

MRI Brain Tumor Classification using Patch-Level Graph Attention Networks

Maheswari KP ^{1, 2, ✉}, Kalaiselvi T ¹

1. Department of Computer Science and Applications, The Gandhigram Rural Institute (Deemed to be University), Dindigul, IND

2. Computer Applications, Fatima College, Madurai, IND

Received: April 23, 2026 | Review began: May 25, 2026 | Review ended: June 25, 2026 | Published: June 30, 2026

© Copyright 2026

This is an open access article distributed under the terms of the Creative Commons Attribution License CC-BY 4.0., which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Abstract

Magnetic Resonance Imaging (MRI) plays a crucial role in the diagnosis and characterization of brain tumors; however, automated classification remains challenging due to tumor heterogeneity, subtle boundary variations, and the complex spatial relationships among tumor regions. Conventional convolutional neural networks (CNNs) primarily focus on local feature extraction and often lack explicit mechanisms to model long-range inter-patch dependencies within medical images. To address these limitations, this work proposes a Patch-Level Graph Attention Network (PL-GAT) framework for MRI brain tumor classification. In the proposed approach, each MRI slice is partitioned into non-overlapping patches, which are independently encoded using a lightweight CNN to extract discriminative local features. These patch embeddings are then represented as nodes in a graph, where spatial adjacency relationships define the graph structure. A graph attention network is employed to selectively weight inter-patch interactions, enabling adaptive aggregation of region-level information. The model is trained end-to-end using multi-class cross-entropy loss and evaluated on a multi-class MRI brain tumor dataset. Experimental results demonstrate that the proposed CNN-GAT hybrid model effectively captures both local texture patterns and global structural dependencies, leading to improved classification performance. Attention heatmaps and Receiver Operating-Area Under the Curve analysis further provide interpretability by highlighting diagnostically relevant image regions, making the approach suitable for clinically explainable brain tumor analysis.

Categories: AI applications, AI/ML-based decision support systems, Deep Learning

Keywords: mri brain tumor classification, medical image analysis, graph attention networks, patch-based image representation, convolutional neural networks (cnn)

Introduction

Brain tumor diagnosis and classification using Magnetic Resonance Imaging (MRI) is a critical task in neuro-oncology, as accurate identification of tumor type directly influences treatment planning and patient prognosis. MRI provides rich structural and contrast information, but the visual complexity of tumor appearance - caused by variations in shape, size, location, and intensity - makes manual interpretation both time-consuming and prone to inter-observer variability. As a result, automated and reliable computer-aided diagnostic systems have become an active area of research.

How to cite this article:

KP M, T K (June 30, 2026) MRI Brain Tumor Classification using Patch-Level Graph Attention Networks. Cureus J Comput Sci 3 : es44389-026-00144-5. DOI <https://doi.org/10.7759/s44389-026-00144-5>

Deep learning techniques, particularly convolutional neural networks (CNNs), have achieved notable success in medical image analysis due to their ability to learn hierarchical feature representations directly from data. CNN-based models have been widely applied to brain tumor classification tasks; however, they inherently operate on local receptive fields and aggregate information through fixed convolutional operations. While this design is effective for capturing local texture patterns, it limits the network's ability to explicitly model relationships between spatially distant yet semantically related regions within an image. In the context of brain tumors, such region-to-region interactions are clinically important, as tumor growth often affects multiple non-adjacent anatomical areas.

To overcome these limitations, graph-based learning has emerged as a promising alternative for structured representation learning. Graph neural networks (GNNs) provide a flexible framework for modeling relational information by representing data as nodes and edges. When applied to images, graph-based approaches enable explicit reasoning over region-level features and their spatial or semantic relationships. Among GNN variants, graph attention networks (GATs) are particularly well-suited for medical imaging tasks, as they incorporate attention mechanisms that dynamically assign importance to neighboring nodes during message passing. This allows the model to focus on diagnostically relevant regions while suppressing less informative ones.

