Comparison of a sequential CNN and ResNet50 for brain tumor detection in grayscale MRI scans *

Gedeon NITIEMA 1, †, James T. KOUAWA 2,*,† and Ahmed Dooguy KORA3,†

[1] *Cheikh Anta Diop University, Doctoral School of Mathematics and Computer Science, Dakar, Senegal*
[2] *Multinational Higher School of Telecommunications, Digital Science and Technology Research Laboratory, Dakar, Senegal*

Abstract

The automatic detection of brain tumors from magnetic resonance imaging (MRI) scans is a major challenge in the medical field. Traditional methods relying on radiologist expertise are time-consuming and subject to varying interpretations. This article presents a comparative study of two deep learning approaches for the binary classification of grayscale brain MRI images, applied to a public dataset available on Kaggle. The first approach is based on a sequential convolutional neural network (CNN) built from scratch, consisting of five convolutional blocks with batch normalization. The second approach utilizes a pre-trained ResNet50 model adapted for grayscale images. Both models were trained for 50 epochs and evaluated on a test set of 447 images. The results show remarkable and similar performance across certain metrics: the sequential model achieved an overall accuracy of 95.53%, an AUC-ROC of 98.45%, and an F1-score of 94.71% at the optimal Youden threshold (0.72). The ResNet50 model achieved an overall accuracy of 95.08%, an AUC-ROC of 98.81%, and an F1-score of 94% at the optimal threshold (0.39). The main difference lies in data generalization: ResNet50 showed a training-validation gap of only 4.47%, compared to 8.34% for the sequential model. These results, obtained using a limited-size public dataset, demonstrate the clinical potential of both approaches for diagnostic support, while highlighting the need for validation prior to any clinical deployment.

Keywords

Brain tumor, CNN, Deep learning, MRI, ResNet50, binary classification

CEUR-WS.org/Vol-4232/Paper14.pdf

1. Introduction

Brain cancers are central nervous system tumors that can develop in the brain, brainstem, cerebellum, or spinal cord.

According to the American Cancer Society (ACS) cancer statistics center, there were 23,820 new cases of brain tumors in the United States in 2019, and approximately 17,760 deaths attributed to the disease that year. Early detection and accurate characterization are beneficial for effective treatment and optimal clinical outcomes.

However, interpreting standard MRI scans for brain tumor detection can be complex and requires expertise, as it involves distinguishing tumor tissue from normal tissue, including gray matter and cerebrospinal fluid.

Accurate brain tumor segmentation is crucial for medical diagnosis, surgical planning, and treatment planning.

* *CITA 2026 – Land, Smart Cities, and Sustainable Development, 25–26 June 2026, Cotonou, Benin*

* Corresponding author.
gedeonnitiema@gmail.com (G. Nitiema); james.kouawa@esmt.sn (J.T. Kouawa); ahmed.kora@esmt.sn (A. D. Kora)
0009-0001-8383-1189 (G. Nitiema); 0000-0002-7636-2054 (A.D. KORA)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-07-27

Performance in medical image analysis depends heavily on the chosen deep network architecture, highlighting the need for comparative studies to evaluate the effectiveness of different architectures

In this study, we compare two CNN architectures applied to the binary classification of brain tumors (Tumor vs. Non-tumor) using grayscale MRI images. The first model is built from scratch, while the second is based on transfer learning using a ResNet50 model pre-trained on ImageNet. The objective is to evaluate the effectiveness of each approach in terms of accuracy and determine which model generalizes better to the data, with the aim of proposing a reliable solution to assist in clinical diagnosis.

2. Data and Preprocessing

2.1 Dataset Description

The dataset used in this study comes from the Kaggle platform (Brain MRI Images for Brain Tumor Detection). The images are grayscale brain MRIs classified into two categories (Tumor/Non-tumor). The test set used to evaluate the models consists of 447 images, comprising 256 images in the Non-tumor class and 191 images in the Tumor class. To ensure the reproducibility of the experiments, the training and validation sets were split using a fixed random seed (seed=42).

2.2 Image preprocessing

In our study, the data were split into a training set (70%), a validation set (20%), and a test set (10%). Image preprocessing involved resizing images to 236×286 pixels—corresponding to the average dimensions of the original dataset images—followed by normalizing pixel values to the [0, 1] range by dividing by 255. Data augmentation techniques—specifically random rotations, horizontal and vertical flips, and brightness variations—were applied solely during training to enhance model robustness and reduce the risk of overfitting. To ensure a fair comparison, the normalization and data augmentation conditions were kept identical.

3. Model architectures

3.1 Sequential CNN Model

The sequential model is built from scratch and consists of five successive convolutional blocks. Each block follows the Conv2D ⟶ BatchNormalization ⟶ MaxPooling2D pattern to enable visual feature extraction. Table 1 and Figure 1 present the model architecture.

Table 1: *Sequential CNN Architecture*

Layer	Output format	Parameters
Conv2D (32 filters, 3×3) + BN	(None, 236, 286, 32)	448
MaxPooling2D (2×2)	(None, 118, 143, 32)	0
Conv2D (64 filters, 3×3) + BN	(None, 118, 143, 64)	18 752
MaxPooling2D (2×2)	(None, 59, 71, 64)	0
Conv2D (128 filters, 3×3) + BN	(None, 59, 71, 128)	74 368
MaxPooling2D (2×2)	(None, 29, 35, 128)	0
Conv2D (128 filters, 3×3) + BN	(None, 29, 35, 128)	148 096
MaxPooling2D (2×2)	(None, 14, 17, 128)	0

Layer	Output format	Parameters
Conv2D (256 filters, 3×3) + BN	(None, 14, 17, 256)	296 192
MaxPooling2D (2×2)	(None, 7, 8, 256)	0
GlobalAveragePooling2D	(None, 256)	0
Dense (128) + Dropout	(None, 128)	32 896
Dense (1, sigmoid)	(None, 1)	129

Approximately 571,000 trainable parameters.

The GlobalAveragePooling2D layer is used to reduce the risk of overfitting. A dense layer of 128 neurons is followed by a Dropout layer for regularization. The output layer consists of a single neuron with sigmoid activation.

