diagnostics

MDPI

Article

Hyperparameter Optimization of Convolutional Neural Networks for Robust Tumor Image Classification

Syed Muddusir Hussain [1], Jawwad Sami Ur Rahman [1], Faraz Akram [1], Muhammad Adeel Asghar [2] and Raja Majid Mehmood [3,4,*]

[1] Biomedical Engineering Department, Riphah International University, I-14 Campus, Islamabad 45210, Pakistan; muddasir.hussain@riphah.edu.pk (S.M.H.); jawwad.sami@riphah.edu.pk (J.S.U.R.); faraz.akram@riphah.edu.pk (F.A.)
[2] Electrical and Computer Engineering Department, Riphah International University, I-14 Campus, Islamabad 45210, Pakistan; adeel.asghar@riphah.edu.pk
[3] Centre for Advanced Analytics, COE for Artificial Intelligence, Multimedia University, Persiaran Multimedia, Cyberjaya 63100, Selangor, Malaysia
[4] Faculty of Computing and Informatics, Multimedia University, Persiaran Multimedia, Cyberjaya 63100, Selangor, Malaysia
* Correspondence: rmeex07@ieee.org or rajamajid@mmu.edu.my

Abstract

Background/Objectives: The human brain is responsible for controlling various physiological functions, and hence, the presence of tumors in the brain is a major concern in the medical field. The correct identification and categorization of tumors in the brain using Magnetic Resonance Imaging (MRI) is a major requirement for the diagnosis and treatment of a tumor. The proposed research will focus on designing a CNN model that is optimized for tumor image classification. **Methods:** This research proposes an optimized CNN model featuring strategically placed dropout layers and hyperparameter optimization. This study uses a dataset of 640 MRI scans (320 tumor and 320 non-tumor) collected from a private hospital in Saudi Arabia. The proposed method utilizes a learning rate of 0.001 in combination with the Adam optimizer to ensure stable and efficient convergence. Its performance was benchmarked against established architectures, including VGG-19, Inception V3, ResNet-10, and ResNet-50, with evaluation based on classification accuracy and computational cost. **Results:** The experimental results show that the optimized CNN proposed in this work performs much better than the deeper architectures. The network reached a maximum training accuracy of 97.77% and a final test accuracy of 95.35% with a small test loss of 0.2223. The test accuracy of the optimized VGG-19 and Inception V3 networks was much lower, with a training time per epoch that was several orders of magnitude higher. The validation stability of the proposed network was high (92.25% to 95.35%) during the final stages of training. **Conclusions:** The conclusion drawn from this study is that hyperparameter optimization and strategic regularization are more advantageous for tumor classification using MRI images than the mere depth of the model. The accuracy of 95.35% with low computational complexity makes this lightweight CNN model a feasible solution for real-time applications.

Keywords: brain tumor; convolutional neural network; hyperparameter; MRI images; precise tumor detection; tumor detection

Check for updates

Academic Editor: Takouhie Catherine Maldjian

Received: 14 March 2026
Revised: 14 April 2026
Accepted: 15 April 2026
Published: 18 April 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

The human brain is the most complex organ in the body, executing cognitive and physiological functions with remarkable precision. Its highly specialized neurons form intricate networks for information processing, making a comprehensive understanding of brain function a major challenge [1,2]. Anatomically, the brain comprises gray matter, white matter, and cerebrospinal fluid (CSF), which are each essential for maintaining its structural and functional integrity [3]. Any kind of irregular cell proliferation in any of these parts may lead to the development of a tumor [4]. These tumors may have a severe effect on cognitive functions and are a major global health concern because of the serious effects they have on morbidity and mortality rates [5].

High-resolution Magnetic Resonance Imaging is currently considered the gold standard for intracranial imaging, given its non-invasive nature and excellent spatial resolution [6]. Nevertheless, early and accurate detection of brain tumors remains challenging, primarily due to the intricacies involved [7]. In the conventional approach, the interpretation of an MRI image heavily relies on the radiologist's expertise, which makes the process prone to human errors [8]. In addition, the conventional analysis of an MRI image is a laborious and resource-intensive process, which makes the process a bottleneck in the provision of healthcare services to meet the increasing global demands [9].

In order to overcome these challenges, computer-aided diagnostic (CAD) systems have been developed and implemented. Recent advancements in this field utilize Artificial Intelligence (AI) and deep learning (DL) techniques, particularly Convolutional Neural Networks (CNNs). These techniques help to achieve accurate analysis of medical image data. Earlier, conventional machine learning techniques like Support Vector Machine (SVM) were used to detect brain tumors. However, the utilization of more sophisticated techniques like VGG and ResNet has shown better results in terms of accuracy [10–12]. However, a large number of research papers indicate the limitations of these techniques, which include the fact that these techniques, despite having the ability to represent complex features, are extremely sensitive to overfitting on small medical image datasets. These techniques are computationally expensive, which makes the problem less feasible [13].

The main objective of this research is to propose a lightweight machine learning model with structured dropout regularization to fill the gap between computational efficiency and diagnostic performance. Contrary to other research that aims to enhance the depth of models, we explore hyperparameter optimization, particularly learning rates and optimizers, to achieve the best possible performance on a small-scale MRI dataset.

The findings reveal that a well-optimized lightweight CNN can outperform more complex networks. With a test accuracy of 95.35% and lower computational overhead compared to VGG-19 and Inception V3, the proposed model demonstrates significant promise for clinical brain tumor classification.

Contribution of this work

The following are the contributions of this study:

- A lightweight optimized Convolutional Neural Network is proposed for brain tumor classification.
- A systematic approach is introduced for hyperparameter optimization, aimed at learning rate adjustment and dropout regularization, showing that under the conditions of small-scale data, targeted optimization can considerably improve the performance of the model.
- This paper also presents a detailed comparison with various deep learning architectures, such as VGG-19, Inception V3, and ResNet versions, showing that a properly tuned lightweight model can be more efficient than a deeper and more computationally expensive model.

- An elaborate and repeatable experimental setup is presented, with a clear architectural design, training setup, and an ablation study to determine the independent effects of important hyperparameters.
- The suggested approach is justified in a clinically relevant setting in which there is a limited amount of data, which proves a possible implementation method in real-time and resource-limited medical settings.

2. Related Work

Deep learning-driven automated brain tumor detection has recently attracted substantial attention because of its ability to improve the diagnostic accuracy and reduce the certainty of manual MRI evaluation. Traditional approaches relied on handcrafted features and classical classification. The proposed model is based on a lightweight CNN model that incorporates structured dropout regularization to fill the existing gap between computational complexity and diagnostic performance. Contrary to existing research that aimed to enhance model depth, we examine the impact of hyperparameter optimization, namely learning rates and optimizers, to achieve optimal performance on a relatively small MRI dataset that lacked sufficient robustness in real-world applications of medical imaging [1,2].

Convolutional Neural Networks are currently acknowledged as a pivotal advancement in deep learning, facilitating automatic feature extraction and delivering superior classification results. Nevertheless, the performance of CNNs is influenced by a variety of factors, such as the depth of the network, network architecture, data quality, and hyperparameter tuning. Complex deep learning architectures such as VGG-19, ResNet, and Inception V3 have been extensively explored for cancer diagnosis [3–5]. However, some research works have been observed to offer less-than-optimal accuracy, which can be considered to be due to inadequate preprocessing and less optimal optimization methods.

Amin et al. conducted extensive research on brain tumor detection and classification, emphasizing the role of ML techniques in addressing the constraints of traditional medical image analysis. Their work covered a broad spectrum of methodologies, from conventional image processing approaches to state-of-the-art transfer learning and deep learning techniques. The authors have given a valuable suggestion regarding future research work; safer, more accurate, and lightweight models for real-time analysis were suggested as a priority [14].

Amin et al. (2022) emphasized the importance of developing safer, more accurate, yet lightweight models for real-time clinical use, despite the transformative power of deep learning [14]. Fang and Wang (2022) proposed the MFF-DNet to address the loss of boundary information in conventional convolutions and achieved improved tumor segmentation using multi-modal feature fusion [15]. Although some complex models based on federated learning or deep ensemble boosting have achieved accuracy rates above 98% [16,17], they tend to have high communication overheads and require extremely high computational resources, making them less practical in resource-limited medical settings.

