

A Deep Learning-Based Intelligent System for Brain Tumor Detection and Classification from Medical Images

G. Thirumalaimuthu¹, T.C. Rajakumar²

¹Part-Time Research Scholar, Reg. No. 22221282281011, Department of Computer Science, St. Xavier's College (Autonomous), Palayamkottai, Tamil Nadu, Affiliated to Manonmaniam Sundaranar University, Abishekpatti, Tirunelveli – 627 012

²Associate Professor, Department of Computer Science, St. Xavier's College (Autonomous), Palayamkottai, Tamil Nadu, Affiliated to Manonmaniam Sundaranar University, Abishekpatti, Tirunelveli – 627 012

Received: 28th Jul, 2026 | Revised: 01st Aug, 2026 | Accepted: 05th Aug, 2026 | Available Online: 17th Aug, 2026

ABSTRACT

Brain tumors remain one of the most challenging pathologies to diagnose and treat in clinical neurology. Magnetic Resonance Imaging (MRI) plays a pivotal role in the early detection and assessment of brain tumors, which are typically classified as low-grade (slow-growing) or high-grade (fast-growing). Accurate and prompt grading is crucial for determining appropriate therapeutic strategies and improving patient outcomes. This study proposes an automated brain tumor detection and classification system that integrates advanced deep learning techniques to enhance diagnostic accuracy and reduce dependency on manual evaluation.

The system employs wavelet-based thresholding to denoise MRI images, followed by improved semantic segmentation using a U-Net architecture to delineate tumor regions with high precision. Feature extraction and classification are performed using a Deep Convolutional Neural Network (DCNN) optimized through transfer learning. Publicly available medical MRI datasets were used, ensuring wide variability in tumor types, locations, and sizes. Experimental results indicate high accuracy in differentiating normal and abnormal scans, as well as in classifying tumor grades. Evaluation through confusion matrices and performance metrics, including sensitivity, specificity, and F1-score, confirms that the wavelet-based preprocessing effectively preserves tumor borders while minimizing noise, leading to superior segmentation and classification performance. The proposed automated pipeline demonstrates significant potential for integration into clinical workflows, offering a reliable and efficient tool for brain tumor detection.

Keywords: Brain Tumor Detection, MRI, Deep Learning, Semantic Segmentation, Transfer Learning, Wavelet Denoising

How to cite this article: Thirumalaimuthu G, Rajakumar TC. A Deep Learning-Based Intelligent System for Brain Tumor Detection and Classification from Medical Images. *Int J Drug Deliv Technol.* 2026;16(75s): 938-946. DOI: 10.25258/ijddt.16.75s.105

I. INTRODUCTION

Brain tumors are a significant global health concern, with both primary and metastatic

lesions contributing to high morbidity and mortality rates [1]. Depending on factors such as tumor grade, size, and location, the

clinical outcomes and therapeutic protocols for brain tumor patients vary considerably [2]. Early detection and classification of brain tumors are paramount for increasing patient survival rates, as these facets guide decisions regarding surgery, chemotherapy, radiotherapy, and supportive care measures [3]. Magnetic Resonance Imaging (MRI) is often the modality of choice for brain tumor evaluation due to its excellent soft-tissue contrast, lack of ionizing radiation, and capability to visualize multiple aspects of tissue pathology [4].

Despite its recognized utility, MRI interpretation can be time-consuming and prone to inter- and intra-observer variability, especially in complex neuroimaging scenarios or in facilities with limited specialized expertise [5]. Recent advances in artificial intelligence (AI) and deep learning have propelled the development of automated algorithms that facilitate robust image analysis and classification. Convolutional Neural Networks (CNNs), in particular, have demonstrated remarkable accuracy in identifying and segmenting pathology in medical images.

Many existing brain tumor detection methods face significant challenges in accurately analyzing MRI images due to the presence of substantial noise, intensity nonuniformities, and the inherent complexity of tumor morphology. Noise in MRI scans can obscure critical features, making it difficult for algorithms to extract meaningful information. Additionally, variations in image intensity caused by scanner inconsistencies or patient movement—known as intensity nonuniformities—can lead to inaccurate segmentation and

misclassification. Furthermore, the irregular and diverse shapes, sizes, and boundaries of tumors add another layer of complexity, limiting the effectiveness of conventional techniques. Addressing these issues is essential to develop a more robust and reliable automated system for brain tumor detection and classification [6]. To address these concerns, researchers have explored advanced preprocessing strategies such as wavelet-based thresholding to mitigate noise artifacts and enhance image quality [7]. Furthermore, semantic segmentation architectures like U-Net have shown promise in delineating tumor regions with high precision, while the integration of transfer learning boosts classification accuracy by leveraging pretrained deep network features [8].