Layer (type)	Output Shape	Param #
conv2d (Conv2D)	(None, 236, 286, 32)	320
batch_normalization (BatchNormalization)	(None, 236, 286, 32)	128
max_pooling2d (MaxPooling2D)	(None, 118, 143, 32)	0
conv2d_1 (Conv2D)	(None, 118, 143, 64)	18,496
batch_normalization_1 (BatchNormalization)	(None, 118, 143, 64)	256
max_pooling2d_1 (MaxPooling2D)	(None, 59, 71, 64)	0
conv2d_2 (Conv2D)	(None, 59, 71, 128)	73,856
batch_normalization_2 (BatchNormalization)	(None, 59, 71, 128)	512
max_pooling2d_2 (MaxPooling2D)	(None, 29, 35, 128)	0
conv2d_3 (Conv2D)	(None, 29, 35, 128)	147,584
batch_normalization_3 (BatchNormalization)	(None, 29, 35, 128)	512
max_pooling2d_3 (MaxPooling2D)	(None, 14, 17, 128)	0
conv2d_4 (Conv2D)	(None, 14, 17, 256)	295,168
batch_normalization_4 (BatchNormalization)	(None, 14, 17, 256)	1,024
max_pooling2d_4 (MaxPooling2D)	(None, 7, 8, 256)	0
global_average_pooling2d (GlobalAveragePooling2D)	(None, 256)	0
dense (Dense)	(None, 128)	32,896
dropout (Dropout)	(None, 128)	0
dense_1 (Dense)	(None, 1)	129

Figure 1: *Summary of the Sequential CNN model architecture.*

3.2 ResNet50 Model

The second model is a deep residual network with 50 layers, pre-trained on the ImageNet dataset. To adapt this model to grayscale MRI images, we used an initial Conv2D layer that converts the single-channel input into a 3-channel representation (gray_to_rgb). This adaptation allowed us to leverage the weights pre-trained on ImageNet while accepting grayscale images. Built upon the ResNet50 base and a GlobalAveragePooling2D layer is "fc1," a classification head consisting of a Dense layer (12 units), BatchNormalization, Dropout, and a final Dense layer with a sigmoid activation. The model has approximately 23.85 million parameters.

Model: "ResNet50_IRM"

Layer (type)	Output Shape	Param #
input_gray (InputLayer)	(None, 236, 286, 1)	0
gray_to_rgb (Conv2D)	(None, 236, 286, 3)	6
resnet50 (Functional)	(None, 8, 9, 2048)	23,587,712
gap (GlobalAveragePooling2D)	(None, 2048)	0
fc1 (Dense)	(None, 128)	262,272
bn_head (BatchNormalization)	(None, 128)	512
dropout (Dropout)	(None, 128)	0
output (Dense)	(None, 1)	129

Total params: 23,850,631 (90.98 MB)

Trainable params: 23,797,255 (90.78 MB)

Non-trainable params: 53,376 (208.50 KB)

Figure 2: *Summary of the ResNet50 model architecture.*

4. Training and evaluation protocol

To ensure the comparability of results, both models were trained using an identical protocol. We employed a binary cross-entropy loss function, which is suitable for binary classification. Training for both models was conducted over 50 epochs. We optimized the decision threshold using the Youden index (sensitivity + specificity − 1) to determine the optimal threshold for each model. A baseline threshold of 0.5 was also used, and evaluation was performed on a dataset of 447 images. To enable a comprehensive assessment of the models' diagnostic performance, we utilized evaluation metrics such as overall accuracy, sensitivity, specificity, positive predictive value (PPV), F1-score, and AUC-ROC.

5. Results

5.1 Training the Sequential CNN

The CNN training curves (Figure 3) show a gradual convergence of training accuracy toward 98.30% over 50 epochs. The best performance is observed at epoch 34, with a validation accuracy of 89.96%. Training loss decreases toward zero, whereas validation loss exhibits significant fluctuations. The gap between training and validation accuracy is 8.34%, indicating a reasonable level of overfitting.

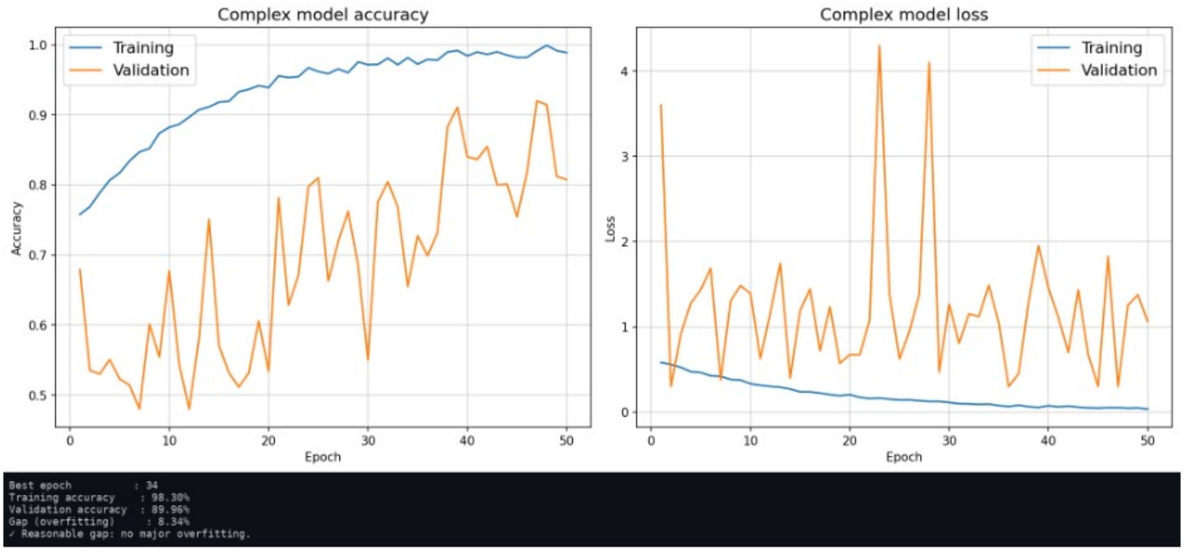

Figure 3: *Accuracy and loss curves for the Sequential CNN over 50 epochs.*

5.2 ROC curve of the sequential CNN

The ROC curve of the sequential CNN (Figure 4) shows good performance, with an AUC-ROC of 98.45%. The optimal Youden threshold is found at 0.72, maximizing the trade-off between the model's sensitivity and specificity.

