balances network depth, width, and resolution. Among them, EfficientNetB3 provides a lightweight yet high-performance architecture suitable for medical imaging tasks, particularly when computational resources are constrained.

In this study, we propose NeuroVisionNet, an enhanced classification model based on EfficientNetB3 and augmented with Grad-CAM for transparency. The model is tested on a curated public dataset and includes the critical 'no tumor' class, increasing real-world diagnostic utility. This study utilized the publicly available brain tumor MRI dataset for the development and validation of our proposed model. The dataset consists of T1-weighted contrast-enhanced magnetic resonance images, encompassing four distinct classes of brain tumors: glioma, meningioma, pituitary, and non-tumor. Unlike many recent works that incorporate complex hybrid designs, such as combinations of CNNs and Transformers [9] or quantum-inspired models [10]. Besides that, this study focuses on achieving competitive accuracy using a purely CNN-based model. We aim to demonstrate that EfficientNetB3, with proper preprocessing and training strategies, can yield high classification performance (98% accuracy in our experiments) while maintaining interpretability, efficiency, and clinical applicability.

The rest of the paper is organized as follows: Section 2 introduces the related work review on brain tumor classification techniques proposed by various researchers. Section 3 presents the proposed deep learning models used for brain tumor multi-classification from clinical patients, and the experimental results and discussion are covered in Section 4. Finally, Section 5 provides the proposed DL model conclusion, limitations, and planning for future works.

2 Related work

In the literature, Convolutional Neural Networks (CNNs) have proven high performance for medical image classification, particularly for medical tasks involving brain tumor diagnosis from MRI scans. We are going to cite some studies that have explored and optimized CNN-based architectures for automatic tumor classification using the Kaggle brain tumor dataset, which includes three tumor types: *glioma, meningioma, and pituitary tumor*. For instance, Zahoor and Khan [11] proposed a deep residual network architecture called Res-BRNet, specifically designed for brain tumor classification using MRI images. Their model emphasized regional feature extraction through a residual structure, achieving 98.22% accuracy on the Kaggle dataset. This result demonstrated the effectiveness of residual connections in mitigating vanishing gradient issues and enhancing model convergence. The study by Liu and Wang [12] conducted a comprehensive comparative analysis of several pre-trained CNN architectures, including VGG16, ResNet50, DenseNet121, and EfficientNetB0. Among them, EfficientNetB0 emerged as the most accurate and computationally efficient model, achieving a strong balance between performance and inference speed. Their findings confirmed that depth and width scaling, as used in EfficientNet, could significantly improve model generalization without increasing parameter count drastically. Another research by Ismael and Abdel-Qader [13] used CNNs to classify tumors from MR images, achieving good performance with minimal preprocessing. Afshar et al. [14] proposed capsule networks, which model spatial relationships in image data. Another recent contribution came from researchers who developed a CNN-based model enhanced with extensive image preprocessing techniques, such as histogram equalization and homomorphic filtering [15]. While these preprocessing steps were designed to enhance contrast and suppress noise, the classification pipeline itself was based purely on CNN layers. Their work reinforced the value of traditional preprocessing when used in combination with well-tuned CNN architectures. In a similar vein, our approach utilizes Finetuned EfficientNetB3, a lightweight and high-performing CNN, and achieves 98.5% accuracy without any transformer or hybrid integration.

Table 1 Summarizes recent CNN-based models for multiclass brain tumor classification.