Motivated by these observations, this work introduces a Patch-Level Graph Attention Network (PL-GAT) for MRI brain tumor classification. Instead of treating the entire image as a single entity, the proposed method decomposes each MRI slice into a set of fixed-size patches. Each patch is encoded using a compact CNN to extract local discriminative features. These patches are then mapped to nodes in a graph, where edges are defined based on spatial adjacency, preserving the inherent anatomical structure of the image. A multi-layer GAT is subsequently applied to learn attention-weighted interactions among patches, enabling effective integration of local and global contextual information.

The proposed framework offers several advantages. First, it explicitly models inter-patch relationships that conventional CNNs cannot capture. Second, the attention mechanism provides inherent interpretability by revealing which regions contribute most to the final decision. Third, the hybrid CNN-GAT design balances computational efficiency with representational power, making it suitable for practical deployment. Comprehensive experimental evaluation, including confusion matrix analysis, Receiver Operating Characteristic Area Under the Curve (ROC-AUC) curves, and attention heatmap visualization, demonstrates the effectiveness and explainability of the proposed approach for multi-class MRI brain tumor classification.

Literature review

Automated brain tumor classification from MRI has attracted sustained research attention due to its critical role in early diagnosis, treatment planning, and prognosis assessment. Early computational approaches primarily relied on handcrafted feature extraction techniques. Researchers such as Zacharaki et al. [1], Chaplot et al. [2], and El-Dahshan et al. [3] employed texture descriptors, wavelet transforms, intensity histograms, and statistical measures combined with classical classifiers like Support Vector Machines (SVMs) and k-Nearest Neighbors (k-NN). While these methods demonstrated moderate success, their performance was highly dependent on domain expertise and handcrafted feature quality, limiting robustness against variations in tumor size, shape, and intensity distribution.

The emergence of deep learning significantly transformed MRI brain tumor analysis. Krizhevsky et al. [4] established the effectiveness of CNNs in visual recognition, which later inspired medical imaging applications. Studies by Pereira et al. [5], Havaei et al. [6], and Kamnitsas et al. [7] demonstrated that CNNs can automatically learn hierarchical spatial features directly from MRI data, outperforming traditional machine learning approaches. Popular architectures such as ResNet [8] and DenseNet [9] have been extensively adopted for brain tumor classification and segmentation tasks.

Despite their strong representational power, CNNs inherently rely on local receptive fields and pooling operations, which restrict their ability to model long-range spatial dependencies. Bengio et al. [10] and Wang et al. [11] noted that CNNs lack explicit mechanisms for capturing region-to-region relationships, an important factor in understanding tumor morphology and spatial progression across the brain. This limitation motivated the exploration of patch-based learning

How to cite this article:

strategies. Patch-based CNN methods divide MRI images into smaller regions to capture fine-grained local characteristics. Nie et al. [12], Faizan et al. [13], and Akkus et al. [14] demonstrated that patch-level analysis improves sensitivity to localized tumor structures and reduces background bias.

Recent studies such as Di et al. [15] have explored global-local representation learning to enhance contextual feature integration in medical image segmentation. Nevertheless, many conventional patch-based frameworks still process patches independently and aggregate predictions using simple pooling or voting mechanisms, limiting their ability to capture spatial dependencies among neighboring regions.

Graph-based learning has emerged as an effective paradigm for modeling relational and structural information. Scarselli et al. [16] introduced GNNs to handle graph-structured data, enabling message passing between connected nodes. In medical imaging, graph representations have been applied to capture anatomical relationships and spatial coherence, as demonstrated by Parisot et al. [17] in brain disease classification and Bronstein et al. [18] in geometric deep learning.

To improve node interaction modeling, Veličković et al. [19] proposed GATs, which incorporate attention mechanisms to assign adaptive weights to neighboring nodes. This innovation allows the network to focus on diagnostically relevant regions while suppressing irrelevant information. Medical imaging studies have shown that graph-based representations improve structural relationship modeling in neuroimaging applications [20-22].

Hybrid CNN-GNN architectures have been introduced to combine local feature extraction with global relational reasoning. Talo et al. [23] demonstrated that CNN-derived features mapped onto graph structures enable effective contextual learning.