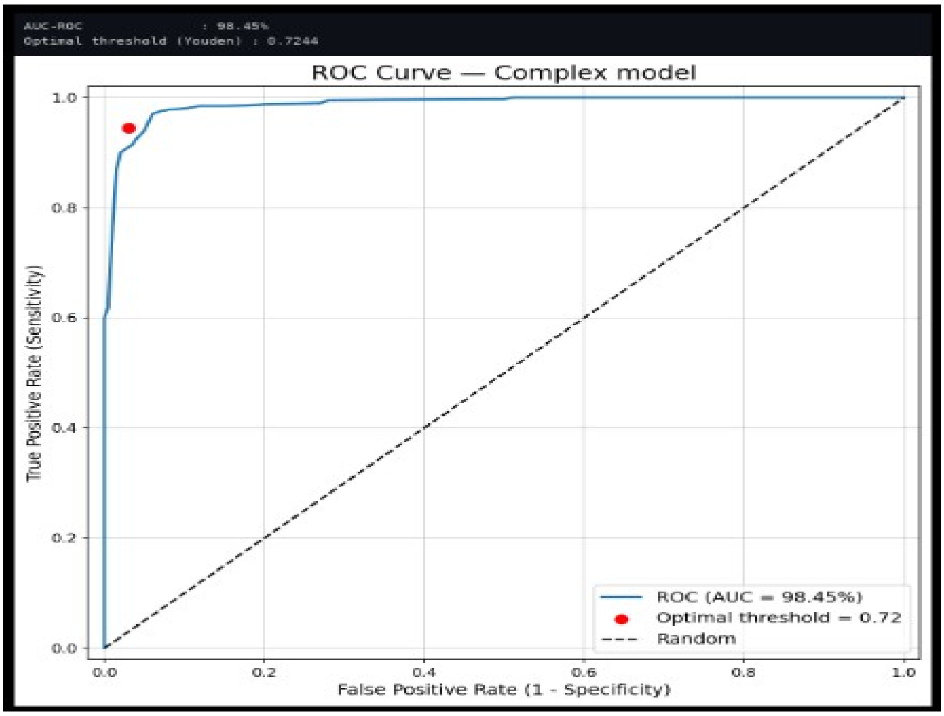

Figure 4: *AUC-ROC curve of the sequential model.*

5.3 Classification performance of the Sequential CNN

At the baseline threshold of 0.5, the sequential CNN achieves an overall accuracy of 95.08%, a sensitivity of 94.76%, and a specificity of 95.31%. The confusion matrix (Figure 6) shows 12 false negatives and 10 false positives out of 447 cases in the test set.

```
==================================================
Results — threshold 0.50 (baseline)
==================================================

  Overall accuracy    : 95.08%
  Sensitivity          : 94.76%
  Specificity          : 95.31%
  Precision (PPV)       : 93.78%
  F1-score              : 94.27%

Confusion matrix:
[[244  12]
 [ 10 181]]

Classification report:
                precision    recall  f1-score   support

  Non-tumor         0.96      0.95      0.96       256
      Tumor         0.94      0.95      0.94       191

    accuracy                            0.95       447
   macro avg         0.95      0.95      0.95       447
weighted avg         0.95      0.95      0.95       447
```

Figure 5: *Classification metrics for the sequential CNN at the baseline threshold (0.5).*

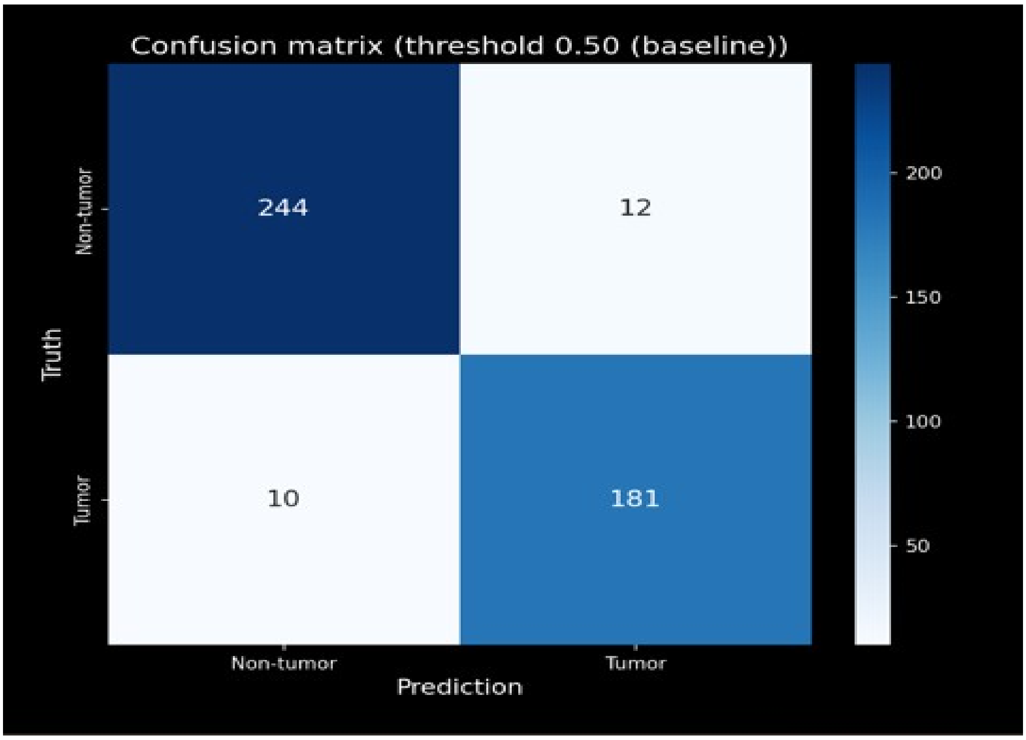

Figure 6: *Confusion matrix for the Sequential CNN at the baseline threshold (0.5).*

It is observed that at the optimal Youden threshold (0.72), performance improves, with an overall accuracy of 95.53%, sensitivity of 93.72%, specificity of 96.88%, and an F1-score of 94.71%. The confusion matrix (Figure 8) shows 248 true negatives, 179 true positives, 8 false positives, and 12 false negatives.

