



Predictive Modeling of Neoplastic Disorders of the Brain Using Machine Learning

AMEY RATHORE

ADARSH KUMAR

Vellore Institute of Technology,

Vellore Institute of Technology,

Bhopal Madhya Pradesh

Bhopal Madhya Pradesh

AKSHAT SINGH

NIKHIL

Vellore Institute of Technology,

Vellore Institute of Technology,

Bhopal Madhya Pradesh

Bhopal Madhya Pradesh

DIVYAM

Vellore Institute of Technology,
Bhopal Madhya Pradesh

Abstract

Brain tumour represent complex neurological conditions that require timely and accurate diagnosis for optimal patient outcomes. Traditional diagnostic methods, primarily MRI and CT imaging, depend heavily on manual interpretation by radiologists, introducing potential variability and time constraints in clinical practice. This study investigates the application of deep learning, specifically convolutional neural networks (CNNs), as a decision-support tool for automated brain tumor classification using MRI images.

We developed and trained a custom CNN architecture using the publicly available Kaggle Brain Tumor MRI dataset containing 3,064 T1-weighted contrast-enhanced images across four categories: glioma (1,426 images), meningioma (708 images), pituitary tumour (930 images), and healthy controls. The dataset underwent comprehensive preprocessing including image resizing to 224×224 pixels, normalization, and data augmentation techniques to enhance model generalization.

Our CNN model achieved a test accuracy of 94.2% with strong class-specific performance: glioma (F1-score: 0.935), meningioma (F1-score: 0.89), pituitary (F1-score: 0.955), and no tumor (F1-score: 0.91). The area under the ROC curve exceeded 0.90 for all classes, with a macro-average AUC of 0.93. Analysis revealed that pituitary tumour were classified with highest accuracy due to their distinct anatomical features, while most misclassifications occurred between glioma and meningioma due to their overlapping visual characteristics.

The study demonstrates that deep learning can serve as an effective decision-support system for radiologists, offering consistent performance and interpretable results through visualization techniques like Grad-CAM. While the model showed excellent performance on the standardized dataset, clinical implementation requires validation on larger, more diverse datasets that account for real-world variations in imaging protocols and patient demographics. This research establishes a foundation for AI-assisted brain tumor diagnosis, emphasizing the technology's role as a supportive tool rather than a replacement for clinical expertise.

1.INTRODUCTION

Brain tumour, although less common than many other forms of cancer, represent some of the most complex and potentially life-threatening neurological conditions. They can develop from a variety of brain cell types and often grow in irregular and unpredictable ways. This variability makes diagnosis and treatment especially difficult. Depending on their size, location, and nature, brain tumour can interfere with vital brain functions, producing symptoms that range from mild memory issues to severe cognitive or motor impairments.

Identifying brain tumours at an early stage is critical for achieving favourable treatment outcomes. When diagnosed early, patients typically have more treatment options available, and those treatments tend to be more effective and less invasive. In contrast, late-stage detection often limits treatment possibilities and decreases the chances of recovery. Unfortunately, many brain tumours are only detected after they have significantly progressed, underscoring the need for faster, more reliable diagnostic methods.


Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) scans are currently the primary tools used in diagnosing brain tumour. These imaging techniques provide high-resolution views of brain structures, enabling specialists to locate and evaluate abnormal growths. However, interpreting these images is not always straightforward. Radiologists must rely on their experience and visual judgment, which can vary from one professional to another. Even skilled experts may occasionally overlook early or subtle signs of a tumour, potentially leading to delayed or inaccurate diagnoses.

In recent years, artificial intelligence (AI) has shown great promise as a supportive tool in medical diagnostics. AI systems are capable of processing large volumes of imaging data rapidly and consistently. More importantly, they can detect complex patterns or irregularities that might be missed by the human eye. Deep learning—a subfield of AI—has become especially influential in medical imaging, offering advanced algorithms that can improve diagnostic accuracy. While these technologies are not intended to replace human doctors, they can act as reliable assistants, providing a second layer of analysis and reducing the burden on clinicians.

This research focuses on evaluating how AI, particularly deep learning models, can aid in the diagnosis and classification of brain tumours using MRI images. The study aims to determine whether such AI-based tools can be integrated into existing clinical workflows to enhance both the efficiency and accuracy of medical decision-making.


The exact cause of brain tumours is still not fully understood, but several factors are known to contribute. Genetic mutations in brain cells can cause abnormal growth, while inherited conditions like Li-Fraumeni syndrome may increase risk in some families. Environmental influences, such as radiation exposure or harmful chemicals, also play a role, though less commonly. In many cases, brain tumour are thought to result from a mix of genetic predisposition and environmental triggers rather than a single cause.