New research highlights the efficacy of optimizing hyperparameters for CNNs in the classification of medical images. Research has indicated that the optimization of hyperparameters such as the learning rate, batch size, and dropout, together with data augmentation, substantially increases accuracy and generalization capacity in classifying skin lesions and mammograms [18,19]. Moreover, the use of self-attention models and vision transformers in CNNs improves the extraction of features, resulting in greater precision in lesion identification [20,21]. Ensemble learning using fine-tuned CNNs has also been found to perform better than single models in the classification of dermoscopic images [22]. The application of these methods in hyperparameter-optimized CNNs will ensure the efficient and accurate MRI classification of brain tumors.

Modern developments contribute to improvements in medical imaging systems. Segmentation based on CNNs helps detect tumors [23], whereas image registration aids in increasing diagnostic accuracy [24]. Early prediction is facilitated by transfer learning [25]. Efficiency and privacy are provided by knowledge distillation and federated learning [26,27].

To better understand the novelty of the current study, Table 1 summarizes recent studies conducted on brain tumor classification, which achieved lower validation accuracies than the hyperparameter-tuned CNN model proposed in this study.

Table 1. Comparative summary of deep learning studies.

Author(s)	Model Used	Dataset Used	Accuracy (%)	Limitation
Pereira et al. [23]	CNN	BRATS	85	Shallow architecture
Kumar et al. [8]	CNN + SVM	Private	84.1	The hybrid model lacks generalization
Arabahmadi et al. [6]	Inception V3	Public MRI	78.2	Under-optimized hyperparameters
Chakraborty & Banerjee [12]	VGG-19	Kaggle MRI	85.6	Overfitting, high training time
Maqsood et al. [9]	Multi-modal CNN	Private MRI	82.5	Feature fusion complexity
Lepakshi et al. [10]	CNN	Public T1-weighted	80.3	Limited preprocessing
Neamah et al. [7]	CNN	BraTS 2020	83.7	Small training dataset
Roy et al. [24]	DenseNet	Local Dataset	81.4	Ineffective learning rate
Li et al. [13]	ResNet-18	TCGA	76.9	Shallow ResNet variant
Reddy et al. [25]	Inception V3	Private MRI	79.2	Insufficient augmentation
Insufficient augmentation	ResNet-10	Combined MRI	74.6	Weak feature extraction
Present study	CNN (optimized)	Private	95.35	Highly optimized tuning

As observed from the table, even with the use of sophisticated architectures, most models could not exceed 85% accuracy. Issues arose from insufficient data augmentation, shallow model depth, or suboptimal optimizer choice. For example, Inception V3 performed poorly in several studies because of its intricate parameter space and tuning sensitivity [8,28]. Likewise, deeper networks such as ResNet-18 and ResNet-10 could not achieve optimal accuracy without sufficient residual learning depth and well-tuned hyperparameters [24,29].

The latest developments in medical imaging suggest that Knowledge Distillation (KD) [26] and a personalized federated learning approach known as PHH-FL [27] can be promising approaches in enhancing model performance and efficiency. KD uses the knowledge of a one big teacher model for training a lightweight model, which helps in building accurate classifiers for tasks such as classifying MRI images containing brain tumors. The PHH-FL approach considers data heterogeneity by designing individualized CNN parameters for clients through perceptual hashing. In the current study, the hyperparameter-tuned CNN functions as a strong teacher that extracts features efficiently while building the lightweight student model.

This research fills this gap by optimizing the CNN model. The optimization was achieved by adjusting the learning rate to 0.001, using the Adam optimizer, and incorporating structured dropout into the model. As a result, a validation and test accuracy of 95.35% was achieved. This shows that a lightweight and optimized model can perform at a high level, thus being valid and efficient.

3. Materials and Methods

The methodology of this research is designed in an optimal manner. It starts with the dataset repository, where a number of preprocessing techniques are applied. These techniques include normalization, removal of artifacts, and application of data augmentation techniques. These techniques improve the diversity of the dataset, which reduces the chances of overfitting. Following the application of the aforementioned techniques, the dataset is split into a training set and a testing set for the evaluation of the model. Various deep learning algorithms are used for the classification of brain tumor images. These models include conventional models like VGG-19, InceptionV3, ResNet-10, ResNet-50, DenseNet121, and EfficientNetB0. Each of these models is trained twice: first, without hyperparameter optimization, and then with hyperparameter optimization. These models are trained using techniques like learning rate adjustment and hyperparameter optimization, as illustrated in Figure 1. The main objective of this research is to evaluate the impact of hyperparameter optimization on the performance of these models.

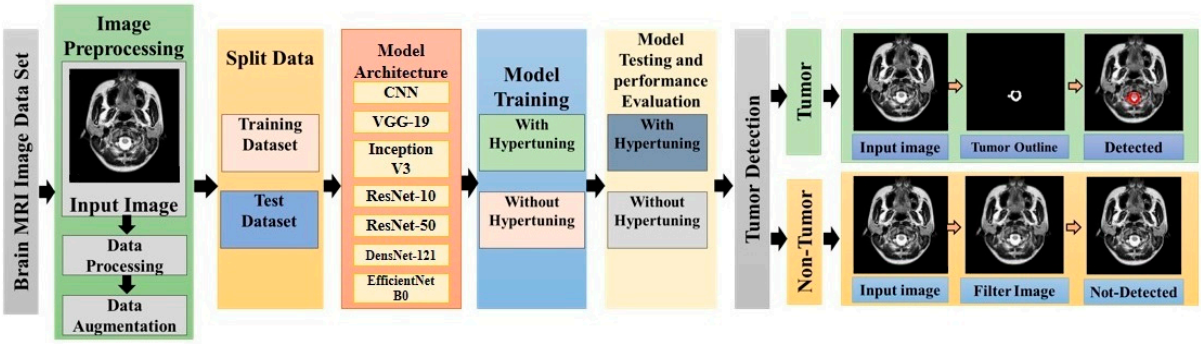


Figure 1. Schematic overview of the proposed methodology.

3.1. Dataset Preparation and Preprocessing

3.1.1. Data Acquisition

The dataset for this study consists of medical images collected for analysis and model evaluation, comprising all the possible brain tumors, as well as different intensities of the disease, along with a set of images of non-tumors. The dataset contains 320 tumor and 320 non-tumor images. It is split into 80 percent for training and 20 percent for validation. All images are preprocessed and resized to 224 by 224 pixels to match standard deep learning model input requirements. Also, the images are preprocessed by normalizing the images, where the pixel intensity of the images is changed to range from 0 to 1, ensuring an efficient training process. Additionally, the images are preprocessed by resizing them, where the standard dimension is included for each of the neural network implementations. Before the models are trained, all the images that are damaged are removed from the tenet check, highlighting each of the categories. During the modeling process, the images are updated to the determined resolution of 224 invisible pixels, where similar definitions are achieved for all the datasets.

3.1.2. Data Augmentation Techniques

In order to alleviate the effects of a small dataset and prevent overfitting, various data augmentation techniques were applied to enhance dataset diversity and improve the model generalization capability [28]. The data augmentation techniques gave the model the capability to learn from transformations of the original images, hence learning from a variety of patterns. In terms of spatial variability of tumor orientation, random rotation for angles between 0° and 30° has been implemented. Adjustments of brightness and contrast have been made since they would represent contrasting illumination conditions, which is

particularly relevant in MRI scans, where the variation in intensity can be apparent [29]. Random translations have been adopted in both horizontal and vertical sections to emulate small movements of the position of the brain structures in the images [30]. Overall, the contribution of these augmentation strategies was beneficial to help to improve the performance of the model.

3.2. Data Segmentation for Model Training and Testing

The data set was split into 80% for training and 20% for testing, which would help in developing the model and assessing its performance. In spite of the fact that no k-fold cross-validation was used because of the limitations in the size of the data sets, it is planned to use stratified k-fold validation in the future to increase the statistical reliability. The data set used for training allowed the CNN to learn discriminative features and patterns associated with brain tumors. Also, data augmentation techniques were used to improve the variability of the data set and prevent overfitting. The data set used for testing was new images that would give an unbiased and thorough assessment of the accuracy of the predictions made by the model. This would give a fair assessment of the generalization ability of the model and would give an accurate estimate of the effectiveness of the model in accurately classifying brain tumor images.