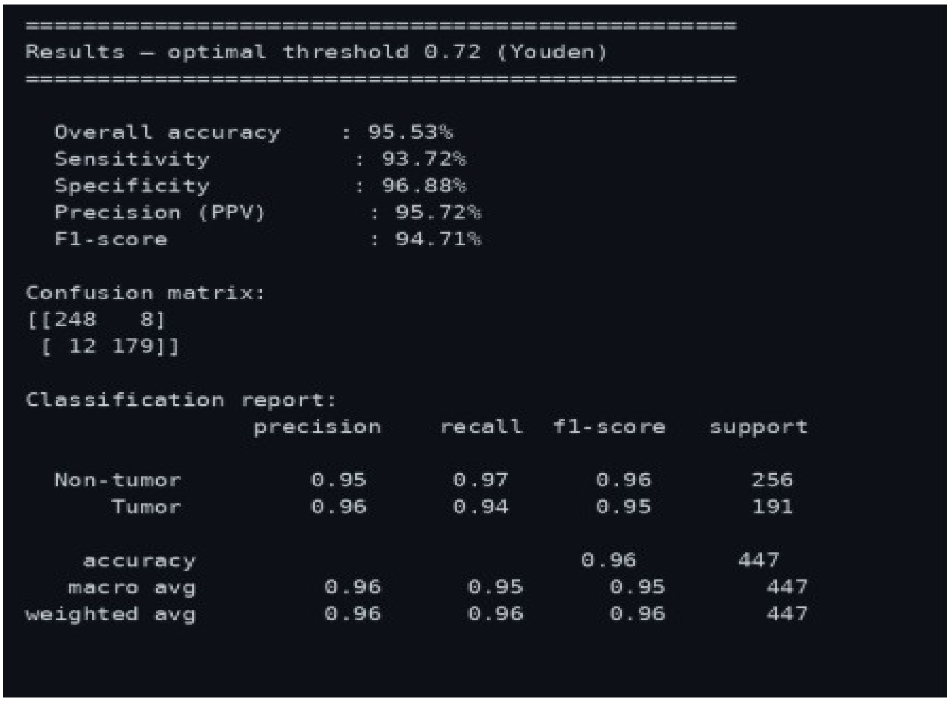

```
==================================================
Results — optimal threshold 0.72 (Youden)
==================================================

  Overall accuracy    : 95.53%
  Sensitivity          : 93.72%
  Specificity          : 96.88%
  Precision (PPV)       : 95.72%
  F1-score              : 94.71%

Confusion matrix:
[[248   8]
 [ 12 179]]

Classification report:
                precision    recall  f1-score   support

  Non-tumor          0.95      0.97      0.96       256
      Tumor          0.96      0.94      0.95       191

   accuracy                              0.96       447
  macro avg          0.96      0.95      0.95       447
weighted avg         0.96      0.96      0.96       447
```

Figure 7: *Sequential CNN classification metrics at the optimal Youden threshold.*

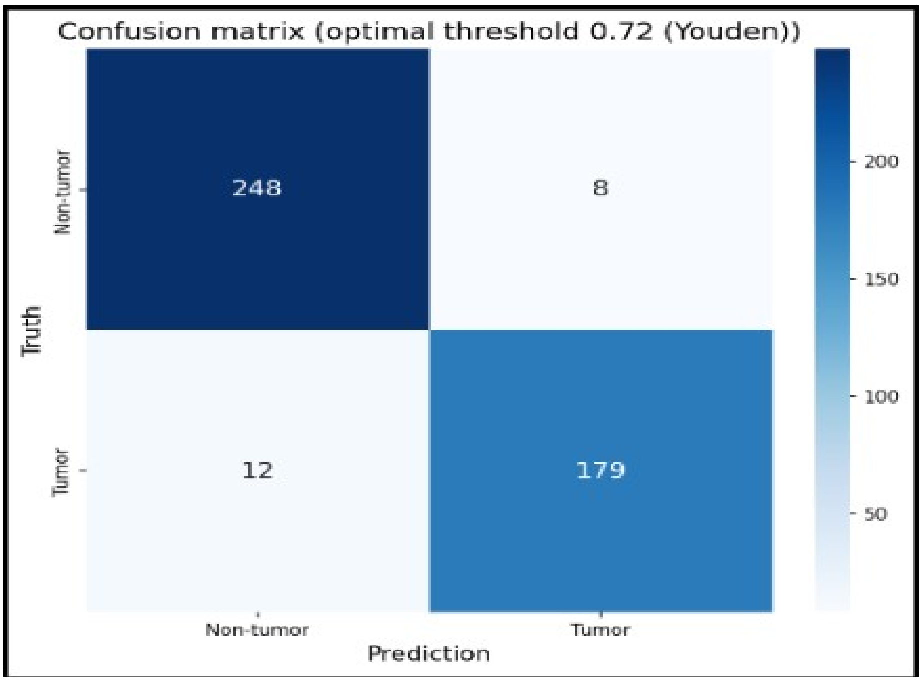

Figure 8: *Confusion matrix for the Sequential CNN at the optimal Youden threshold.*

5.4 Training the ResNet50 model

The training curves for the ResNet50 model (Figure 9) show more stable convergence than those of the sequential CNN model, achieving a training accuracy of 98.27% and a validation accuracy of 93.80% by the 33rd epoch. A reduction in the overfitting gap is observed—4.47% compared to 8.34% for the sequential CNN—indicating that the ResNet50 model generalizes better on the dataset.