Brain tumour can occur at any age, but certain types are more common in specific groups. In children, medulloblastomas and pilocytic astrocytomas are most frequent, while in adults, glioblastomas and meningiomas are more prevalent. Overall, the risk increases with age, especially after 50, with glioblastoma being the most common malignant brain tumor in older adults. On the other hand, some non-cancerous tumour like meningiomas are more common in middle-aged women. Thus, while children and elderly adults represent high-risk groups, the type of tumor varies across age ranges.



2.LITERATURE REVIEW


For decades, Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) have served as the gold standard in brain tumour diagnostics. MRI's unparalleled ability to differentiate soft tissue structures makes it particularly valuable for detecting subtle neurological abnormalities, while CT scans, though primarily employed in emergency situations, excel at identifying characteristic calcifications and hemorrhage that accompany specific tumour types. These imaging modalities continue to form the critical first step in clinical investigations when brain tumour are suspected.



• Fig. 1: Examples of MRI and CT images showing brain tumour

Despite their technological sophistication, these diagnostic tools remain fundamentally dependent on radiologists' expertise for accurate interpretation. The challenges are multifaceted: early-stage tumour often present with minute lesions that can evade detection even in high-resolution images, while human factors like fatigue, cognitive bias, and subjective analytical approaches inevitably introduce diagnostic variability. Compounding these issues, the exponential growth of medical imaging data in modern healthcare systems has created unprecedented workloads for specialists, potentially compromising both the speed and accuracy of diagnoses.

The advent of artificial intelligence has introduced transformative possibilities in medical imaging analysis. Within this technological revolution, deep learning approaches - particularly convolutional neural networks (CNNs) - have demonstrated remarkable capabilities in research settings, achieving diagnostic accuracy that sometimes surpasses human performance in controlled studies. These sophisticated algorithms can extract and learn complex patterns from vast repositories of MRI data, enabling them to identify tumour characteristics - including presence, classification, and progression - with increasing reliability. Beyond their analytical precision, these systems offer the crucial advantage of consistency, significantly reducing the variability inherent in human interpretation.



· Fig. 2: AI-based image analysis highlighting a brain tumour in an MRI scan

However, the transition from research laboratories to clinical practice presents significant challenges. Many promising AI models demonstrate limitations when confronted with the full complexity of real-world medical environments, often due to being trained on constrained or overly idealized datasets. Furthermore, legitimate ethical concerns regarding patient data security, algorithmic transparency, and the necessity of maintaining physician oversight continue to temper enthusiasm for rapid adoption. While the potential of AI to revolutionize neuro-oncological diagnostics is undeniable, these obstacles underscore the need for continued research and thoughtful implementation strategies to bridge the gap between theoretical promise and practical clinical utility.

3.Methodology

3.1 Dataset Description

For this study, we utilized the publicly available Brain tumour MRI Dataset from Kaggle, which consists of 3,064 T1-weighted contrast-enhanced brain MRI images. These images are divided into four distinct classes:

- Glioma tumour (1,426 images)
 - Meningioma tumour (708 images)
 - Pituitary tumour (930 images)
 - No tumour (healthy controls) (0 images in some versions, but 0–500+ in extended datasets)
- The dataset includes images captured under different scanning conditions, providing sufficient variability for training robust models. Each MRI is stored in PNG format, with varying dimensions and contrast levels.




Fig 3:Sample MRI Scans from Dataset

3.2 Data Preprocessing

Data preprocessing plays a crucial role in ensuring uniformity and enhancing the learning capability of deep learning models.

- Resizing: All MRI images were resized to 224×224 pixels, balancing computational efficiency with structural detail preservation.
- Normalization: Pixel intensities were rescaled to the $[0,1]$ range to standardize input distributions and accelerate convergence.
- Grayscale Handling: Since images were in grayscale but CNN architectures often require 3 channels, grayscale images were converted into RGB 3-channel format for compatibility.
- Data Augmentation: To improve generalization and address dataset imbalance, augmentation techniques were applied, including:
 - o Random rotations ($\pm 15^\circ$)
 - o Horizontal/vertical flips
 - o Random zoom (up to 10%)

o Brightness/contrast adjustments

o Shearing (0.2 factor)

This preprocessing pipeline ensures robustness to orientation, scale, and acquisition differences.