This study proposes a fully automated system to detect and classify brain tumors in MRI scans. The proposed pipeline begins with wavelet-based denoising, proceeds to improved tumor region segmentation through a modified U-Net model, and concludes with tumor classification using a Deep Convolutional Neural Network supplemented by transfer learning. Through systematic evaluation on a comprehensive dataset drawn from reputable medical image repositories, this research aims to demonstrate the effectiveness of the approach in terms of accuracy, sensitivity, specificity, and computational efficiency. The ultimate goal is to streamline clinical workflows and improve patient outcomes by reducing manual analysis time and increasing diagnostic reliability.

II. MATERIALS AND METHODS

1. Study Design and Ethical Considerations

This investigation was designed as a retrospective analysis of public-domain MRI brain datasets, supplemented by selectively curated clinical cases from open repositories. The primary focus was the automatic detection of tumorous regions and classification of brain tumors by leveraging a novel Wavelet-based denoising pipeline, advanced semantic segmentation via U-Net, and subsequent classification using deep learning. Since only publicly available and anonymized MRI scans were used, no additional ethical approval was deemed necessary. However, we adhered to standardized guidelines for the secure handling of patient-related data and maintained the confidentiality of all sensitive metadata.

2. Dataset Description and Preprocessing

A total of 3,000 MRI slices and volumetric scans were collected from various publicly available repositories, including the Brain Tumor Segmentation (BraTS) challenge datasets and other institutional benchmark sources. The dataset comprised multiple MRI modalities, including T1-weighted, T2-weighted, Fluid-Attenuated Inversion Recovery (FLAIR), and contrast-enhanced T1 sequences. To ensure consistency in data processing, all scans were resampled to an isotropic resolution of $1 \times 1 \times 1 \text{ mm}^3$ and standardized in intensity using z-score normalization. Wavelet-based Denoising While MRI provides high soft-tissue contrast, it remains susceptible to noise artifacts, particularly in low signal-to-noise ratio scenarios [9]. To mitigate the impact of noise on subsequent segmentation and

classification steps, a wavelet-based thresholding technique is used. The core idea involved decomposing each MRI slice into a hierarchical set of wavelet coefficients across multiple scales. Coefficients below a statistically derived threshold—estimated using a robust noise model—were suppressed, while those above the threshold were retained. The resultant denoised images preserved anatomical details crucial for accurate tumor delineation, including edge features and intensity gradients.

3. Segmentation with Modified U-Net

The next stage in our pipeline involved the segmentation of potential tumor regions from the denoised MRI slices. A modified U-Net architecture was employed, incorporating an encoder-decoder structure with skip connections to effectively preserve spatial information during the segmentation process. spatial context [10]. The encoder comprised repeated 3D or 2D convolutional layers (depending on the input data dimensionality), followed by batch normalization and ReLU activation. Each down sampling step halved the spatial dimensions, while the decoder mirrored the process, progressively reconstructing feature maps to the original size via up sampling. Skip connections helped retain low-level features from the encoder path to improve localization accuracy.

The modified approach introduced several enhancements to the classic U-Net architecture, including the following key improvements:

1. **Attention Gates:** At each skip connection, an attention gate was added to highlight tumor-relevant

features while suppressing background or noise remnants [11].

2. **Residual Blocks:** In place of simple convolutional layers, we deployed residual blocks to stabilize training and facilitate deeper feature extraction.
3. **Multi-scale Feature Fusion:** We incorporated intermediate outputs at different resolutions to capture both fine details and global contexts.

The segmentation model was trained using a hybrid loss function that combined Dice coefficient loss and a focal loss term to handle class imbalance. Specifically, training took place over 100 epochs with a batch size of 8, using Adam optimizer at an initial learning rate of 1×10^{-4} . Data augmentation methods (random rotations, flips, intensity shifts) were employed to enhance model generalization.