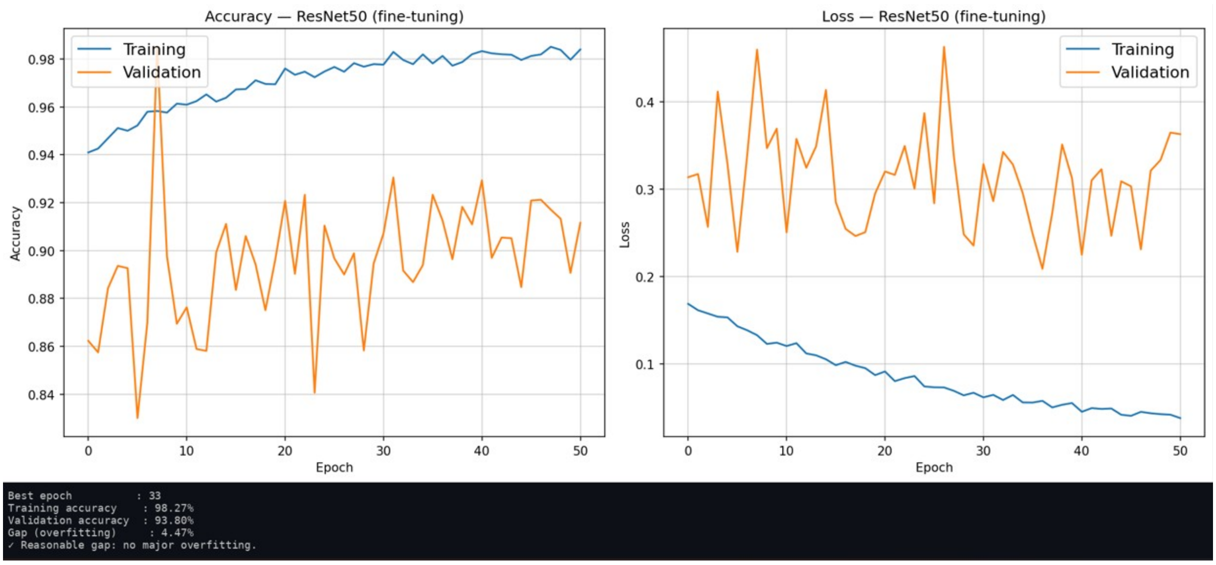

Figure 9: *Accuracy and loss curves for the ResNet50 model.*

5.5 ResNet50 AUC-ROC curve

The AUC-ROC curve (Figure 10) for the pre-trained model shows a value of 98.81%, which is slightly higher than that of the sequential CNN (98.45%). The optimal Youden threshold achieved its best balance at 0.39.

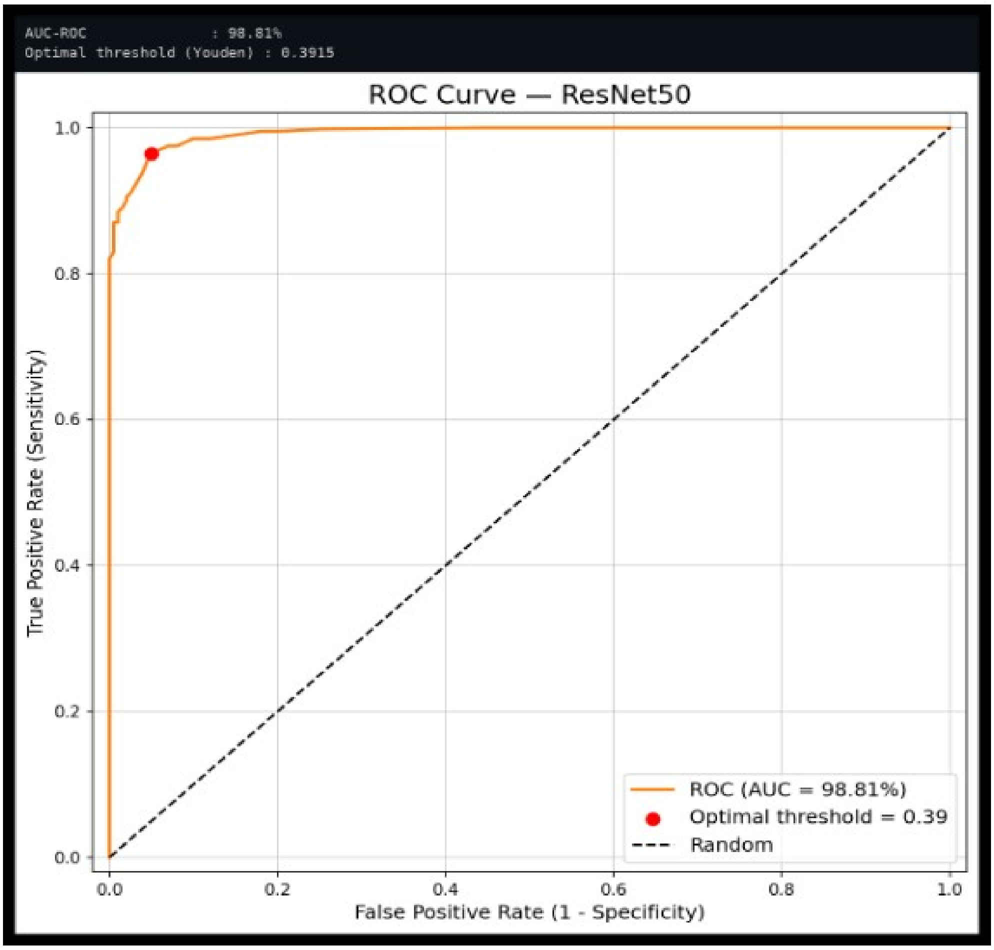

Figure 10: *ResNet50 AUC-ROC curve at 98.81%.*

5.6 ResNet50 classification performance

At the baseline threshold of 0.5, ResNet50 achieved an overall accuracy of 95.30%, a sensitivity of 94.76%, and a specificity of 95.70%. The confusion matrix (Figure 12) shows 11 false negatives and 10 false positives.

```
==================================================
Results ResNet50 – threshold 0.50
==================================================

  Overall accuracy    : 95.30%
  Sensitivity          : 94.76%
  Specificity          : 95.70%

Classification report:
               precision    recall  f1-score   support

  Non-tumor       0.96      0.96      0.96       256
      Tumor       0.94      0.95      0.95       191

    accuracy                          0.95       447
   macro avg       0.95      0.95      0.95       447
weighted avg       0.95      0.95      0.95       447
```

