

Deep Learning–Driven Brain Tumour Diagnosis: Training ResNet18 from Scratch on MRI for Real-Time Clinical Use

Affan Osman¹, R Naveen Kumar²

¹24E31D5801, Dept. Of CSE, MIST, Bandlaguda Jagir, Hyderabad, Telangana.

² Assistant Professor, Dept. Of CSE, MIST, Bandlaguda Jagir, Hyderabad, Telangana.

Abstract

Highly accurate results of brain tumor detection from MRI scans are a vital step toward early diagnosis and clinical decision support. In this study, we trained a ResNet18 deep convolutional neural network from-scratch with a custom architecture on a dataset with 3,762 of brain MRIs of two-class nomenclature tumor and no-tumor. After standard preprocessing and stratified train-test splitting, the network was trained for 15 epochs with a variable learning-rate schedule and checkpointing based on performance. The model had good convergence properties, and the training accuracy improved consistently while the validation accuracy has stabilized. In an out-of-sample validation set ($n = 753$) the model yielded an overall accuracy of 95.22%, precision of 0.9822 for no-tumor and 0.9192 for tumor, and recall of 0.9303 and 0.9792 on this data, respectively. The resulting confusion matrix: $TN = 387$, $FP = 29$, $FN = 7$, $TP = 330$ shows the sensitivity to tumor detection, minimizing the risk of missed diagnosis. The ResNet18 system continued to show excellent discriminative capability with a ROC-AUC of 0.9916 (95% CI 0.9836 to 0.9965) and a PR-AUC of 0.9924 (95% CI 0.9828 to 0.9949), despite class imbalance. The prediction confidence distributions for the classes were very well separated, as visualized in figure??. The final ranked ResNet18 model enables high-accuracy and reliable brain tumor classification and is feasible for seamless integration into real-time clinical screening workflows, but external validation with a focus on detection of multiple tumor subtypes are requirements of future work.

Keywords: *Brain Tumor Classification, Magnetic Resonance Imaging (MRI), Deep Convolutional Neural Network, ResNet18 Architecture*

1. Introduction

Brain tumor is still one of the most lethal nervous system diseases; proven that a diagnostic delay cannot only affect treatment results but also prognosis and overall survival [2]. Due to its soft-tissue contrast and non-invasiveness, magnetic resonance imaging (MRI) is the most common modality for screening brain tumors. Nevertheless, MRI scans are interpreted manually, a time-consuming process subject to a considerable range of interpretive variability, and consequently misdiagnosis. Automated systems based on deep learning have emerged as a promising solution to aid the radiologist in the identification of tumors avoiding laborious and time-consuming work while providing fast, reproducible, and highly accurate results. Several recent studies have exhibited improved performance for MRI-based tumor detection using residual networks and other deep learning architectures in comparison to traditional handcrafted feature methods [1]–[4]. In addition, new strides in explainable deep learning [5] emphasize that transparency and decision interpretability are fundamental characteristics that need to be learned in order to implement used solutions in the clinic. Deep learning based brain tumor classification approaches have made tremendous improvement, however, most of the existing systems follow computationally heavy architectures, transfer-learning pipelines or hybrid multi-stage frameworks posing a bottleneck for translation into clinical practice. Hence, a deep learning model trained from scratch lightweight in size but high in accuracy will be ideal for efficient tumor/no-tumor categorization but also clinically meaningful.

2. Related Work

A number of deep learning methods have been applied to facilitate MRI-based brain tumor analysis. Ashimgaliyev et al. implemented MRI segmentation in

deep learners to improve identification and lower false positives [6]. This includes the fine-tuned networks proposed by Rastogi et al. to be widely adopted to work on common approaches of transfer learning. for tumor prediction [7]. Recent work by Gorenshtein et al. [8], while Awad et al. built a small hybrid model for scalable, real-time applications. [9] Arabahmadi et al. tumor detection approaches based on machine learning and deep learning for smart healthcare systems, as well as providing an overview of current trends and some important technical challenges [10] Although there are a few papers that focus on deep learning-based tumor detection methods for smart healthcare systems [10], there remains a lack of attention directed towards summarizing the state-of-the-art in automatic tumor detection systems.