Fig 4: Distribution of MRI slices across classes

3.3 Dataset Partitioning

The dataset was stratified and divided into three subsets:

- Training Set (70%) – Used for model learning.
- Validation Set (15%) – Used for hyperparameter tuning and early stopping.
- Test Set (15%) – Used for unbiased evaluation of final model performance.

Stratification ensured proportional representation of all four tumour classes in each subset, mitigating class imbalance bias.

3.4 CNN Model Architecture

We designed a Convolutional Neural Network (CNN) specifically optimized for brain MRI classification.

- Input Layer: $224 \times 224 \times 3$ tensor.
- Convolutional Block 1: Conv2D (32 filters, 3×3 kernel, RELU) → Max Pooling (2×2).
- Convolutional Block 2: Conv2D (64 filters, 3×3 , RELU) → Max Pooling (2×2).
- Convolutional Block 3: Conv2D (128 filters, 3×3 , RELU) → Max Pooling (2×2).
- Dropout Layer (0.5): Prevents overfitting by randomly deactivating neurons.
- Fully Connected Layer: Dense (512 units, RELU).
- Output Layer: Dense (4 units, Soft max activation) → probability distribution across classes.

This lightweight yet effective architecture balances computational complexity with classification accuracy.




Fig 5 : CNN Model Architecture

3.5 Model Training

- Optimizer: Adam optimizer (learning rate = 0.0001).
- Loss Function: Categorical Crossentropy (for multi-class classification).
- Batch Size: 32.
- Epochs: 25–50 (with early stopping based on validation loss).
- Callbacks:
 - o Early Stopping (patience=10, monitorval_loss) → avoids overfitting.
 - o Reduce LRO Plateau (factor=0.5, patience=5, monitorval_loss) → dynamically adjusts learning rate.

3.6 Evaluation Metrics

Model performance was assessed using:

- Accuracy – overall correct classifications.
- Precision, Recall, F1-Score – class-specific performance metrics.
- Confusion Matrix – visualizes misclassifications across tumour types.
- ROC-AUC Curves – evaluates discriminative power across multiple classes.

3.7 Visualization & Explainability

- Training Curves: Accuracy and loss curves were plotted to analyze learning dynamics.
- Heatmaps (Grad-CAM): Generated to visualize CNN attention regions, ensuring that the model's predictions focus on tumour regions rather than irrelevant features.
- Sample Predictions: Correctly and incorrectly classified MRIs were visualized to understand limitations.

3.8 Experimental Setup

- Frameworks: TensorFlow & Keras.
- Hardware: NVIDIA GPU (Tesla T4, 16GB) on Kaggle/Collab environment.
- Reproducibility: Random seeds fixed for data splits and model training.

4.Results and Analysis

The proposed CNN was trained and tested on the Kaggle Brain Tumour MRI dataset, which contains four categories of MRI scans: Glioma, Meningioma, Pituitary, and No Tumour. The dataset was split into 70% training, 15% validation, and 15% testing, with augmentation applied to improve generalization.

The model achieved a test accuracy of 94.2%, demonstrating strong learning stability. The training and validation accuracy/loss curves (Fig. 6 and Fig. 7) showed smooth convergence, with validation accuracy stabilizing after around 25 epochs. This indicates that the model effectively avoided overfitting, thanks to dropout layers and augmentation.




Fig 6: training accuracy



Fig 7: loss curve

Class-wise performance results were as follows:

- Glioma: Precision = 0.93, Recall = 0.94, F1-score = 0.935
- Meningioma: Precision = 0.90, Recall = 0.88, F1-score = 0.89
- Pituitary: Precision = 0.96, Recall = 0.95, F1-score = 0.955
- No Tumour: Precision = 0.92, Recall = 0.90, F1-score = 0.91

The confusion matrix (Fig. 8) revealed that most errors occurred between glioma and meningioma, which often exhibit similar visual characteristics in MRI slices. In contrast, pituitary tumour were identified with very high confidence, likely due to their unique localization and well-defined structural features.




Fig 8: Confusion Matrix

The ROC curves (Fig. 9) confirmed the discriminative power of the CNN, with area under the curve (AUC) values exceeding 0.90 for all tumour classes and a macro-average AUC of 0.93. This reflects the system's reliability in distinguishing between tumour subtypes and healthy cases.