Explainability has become a fundamental requirement for clinical adoption of artificial intelligence (AI) systems. Regulatory perspectives discussed by Holzinger et al. [24] emphasize transparency and interpretability in medical AI. Attention-based graph models inherently support explainability by visualizing attention weights that highlight critical regions and inter-region interactions. Studies by Montavon et al. [25] and Samek et al. [26] further support the importance of interpretable deep learning in healthcare applications.

Motivated by these research gaps, the present work focuses on a PL-GAT framework that integrates CNN-based patch encoding with graph attention-based relational learning. By explicitly modeling spatial interactions among MRI patches, the proposed approach addresses the limitations of conventional CNNs and independent patch-based methods, offering an accurate, robust, and interpretable solution for MRI brain tumor classification.

Research gap and contribution

Despite the strong performance of CNNs in MRI brain tumor classification, most existing approaches implicitly aggregate local features and lack explicit modeling of spatial interactions among anatomically adjacent regions. Patch-based CNN methods improve local sensitivity but typically encode patches independently and combine them using global pooling or voting, thereby ignoring structured inter-patch dependencies that are important for characterizing tumor heterogeneity. In addition, graph-based learning and attention mechanisms remain underexplored at the patch level in MRI classification, and interpretability is often addressed using post hoc techniques detached from the learning process. To address these limitations, this work presents a patch-level graph attention framework in which MRI images are partitioned into spatially ordered patches represented as graph nodes with explicitly defined neighborhood connections. Patch features are extracted using a convolutional encoder and refined through graph attention-based message passing, enabling adaptive weighting of inter-patch relationships and improved global contextual representation. The resulting graph-level embeddings are used for end-to-end classification, while the learned attention weights provide intrinsic interpretability through patch-level importance visualization, thereby offering a unified and reproducible solution for MRI brain tumor classification.

Materials And Methods

The development phases of the proposed work are given in Figure 1.

How to cite this article:

KP M, T K (June 30, 2026) MRI Brain Tumor Classification using Patch-Level Graph Attention Networks. Cureus J Comput Sci 3 : es44389-026-00144-5. DOI <https://doi.org/10.7759/s44389-026-00144-5>




FIGURE 1: Steps for patch-level graph attention network.

CNN, Convolutional Neural Network; MRI, Magnetic Resonance Imaging

Each MRI slice is first divided into fixed-size, non-overlapping patches to preserve localized anatomical details and address intra-tumoral heterogeneity. These patches are independently processed using a lightweight CNN to extract compact and discriminative feature representations. The resulting patch embeddings are treated as nodes in a graph, where edges are defined based on spatial neighborhood relationships to retain the structural layout of the image.

A GAT is then applied to the constructed patch graph to learn adaptive inter-patch interactions. By assigning attention weights to neighboring nodes, the model selectively emphasizes diagnostically relevant regions and captures both local and long-range spatial dependencies. Finally, node-level features are aggregated using a global pooling operation to form a graph-level representation, which is passed to a fully connected classifier to predict the tumor class. The entire framework is trained end-to-end using a multi-class cross-entropy loss function.




FIGURE 2: Patch CNN encoder for local feature representation.

CNN, Convolutional Neural Network; MRI, Magnetic Resonance Imaging

The patch CNN encoder is designed as in Figure 2 to extract compact and discriminative feature representations from individual MRI image patches prior to graph-based learning. Each input patch corresponds to a single-channel grayscale MRI region, which serves as the input to the convolutional feature extractor. The encoder begins with an initial convolutional layer employing 16 filters with a kernel size of 3×3, enabling the capture of low-level texture patterns such as intensity gradients and edge-related structures commonly observed in brain MRI scans.

Subsequently, a second convolutional layer with 32 filters and a 3×3 kernel is applied to learn higher-level semantic features that encode more complex tissue characteristics and localized tumor-related patterns. Rectified Linear Unit (ReLU) activation functions are used after each convolution to introduce non-linearity and enhance the representational capacity of the encoder.

To obtain a fixed-dimensional feature vector suitable for graph construction, the convolutional outputs are processed through a combination of max pooling and adaptive average pooling operations. Max pooling emphasizes salient local activations, while adaptive average pooling ensures scale-invariant feature aggregation and produces a consistent output size irrespective of patch resolution. The final output of the encoder is a 32-dimensional feature vector, which serves as the node representation for each patch in the subsequent GAT. This lightweight yet effective encoding strategy balances feature expressiveness and computational efficiency, making it well-suited for patch-level graph-based MRI analysis.