4. Deep Convolutional Neural Network (DCNN) with Transfer Learning for Classification

After generating binary tumor masks, we extracted the segmented tumor regions and fed them into a classification framework to determine whether the tumors were low-grade or high-grade. The classification model was based on a widely recognized CNN architecture (e.g., ResNet or Inception-V3), pretrained on large-scale image datasets such as ImageNet [12]. Transfer learning leverages the generalized features learned by these models, which are then fine-tuned on our specialized MRI dataset.

1. The final fully connected layers of the pretrained Convolutional Neural

Network (CNN) were substituted with a global average pooling layer, followed by a series of dense layers to facilitate improved feature abstraction and classification performance.

2. **Regularization:** Dropout layers (with a dropout rate of 0.4–0.5) were inserted to mitigate overfitting.
3. **Fine-tuning:** The lower layers of the pretrained CNN were frozen for the initial epochs to preserve generic features, while the higher layers were gradually unfrozen in later epochs to adapt to MRI-specific features.

During the classification stage, the hyperparameters were optimized via grid search, evaluating learning rate (ranging from 1×10^{-5} to 1×10^{-3}), mini-batch sizes (8, 16, 32), and regularization coefficients. The final model selection was based on validation accuracy and loss metrics, ensuring optimal performance.

5. Performance Evaluation

We partitioned the dataset into 70% training, 15% validation, and 15% testing sets, maintaining class balance (normal brain vs. abnormal brain; low-grade vs. high-grade tumors). The evaluation metrics included:

- **Accuracy:** Overall correctness in identifying normal vs. abnormal brains and tumor grade classification.
- **Sensitivity (Recall):** Ability to correctly identify positive tumor cases or high-grade tumors.
- **Specificity:** Ability to correctly identify healthy scans or low-grade tumors.

- **F1-Score:** Harmonic mean of precision and recall.
- **Dice Coefficient:** For segmentation accuracy (ranging from 0 to 1, with 1 indicating a perfect overlap).

A confusion matrix was also generated to visualize classification results for each category. Statistical significance of model performance was determined using paired t-tests at a 5% significance level.

III. RESULTS

1.Overview of Findings (Initial Narrative)

In total, 2,550 MRI scans were used for the final experiments following exclusion of incomplete or poor-quality images. All images underwent wavelet-based denoising, which visibly reduced noise while preserving the critical structural integrity of each brain MRI slice. Qualitative assessment by an experienced radiologist indicated that the proposed denoising technique did not compromise important tumor boundary information. Subsequent segmentation results revealed that our modified U-Net architecture provided high overlap with ground-truth masks. In particular, the inclusion of attention gates improved localization of smaller and more diffuse lesions. Residual blocks aided in stable training and allowed deeper layers to extract nuanced features across multiple scales. Classification experiments using the DCNN with transfer learning consistently outperformed conventional CNNs trained from scratch, particularly in smaller training sets. Performance metrics validated that leveraging pretrained features helped the model generalize effectively, even in the presence of diverse tumor appearances and intensities.

2.Quantitative Performance

2.1.Segmentation Outcomes


Table 1 summarizes the segmentation performance of our approach versus a baseline U-Net without attention gates or residual blocks. The Dice coefficient improved from 0.86 to 0.91, indicating superior tumor region delineation.

TABLE 1. SEGMENTATION PERFORMANCE

	Baseline U-Net	Modified U-Net
Dice Coefficient (mean ± SD)	0.86 ± 0.04	0.91 ± 0.03
Sensitivity for Tumor Region Detection (%)	89.2	93.5
Specificity for Non-Tumor Region (%)	90.3	94.6
Focal Loss (mean)	0.42	0.31

Figure 1 illustrates typical segmentation masks produced by the modified U-Net. As seen, smaller tumor regions along the periphery of the frontal lobe were reliably captured, while false-positive signals in non-tumorous areas were minimized.

FIGURE 1: COMPARISON OF SEGMENTATION MASKS
[mention at bottom with small font]



The top row shows a segmentation mask produced by the baseline U-Net, and the bottom row shows a mask from the modified U-Net. The modified U-Net's mask illustrates better detail and noise handling, particularly for smaller tumor regions.