There are also architectural improvements that have been considered. A skip connected ResNet18 encoder-decoder model was also used by Pravallika and Reddy to achieve better segmentation performance with the help of Dice loss [11]. Nassar et al. hybrid deep Net: A fully integrated hybrid deep framework for robust tumor classification that merges several learning modules [12]. Ben Brahim et al. showed their work on CNN-driven detection system of MRI based tumors [13], and Kibriya et al. combined deep features with conventional ML classifiers to improve classification performance [14]. Anagun [15] explored smart CNN-powered tumor diagnosis systems; Saurav et al. proposed CNNs that emphasize specific features using attention guidance [16]. Ullah et al. TumorDetNet—a unified deep learning model for simultaneous detection and classification [17];

More progress includes the work of Pacal et al. on the global and channel attention mechanism strengthened EfficientNetV2 variation. Lastly, challenges, limitations, and future directions in MRI-based tumor classification: A review by Ottoni et al [18]. [19]. Kumar et al. provided a summary of recent neuroimaging based ML solutions for brain tumour and Alzheimer detection [20]. Then, Karacı and Akyol presented a hybrid model architecture combining YOLO + DenseNet + Bi-LSTM to obtain high accuracy for multi-class brain tumor classification [21].

3. Methodology:

The five-phase deep learning workflow for brain tumor classification is organized and will be carried out in this way. The dataset containing 3,762 MRI images was first prepared by reading a CSV file listing image filenames and corresponding labels, converting each image to RGB, resizing all images to the same dimensions of 224×224 pixels, and splitting the dataset with stratified sampling into 80% training set (3,009 images) and 20% validation sets (753 images) to maintain class balance. The second stage applies preprocessing and augmentation on images to ensure better robustness of the model: random horizontal and vertical flips, small rotations, affine translations, color jittering, and normalizing with ImageNet-based statistics were performed on the training set, while images in the validation set were resized and normalized without performing other transformations to avoid any biased evaluation. Stage three: Implemented a full ResNet18 architecture from scratch in PyTorch as shown in figure 1. Each layer of the architecture was expressed in terms of primitive PyTorch operations starting with a 7×7 convolution layer then batch normalization, ReLU activation, and max pooling followed by four residual stages with two BasicBlocks in each that included skip connections and batch normalization after each 3×3 convolution; Added global average pooling with a fully connected classifier with dropout layers giving a model with 11.3 million trainable parameters. All parameters were initialized using He/Kaiming normalization. In the fourth stage, the entire 2D ResNet50k architecture was trained for 15 epochs on all image patches using the AdamW optimizer (with a learning rate of 0.001 and weight decay $1e-4$), where the objective function was cross-entropy loss, and a ReduceLROnPlateau scheduler was used to adaptively reduce the learning rate whenever the validation accuracy plateaued with no improvement; per-epoch training and validation accuracy and loss were collected each epoch and the version giving the best validation accuracy was kept. The last step included evaluation of the model with the best performance on the validation set in terms of accuracy, precision, recall, F1-score, ROC-AUC, PR-AUC, sensitivity, specificity, and confusion matrices along with visualizations including ROC and precision-recall curves, training curves, probability distributions, and sample prediction results. The pipeline follows a structured approach, making it a reproducible and well-validated deep learning system for brain tumor detection.




Figure 1: *Proposed model architecture*

4. Results and Analysis

Our custom-built ResNet18 model performed well in classifying the brain MRI images into tumor and no-tumor categories. The training was non-problematic across the 15 epochs with both training and validation losses decreasing steadily validating the learned features and the generalization. The learning-rate scheduler also provided additional stability by automatically reducing the learning rate whenever validation accuracy plateaued.

The model produced a global accuracy of 95.22% on the validation set composed of 753 images, indicating an excellent ability to generalize. Print (cr) Output: 0.9822 for No-Tumor class Precision: 0.9303 for No-Tumor class Recall: 0.9192 for Tumor class Precision: 0.9792 for Tumor class Recall: Results from the

classification report, we can say that both classes performed well. These results underscore that while detection of tumor cases are highly sensitive under the model (reducing the risk of missing tumors), healthy cases are highly specific. Table 1 also highlights balanced performance by providing the corresponding F1-scores of 0.9556 (no tumor) and 0.9483 (tumor)

Looking under the hood (Figure 3) weirdly explains a lot, 387 true negatives and 330 true positives but only 29 false positives and 7 false negatives — a confusion matrix. This means that there would be very few misclassifications or a few tumorcases misclassified as healthy. This is very important for medical screening systems, where any missed tumor would have potentially fatal consequences.