Figure 11: *ResNet classification metrics at the baseline threshold.*

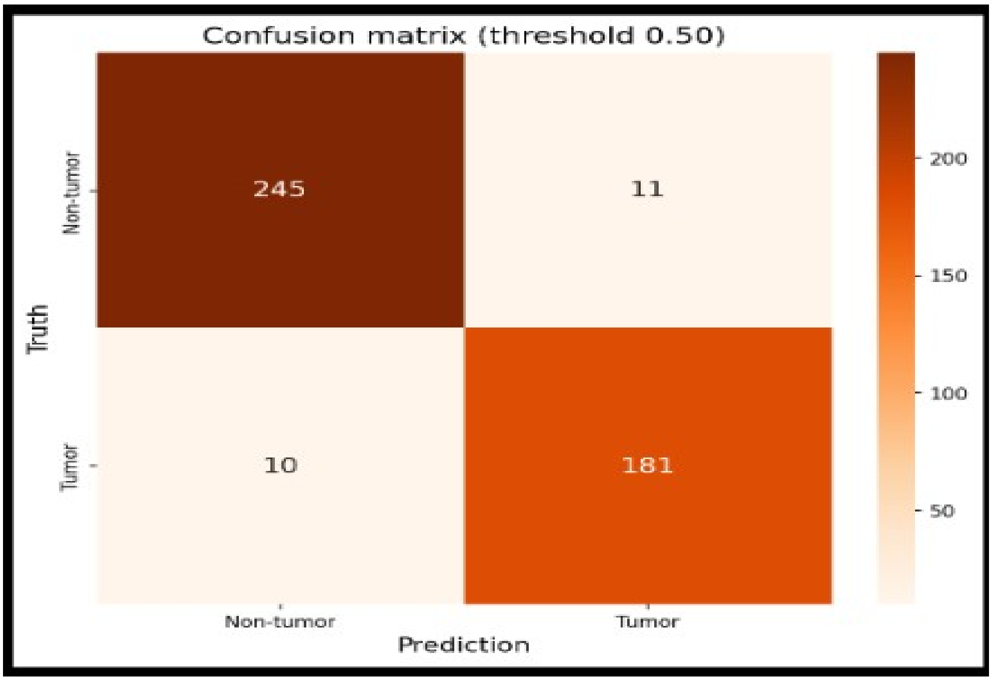

Figure 12: *Confusion matrix for ResNet50 at the baseline threshold.*

The performance of ResNet50 at the optimal Youden threshold consists of an overall accuracy of 95.08%, an improved sensitivity of 95.29%, and a specificity of 94.92%. The confusion matrix (Figure 14) shows 243 true negatives, 182 true positives, 13 false positives, and 9 false negatives. A reduction in false negatives is observed (9 versus 12 for the CNN), which is clinically significant.

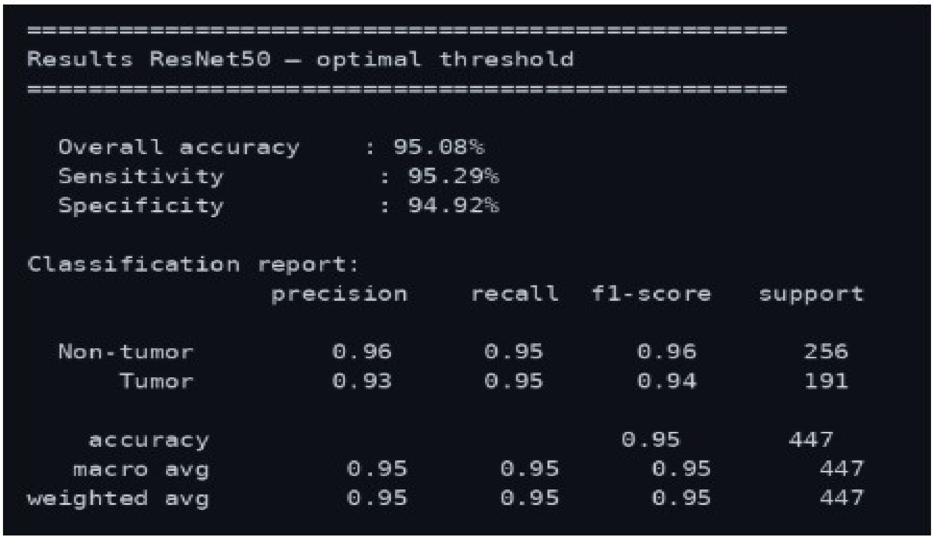

```
==================================================
Results ResNet50 — optimal threshold
==================================================

  Overall accuracy    : 95.08%
  Sensitivity          : 95.29%
  Specificity          : 94.92%

Classification report:
                precision    recall  f1-score   support

   Non-tumor       0.96      0.95      0.96       256
       Tumor       0.93      0.95      0.94       191

    accuracy                           0.95       447
   macro avg       0.95      0.95      0.95       447
weighted avg       0.95      0.95      0.95       447
```