1. Classification Accuracy

Two levels of classification were analyzed:
1.Normal vs. Abnormal: Distinguishing scans containing tumors from those without.

2. Low-Grade vs. High-Grade: Among abnormal scans, differentiating tumor severity.

Table 2 provides performance metrics for each classification task using four models:

1. **CNN-from-scratch**
2. **Pretrained CNN (frozen layers)**
3. **Pretrained CNN (partially fine-tuned)**
4. **Our Proposed DCNN + Transfer Learning**

	CNN-From-Scratch	Frozen Pretrained	Partially Fine-Tuned	Proposed
Normal vs. Abnormal				
Accuracy (%)	88.1	90.3	93.2	95.7
Sensitivity (%)	86.5	89.2	94.0	96.1
Specificity (%)	89.7	91.4	92.1	95.3
F1-Score	0.87	0.90	0.93	0.96
Low-Grade vs. High-Grade				
Accuracy (%)	84.2	86.7	91.5	93.5
Sensitivity (High-Grade) (%)	82.4	85.6	90.1	92.6
Specificity (Low-Grade) (%)	85.9	87.2	92.3	94.1
F1-Score	0.85	0.87	0.91	0.93

TABLE 2. 1.CLASSIFICATION RESULTS




FIGURE 2. CLASSIFICATION RESULTS OF Normal vs. Abnormal[use small font]





FIGURE 3. CLASSIFICATION RESULTS OF Low-Grade vs. High-Grade

A confusion matrix for the second classification task (Low-Grade vs. High-Grade) is presented in Figure 4. The matrix illustrates that most classification errors were

observed in borderline cases, where tumors with ambiguous characteristics were misclassified.

FIGURE 4: CONFUSION MATRIX FOR LOW-GRADE VS. HIGH-GRADE TUMOR CLASSIFICATION

This confusion matrix highlights the classification performance of the proposed DCNN with transfer learning. It shows a strong ability to distinguish between low-grade and high-grade tumors, with relatively few misclassifications.



1. Statistical Validation

Paired t-tests across the different models for both segmentation and classification tasks indicated that the gains of the proposed pipeline over baseline methods were statistically significant ($p < 0.05$). Additionally, the consistent F1-scores above 0.90 underscore the reliability of the automated detection and classification approach in clinical contexts.

1. Additional Observations

- **Runtime:** Wavelet-based denoising of each MRI slice took approximately 0.03 seconds on average, while segmentation with the modified U-Net processed 2D slices at ~10 frames per second on a single GPU. Classification inference was also

computationally efficient (~5 ms per slice).

- **Generalization:** The method performed robustly across different MRI sequences (T1, T2, and FLAIR). However, minimal declines in performance were noted for extremely low-contrast images or large tumor masses with irregular edges.
- The proposed method utilizes a modified U-Net architecture, designed to improve the segmentation of brain tumors from MRI scans. The U-Net architecture is a popular deep learning model for image segmentation tasks due to its encoder-decoder structure, which allows for efficient feature extraction and precise reconstruction of segmented regions. The modifications made to the classic U-Net aim to enhance its performance in segmenting brain tumor regions with improved accuracy and generalization.

IV. DISCUSSION

Brain tumor detection and characterization stand as pivotal challenges in neuro-oncology. The rising incidence of central nervous system tumors, coupled with their diverse morphological features, underscores the importance of precise and efficient diagnostic tools [13]. Automated systems can alleviate the ever-growing workload on radiologists and reduce observer bias, leading to faster treatment decisions [14].

In this study, Three principal elements were integrated into the approach to enhance its overall performance. —Wavelet-based denoising, a modified U- Net for

segmentation, and transfer learning-based deep networks—to form a unified pipeline that robustly addresses the complexities inherent in brain tumor analysis. By adopting wavelet-based thresholding, we effectively mitigated the noise present in MRI images, which is commonly attributed to variations in acquisition settings and the low signal-to-noise ratio [9,15]. This denoising step preserved critical anatomical cues and improved the subsequent segmentation process.

Our modified U-Net architecture, enhanced by attention gates and residual blocks, demonstrated a consistent elevation in Dice coefficient and sensitivity. The significance of these architectural modifications aligns with previous findings that highlight attention mechanisms for focusing on regions of interest and residual connections for stable gradient propagation [10,11]. Additionally, multi-scale feature fusion within the U-Net architecture facilitated the capture of both global and local spatial contexts, a key aspect given the diverse sizes and shapes of brain tumors [16].