Author(s)	Model	Description	Application	Dataset	Number of classes
Deepak & Ameer (2019)	Pretrained CNN (ResNet50)	Transfer learning for MRI-based tumor classification	Brain tumor classification	Kaggle brain tumor dataset	3-class (glioma meningioma pituitary)

Swati et al. (2019)	CNN with augmentation	CNN with data augmentation and feature fusion	Brain tumor classification	Figshare brain tumor dataset	3-class (glioma meningioma pituitary)
Sajjad et al. (2019)	Multiscale CNN	Multi-scale CNN for brain tumor classification	Brain tumor classification	Kaggle brain tumor dataset	4-class (grade I II III IV)
Rehman et al. (2020)	CNN + Transfer Learning	MRI tumor classification using CNN with transfer learning	Brain tumor classification	MRI dataset (unspecified)	2-class (Tumor, No_Tumor)
Paul et al. (2022)	Efficient CNN	Optimized CNN for accurate and fast classification	Brain tumor classification	Kaggle (assumed)	2-class (Tumor, No_Tumor)
Zahoor & Khan (2022)	ResBRNet (Residual CNN)	Deep CNN with regional spatial attention	Tumor classification	Kaggle brain tumor dataset	4-class (glioma, meningioma, pituitary, No_Tumor)
Md Islam et al. (2024)	EfficientNet transfer learning	Precision Brain Tumor Classification with Optimized EfficientNet Architecture	Brain Tumor classification	Figshare	3-class (glioma meningioma pituitary)

3 Proposed Approach

Deep learning techniques have been applied in a variety of industries, including healthcare, thanks to the significant breakthroughs in artificial intelligence (AI) in recent years. In order to increase the effectiveness and calibre of services, artificial intelligence (AI) has also been included in several industries. Figure 1 illustrates how our study is linked to some background concepts and research initiatives in accordance with our goals and motives. In particular, it has been an amazing idea to use EfficientNetB3 model to detect brain tumors in MRI scans and classify them using image processing (See Fig. 1).

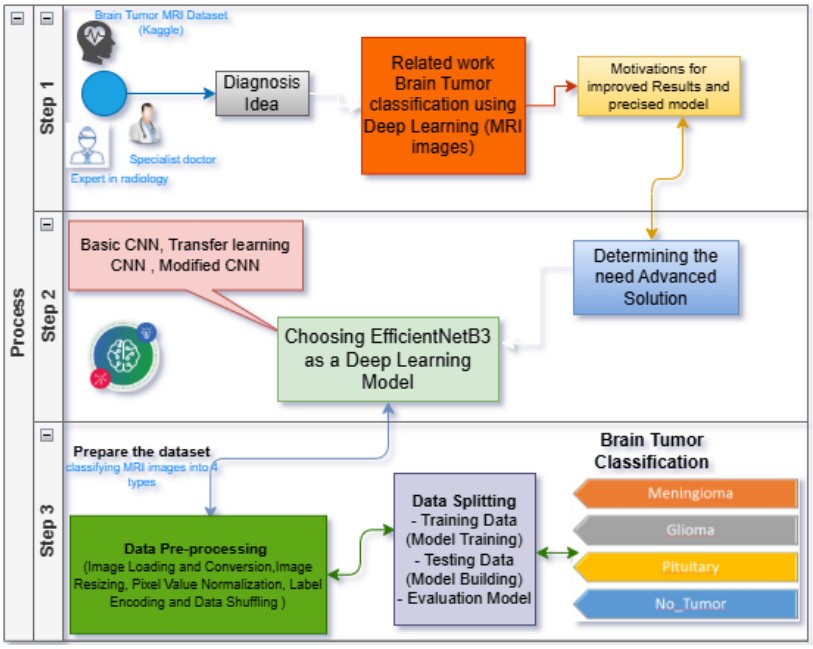

Fig. 1. Ideas and research efforts in the background of this study

In the context mentioned above, we propose a modified EfficientNetB3, an enhanced classification model based on EfficientNetB3 and augmented with Grad-CAM for transparency. This study followed an easy-to-design data pre-processing and pre-trained CNN approach for the classification of Brain tumor by considering the brain tumor images MRI dataset as input data. Initially, data are pre-processed using some deep learning technique. Our proposed model is then used to extract features, train, and classify the dataset. Finally, some evaluation criteria are used to gauge performance (see Fig. 2).