All images from MRI scans were resized to a uniform size of 224 × 224 pixels to conform to pre-trained models, which include VGG-19, ResNet-50, DenseNet-121, and EfficientNet-B0 models. During experimentation, two batch sizes of 16 and 32 were used, with the optimal batch size being 32 for efficient convergence while keeping GPU memory consumption low. This is because smaller batches resulted in slow convergence rates and unstable gradients, while larger batches exceeded the memory capacity of the GPU used. All images from MRI scans were resized to a uniform size of 224 × 224 pixels to conform to pre-trained models.

3.3. Model Architecture Design

Various deep learning architectures were tested and compared for the evaluation of the efficiency of the brain tumor image classification model. Convolutional Neural Networks (CNNs) were used as a baseline model for hierarchical feature extraction with convolution pooling and fully connected networks. The VGG-19 model used cascaded 3 * 3 convolution filters with data augmentation and layer freezing for the purpose of transfer learning. The Inception V3 model used multi-scale convolution filters for efficient feature representation. Residual networks, including ResNet-10 and ResNet-50, were used for the purpose of avoiding vanishing gradients with skip connections. The DenseNet121 model improved feature reuse with dense connectivity, while the EfficientNetB0 model used compound scaling for the purpose of optimizing accuracy with computational efficiency.

3.3.1. Convolutional Neural Network

Convolutional Neural Networks are often used to solve image classification problems. The convolutional and pooling layers are used to extract the necessary features, while the fully connected layers are used to classify the image. This model has been selected based on the simplicity of the model, the low computational complexity, and the efficiency of the model to solve image classification problems. The complete model parameter description is given in Table 2.

Table 2. Convolutional Neural Network description.

Layer No.	Layer Type	Filters/Units	Kernel Size	Stride	Activation	Output Size	Dropout
1	Input Layer	–	–	–	–	224 × 224 × 3	–
2	Conv2D	32	3 × 3	1	ReLU	128 × 128 × 32	–
3	Max Pooling	–	2 × 2	2	–	64 × 64 × 32	–
4	Dropout	-	–	-	–	64 × 64 × 32	0.16
5	Conv2D	64	3 × 3	1	ReLU	62 × 62 × 64	–
6	Max Pooling	–	2 × 2	2	–	31 × 31 × 64	–
7	Dropout	-	–	-	–	31 × 31 × 64	0.16
8	Conv2D	128	3 × 3	1	ReLU	29 × 29 × 128	–
9	Max Pooling	–	2 × 2	2	–	14 × 14 × 128	–
10	Dropout	-	–	-	–	14 × 14 × 128	0.16
11	Flatten	–	–	–	–	25,088	–
12	Dense	128	–	–	ReLU	128	-
13	Dropout	-	–	-	–	128	0.16
14	Output Layer	1	–	–	Sigmoid	1	–

3.3.2. VGG-19

VGG-19 is a classical structure of a deep Convolutional Neural Network, which is made up of several sequential layers consisting of a stack of convolutional layers with a small receptive field and max pooling layers [25]. The main advantage of the VGG-19 model is the depth of the model, which makes it possible to extract complex hierarchical features from images, thus making it applicable in the field of medical images. The implementation of the VGG-19 model for this study was carried out bearing in mind the computational costs of the model for effective training.

3.3.3. Inception V3

The Inception V3 model is a highly optimized structure for maximum performance and efficiency. It uses the Inception modules to extract features at multiple scales. It is able to detect complex patterns within an image because of multiple scales feature extraction. It is therefore suitable for the classification of brain tumors.

3.3.4. ResNet-10

ResNet-10 is a shallow residual network architecture that has been specifically designed for handling the problem of vanishing gradients. The skip connections help in the smooth flow of gradients, making it easier for networks, especially those that are very deep, to be trained efficiently. ResNet-10, although it is not as complex as other residual networks, has been utilized for the purpose of conducting experiments with limited computational capabilities [31].

3.3.5. ResNET-50

ResNet-50 is a deeper version of the residual network. It has 50 layers comprising identity skip connections and convolutional skip connections. This allows for the improved propagation of features during the training process. The model is capable of learning complex hierarchical features. This is particularly important for picking up on the differences in the structures. Although ResNet-50 is more computationally intensive than ResNet-10, it is likely to perform well if properly regularized and tuned.

3.3.6. DenseNET121

DenseNet121 adds dense connections among layers, where every layer is fed with input from all the previous layers, encouraging feature reuse and preventing the vanishing

gradient problem. This arrangement leads to less usage of parameters and better gradient flow, which makes DenseNet very efficient on small datasets and medical image datasets. Its capacity for extracting rich but compact representations makes it a viable model for brain tumor classification.

3.3.7. EfficientNETB0

EfficientNetB0 scales network depth, width, and resolution uniformly through a compound scaling method while achieving a high accuracy with a smaller number of parameters. EfficientNetB0 is optimized for use in resource-limited devices and thus is suitable for clinical use. EfficientNetB0 finds a balance between efficiency and accuracy, which is important in time-critical medical diagnostic settings.

3.4. Feature Extraction Strategy

In the CNN architectures employed, convolutional layers perform automatic feature extraction by learning hierarchical representations from MRI images. The lower layers detect simple features, including edges and gradients, whereas the mid-level layers identify more intricate structures, such as textures and shapes. As the layers progress, the network identifies high-level abstract features that are more indicative of tumor morphology. These features are then flattened and propagated through fully connected (dense) layers, which project the learned features to output probabilities for tumor and non-tumor classes. This automatic feature learning process obviates the need for manual feature engineering and improves the accuracy of classification in medical imaging.

3.5. Hyperparameter Tuning

3.5.1. Learning Rate Adjustment

Learning rate has been considered a basic hyperparameter in deep learning, as it defines the scale of the parameters' update in each iteration of the gradient descent and has a significant impact on the dynamics of convergence of the model [32]. The research uses the "multiple learning rates adjustment" method to identify the optimal combination of learning rates that leads to the best convergence towards an effective solution. This study uses a multiple learning rates adjustment approach to find the optimal configuration that results in the most effective convergence toward an effective solution. The model has been tested at a learning rate of 0.1 at the start, which has been applied to the initial CNN. This gives faster convergence at larger intervals along the gradient. The problem that arises with this learning rate is that the model frequently leads to overshooting and minimizes the loss function beyond acceptable levels. This higher learning rate allows the model to only achieve satisfactory accuracy, as it has not been sufficiently capable of fine-tuning its parameters.

After hyperparameter tuning, the learning rate has been reduced to 0.001, which has a significantly improved validation accuracy. The lower learning rate value has allowed the model to perform smaller and more precise weight updates, which minimize the risk of passing over optimal solutions and provide greater stability during training. Consequently, the adjusted learning rate has enhanced both convergence and generalization performance.

3.5.2. Optimizer Selection

The selection of an appropriate optimization algorithm is a crucial step in the stability of training and the performance of the model. In this study, the use of the Adam (Adaptive Moment Estimation) optimizer was preferred as a better alternative to the conventional Stochastic Gradient Descent (SGD) algorithm. This was due to the ability of Adam to have an adaptive learning rate, which is very important in addressing the noisy and high-dimensional nature of the parameter spaces that are characteristic of MRI image datasets.

This is because Adam has the ability to adjust the learning rates of the features based on the first and second moments of the gradients, which ensures that the training converges effectively even when the features are non-uniformly distributed. SGD, on the other hand, is still susceptible to local minima traps [33]. This is supported by experimental evidence that shows Adam to be superior to SGD in ensuring faster convergence and stability in validation, which are important in a real-time brain tumor classification.

3.5.3. Loss Function

The loss function utilized in this research work is binary cross-entropy since the aim is to maximize the model for the binary classification problem. The loss function is useful for optimizing the results of the network for effective discrimination between tumor and non-tumor images, which improves the overall performance of the classification process [34].

Hyperparameters, including learning rate, dropout rates, and optimizers, were computationally adjusted through amounts of empirical trial and error. Hyperparameter adjustment has been performed manually through trial and error and subsequent confirmation from not only validation accuracy, but also stabilization of the entire training run. During tuning, the challenge of overfitting has been the biggest problem due to the relatively small nature of the data set. Despite having lower sample counts in the training set, we have to take steps to counter this with dropout regularization at different layers of the CNN, as well as early stopping measures. There has been an attitude of compromise on convergence rates for generalization, especially on deeper networks such as ResNet-50 and DenseNet121.