Additionally, the model is further reinforced by threshold-independent metrics, thus proving its robustness. An ROC curve resulted in an AUC of 0.9916 and a precision–recall curve resulted in an AUC of 0.9924, both suggesting excellent discriminative ability even in the presence of class imbalance. Figure 5: Visualisation of the predicted probability distribution $\times f(x)$ for Tumour vs No Tumour class as it is likely to be highly normalised and therefore demonstrates clear separation. Parse error Protocol description: The Motor Sequence Task is designed to assess aspects of learning as reflected in general speed, habitual responses and cognitive novelty. It has been successfully used with animals and humans.

Qualitative inspection of sample predictions show the model has considered clearly a tumor region and a very high confident prediction is made in separating the tumor from all other normal brain structure. The top images that are correctly predicted show high activations of the features, while the very few misclassified samples are the ones with low contrast, motion blur or overlapping structures in the tissue, thereby mimicking the natural challenges when interpreting MRIs.

In summary, the findings show that a ResNet18 model that trained from scratch is an efficient model for binary brain tumor detection while achieving high accuracy, sensitivity, specificity, and high AUC score. The model reliability and robustness are due in part to: deep residual learning, extensive data augmentation, and adaptive learning-rate scheduling. These results

demonstrate the potential of the model for use in real-world clinical screening and computer-aided diagnosis.

Figure 2 below displays 16 example MRI predictions by ResNet18. The true label, model predicted label/text and the prediction confidence are shown for each image. The majority of correct predictions —right (tumor) or wrong (non-tumor) — have a high confidence score, which indicates a good visual understanding of both sides. Examples of Misclassified occur in low-contrast or ambiguous images.




Figure 2: predictions of proposed model




Figure 3: Confusion Matrix and Normalized Confusion Matrix for the ResNet18 Model




Figure 4: ROC curve of proposed model




Figure 5: Training and Validation Loss and Accuracy Curves of the ResNet18 Model

Figure 5: The training and validation performance of ResNet18 across 15 epochs. The loss graph looks good, with a consistent trend of decreasing training loss, and a fluctuating but generally downward trend of validation loss, which tells the model is learning some meaningful features. From the accuracy curve, we see a steady increase in training accuracy while validation accuracy graphs also periodically rise but become stable at a value greater than 90%. In general, the curves suggest the models were able to train and generalize well to unseen validation data.

Table 1: Classification Report for the ResNet18 Model

Class	Precision	Recall	F1-Score	Support
No Tumor	0.9822	0.9303	0.9556	416
Tumor	0.9192	0.9792	0.9483	337
Accuracy	—	—	0.9522	753
Macro Avg	0.9507	0.9548	0.9519	753
Weighted Avg	0.9540	0.9522	0.9523	753

5. Conclusion

This paper then provided a deep learning pipeline for brain classification that was end-to-end. It introduced a different model, but trained it end-to-end for the whole task, using the resNet18 implementation by Asif et al. By properly preprocessing the data, making clever augmentation, and training the model with attention to detail, the model managed to achieve near-perfect accuracy and generalization score, and by only misclassifying a few of the tumor and human ones. The results show that the proposed model is accurate and speedy, with high-sensitivity in identifying the tumor and strong specificity for the healthy cases. The implications of these findings suggest that deep residual learning has the potential to be a useful tool in the hands of clinicians for the early and accurate diagnosis of brain tumors. Our ongoing and future work here could involve extension of the approach to multi-class tumor grading, larger dataset cohorts, and real-time deployment within a clinical context.