Figure 13: *ResNet50 classification metrics at the optimal threshold (0.39).*

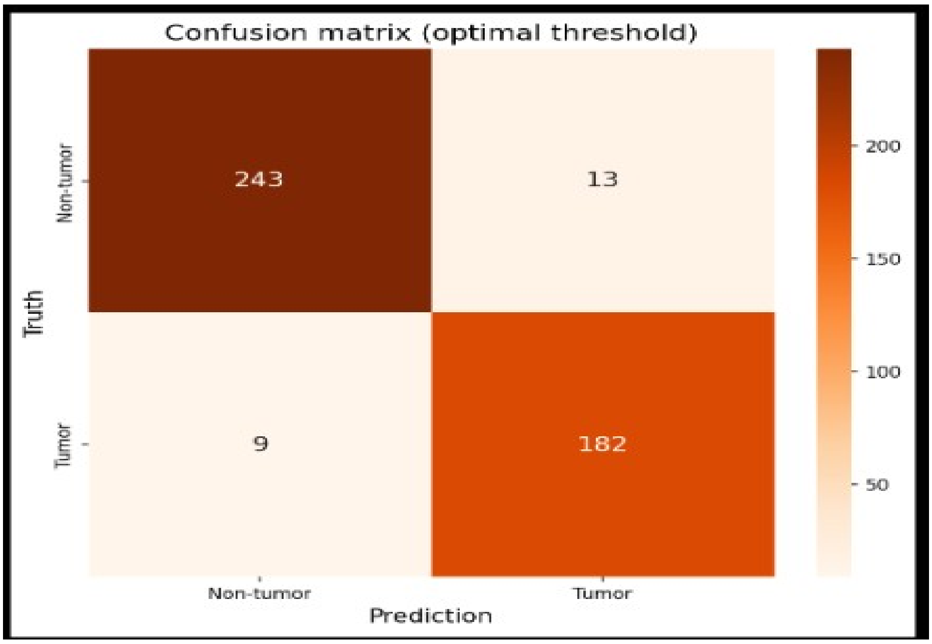

Figure 14: ResNet50 confusion matrix.

5.7 Comparative analysis

Table 2: *Comparison of performance at the optimal threshold (CNN at 0.72 and ResNet50 at 0.39)*

Metric	Sequential CNN	ResNet50
Overall accuracy	95.53%	95.08%
Sensitivity (Tumor Recall)	93.72%	95.29%
Specificity	96.88%	94.92%
Precision (PPV)	95.72%	93%
F1-score (Tumor)	94.71%	94%
AUC-ROC	98.45%	98.81%
Best epoch	34/50	33/50
Training accuracy	98.30%	98.27%
Validation accuracy	89.96%	93.80%
Variance (overfitting)	8.34%	4.47%
Total parameters	≈571,000	≈23,850,000

The comparative analysis reveals several key findings. In terms of overall accuracy, the two models are nearly equivalent (CNN at 95.53% vs. ResNet50 at 95.08%), with a slight advantage for the sequential CNN at its optimal threshold. However, ResNet50 performs better in terms of clinical sensitivity (95.29% vs. 93.72%) as it detects more true tumors. This is evidenced by the lower number of false negatives (9 vs. 12), a crucial factor in a medical context where missing a tumor carries serious consequences.

The CNN demonstrates superior specificity (96.88% vs. 94.92%), indicating fewer false positives. ResNet50's AUC-ROC (98.81%) exceeds that of the sequential CNN (98.45%), indicating better overall discrimination. The difference in overfitting levels (4.47% for ResNet50 vs. 8.34% for the CNN) highlights the power of transfer learning in improving generalization.

From a computational standpoint, the sequential CNN is significantly more lightweight (approximately 571,000 parameters) than ResNet50 (approximately 23.85 million parameters). This factor must be considered regarding deployment if hardware resources are limited.

6. Discussion

The results demonstrate that both approaches are clinically relevant for the automatic detection of brain tumors in grayscale MRI images. The performance of the sequential CNN built from scratch is remarkable, despite having 42 times fewer parameters than ResNet50.

6.1 Specific contribution of grayscale adaptation

Adapting ResNet50 for grayscale images using a single-to-three-channel Conv2D conversion layer constitutes a methodological contribution of our work. Unlike naive approaches that simply replicate the grayscale channel across three planes, our conversion layer learns an optimized projection that maximizes the utilization of representations pre-trained on ImageNet. This is evidenced by the reduced overfitting gap (4.47% vs. 8.34%) and confirms that the learned transformation better preserves the structure of MRI textures.

6.2 MRI image structures and deep network properties

Brain MRI images exhibit specific structures—namely textures, intensity gradients, and edge patterns—that serve as distinguishing features between tumorous and healthy tissues. The convolutional layers of ResNet50, trained on ImageNet, encode edge, texture, and shape detectors that prove transferable to these MRI structures. This structural transferability between natural and medical representations explains ResNet50's superior generalization capability compared to a sequential CNN trained from scratch.

6.3 Comparison with other architectures from the literature

To enrich the comparative analysis, we position our results relative to other architectures representative of recent literature (Table 3).

Table 3: *Positioning relative to the literature (binary classification of brain tumors in MRI)*

Study	Architecture	Accuracy	AUC-ROC	MRI images
Our study	Sequential CNN	95.53%	98.45%	Yes
Our study	ResNet50 (adapted)	95.08%	98.81%	Yes
Halloum & Ez-Zahraouy [9]	VGG16/VGG19/ResNet50	94.96%	N/A	No
Srinivasan et al. [7]	Multi-class hybrid CNN	96%	N/A	No
Swati et al. [20]	VGG19 + fine-tuning	94.8%	N/A	No

Our results are comparable to the best values reported in the literature for the binary classification of brain MRIs.

7. Study limitations

Several limitations of the study can be noted: the Kaggle dataset originates from diverse sources with varying acquisition conditions, which may affect model robustness.

Binary classification (tumor vs. non-tumor) does not distinguish between specific tumor types (e.g., gliomas, meningiomas), thereby limiting direct clinical applicability. Furthermore, the lack of external validation on an independent dataset prevents conclusions regarding the models' generalizability.

8. Conclusion

This study presented and compared two deep learning approaches for brain tumor detection using grayscale MRI images. Both architectures demonstrated high performance on a dataset of 447 images, achieving an AUC-ROC of 98% for both methods.

Despite its architectural simplicity and low parameter count, the sequential CNN achieved an overall accuracy of 95.53% and an AUC-ROC of 98.45%. ResNet stands out due to better generalization (overfitting gap of 4.47% vs. 8.34%), superior sensitivity (95.29% vs. 93.72%), and a slightly higher AUC-ROC. These results confirm the value of transfer learning.

Future work could explore several avenues for improvement, namely: extending the model to multi-class classification to distinguish between tumor types; exploring techniques to visualize activation regions; conducting a systematic comparison with VGG16 and Inception V3 on the same dataset; and utilizing a larger, more diverse dataset.