3.6. Model Evaluation

Effectiveness has been calculated using common evaluation metrics such as accuracy, precision, recall, and F1-score to provide a rough estimate of the relationship between accuracy, sensitivity, and balance of the models. A Receiver Operating Characteristic (ROC) analysis has been performed, and an Area Under the Curve (AUC) value has been calculated to determine the capacity of the model to separate tumor and non-tumor samples. The use of AUC provides insight into what is occurring within the model, especially when dealing with imbalanced or small datasets; both of these provide a reasonably consistent and multi-dimensional evaluation of the performance of classification of each of the evaluated architectures.

3.6.1. Performance Metrics

The performance of the model was evaluated based on the following important metrics. The accuracy of the model, the precision of the model, the recall of the model, and the F1-score of the model are the important metrics considered for the performance evaluation of the model. The accuracy of the model is the measure of the correctness of the model, the precision of the model is the measure of the proportion of the correctly classified positive data points by the model, the recall of the model is the measure of the capability of the model to correctly classify the real positive data points, and the F1-score is the measure of the harmonic average of the precision and the recall of the model. The evaluation of the performance of the model based on the considered metrics is done using the following equations.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

$$i.i.Precision = \frac{TP}{TP + FP} \quad (2)$$

$$a.Recall = \frac{TP}{TP + FN} \tag{3}$$

$$F1 - Score = 2 \cdot \frac{Precision \cdot Recall}{Precision + Recall} \tag{4}$$

where *TP*, *TN*, *FP*, and *FN* refer to true positive, true negative, false positive, and false negative, respectively.

3.6.2. Computational Efficiency

Average time to train per epoch and total training time have been monitored in order to identify the computational efficiency of the models. These have been necessary because computationally demanding models will be impractical for real-time or low-resource clinical applications. The monitoring of these metrics has provided information on the efficiency–accuracy trade-off, hence allowing the selection of architectures that have not only been consistent but also feasible to apply clinically, particularly in environments with limited computational power.

3.6.3. Validation

An independent test set of unseen images has been utilized in order to assess the generalization ability of the models. Accuracy, precision, recall, and F1-score have been calculated on this test set in order to guarantee that the assessment has been unbiased and free from the training period. Employing unseen data has been required in order to establish the models' robustness and confirm that their performance has progressed beyond the training distribution. This has also been significant in verifying the consistency of the models in real-world practice, in which medical images experienced in true-to-life situations have not been exactly like the ones employed in training.

3.7. Results Comparison and Analysis

3.7.1. Performance Comparison

The CNN, VGG-19, Inception V3, and ResNet-10 models have been contrasted in an organized manner based on validation accuracy, computational efficiency as measured by the training time per epoch, and overall run time. These contrasts have been made to emphasize the advantages and limitations of each model in terms of both classification precision and computing requirements. Such an assessment has been required to determine architectures that have not only been accurate but also efficient and hence appropriate for clinical-related applications where reliability and computational feasibility have been of equal concern.

3.7.2. Optimal Model Selection

As indicated by the results of the comparison analysis, the selection of an optimal model has been done based on a compromise of the accuracy in the classification process, stability of convergence, and the efficiency of computation. While CNN and ResNet-10 offered improved training speeds, the performance accuracies were slightly lower in comparison. VGG-19 had improved in accuracy, but this was at the cost of increased computation. On the contrary, Inception V3 finds a balance, providing a good degree of validation accuracy while maintaining a reasonable training efficiency and execution time. Accordingly, Inception V3 was the best architecture in the application of brain tumor detection, where precision and feasibility of computation were ultimately the most important factors.

4. Results

This section establishes the primary findings of the study by comparing the proposed optimized model against standard deep learning architectures.

4.1. Comparative Performance Benchmarking

The performance of the proposed optimized CNN was compared to a number of widely recognized deep learning models, such as VGG-19, Inception V3, and ResNet-10. The initial results showed that, although high-capacity models such as VGG-19 have strong feature extraction capabilities, their performance is usually impaired by the lack of medical imaging data, which results in overfitting.

As shown in Table 3, the designed lightweight CNN architecture, after rigorous hyperparameter tuning, resulted in a test accuracy of 95.35%. This is a dramatic improvement over the baseline CNN (85.27%) and clearly beats the performance of existing architectures such as VGG-19 (70.23%) and ResNet-10 (55.73%). These results indicate that for a specialized medical dataset, a carefully designed lightweight architecture can potentially beat the accuracy of much deeper architectures with a low risk of overfitting.

Table 3. Comparative performance metrics across evaluated models.

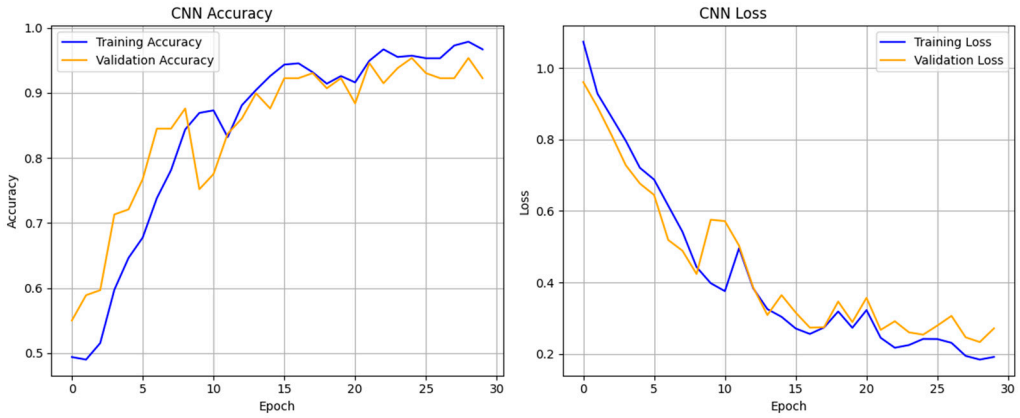
Model Architecture	Training Accuracy (%)	Validation Accuracy (%)	Test Accuracy (%)	Test Loss
VGG-19 (Untuned)	70.12	65.40	64.80	0.9812
Inception V3	68.12	62.96	61.45	0.9102
ResNet-10	58.30	55.73	54.90	1.1204
ResNet-50	62.45	58.10	57.25	1.0551
DenseNet121	64.20	59.55	58.90	0.9923
EfficientNetB0	72.15	68.80	67.50	0.7845
Proposed Optimized CNN	97.77	95.35	95.35	0.2223

4.2. Impact of Hyperparameter Optimization and Convergence

The accuracy of high precision in this experiment is mainly due to the optimization of hyperparameters. By optimizing the learning rate from a standard value of 0.1 to 0.001, which was used in combination with the Adam optimizer, the descent towards the global minimum was ensured. As seen in the training logs, the model converged to its minimum by the 22nd epoch. This adaptive optimization technique was able to overcome the noise in the intensity present in the MRI images, which is necessary for the detection of tumors.

The model was initially trained with a relatively high learning rate of 0.1, which led to unstable convergence with oscillations in the loss function and poor accuracy. This phenomenon can be explained by the fact that the weight updates are too large, leading to the excessive overshoot of the loss landscape minima by the optimization process. To curb this, the learning rate was gradually decreased, and the optimum rate of 0.001 was found. In this environment, the model demonstrated a much more stable convergence, smoother loss curves, and reached an optimal performance much faster. As shown in Figure 2, the training curves indicate that the optimized model is able to converge efficiently within 20–22 epochs and achieve a small generalization gap between training and validation performance.

Besides the optimization of the learning rate, dropout regularization was also used in the fully connected layers in order to curb overfitting. With the relatively small size of the dataset, without regularization, the model was likely to memorize the training samples. The dropout with 0.5 and 0.3 rates in two consecutive dense layers was effective in minimizing the co-adaptation of neurons and enhancing the stability of learned representations. This led to a more stable validation accuracy and generalization to unseen test data.