This study's dataset was accessed from the Kaggle Brain Tumor Classification repository, which is publicly available. It consists of 3,094 structural brain MR images acquired with high resolution, size $512 \times 512 \times 3$ pixels. The imaging modalities include a wide range of T1-weighted, T2-weighted, and Fluid-Attenuated Inversion Recovery (FLAIR) sequences. The dataset is split into four target categories: Glioma Tumor ($n = 777$); Meningioma Tumor ($n = 773$); Pituitary Tumor ($n = 778$); Normal control cases ($n = 372$), for the purpose of evaluating the clinical generalizability of the proposed PL-GAT framework to deep learning baselines. The baseline architecture was tested with an 80% training configuration and 20% independent testing configuration.

To test the performance of the proposed PL-GAT framework, we have devised a standardized training pipeline for comparison with our standard deep learning baselines. All the models were written in Python and PyTorch, and trained in Google Colab. All the architectural parameters, optimization parameters and regularization parameters used in the baseline models and the proposed framework are described in detail in Table 1. The differences reported in our study can be confidentially attributed to basic architectural differences by keeping the optimization parameters unchanged, e.g. using the Adam optimizer, an identical baseline learning rate of 1×10^{-4} , and a uniform batch size of 32. The proposed PL-GAT method uses patch embedding parameters to capture non-Euclidean anatomical relationships via a fixed grid of patch sizes, with the introduction of multi-head graph attention layers, and LeakyReLU functions, in a deeply stacked manner.

Configuration parameter	Baseline 1: ResNet-50	Baseline 2: DenseNet-121	Proposed PL-GAT framework
Core Structural Archetype	Residual Deep Grid Convolution	Dense Feature Inter-layer Re-use	Hybrid Patch Encoding + Graph Attention Network
Input Feature Extraction	50-layer Bottleneck ResNet Blocks	121-layer Dense Block Structures	Lightweight CNN (Conv + MaxPool)
Feature/Node Dimension	2048 channels (Global Pooling)	1024 channels (Global Pooling)	64 channels per patch (Structural Node Embedding)
Spatial Processing Block	2D Grid Kernels (3×3, 1×1)	2D Grid Kernels (3×3, 1×1)	2-Layer Stacked Multi-Head Graph Attention (GAT)
Attention Configuration	N/A	N/A	4/8 Attention Heads (K)
Intermediate Activation	ReLU	ReLU	LeakyReLU (α = 0.2)
Output/Classification Head	Linear Layer + Softmax	Linear Layer + Softmax	Multi-Head Attention Pool + Softmax
Regularization (Dropout)	0.2 (Fully Connected Layer)	0.2 (Classification Head)	0.2 (Hidden)/0.5 (Attention Weights)
Optimization Algorithm	Adam	Adam	Adam
Initial Learning Rate (η)	1×10^{-4}	1×10^{-4}	1×10^{-4} (with StepLR Decay)
Batch Size/Max Epochs	32 Slices/100 Epochs	32 Slices/100 Epochs	32 Slices/100 Epochs (Early Stopping Δ = 10)

TABLE 1: Experimental setup and training configuration.

ReLU, Rectified Linear Unit

Results And Discussion

Figure 3 illustrates representative image patches extracted from MRI brain scans as part of the proposed patch-level learning framework. The visual diversity observed across the patches highlights substantial variations in intensity distribution, texture patterns, and structural composition within a single MRI slice. Some patches predominantly capture homogeneous background regions, while others contain heterogeneous tissue characteristics that are indicative of tumor presence or boundary transitions. This variability confirms the inherent intra-image heterogeneity commonly observed in brain tumor MRI data.

How to cite this article:




FIGURE 3: MRI image patches.