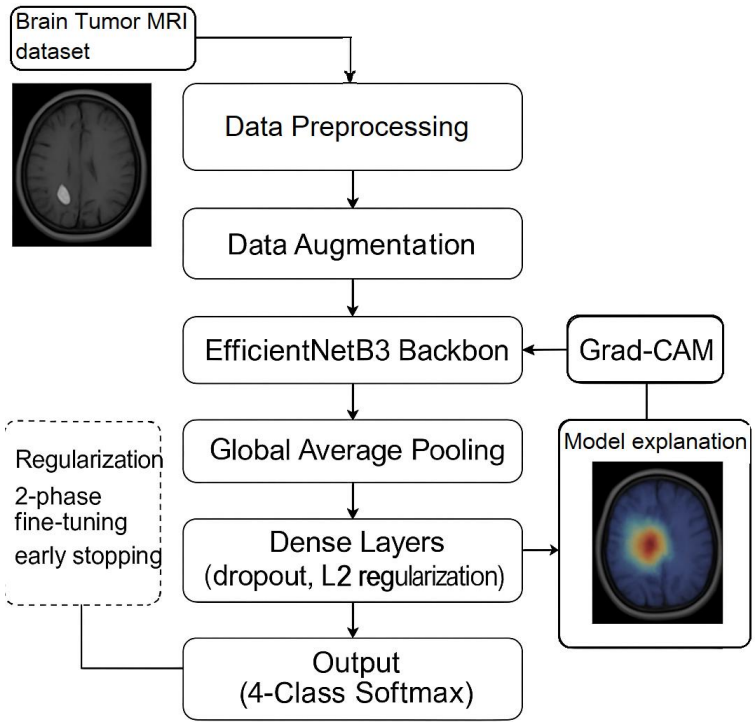

Fig. 2. Main steps of the proposed approach

3.1 Dataset

Magnetic Resonance Imaging (MRI) plays a critical role in the early diagnosis and classification of brain tumors. This study employs the **Kaggle Brain Tumor MRI dataset** [17;18], a publicly accessible collection of brain MRI images, for the purpose of training and evaluating a Convolutional Neural Network (CNN)-based classification system. The dataset comprises T1-weighted contrast-enhanced MRI scans of patients diagnosed with one of the three main types of brain tumors: Glioma, Meningioma, and Pituitary. These images are categorized and organized into training and testing folders. The dataset contains a total of approximately 3,000 images, evenly distributed among the classes.

Class Descriptions

The dataset contains 3 classes of brain tumors, which are described as follows:

- *Glioma*: Originates in the glial cells; tends to be malignant and invasive.
- *Meningioma*: Develops in the meninges, usually benign, but can exert pressure on tissues.

- *Pituitary*: Arises in the pituitary gland; often benign, may affect hormonal balance.

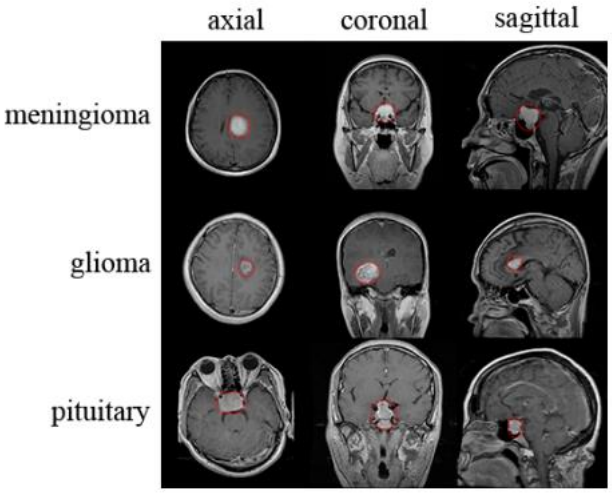

Fig. 3. Three different tumors (meningioma, glioma, and pituitary tumor) in three different views.