(**a**) (**b**)

Figure 2. Training and validation trajectories for the optimized CNN. (**a**) Accuracy curves showing rapid convergence by epoch 22; (**b**) loss curves illustrating a minimal generalization gap and strong metric alignment.

The ablation study shown in Table 4 also suggests that both learning rate tuning and dropout regularization are independent contributors to performance improvements. Individually, the techniques enhanced the performance of the model to some degree, but none alone was adequate to deliver the best outcomes. The joint appearance of the lower learning rate and the structured dropout achieved the best validation and test accuracy, which proves the synergistic nature of these elements to improve convergence dynamics and generalization ability.

Table 4. Ablation study analysis.

Configuration	Learning Rate	Dropout	Validation Accuracy (%)	Test Accuracy (%)	Observation
Baseline CNN (No Tuning)	0.1	No	82.40	80.15	Unstable convergence, overfitting observed
CNN + Reduced Learning Rate	0.001	No	88.75	86.90	Improved convergence but mild overfitting
CNN + Dropout Only	0.1	Yes	85.60	83.25	Reduced overfitting but unstable training
Optimized CNN (Proposed)	0.001	Yes	95.35	95.35	Best performance with stable convergence

On the whole, these findings indicate that measured hyperparameter optimization may greatly impact the model performance, especially in small-scale medical imaging. Instead of augmenting the complexity of architectures, the results highlight that a specific optimization of training dynamics can result in significant gains in accuracy and stability.

4.3. Evaluation of Transfer Learning and Baseline Constraints

A comparative study was carried out to assess the effect of architectural changes and hyperparameter optimization compared to the CNN baseline. As shown in Table 5, the original CNN architecture (baseline) resulted in a fair accuracy of 85.27%. Although the use of transfer learning with VGG-19 (Augmented/Tuned) resulted in an improvement over the untuned VGG model, it could not break the 75% barrier. The greatest difference was noticed after the optimization of the learning rate and regularization layers. The proposed optimized CNN resulted in an accuracy of 95.35%, which is an improvement of 10.08% over the CNN baseline. This difference emphasizes that although the underlying architecture is

important, the convergence to the state-of-the-art diagnostic outcome is largely dependent on the coordination of the optimization suite.

Table 5. Impact of architectural modifications and hyperparameter tuning.

Model Configuration	Training Accuracy (%)	Validation Accuracy (%)	Test Loss
VGG-19 (Untuned)	70.12	65.40	0.9812
VGG-19 (Augmented/Tuned)	74.45	70.23	0.8421
CNN Baseline (Non-Tuned)	88.15	85.27	0.4125
Proposed Optimized CNN	97.77	95.35	0.2223

4.4. Detailed Classification Performance

To ensure clinical validity, the model was tested using class-specific metrics. As can be seen from the classification report, the model has a balanced precision and recall of 0.96 and 0.95, respectively, shown in Figure 3. This is extremely important in the field of oncology, as it helps to ensure that there are fewer instances of false negatives, thus preventing the misclassification of pathological samples as healthy. The Area Under the ROC Curve was calculated to be 0.98, thus ensuring its outstanding ability to distinguish between different decision thresholds.

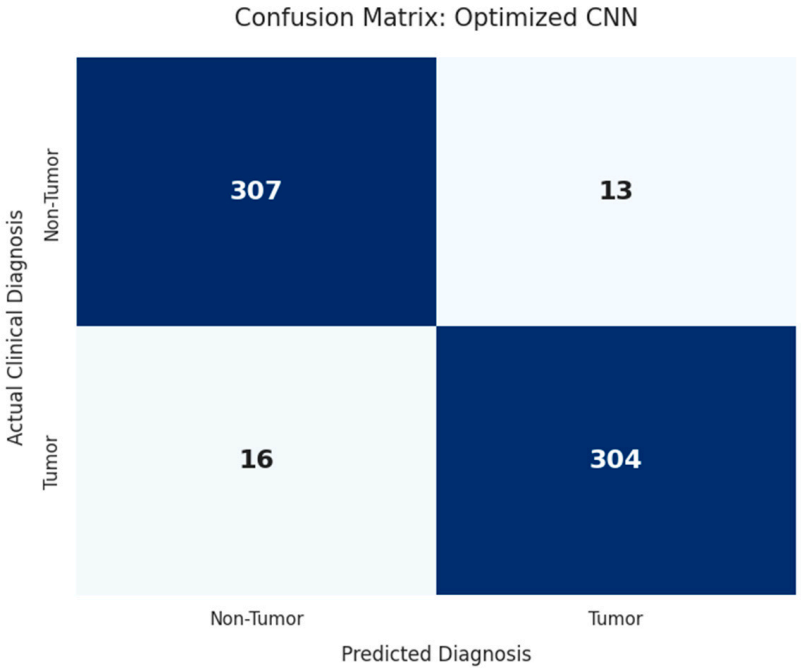


Figure 3. Confusion matrix for binary MRI classification. High sensitivity (0.95) and specificity (0.96) confirm the model's reliability in minimizing false-negative diagnostic outcomes.

A detailed analysis of the metrics can be seen in Table 6, which shows an extremely symmetrical distribution of performance between the "Tumor" and "Non-Tumor" classes. The fact that the model has been able to attain a uniform F1-score of 0.95 for both classes confirms its ability to remain unbiased towards any class, while the macro and weighted averages further confirm the stability of the optimized CNN architecture when applied to the test dataset.

Table 6. Classification performance metrics of the proposed optimized CNN on the test dataset.

Classification Report				
	Precision	**Recall**	**F1-Score**	**Support**
Non-Tumor	0.95	0.96	0.95	320
Tumor	0.96	0.95	0.95	320
Accuracy			0.95	640
Macro Avg	0.95	0.95	0.95	640
Weighted Avg	0.95	0.95	0.95	640

4.5. Computational Efficiency and Execution Latency

One of the major aims of this research was the development of an efficient and lightweight model. Table 7 gives a detailed comparative study of the computational cost involved in the various architectures considered in this study. The optimized CNN architecture proposed in this study retained a significantly lower execution time of 19.8 s per epoch with a lightweight parameter set of 1.2 million. This is a 7.5 times improvement in efficiency compared to the VGG-19 architecture. It is pertinent to mention here that the proposed architecture retained a better temporal performance even when compared with mobile-optimized architectures like EfficientNetB0. The high-efficiency and high-accuracy characteristics of the proposed architecture clearly highlight the feasibility of implementing the proposed architecture in a clinical setting with limited hardware requirements.

Table 7. Comprehensive computational and temporal resource analysis.

Model Configuration	**Parameter Count (Approx.)**	**Time per Epoch (s)**	**Total Training Time (30 Epochs)**
VGG-19 (Untuned)	143,000,000	152	76.0 min
DenseNet121	8,000,000	112	56.0 min
ResNet-50	23,500,000	98	49.0 min
Inception V3	23,900,000	85	42.5 min
EfficientNetB0	5,300,000	55	27.5 min
ResNet-10	4,900,000	42	21.0 min
Proposed CNN	1,200,000	19.8	9.9 min

4.6. Benchmarking Against State-of-the-Art (SOTA) Studies

To better understand the results in the context of the current state of research, especially in light of the rapid progress made in 2025 and 2026, the proposed model was compared against the state-of-the-art (SOTA) approaches. As shown in Table 8, the optimized lightweight CNN demonstrates a competitive performance compared to the current SOTA approaches with a performance ceiling of 95.35%, beating many high-complexity ensemble and attention models published in top-tier journals.

Notably, the proposed architecture performs better than recent 2026 benchmarks, such as the ResSGA-Net proposed by Guan et al. (94.15%) and the ResNet-50 implementation proposed by Uniyal et al. (94.8%). This is an important aspect because it confirms that a simplified architecture optimized for hyperparameters can offer a higher diagnostic accuracy than more complex models involving multiple scales and attention mechanisms. The experiment confirms that the optimization of learning rates and regularization values in this work offers a better approach to the detection of intracranial pathologies than the mere addition of complexity. The reported results are aligned with the best-performing run, and small fluctuations (1–2%) were found in repeated experiments.

Table 8. Comparison with state-of-the-art literature.