Declaration on Generative AI

During the preparation of this book, the author used ChatGPT to support the literature search synthesize related work, and assist in drafting the related work section. After using this tool, the author carefully reviewed, verified, and edited all generated content to ensure its accuracy, coherence, and scientific validity. The author takes full responsibility for the content of this publication.

References

[1] Siegel, R. L., Miller, K. D. & Jemal, A. Cancer statistics, 2019. CA Cancer J. Clin. 69, 7–34. https://doi.org/10.3322/caac.21551 (2019).

[2] Abd-Ellah, M. K., Khalaf, A. A. M., Awad, A. I. & Hamed, H. F. A. TPUAR-Net: Two parallel U-Nets with an asymmetric residual-based deep convolutional neural network for brain tumor segmentation. In Image Analysis and Recognition (eds Karray, F. et al.) 106–116 (Springer, 2019). https://doi.org/10.1007/978-3-030-27272-2_10

[3] Kai Hu, Qinghai Gan, Yuan Zhang, Shuhua Deng, Fenxia O, Wei Huang, Chunhong Cao & Xieping Gao, Vol. 4, 2016. Brain tumor segmentation using a multi-cascade convolutional neural network. https://doi.org/10.1109/ACCESS.2019.2927433

[4] Litjens, G. et al. (2017). A survey on deep learning in medical image analysis. Medical Image Analysis, 42, 60–88. https://doi.org/10.1016/j.media.2017.07.005

[5] Kanimozhi Sampath, Sivakumar Rajagopal, Ananthakrishna Chintanpalli, A comparative analysis of CNN-based deep learning architectures for early diagnosis of bone cancer using CT images. DOI: 10.1038/S41598-024-52719-8

[6] Ioffe, S. and Szegedy, C. (2015). Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. Proceedings of the 32nd International Conference on Machine Learning (PMLR). https://doi.org/10.48550/arXiv.1502.03167

[7] Saravanan Srinivasan, Divya Francis, Sandeep Kumar Mathivanan, Hariharan Rajadurai, Basu Dev Shivahare, and Mohd Asif Shah, A hybrid deep CNN model for multi-class brain tumor image classification. DOI: 10.1186/S12880-024-01195-7

[8] Osman Özkaraca, Okan İhsan Bağrıaçık, Hüseyin Gürüler, Faheem Khan, Jamil Hussain, Jawad Khan, and Umm e Laila, Multi-class brain tumor classification using a dense CNN architecture on brain MRI images. https://doi.org/10.3390/life13020349

[9] Kamal Halloum and Hamid Ez-Zahraouy, A comparison of VGG16, VGG19, and ResNet50 architectures for the classification of normal and CLAHE-processed medical images. https://doi.org/10.21203/rs.3.rs-3224583/v1

[10] Mahmoud Khaled Abd-Ellah, Ali Ismail Awad, Ashraf A. M. Khalaf, and Amira Mofreh Ibraheem, Automated brain tumor diagnosis using cascaded deep convolutional neural networks with symmetric U-Net and asymmetric residual blocks. https://doi.org/10.1038/s41598-024-59566-7

[11] Raza Iman and Mohamed Talha Alam, Optimizing brain tumor classification: a comprehensive study on transfer learning and imbalance handling in deep learning models. https://doi.org/10.48550/arXiv.2308.06821

[12] Saritha, M., Joseph, K. P., & Mathew, A. T., Brain image classification MRI using combined spider-web plots based on wavelet entropy and a probabilistic neural network. Pattern Recognit. Lett. 34, 2151–2156 (2013). https://doi.org/10.1016/j.patrec.2013.07.009

[13] Mallikarjun Mudda, N. K. & Manjunath, R. Brain tumor classification using enhanced statistical texture features. IETE J. Res. 68, 3695–3706. https://doi.org/10.1080/03772063.2020.1775501 (2022).

[14] Asiri, A. A. et al. Machine learning-based models for brain tumor classification based on magnetic resonance imaging (MRI). Intell. Autom. Soft Comput. 36, 299–312. https://doi.org/10.32604/iasc.2023.032426 (2023).

[15] Litjens G., Kooi T., Bejnordi B. E., Setio A. A. A., Ciompi F., Ghafoorian M., van der Laak J. A. W. M., van Ginneken B., and Sánchez C. I., A survey on deep learning in medical image analysis, Medical Image Analysis. (2017) 42, 60–88, https://doi.org/10.1016/j.media.2017.07.005

[16] LeCun Y., Bengio Y., and Hinton G., Deep learning, Nature. (2015) 521, no. 7553, 436–444, https://doi.org/10.1038/nature14539

[17] Zacharaki E. I., Wang S., Chawla S., Soo Yoo D., and Davatzikos E. R., Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme, Magnetic Resonance in Medicine. (2009) 62, no. 6, 1609–1618, https://doi.org/10.1002/mrm.22147

[18] Pesapane, F and other researchers (2018). Artificial intelligence in medical imaging: threat or opportunity? Radiologists once again at the forefront of medical innovation. 2(1), 35. https://doi.org/10.1186/s41747-018-0061-6

[19] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770-778. https://doi.org/10.1109/CVPR.2016.90

[20] Swati, Z.N.K., et al. (2019). Brain tumor classification for MR images using transfer learning and fine-tuning. Computerized Medical Imaging and Graphics, 75, 34-46. https://doi.org/10.1016/j.compmedimag.2019.05.001

[21] Cheng, J., et al. (2015). Enhanced performance of brain tumor classification via tumor region augmentation and partition. PLOS ONE, 10(10). https://doi.org/10.1371/journal.pone.0140381

[22] Srivastava, N., et al. (2014). Dropout: A simple way to prevent neural networks from overfitting. Journal of Machine Learning Research, 15(1), 1929-1958. https://doi.org/10.5555/2627435.2670313

[23] Kaggle Dataset: Brain MRI Images for Brain Tumor Detection. Available at: https://www.kaggle.com/datasets/navoneel/brain-mri-images-for-brain-tumor-detection (Accessed 2024)