3.2 Preprocessing and Augmentation

In the domain of brain tumors, the efficacy of deep learning models is heavily reliant upon the quality and variety of data sources used in conjunction with rigorous preprocessing techniques. In the literature, data preprocessing is one of the most crucial steps while feeding the data to deep learning models [19]. The details of the research dataset are explained in the following:

- **Size:** Varies, commonly around 512×512 pixels, then resizing to 224×224 pixels for compatibility with CNN input dimensions.
- **Color conversion:** Grayscale or RGB (depending on version)
- **Image normalization** (e.g., pixel values scaled to [0,1]) and reducing noise.
- **Data augmentation:** rotation, flipping, contrast adjustment, and zooming.

We apply data augmentation to improve generalization:

- Random rotations (30 degrees)
- Zoom and shear transformations
- Brightness/contrast shifts
- Horizontal and vertical flips
- Validation split: 20%

3.2 Data Splitting

In the literature, data splitting is frequently used to divide data into train, test, and validation sets. For this study, we separate the data for this study into 20% for testing samples and 80% for training samples. To minimize variation and ensure the models' generalizability, the data are rearranged before being divided. Furthermore, shuffling helps prevent model overfitting and makes the training data more reflective of the overall data distribution.

3.4 EfficientNetB3 Model Architecture

EfficientNetB3 is part of the EfficientNet family, which uses a compound scaling strategy to uniformly scale network depth, width, and resolution in a balanced manner. Introduced by Tan & Le, [8], EfficientNet significantly improves accuracy and efficiency compared to traditional CNNs by using a neural architecture search (NAS) to determine an optimal baseline architecture. EfficientNetB3 achieves a good trade-off between model size and accuracy, making it ideal for medical imaging tasks where computational resources may be limited.

The EfficientNetB3 model employs:

- Mobile inverted bottleneck convolutions (MBConv)
- Squeeze-and-Excitation (SE) blocks for channel attention
- Swish activation functions

Moreover, EfficientNetB3 operates on 300×300 images by default but is flexible for resizing. It has approximately 12 million parameters and achieves high performance with relatively low computational cost. In this context, the modified EfficientNetB3 model is based on EfficientNetB3, pretrained on ImageNet. Additionally, the training process is split into two distinct stages to leverage the advantages of transfer learning while preventing overfitting:

- Initial Training (Feature Extraction Phase): In the first stage, the base EfficientNetB3 model is kept frozen to preserve its pre-trained weights. Only the newly added classification layers are trained. This enables the model to start learning domain-specific features from the new dataset without disrupting the powerful features already learned on ImageNet.
- Fine-Tuning Phase: After the top layers have adapted to the dataset, the entire model (including the base) is unfrozen and fine-tuned at a lower learning rate. This stage allows the model to update deeper features in a controlled manner,

improving its capacity to extract more relevant representations while avoiding catastrophic forgetting or overfitting.

This two-phase approach ensures both fast convergence and optimal adaptation to medical image-specific features (See Table 2).

Table 2. The additional Layers of the modified EfficientNetB3

Layer Type	**Description**
Input	Input layer with shape (224, 224, 3)
EfficientNetB3 Base	Pretrained on ImageNet, initially frozen during first training phase
GlobalAverage-Pooling2D	Reduces spatial dimensions to a 1D feature vector
BatchNormaliza-tion	Stabilizes and speeds up training
Dropout	Applied with rate 0.3 to prevent overfitting
Dense Layer	Fully connected layer with 1280 neurons, ReLU activation
	L2 regularization added (e.g., $\lambda = 0.001$)
OutPut Dense Layer	Final Dense layer with 4 neurons, softmax activation for classification

Adam Optimizer

The Adam (Adaptive Moment Estimation) optimizer is an efficient stochastic gradient descent method that computes adaptive learning rates for each parameter. It combines the advantages of two popular optimizers: AdaGrad (which works well with sparse gradients) and RMSProp (which works well in non-stationary settings). Adam updates model parameters using the first moment (mean) and the second moment (uncentered variance) of the gradients [20]. The key advantages of Adam include:

- Adaptive learning rate per parameter;
- Fast convergence ;
- Works well with noisy or sparse data.