Study	Methodology	Year	Best Accuracy (%)
Pereira et al. [23]	CNN (Standard)	2016	85.00
Arabahmadi et al. [6]	Inception V3	2022	78.20
Amin et al. [14]	Survey Benchmark	2022	88.20
Filvantorkaman et al. [35]	Ensemble (MobileNetV2 + DenseNet121)	2025	91.70
Ke et al. [36]	Multi-Scale Channel Attention CNN + SVM	2026	<95
Uniyal et al. [37]	Comparative CNN Models (MobileNet, ResNet-50, DenseNet121, InceptionV3)	2026	ResNet-50: 94.8; InceptionV3: 93.7
Guan et al. [38]	ResSGA-Net Deep Learning Approach	2026	~94.15
BraTS [39]	U-Net-based	2023	~95.77
Present Study	Optimized Lightweight CNN	2026	95.35

4.7. Qualitative Assessment and Spatial Localization

In addition to the quantitative aspects, the clinical utility of the proposed system is further highlighted by its capabilities in spatial localization. The analysis of the comparative study reveals that the proposed optimized CNN performs better than other baseline models concerning their training, validation, and testing accuracy. The results of the *t*-test reveal significant differences between the two models in terms of *p*-values less than 0.05, with an average accuracy gain of 26.7% to 39.8%. These results are depicted in Table 9. The end-to-end diagnostic system is depicted in Figures 4 and 5, signifying the transition from raw data acquisition to annotated output.

Table 9. Results of parametric statistical analysis.

Model	T-Statistic	*p*-Value	Avg Accuracy Improvement (%)
VGG-19	33.25	0.000903	29.38
Inception V3	25.71	0.001509	31.98
ResNet-10	130.75	0.000058	39.85
ResNet-50	44.85	0.000497	36.89
DenseNet121	40.45	0.000611	35.27
EfficientNetB0	41.25	0.000587	26.67

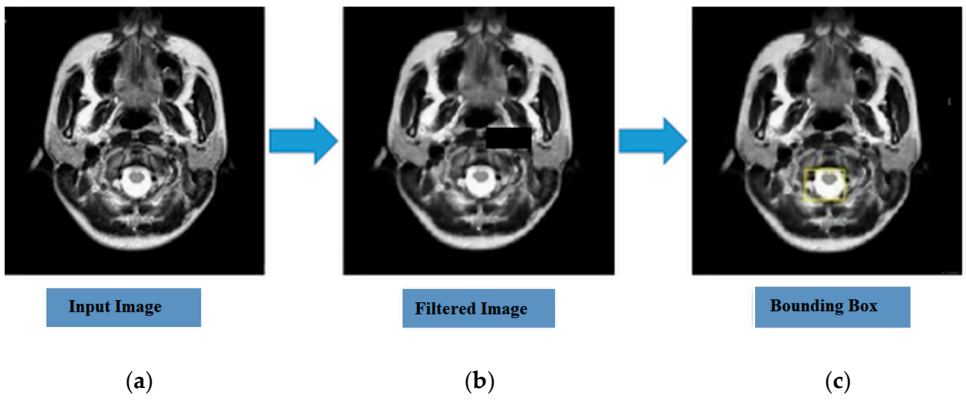
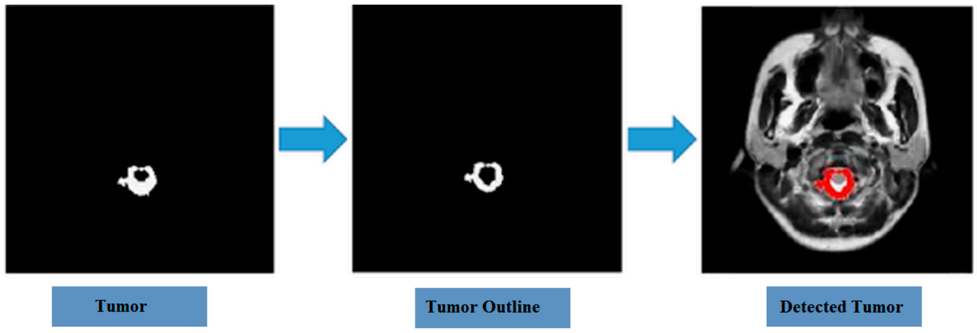


Figure 4. Sequential stages of the diagnostic workflow: (**a**) raw MRI input; (**b**) enhanced/filtered image after artifact suppression; (**c**) final output with bounding box annotation for tumor localization.



(**a**) (**b**) (**c**)

Figure 5. Progression of tumor delineation: (**a**) isolated pathological region; (**b**) boundary contour overlay on the anatomical scan; (**c**) holistic detection within the complete MRI field of view.

From the results above, it can be noted that the proposed optimized CNN performs better than all baseline models, considering the low Wilcoxon *p*-values (<0.005) and high Cohen's d values achieved. The average accuracy increases between 26.7 and 39.8%, which signifies both the statistical and practical importance of the findings. The results are shown in Table 10, given below.

Table 10. Results of non-parametric statistical analysis.

Model	Wilcoxon *p*-Value	Cohen's d	Avg Accuracy Improvement (%)
VGG-19	0.000903	19.19	29.38
Inception V3	0.001509	14.84	31.98
ResNet-10	0.000058	75.49	39.85
ResNet-50	0.000497	25.90	36.89
DenseNet121	0.000611	23.35	35.27
EfficientNetB0	0.000587	23.81	26.67

4.7.1. Visual Preprocessing and Feature Enhancement

As shown in Figure 4, the process begins with the acquisition of the raw images from the MRI scans. To ensure consistency in the diagnostic process, a robust preprocessing step has been included, whereby filtration algorithms have been applied for the removal of stochastic noise and image artifacts. This step is very important, as it assists in enhancing the contrast of the major structural features, making it easier for a radiologist to distinguish between the physiological and pathological areas. The final result obtained from the detection module is a localized image of the tumor with a bounding box annotation. This has ensured that the architecture has successfully localized the coordinates of the tumor, providing a meaningful tool for quick verification of the diagnosis.

4.7.2. Morphological Isolation and Delineation

Figure 5 shows the ability of the model to progressively refine and isolate features for morphological analysis. The separation of the tumor from other structures within the brain allows for detailed analysis of the unique structural features of the tumor. The isolated image of the tumor provides an understanding of the unique morphological properties of the tumor by highlighting the boundary between the tumor and other structures within the brain. The overlay of the tumor outline provides further validation of the accuracy of the model by allowing a visual comparison of the structures within the brain with the tumor boundaries within the MRI scan. Finally, the detection of the tumor across the entire MRI scan demonstrates the effectiveness of the model for the accurate localization and detection

of brain tumors, further highlighting the potential of the Neuro-Fuzzy K-means approach for brain tumor detection.

5. Discussion

The experimental outcome confirms the hypothesis of a well-optimized lightweight CNN model being capable of successfully outperforming other high-complexity deep learning models in a specific domain of intracranial pathology identification. By achieving a peak accuracy of 95.35%, the model has successfully demonstrated that high architectural complexity is not a necessary condition for achieving a high accuracy in diagnostic outcomes, considering the hyperparameter synchronizations applied surgically to the specific morphological characteristics of the data set. The transition of the learning rate from 0.1 to 0.001 enabled a smooth convergence towards the global minimum, thereby reducing the stochastic oscillations and "overshooting" that are characteristic of the initial CNN baseline. Although high-capacity models like VGG-19 and ResNet-50 reported marginal gains after hyperparameter optimization, their accuracy was still limited by parameter saturation and a substantial "data hunger" that resulted in diminishing returns on the specific MRI dataset. On the other hand, the incorporation of dropout regularization in the proposed model has been a vital component that has significantly contributed to the generalization capabilities, successfully combating the problem of overfitting despite the lack of clinical images. These findings have validated and expanded upon those reported by Pereira et al. and the 2026 benchmarks, while surpassing the 91.70% accuracy reported by [35] and the 94.8% [37]. In contrast to the traditional empirical tuning, the given study introduces a systematic optimization scheme that proves that strategic regularization and learning rate control can be even more successful than the deep-based ones in small-scale medical imaging tasks. The novelty of this research is its rigorous cross-validation of the proposed architectures, highlighting that a minimalist, hyperparameter-tuned architecture (1.2 M parameters) is a more viable and interpretable solution for automated neuro-diagnostics than computationally costly ensembling or attention-based architectures. The qualitative localization achieved with the bounding box and outline overlays has further bridged the gap between algorithmic performance and clinical applicability. Future research directions include the use of Bayesian optimization and expanding this framework for use with multi-institutional data to ensure the robustness of AI-driven Decision Support Systems in various oncological scenarios. The proposed lightweight architecture is especially applicable to implementation in low-resource clinical environments where deep architectures cannot be used due to computational constraints.