6. References:

1. Neamah, K., Mohamed, F., Waheed, S. R., Kurdi, W. H. M., Taha, A. Y., & Kadhim, K. A. (2024). Utilizing deep improved resnet50 for brain tumor classification based MRI. *IEEE Open Journal of the Computer Society*, 5, 446-456.
2. Abdusalomov, A. B., Mukhiddinov, M., & Whangbo, T. K. (2023). Brain tumor detection based on deep learning approaches and magnetic resonance imaging. *Cancers*, 15(16), 4172.
3. Abdulrahman, S. (2024). ENHANCING BRAIN TUMOUR DIAGNOSIS WITH AI: A COMPARATIVE ANALYSIS OF RESNET AND YOLO ALGORITHM FOR TUMOUR CLASSIFICATION IN MRI SCANS.
4. Rath, A., Mishra, B. S. P., & Bagal, D. K. (2025). ResNet50-based Deep Learning model for accurate brain tumor detection in MRI scans. *Next Research*, 2(1), 100104.
5. Li, Z., & Dib, O. (2024). Empowering brain tumor diagnosis through explainable deep learning. *Machine Learning and Knowledge Extraction*, 6(4), 2248-2281.
6. Ashimgaliyev, M., Matkarimov, B., Barlybayev, A., Li, R. Y. M., & Zhumadillayeva, A. (2024). Accurate MRI-based brain tumor diagnosis: integrating segmentation and deep learning approaches. *Applied Sciences*, 14(16), 7281.
7. Rastogi, D., Johri, P., Donelli, M., Kumar, L., Bindewari, S., Raghav, A., & Khatri, S. K. (2025). Brain tumor detection and prediction in MRI images utilizing a Fine-Tuned transfer learning model integrated within deep learning frameworks. *Life*, 15(3), 327.
8. Gorenstein, A., Liba, T., & Goren, A. (2025). Lightweight Transfer Learning Models for Multi-Class Brain Tumor Classification: Glioma, Meningioma, Pituitary Tumors, and No Tumor MRI Screening. *Journal of Imaging Informatics in Medicine*, 1-10.
9. Awad, S. R., Daood, A. I., & Dawood, A. A. (2025). Scalable and Efficient Multi-Class Brain Tumor Classification with a Compact Hybrid Deep Learning Model for Real-Time Applications. *AROTHE SCIENTIFIC JOURNAL OF KOYA UNIVERSITY*, 13(1), 162-174.
10. Arabahmadi, M., Farahbakhsh, R., & Rezazadeh, J. (2022). Deep learning for smart Healthcare—A survey on brain tumor detection from medical imaging. *Sensors*, 22(5), 1960.
11. Pravallika, R. L., & Reddy, R. P. K. (2025). ENHANCED BRAIN TUMOR SEGMENTATION WITH SKIP CONNECTED RESNET-18 ENCODER-DECODER FRAME WORK WITH DICE LOSS. *International Journal of Applied Mathematics*, 38(2s), 743-769.
12. Nassar, S. E., Yasser, I., Amer, H. M., & Mohamed, M. A. (2024). A robust MRI-based brain tumor classification via a hybrid deep learning technique. *The Journal of Supercomputing*, 80(2), 2403-2427.
13. Ben Brahim, S., Dardouri, S., & Bouallegue, R. (2024). Brain tumor detection using a deep cnn model. *Applied Computational Intelligence and Soft Computing*, 2024(1), 7634426.
14. Kibriya, H., Amin, R., Alshehri, A. H., Masood, M., Alshamrani, S. S., & Alshehri, A. (2022). A novel and effective brain tumor classification model using deep feature fusion and famous machine learning classifiers. *Computational Intelligence and Neuroscience*, 2022(1), 7897669.
15. Anagun, Y. (2023). Smart brain tumor diagnosis system utilizing deep convolutional neural networks. *Multimedia tools and applications*, 82(28), 44527-44553.
16. Saurav, S., Sharma, A., Saini, R., & Singh, S. (2023). An attention-guided convolutional neural network for automated classification of brain tumor from MRI. *Neural Computing and Applications*, 35(3), 2541-2560.

17. Ullah, N., Javed, A., Alhazmi, A., Hasnain, S. M., Tahir, A., & Ashraf, R. (2023). TumorDetNet: A unified deep learning model for brain tumor detection and classification. *Plos one*, 18(9), e0291200.
18. Pacal, I., Celik, O., Bayram, B., & Cunha, A. (2024). Enhancing EfficientNetv2 with global and efficient channel attention mechanisms for accurate MRI-Based brain tumor classification. *Cluster Computing*, 27(8), 11187-11212.
19. Ottoni, M., Kasperczuk, A., & Tavora, L. M. (2025). Machine Learning in MRI Brain Imaging: A Review of Methods, Challenges, and Future Directions. *Diagnostics*, 15(21), 2692.
20. Kumar, N., Narawade, V., Chheda, K., Patkar, H., & Mishra, A. (2024). Advanced Machine Learning for Brain Tumor and Alzheimer's Disease Detection: A Comprehensive Review of Neuroimaging-based Classification Techniques. *Advances in Artificial Intelligence Research*, 4(2), 62-68.
21. Karacı, A., & Akyol, K. (2023). YoDenBi-NET: YOLO+ DenseNet+ Bi-LSTM-based hybrid deep learning model for brain tumor classification. *Neural Computing and Applications*, 35(17), 12583-12598.