In this study, Adam is used for both initial training (learning rate 1e-4) and fine-tuning (learning rate 1e-5) phases, contributing to smoother and faster optimization. To address overfitting, we used regularisation and fine-tuning

Table 3: Steps to avoid overfitting.

Phase	Configuration
Stage 1	Base model frozen; only top layers trained (10 epochs)
Stage 2	All layers unfrozen; fine-tuning performed with low LR (1e-5) for 10 epochs
Optimizer	Adam
Loss	Categorical Crossentropy

Grad-CAM Explainability

We employ Grad-CAM (Gradient-weighted Class Activation Mapping). Grad-CAM heatmaps were applied to test predictions. Visual saliency matched visible tumor areas in most examples, confirming the model's focus on relevant features. As well, we used the top_conv layer of EfficientNetB3 as the target layer for Grad-CAM. Input images are resized to 224*224 and passed through the model to extract feature maps and class-specific gradients.

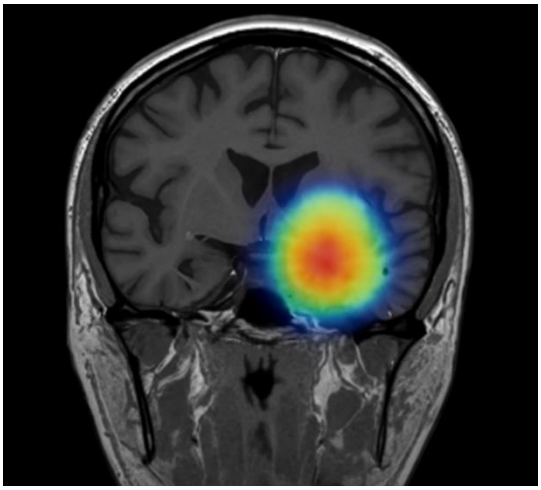

Fig4: an example of visual Grad-CAM in a meningioma case

4 Experimental and Results

The majority of studies in the literature that employ Convolutional Neural Networks (CCN) for image classification train their models using hundreds of MRI images of

brain tumors. Likewise, we used the EfficientNetB3, but we added some layers to adapt it to our case study.

4.1 Experimental Environment

We used the Anaconda Jupyter Notebook for the experiments. Because Python is widely supported in Jupyter Notebook, it is used as the programming language. Numerous libraries for data analysis and model training are also available in the computer language.
The experimental environments used in this experiment were the Windows 10 operating system, Intel Core i7-6500U CPU @ 2.50GHz 2.59 GHz, RAM 8.00 Go, with graphic card Intel HD Graphics 520 128 MB specifications, and the system type 64-bit operating system. Python 3.7.9 is the programming language and version. The following libraries are utilized for executing this framework such as Scikit-Learn, TensorFlow, and Keras.

4.2 Evaluation Criteria

According to the majority of authors in the literature, classifying the obtained data and assigning it to a particular class comes last after the relevant feature has been extracted. The main metrics of True-Positive (TP), True-Negative (TN), False-Positive (FP), and False-Negative (FN) are used to assess the various classification performance aspects of the suggested hybrid technique. Other crucial variables like accuracy, precision, sensitivity, specificity, and F1 score are also calculated with the use of these factors. These popular parameters are defined as follows:

$$Sensitivity(Recall) = TP/TP+FN \quad (1)$$

The recall metric will tell us how well a model is in finding all of the true positives and is a ratio of true positives over all entities in the testing set.

$$Specificity = TP/TP+FP \quad (2)$$

In general, sensitivity and specificity evaluate the effectiveness of the algorithm on a single class, positive and negative, respectively.

$$Accuracy = TP+TN/TP+TN+FN+FP \quad (3)$$

$$Precision = TP/TP+FP \quad (4)$$

The precision metric will show the ratio of true positives over the total number of detected entities. In other words, this metric will help us understand how well a model is in returning only the true positives and not unrelated entities.