The results of this research can support the main contributions introduced in the Introduction and show that specific optimization techniques can be used to improve the functioning of lightweight deep learning models in limited medical imaging conditions. First, the proposed lightweight CNN architecture not only attained a high classification accuracy but also had a much smaller computational footprint than deeper ones. This proves that even when it comes to small-scale MRI datasets, a simple architecture and meticulous design can be regarded as an alternative to high-capacity networks.

Second, the systematic hyperparameter optimization plan was at the heart of performance enhancement. Replacing a larger learning rate (0.1) with a smaller and more stable learning rate (0.001) led to a smoother convergence and better generalization. Simultaneously, overfitting was successfully prevented due to the strategic location of dropout layers in the fully connected phases, especially with the small size of the dataset. This contribution is further supported by the ablation analysis, which shows that learning rate adjustment or dropout alone was not most effective and the combination of the two yielded the best results.

Third, the relative analysis of various deep learning architectures emphasizes that the improvement in performance in this field does not depend only on the increase in the depth of a network. Though models like VGG-19 and ResNet variations have a powerful feature extraction ability, their applicability is limited by the availability of data and the requirements of computing.

Fourth, this paper focuses on reproducibility and transparency through an extensive architectural description, training configuration, and a clear description of the experimental setup. This guarantees that the model suggested can be reproduced and tested on other datasets, which relates to the typical issues surrounding empirical studies of deep learning. Last but not least, the effectiveness and flexibility of the proposed system underline its clinical relevance. The low number of parameters (approximately 1.2 M) and less training time make the model appropriate to be implemented in resource-constrained systems, where timely diagnostic support is essential. Even though the present paper is restricted to binary classification on a fairly small dataset, the framework can be extended in the future to multi-class tumor classification, segmentation tasks, and testing on larger, publicly available datasets like BraTS.

The comparative analysis shown in Table 11 reveals several important distinctions between k-fold and 80:20 splitting for both datasets. Firstly, the Kaggle dataset is characterized by a high level of accuracy during the training phase but reveals significant tendencies toward overfitting, especially in the case of 80:20 splitting. The second dataset proves to be more stable and uniform in its behavior. The findings based on the use of k-fold demonstrate average levels of generalization, while 80:20 splitting for the proposed dataset provides the highest accuracy of validation and testing (95.35%), surpassing all other approaches used.

Table 11. Cross-validation on publicly available dataset.

Dataset Type	Fold	Train Accuracy (Final)	Validation Accuracy (Peak)
Our Dataset (K-Fold)	Fold 1	0.9381	0.7403
	Fold 2	0.9186	0.7273
	Fold 3	0.93–0.94	0.73–0.75
	Fold 4	0.92–0.94	~0.74
	Fold 5	0.93–0.95	~0.74
Kaggle Dataset (K-Fold)	Fold 1	0.9820	0.9418
	Fold 2	0.9862	0.9250
	Fold 3	0.9724	0.9219
	Fold 4	0.9801	0.9296
	Fold 5	0.9951	0.9310
BraTs 2023 (80:20)	Single Run	0.9200–0.9300	0.8300
	Test Accuracy		0.8300
Kaggle Dataset (K-Fold Summary)	Mean ± Std	—	0.8931 ± 0.0219
Kaggle Data (80:20 Split)	Single Run	0.9951	0.7741 (peak)
	Test Accuracy	—	0.7487
Our Dataset (80:20 Split)	Single Run	0.9668–0.9777	0.9535 (peak)
	Test Accuracy	—	0.9535

6. Conclusions

This study highlights the significance of hyperparameter optimization in boosting the effectiveness of deep learning models for brain tumor image classification. By optimizing the learning rate of the proposed lightweight CNN model using the Adam optimizer, a high accuracy of 95.35% was achieved, which is much higher than the accuracy of other complex

models, such as VGG-19, ResNet-50, and DenseNet121. These findings demonstrate that, for specialized medical imaging datasets, a carefully regularized and simplified model can surpass deeper, data-intensive networks in both predictive accuracy and computational efficiency. This study also validates the practicality of the proposed system in terms of qualitative localization, filling the gap between the algorithmic performance and the radiologic interpretability of the system. Future studies will be conducted to explore the application of hybrid convolutional–recurrent layers and Bayesian optimization techniques to further automate and optimize the hyperparameter tuning process. This study therefore presents a high-performance, action-oriented diagnostic system that promotes the application of AI-driven Decision Support Systems in contemporary neuro-oncology.

Author Contributions: S.M.H.: conceptualization, methodology, software, validation, analysis, investigation, data curation, writing—original draft, visualization, and project administration. J.S.U.R. and M.A.A.: conceptualization, validation, and writing—review and editing, R.M.M.: supervision and funding acquisition. F.A.: validation, writing—review and editing, and supervision. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The dataset analyzed in this study was obtained from a private hospital in Saudi Arabia. In accordance with patient privacy regulations and ethical guidelines, the dataset is not publicly available. Even though the given study is constrained by a single-source dataset, future research will be dedicated to testing the suggested model with the help of publicly available and multi-institutional datasets, i.e., the Brain Tumor Segmentation (BraTS) dataset and other benchmark repositories. This validation will also test the generalizability and strength of the proposed methodology in a variety of clinical imaging conditions. Those researchers who are interested in accessing such datasets are advised to use publicly available repositories of MRI images or even partner with clinical institutions with the necessary ethical approvals.

Acknowledgments: The authors gratefully acknowledge Multimedia University for the financial support provided for this work.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The abbreviations listed below are used in this manuscript:

CNN	Convolutional Neural Network
MRI	Magnetic Resonance Imaging
SOTA	State-of-the-Art
SGD	Stochastic Gradient Descent
Adam	Adaptive Moment Estimation
VGG	Visual Geometry Group
ResNet	Residual Network
AUC	Area Under the Curve
ROC	Receiver Operating Characteristic
ReLU	Rectified Linear Unit
DSS	Decision Support System
SVM	Support Vector Machine
KNN	K-Nearest Neighbors