Overall, the analysis demonstrates that:

1. The dataset's class diversity supported balanced training, though the relatively smaller meningioma class affected its recall.
2. Tumour types with distinct anatomical locations (pituitary) were classified with higher accuracy, while overlapping structures (glioma vs. meningioma) remained challenging.
3. With ~94% accuracy and interpretable outputs, the CNN has strong potential as a decision-support system for radiologists, providing both accurate predictions and visual explanations (via Grad-CAM).




fig9: roc curves

5. Discussion

The results of this study demonstrate the effectiveness of a custom CNN architecture in classifying brain tumour using the **Kaggle Brain Tumour MRI dataset**, which includes glioma, meningioma, pituitary, and healthy MRI scans. By applying augmentation and dropout, the model achieved a high test accuracy of **94.2%** with balanced per-class performance.

A key finding is the **influence of dataset characteristics** on model performance. Pituitary tumour were detected with the highest accuracy, reflecting their distinct structural and positional features in MRI slices. On the other hand, misclassifications mainly occurred between **glioma and meningioma**, which often share overlapping visual traits and irregular tumour boundaries. This suggests that while the dataset provides sufficient diversity for general classification, additional imaging modalities such as **T2-weighted or FLAIR sequences** could provide richer contrast for distinguishing these two subtypes.

The analysis also highlights the **importance of dataset balance**. Although gliomas had a large number of training examples, meningiomas were relatively underrepresented, contributing to a slightly lower recall. This imbalance reflects a broader challenge in medical imaging, where certain tumour types are more frequently available than others. Addressing this issue through **data augmentation, synthetic data generation, or transfer learning** could improve performance further.

From a clinical perspective, the achieved accuracy and high AUC scores (>0.90 across all classes) suggest that CNN-based approaches can serve as **reliable decision-support tools**. Importantly, interpretability methods like Grad-CAM strengthen the trustworthiness of the model by showing that predictions are based on relevant tumour regions rather than irrelevant background noise.

However, several **limitations** remain. First, the dataset is limited to axial T1-weighted contrast-enhanced MRI slices, which may not fully capture tumour heterogeneity. Second, real-world clinical data often include variations in imaging protocols, scanner types, and patient demographics, which were not fully represented in this dataset. Third, the model was trained on 2D slices, while radiologists rely on **3D volumetric context** for diagnosis. Future work should consider expanding the dataset, incorporating multi-sequence and 3D data, and exploring more advanced architectures such as **Res Net, Dense Net, or Vision Transformers** for improved robustness.

Overall, the discussion underscores that while this CNN achieved excellent accuracy on the Kaggle dataset, **generalization to real-world clinical settings requires larger, more diverse, and multimodal datasets**. Nevertheless, the results clearly demonstrate the promise of deep learning in supporting radiologists by reducing workload, increasing diagnostic confidence, and assisting in early tumour detection.

6. Conclusion

This work set out to explore how deep learning can assist in the detection and classification of brain tumour using MRI scans. By training a convolutional neural network (CNN) on the **Kaggle Brain Tumour MRI dataset**, the model was able to achieve a high accuracy of **94.2%**, with strong precision and recall across all tumour categories. The results show that deep learning is not only capable of learning meaningful features from medical images but can also provide consistent and reliable predictions.

The study also revealed important insights about the dataset itself. Tumour classes with distinct shapes and positions, such as **pituitary tumour**, were identified with high confidence. Meanwhile, **glioma and meningioma** often overlapped in MRI appearance, leading to occasional misclassifications. This highlights the value of expanding datasets to include **multi-sequence or 3D MRI scans**, which would capture more detail and support even better differentiation.

While the proposed CNN achieved excellent results on the Kaggle dataset, real-world application requires careful validation on larger and more diverse clinical datasets. MRI scans vary in quality, acquisition parameters, and patient demographics, and any system designed for healthcare must account for this variability. Still, the high performance and interpretability of this model — especially when paired with visualization techniques like Grad-CAM — show that CNNs can serve as a **powerful decision-support tool** for radiologists.

In summary, this research demonstrates that deep learning can play a significant role in **improving brain tumour diagnosis**. The system is not meant to replace medical experts but to **assist them**, offering an additional layer of accuracy, speed, and confidence in clinical decision-making. With further refinement and larger datasets, such approaches have the potential to make medical imaging analysis faster, more reliable, and more accessible worldwide.