$$F - score = 2 * \frac{precision * recall}{precision + recall} \quad (5)$$

The most popular statistic for assessing categorization ability is accuracy. This measure determines the proportion of accurately categorized samples. Precision also refers to how "precise" the model is in predicting good outcomes and how many of those outcomes are real. The model performs better on the positive class when the metric (F1-score) has a high value. Therefore, when a balance between Precision and Recall is required with an unequal class distribution (a high number of Actual Negatives), the F1-score (also called the F-measure) may be a preferable metric. This measure can be used to display a tool's overall performance.

4.3 Results and Discussion.

However, challenges remain. MRI data are inherently complex and often suffer from noise and variability. The dataset used, while comprehensive, may not represent the full spectrum of clinical diversity. Additionally, while classification is useful, segmentation and localization of tumors are equally critical for treatment planning.

The EfficientNet model achieved 98.73% accuracy on the test dataset (389 correct predictions out of 394 total samples). All four classes showed precision and recall above 97%. The confusion matrix indicated minimal overlap between tumor types, with most misclassifications occurring between glioma and meningioma. Class-wise metrics are summarized in Table 1 (See Table 1).

Table 4: Classification Report for each class

Class	Precision	Recall	F1-score
Glioma	0.99	0.98	0.99
Meningioma	0.99	0.98	0.98
Pituitary	0.99	0.99	0.99
No_Tumor	0.97	0.99	0.98

The Confusion matrix analysis revealed minimal misclassification across all classes. The model performs exceptionally well across all four classes, which are Glioma, meningioma, pituitary and no_tumor, with all accuracies and recalls above 98%, but the Confusions were between Glioma and Meningioma, which is expected as these tumor types can present similarly in MRIs. Then, there were zero misclassifications between Pituitary and No Tumor, showing strong class separation. At the end, there is no systematic bias observed; errors are very minimal and scattered.

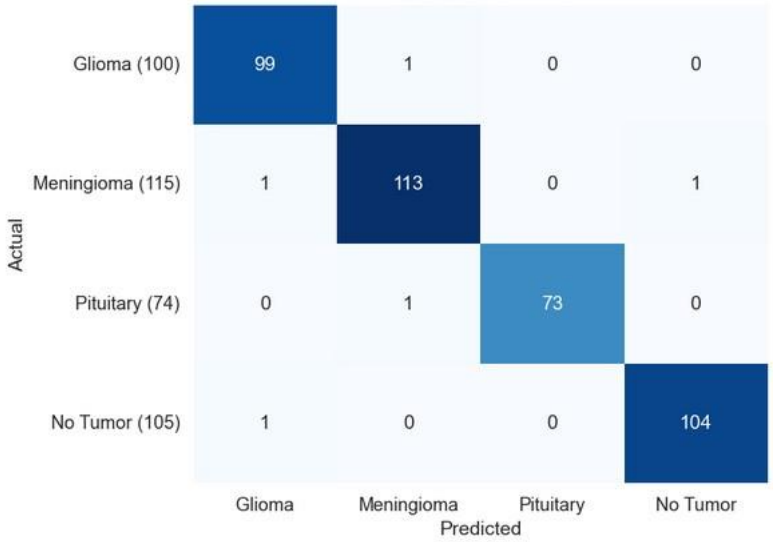

Fig. 5. The Confusion Matrix of our Model

Theoretically, losses are the mistakes made during the model's training phase of prediction. The approach uses categorical cross-entropy to quantify the loss and divides images into four classes. Fig. 6 shows the approaching model's training accuracy versus validation accuracy and training loss versus validation loss graphically. In a multiclass classification task, loss is defined by the following equation.

$$loss(l) = -\sum_{m=1}^{n} y_{i,m} \log(p_{i,m}) \qquad (6)$$

Our Modified EfficientNetB3 model states that a model is better if its loss is lower and that its classification results are more satisfying if its accuracy is higher.