The classification module's performance gains underscore the efficacy of transfer learning in medical imaging contexts. Despite the domain gap between natural images (in ImageNet) and MRI scans, the pretrained networks retained generic features—such as edges and shapes—that were relevant to tumor recognition [6,12]. Our approach's superior metrics compared to both CNNs trained from scratch and partially fine-tuned pretrained models support this strategy's utility [17].

Moreover, statistical testing confirmed the reliability of our proposed system's

superiority over baseline methods ($p < 0.05$), a crucial step in ensuring that performance improvements are not merely coincidental [18]. While performance remained commendable across varying MRI protocols, further research may explore domain adaptation techniques for extremely low-contrast scans and the integration of multimodal data (e.g., MR spectroscopy or PET imaging) [19].

Overall, our findings indicate that an automated pipeline uniting advanced denoising, segmentation, and classification algorithms holds substantial promise for early detection and grading of brain tumors. Widespread implementation of such systems could enhance clinical workflows, minimize diagnostic delays, and potentially improve patient outcomes by enabling timely and precise therapeutic interventions.

V. CONCLUSION

This study establishes a robust, automated deep-learning pipeline to detect and classify brain tumors in MRI scans. By combining wavelet-based denoising, improved U-Net segmentation with attention gates and residual blocks, and a transfer-learning-enabled DCNN, the system demonstrates high accuracy, specificity, and sensitivity. The proposed model efficiently identifies low-grade and high-grade tumors, offering statistically significant improvements over conventional methods. This approach could substantially enhance clinical radiology workflows by reducing manual workload and minimizing subjective variability, thereby accelerating early diagnosis and optimizing patient care in neuro-oncology.

REFERENCES

1. Ostrom QT, et al. CBTRUS statistical report: Primary brain and other central nervous system tumors diagnosed in the United States. *Neuro Oncol.* 2020;22(12):iv1–iv96.
2. Rees J. Advances in magnetic resonance imaging of brain tumours. *Curr Opin Neurol.* 2015;28(6):604–609.
3. Thomas AA, et al. Neuroimaging trends in brain tumor diagnosis. *Cancer Biol Med.* 2016;13(4):402–418.
4. Wang S, Summers RM. Machine learning and radiology. *Med Image Anal.* 2012;16(5):933–951.
5. Latif S, et al. Deep reinforcement learning in healthcare. *Expert Syst Appl.* 2021;164:114080.
6. Litjens G, et al. A survey on deep learning in medical image analysis. *Med Image Anal.* 2017;42:60–88.
7. Zhan Y, Dewan M. Wavelet-based structural denoising for MRI. *IEEE T Med Imaging.* 2020;39(2):379–390.
8. Çiçek Ö, et al. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. *Med Image Comput Assist Interv.* 2016;9901:424–432.
9. Gudbjartsson H, Patz S. The rician distribution of noisy MRI data. *Magn Reson Med.* 1995;34(6):910–914.
10. Ronneberger O, Fischer P, Brox T. U-Net: Convolutional networks for biomedical image segmentation. *Med Image Comput*
Comput Assist Interv. 2015;9351:234–241.
11. Oktay O, et al. Attention U-Net: Learning where to look for the pancreas. *arXiv preprint.* 2018;arXiv:1804.03999.
12. 2015:1
13. RL, et al. Cancer statistics. *CA Cancer J Clin.* 2023;73(2):17–48.
14. 14)Choi YR, et al. Cognitive burden in neuroimaging. *Front Psychol.* 2015;6:1480.
15. 15)Nowinski WL, et al. Brain segmentation and classification. *Neuroinformatics.* 2006;4(3):243–265.
16. Programs Biomed. 2019;177:49–61.
17. Vis. 2014;8689:818–833.
18. Bland JM, Altman DG. Statistical methods for assessing agreement between two methods of clinical measurement. *Lancet.* 1986;327(8476):307–310.
19. Kickingereder P, et al. Automated quantitative tumor response assessment of MRI in neuro-oncology with artificial neural networks. *Eur Radiol.* 2018;28(10):4179–4189.