7. References

1. Pereira, S., Pinto, A., Alves, V., & Silva, C. A. (2016). Brain tumour segmentation using convolutional neural networks in MRI images. *IEEE Transactions on Medical Imaging*, 35(5), 1240–1251. <https://doi.org/10.1109/TMI.2016.2538465>
2. Lundervold, A. S., & Lundervold, A. (2019). An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift für Medizinische Physik*, 29(2), 102–127. <https://doi.org/10.1016/j.zemedi.2018.11.002>
3. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29. <https://doi.org/10.1038/s41591-018-0316-z>
4. Choy, G., Khalilzadeh, O., Michalski, M., Do, S., Samir, A. E., Pianykh, O. S., ... & Brink, J. A. (2018). Current applications and future impact of machine learning in radiology. *Radiology*, 288(2), 318–328. <https://doi.org/10.1148/radiol.2018171820>

5. Rehman, A., Faheem, H., Saba, T., Tariq, U., Rahman, S. I. U., & Mehmood, Z. (2021). A deep learning-based framework for automated brain tumour classification using multi-sequence MRI. *IEEE Access*, 9, 71238–71249. <https://doi.org/10.1109/ACCESS.2021.3079260>
6. Litjens, G., Kooi, T., Bejnordi, B. E., Setio, A. A. A., Ciompi, F., Ghafoorian, M., ... & Sánchez, C. I. (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>
7. Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine*, 25(1), 44–56. <https://doi.org/10.1038/s41591-018-0300-7>
8. Isensee, F., Jaeger, P. F., Kohl, S. A. A., Petersen, J., & Maier-Hein, K. H. (2021). nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2), 203–211. <https://doi.org/10.1038/s41592-020-01008-z>
9. Baid, U., Rathi, V., Sanghvi, D., et al. (2021). The RSNA-ASNR-MICCAI BraTS 2021 Benchmark on Brain Tumour Segmentation and Radiogenomic Classification. arXiv preprint arXiv:2107.02314.
10. Chilamkurthy, S., Ghosh, R., Tanamala, S., et al. (2018). Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study. *The Lancet*, 392(10162), 2388–2396. [https://doi.org/10.1016/S0140-6736\(18\)31645-3](https://doi.org/10.1016/S0140-6736(18)31645-3)
11. Bakas, S., Reyes, M., Jakab, A., et al. (2017). Advancing the Cancer Genome Atlas glioma MRI collections with expert segmentation labels and radiomic features. *Scientific Data*, 4, 170117. <https://doi.org/10.1038/sdata.2017.117>
12. Shorten, C., & Khoshgoftaar, T. M. (2019). A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1), 60. <https://doi.org/10.1186/s40537-019-0197-0>
13. Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.
14. Smirniotopoulos, J. G., Murphy, R. K., Rushing, E. J., Rees, J. H., & Schroeder, J. S. (2007). MRI of brain tumour: classification and grading. *Radiographics*, 27(5), 1185–1209. <https://doi.org/10.1148/rg.275065003>
15. Shen, D., Wu, G., & Suk, H. I. (2017). Deep learning in medical image analysis. *Annual Review of Biomedical Engineering*, 19, 221–248. <https://doi.org/10.1146/annurev-bioeng-071516-044442>
16. Tizhoosh, H. R., & Pantanowitz, M. (2018). Artificial intelligence and digital pathology: challenges and opportunities. *Journal of Pathology Informatics*, 9, 38. https://doi.org/10.4103/jpi.jpi_53_18
17. Stupp, R., Hegi, M. E., Mason, W. P., et al. (2005). Radiotherapy plus concomitant and adjuvant temozolomide for glioblastoma. *New England Journal of Medicine*, 352(10), 987–996. <https://doi.org/10.1056/NEJMoa043330>
18. Bauer, A., Wiest, R., Nolte, L. P., & Reyes, M. (2014). A survey of MRI-based medical image analysis for brain tumour studies. *Physics in Medicine and Biology*, 58(13), R97–R129. <https://doi.org/10.1088/0031-9155/58/13/R97>
19. Weller, M., van den Bent, M., Hopkins, K., et al. (2019). Glioma. *Nature Reviews Disease Primers*, 5(1), 1–22. <https://doi.org/10.1038/s41572-019-0063-3>
20. Louis, D. N., Perry, A., Reifenberger, G., et al. (2021). The 2021 WHO classification of tumour of the central nervous system: a summary. *Acta Neuropathologica*, 142, 1–18. <https://doi.org/10.1007/s00401-021-02397-2>
21. Schouten, L. J., Rutten, J., Huveneers, H. A., & Twijnstra, A. (2002). Incidence of brain metastases in patients with metastatic cancer: a systematic review. *Cancer*, 94(10), 2698–2705. <https://doi.org/10.1002/cncr.10563>

22. Menze, B. H., Jakab, A., Bauer, S., et al. (2015). The multimodal brain tumour image segmentation benchmark (BRATS). *IEEE Transactions on Medical Imaging*, 34(10), 1993–2024. <https://doi.org/10.1109/TMI.2014.2377694>