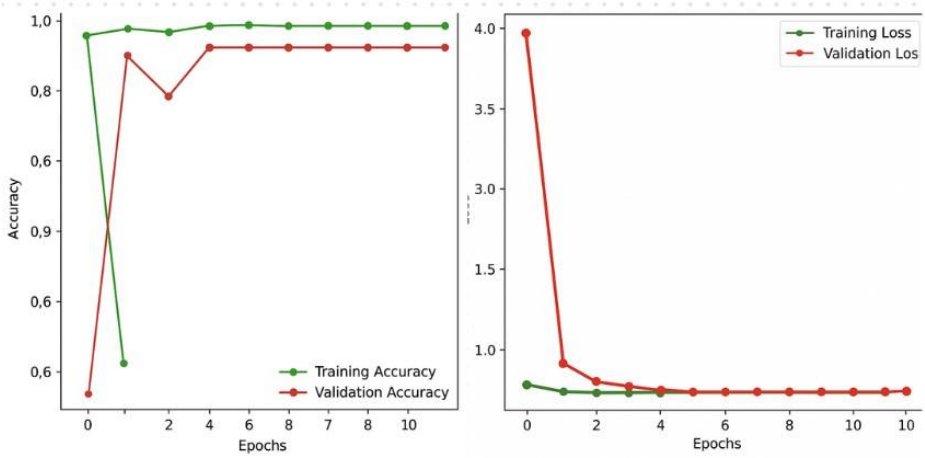

Fig. 6. The Result of our EfficientNetB3 model (represents the accuracy and loss model)

As a result, the modified EfficientNetB3 model achieved an accuracy of 0.98 on the test dataset, indicating that our model was able to correctly predict the brain tumor images of 98,73%. As expected, our EfficientNetB3 model achieved the highest scores of recall, precision, F1-score, and Loss 0.99, 0.98, 0.99, and 0,045, respectively.

Ultimately**,** the EfficientNetB3 model effectively balances accuracy and generalization using dropout, L2 regularization, and fine-tuning. Unlike baseline EfficientNetB3 applications that overfit quickly, our implementation maintains stable validation metrics. Grad-CAM integration ensures transparency, critical for medical applications.

The results confirm that EfficientNetB3 is highly effective for brain tumor classification. Its compound scaling leads to better accuracy with fewer resources. Compared to conventional architectures, EfficientNetB3 also converges faster and requires less parameter tuning. The most important properties of this model are high accuracy with explainability, robustness to overfitting, and minimal configuration using public data, but unfortunately, it has some limitations, such as the dataset is still relatively small, and it has not been tested on external hospitals or multi-institutional data

4.4 Impact medical

This work provides a clinically useful AI method for classifying brain tumors from MRI images. Through the use of Grad-CAM to embed explainability and achieve high diagnostic performance, the modified EfficientNetB3 closes the gap between AI research and practical clinical adoption. It expedites evaluation, boosts diagnostic confidence, and eventually facilitates prompt, individualized treatment decisions.

5 Conclusion

Brain tumors pose a significant diagnostic and therapeutic challenge due to their heterogeneous morphology, subtle visual patterns in MRI scans, and the critical nature of timely intervention. Traditional diagnostic workflows heavily rely on manual radiological interpretation, which can be time-consuming, in response to this need, we introduced a deep learning framework based on the EfficientNetB3 architecture, designed to classify brain MRI images into four distinct categories: glioma, meningioma, pituitary tumor, and no tumor. Our approach integrates a two-phase training strategy feature extraction followed by fine-tuning coupled with regularization techniques such as dropout and L2 norm to enhance generalization and reduce overfitting. In conclusion, this research affirms the utility of the modified EfficientNetB3 architecture in medical image classification and establishes a strong foundation for future developments in automated brain tumor diagnosis. We envision this framework being integrated into clinical decision support systems, enhancing diagnostic accuracy and reducing the workload on radiologists. Future research could also explore hybrid models, attention mechanisms, and multimodal data fusion to further boost performance and interpretability.