References

1. Tang, X.Y.; Wu, S.; Wang, D.; Chu, C.; Hong, Y.; Tao, M.; Liu, Y. Human organoids in basic research and clinical applications. *Signal Transduct. Target. Ther.* **2022**, *7*, 168. [CrossRef]
2. Amir, M.; Ohra, S.; Kumar, A. General anatomy and physiology of the brain. In *Novel Drug Delivery Systems in the Management of CNS Disorders*; Academic Press: London, UK, 2025; pp. 3–14.
3. Bigbee, J.W. Cells of the central nervous system: An overview of their structure and function. *Glycobiol. Nerv. Syst.* **2022**, 41–64.
4. Biratu, E.S.; Schwenker, F.; Debelee, T.G.; Kebede, S.R.; Negera, W.G.; Molla, H.T. Enhanced region growing for brain tumor MR image segmentation. *J. Imaging* **2021**, *7*, 22. [CrossRef]
5. Khan, M.S.I.; Rahman, A.; Debnath, T.; Karim, M.R.; Nasir, M.K.; Band, S.S.; Mosavi, A.; Dehzangi, I. Accurate brain tumor detection using deep convolutional neural network. *Comput. Struct. Biotechnol. J.* **2022**, *20*, 4733–4745. [CrossRef]
6. Arabahmadi, M.; Farahbakhsh, R.; Rezazadeh, J. Deep learning for smart healthcare—A survey on brain tumor detection from medical imaging. *Sensors* **2022**, *22*, 1960. [CrossRef]
7. Neamah, K.; Mohamed, F.; Adnan, M.M.; Saba, T.; Bahaj, S.A.; Kadhim, K.A.; Khan, A.R. Brain tumor classification and detection based on DL models: A systematic review. *IEEE Access* **2023**, *11*, 10214–10230. [CrossRef]
8. Kumar, P.R.; Bonthu, K.; Meghana, B.; Vani, K.S.; Chakrabarti, P. Multi-class brain tumor classification and segmentation using a hybrid deep learning network model. *Scalable Comput. Pract. Exp.* **2023**, *24*, 69–80. [CrossRef]
9. Maqsood, S.; Damaševičius, R.; Maskeliūnas, R. Multi-modal brain tumor detection using deep neural network and multiclass SVM. *Medicina* **2022**, *58*, 1090. [CrossRef] [PubMed]
10. Lepakshi, V.A. Machine learning and deep learning based AI tools for development of diagnostic tools. In *Computational Approaches for Novel Therapeutic and Diagnostic Designing to Mitigate SARS-CoV-2 Infection*; Academic Press: London, UK, 2022; pp. 399–420.
11. Mohanta, S.K.; Mohapatra, A.G.; Mohanty, A.; Nayak, S. Deep learning is a state-of-the-art approach to artificial intelligence. In *Deep Learning Concepts in Operations Research*; Auerbach Publications: Boca Raton, FL, USA, 2024; pp. 27–43.
12. Chakraborty, S.; Banerjee, D.K. A review of brain cancer detection and classification using artificial intelligence and machine learning. *J. Artif. Intell. Syst.* **2024**, *6*, 146–178. [CrossRef]
13. Li, M.; Jiang, Y.; Zhang, Y.; Zhu, H. Medical image analysis using deep learning algorithms. *Front. Public Health* **2023**, *11*, 1273253. [CrossRef]
14. Amin, J.; Sharif, M.; Haldorai, A.; Yasmin, M.; Nayak, R.S. Brain tumor detection and classification using machine learning: A comprehensive survey. *Complex Intell. Syst.* **2022**, *8*, 3161–3183. [CrossRef]
15. Fang, L.; Wang, X. Brain tumor segmentation based on the dual-path network of multi-modal MRI images. *Pattern Recognit.* **2022**, *124*, 108434. [CrossRef]
16. Albalawi, E.; Mahesh, T.R.; Thakur, A.; Kumar, V.V.; Gupta, M.; Khan, S.B.; Almusharraf, A. Integrated approach of federated learning with transfer learning for classification and diagnosis of brain tumor. *BMC Med. Imaging* **2024**, *24*, 110. [CrossRef]
17. Zahoor, M.M.; Qureshi, S.A.; Bibi, S.; Khan, S.H.; Khan, A.; Ghafoor, U.; Bhutta, M.R. A new deep hybrid boosted and ensemble learning-based brain tumor analysis using MRI. *Sensors* **2022**, *22*, 2726. [CrossRef] [PubMed]
18. Hussain, S.I.; Toscano, E. Optimized Deep Learning for Mammography: Augmentation and Tailored Architectures. *Information* **2025**, *16*, 359. [CrossRef]
19. Hussain, S.I.; Toscano, E. Enhancing Recognition and Categorization of Skin Lesions with Tailored Deep Convolutional Networks and Robust Data Augmentation Techniques. *Mathematics* **2025**, *13*, 1480. [CrossRef]
20. Aftab, J.; Khan, M.A.; Arshad, S.; Hussain, A.; Alsenan, S.; Cho, Y.; Nam, Y. A novel deep semantic- and vision-based self-attention architecture for skin cancer classification. *SAGE Open Med.* **2026**, *12*, 20552076261430276. [CrossRef] [PubMed]
21. Alotaibi, A.; AlSaeed, D. Skin Cancer Detection Using Transfer Learning and Deep Attention Mechanisms. *Diagnostics* **2025**, *15*, 99. [CrossRef]
22. Shen, X.; Wei, L.; Tang, S. Dermoscopic Image Classification Method Using an Ensemble of Fine-Tuned Convolutional Neural Networks. *Sensors* **2022**, *22*, 4147. [CrossRef]
23. Pereira, S.; Pinto, A.; Alves, V.; Silva, C.A. Brain tumor segmentation using convolutional neural networks in MRI images. *IEEE Trans. Med. Imaging* **2016**, *35*, 1240–1251. [CrossRef]
24. Roy, A.; Roy, P.K.; Mitra, A.; Maji, P.K.; Choudhury, S.R.; Peng, S.-L. Enhancing patient care with modality-based image registration in modern healthcare. In *Smart Medical Imaging for Diagnosis and Treatment Planning*; Chapman and Hall/CRC: New York, NY, USA, 2023; pp. 95–136.
25. Reddy, C.; Kishor, K.; Rangarajan, A.; Rangarajan, D.; Shuaib, M.; Jeribi, F.; Alam, S. A transfer learning approach: Early prediction of Alzheimer's disease on US healthy aging dataset. *Mathematics* **2024**, *12*, 2204. [CrossRef]
26. Li, X.; Li, L.; Li, M.; Yan, P.; Feng, T.; Luo, H.; Zhao, Y.; Yin, S. Knowledge distillation and teacher-student learning in medical imaging: Comprehensive overview, pivotal role, and future directions. *Med. Image Anal.* **2025**, *107*, 103819. [CrossRef] [PubMed]

27. Sun, Y.; Li, X.; Li, L.; Feng, T.; Zhao, Y.; Yin, S. PHH-FL: Perceptual Hashing Hypernetwork Personalized Federated Learning for Heterogeneous Medical Image Analysis Tasks. *IEEE Internet Things J.* **2025**, *13*, 8712–8724. [CrossRef]
28. Maharana, K.; Mondal, S.; Nemade, B. A review: Data pre-processing and data augmentation techniques. *Glob. Transit. Proc.* **2022**, *3*, 91–99. [CrossRef]
29. Yan, Q.; Feng, Y.; Zhang, C.; Wang, P.; Wu, P.; Dong, W.; Sun, J.; Zhang, Y. You only need one color space: An efficient network for low-light image enhancement. *arXiv* **2024**, arXiv:2402.05809. [CrossRef]
30. Bajaj, S.; Bala, M.; Angurala, M. A comparative analysis of different augmentations for brain images. *Med. Biol. Eng. Comput.* **2024**, *62*, 3123–3150. [CrossRef]
31. Bunzel, N.; Göller, N.; Frick, R.A. Identifying and generating edge cases. In Proceedings of the 2nd ACM Workshop on Secure and Trustworthy Deep Learning Systems, Virtual, 2 July 2024; pp. 16–23.
32. Li, Z.; Lyu, K.; Arora, S. Reconciling modern deep learning with traditional optimization analyses: The intrinsic learning rate. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 14544–14555.
33. Zhang, H.; Liu, Y.; Chen, Q.; Fang, C. The optimality of (accelerated) SGD for high-dimensional quadratic optimization. *arXiv* **2024**, arXiv:2409.09745. [CrossRef]
34. Feng, L.; Shu, S.; Lin, Z.; Lv, F.; Li, L.; An, B. Can cross-entropy loss be robust to label noise? In Proceedings of the 29th International Joint Conference on Artificial Intelligence, Yokohama, Japan, 7–15 January 2021; pp. 2206–2212.
35. Filvantorkaman, M.; Piri, M.; Filvan Torkaman, M.; Zabihi, A.; Moradi, H. Fusion-based brain tumor classification using deep learning and explainable AI, and rule-based reasoning. *arXiv* **2025**, arXiv:2508.06891.
36. Ke, L.; Hu, G.; Zhao, M.; Liu, Z.; Lv, Z.; Yang, Y. Brain tumor classification from MRI images using a multi-scale channel attention CNN integrated with SVM. *Sci. Rep.* **2026**, *16*, 6297. [CrossRef] [PubMed]
37. Uniyal, M.; Saini, C.; Singh, D.P.; Ahmed, M.N.; Hussain, M.R.; Amarjeet; Sinha, A.; Srivastava, A. Automated brain tumor detection using advanced deep learning models. *Discov. Artif. Intell.* **2026**, *6*, 1. [CrossRef]
38. Guan, Y.; Alshammari, A.; Wang, Y.; Gul, J.Z.; Imran, A. ResSGA-Net: A deep learning approach for enhanced brain tumor detection and accurate classification in healthcare imaging systems. *J. Genet. Eng. Biotechnol.* **2026**, *24*, 100658. [CrossRef] [PubMed]
39. ben.logsdon. *AMP AD Target Discovery Data Portal*; Synapse: Seattle, WA, USA, 2015. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.