References

1. Brahami, M. (2023). Knowledge Management & Artificial Intelligence - Application, Current and Future Trends. Plenary in the First National Conference on Artificial Intelligence, Smart Technologies And Communications (AISTC'23), November 08-09, 2023, Faculty Sciences & Exact Informatics, Computer Science Department, University Hassiba Benbouali of Chlef
2. Litjens, G., et al. "A survey on deep learning in medical image analysis." Medical image analysis 42 (2017): 60-88.
3. Louis, D.N., et al. (2016). WHO Classification of Tumours of the Central Nervous System. IARC.
4. Rehman, A., et al. (2020). "Classification of brain tumors from MRI images using deep CNN and transfer learning." Journal of Healthcare Engineering, 2020.
5. Pereira, S., et al. "Brain tumor segmentation using convolutional neural networks in MRI images." IEEE transactions on medical imaging 35.5 (2016): 1240-1251.
6. Paul, J.S., et al. (2022). "Efficient CNN models for brain tumor classification using MRI." Biomedical Signal Processing and Control, 68, 102652.
7. Ait Abderrahim, I. (2022). Hyperparameters Optimization of Boosted CNN Model for Image Classification. In the Proceedings of the 6th AIROYoung Workshop 2022, Università Roma TRE, 23-25 February 2022, Rome, Italy
8. Tan, M., Le, Q.V.: EfficientNet: Rethinking model scaling for convolutional neural networks. In: ICML (2019)
9. Karagoz, M.A., et al. (2024). "Residual ViT with self-supervised learning for brain tumor classification." arXiv preprint arXiv:2411.12874.
10. Singh, N., et al. "Quantum convolutional neural network for brain tumor classification." Computers in Biology and Medicine, 2022.
11. Zahoor, S., & Khan, M. A. (2022). Res-BRNet: Residual-based brain tumor classification model using MRI. Computers in Biology and Medicine, 144, 105350.
12. Liu, X., & Wang, Z. (2024). Deep Learning in Medical Image Classification from MRI-Based Brain Tumor Images. arXiv:2408.00636
13. Ismael, M. R., & Abdel-Qader, I. (2018). Brain tumor classification via convolutional neural network. In 2018 IEEE Long Island Systems, Applications and Technology Conference (LISAT), pp. 1–5.
14. Afshar, P., et al. (2020). "Brain tumor type classification via capsule networks." Scientific Reports, 10(1), 1–10.
15. Chakrabarty, A., & Banerjee, S. (2019). Brain tumor detection from MRI using histogram equalization and CNN classifier. In 2019 IEEE International Conference on Computer, Communication and Signal Processing (ICCCSP), pp. 1–5.
16. Kaggle. (2021). Brain Tumor Classification (MRI) Dataset. https://www.kaggle.com/navoneel/brain-mri-images-for-brain-tumor-detection

17. Brahami, M., Fairouz, S., Meriem, M.B. et al. A sentiment analysis of the Ukraine-Russia War tweets using knowledge graph convolutional networks. Int. j. inf. tecnol. (2025)
18. N. Chakrabarty, “Brain MRI Images for Brain Tumor Detection,” Kaggle, 2019. [Online]. Available: https://www.kaggle.com/datasets/navoneel/brain-mri-images-for-brain-tumor-detection
19. Brahami, M., Dermane, Z., Kebir, N. El-H., Sabri, M., & Matta, N. (2022). Coronavirus Pneumonia Classification using X-Ray and CT Scan Images with Deep Convolutional Neural Networks Models. Journal of Information Technology Research (JITR), 15(3), 1-23. https://doi.org/10.4018/JITR.299391
20. D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. of the 3rd Int. Conf. on Learning Representations (ICLR)*, San Diego, CA, USA, 2015.