MRI, Magnetic Resonance Imaging

The patch-wise representation enables the model to isolate localized anatomical information that may be overlooked when processing the entire image holistically. Notably, patches corresponding to tumor-affected regions exhibit irregular texture patterns and contrast variations, whereas non-tumor patches tend to display smoother intensity profiles. Such distinctions support the effectiveness of patch-level feature learning in preserving fine-grained spatial cues essential for accurate tumor characterization.

By treating each patch as an independent unit before relational modeling, the proposed approach facilitates region-specific analysis and enhances robustness to noise and irrelevant background information. These observations justify the adoption of patch-based graph modeling, as it allows subsequent graph attention mechanisms to selectively emphasize

How to cite this article:

KP M, T K (June 30, 2026) MRI Brain Tumor Classification using Patch-Level Graph Attention Networks. Cureus J Comput Sci 3 : es44389-026-00144-5. DOI <https://doi.org/10.7759/s44389-026-00144-5>

clinically informative regions during inter-patch interaction learning. Overall, the qualitative visualization of MRI patches in Figure 3 reinforces the motivation for patch-level graph construction and demonstrates its suitability for capturing both local and contextual variations in brain tumor imaging.




FIGURE 4: Patch graph structure.

Figure 4 presents the patch graph representation, in which individual nodes correspond to localized image patches and edges encode spatial or feature-based neighborhood connections. The graph topology preserves the intrinsic geometric arrangement of patches, enabling the model to capture both local continuity and broader structural dependencies across the image.

The irregular yet smoothly connected layout demonstrates the ability of the graph formulation to adapt to non-uniform spatial distributions, which is particularly important in medical imaging scenarios where anatomical structures exhibit complex and curved boundaries. By maintaining connectivity among adjacent patches, the graph facilitates effective information propagation during graph-based learning, allowing contextual features from neighboring regions to influence each node representation.

From a results perspective, this structured representation contributes to improved feature discrimination by integrating spatial coherence with learned patch embeddings. The patch graph enables the model to move beyond independent patch analysis, supporting relational reasoning across regions and enhancing robustness to local noise or intensity variations. Overall, the graph structure shown in Figure 4 plays a critical role in preserving anatomical consistency and improving downstream performance in the proposed framework.




FIGURE 5: Distribution of graph attention weights.

In Figure 5, the empirical distribution of attention weights learned by the graph attention mechanism. The histogram shows that the majority of attention values are concentrated within a moderate range, indicating that the model assigns balanced importance to most neighboring nodes rather than relying on a small subset. This behavior reflects stable and well-regularized attention learning, avoiding excessive dominance of individual connections.

How to cite this article:

A smaller number of higher attention weights can also be observed, corresponding to particularly informative or structurally relevant patch relationships. These higher values suggest that the model successfully identifies and emphasizes critical interactions between specific regions, which is essential for capturing discriminative patterns in complex image structures.

From a performance standpoint, this distribution highlights the effectiveness of the attention mechanism in adaptively weighting node contributions based on contextual relevance. By maintaining a smooth attention profile with selective emphasis, the GAT enhances feature aggregation while preserving robustness against noise and redundant connections. Overall, the learned attention distribution supports improved relational modeling and contributes to the observed gains in classification accuracy.




FIGURE 6: Patch-level attention heatmap.

How to cite this article:

Figure 6 illustrates the patch-level attention heatmap generated by the proposed graph attention framework. Each cell corresponds to the attention score assigned to an individual image patch, providing a spatial visualization of how the model distributes focus across the input. Brighter regions indicate patches that contribute more strongly to the final representation, while darker regions denote comparatively lower influence.

The heterogeneous distribution observed in the heatmap demonstrates that the model selectively emphasizes informative regions rather than uniformly weighting all patches. This selective behavior is predominantly important in medical imaging, where diagnostically relevant patterns are often localized and subtle. The presence of high-attention clusters suggests that the network effectively captures salient structural or textural cues within specific spatial neighborhood.

From an interpretability perspective, the heatmap offers insight into the process of decision-making process of the model by highlighting regions that drive the learning outcome. The smooth transitions between attention values further indicate stable attention assignment, reducing sensitivity to local noise. Overall, the patch-level attention visualization confirms that the proposed approach successfully integrates spatial context and learned feature relevance, contributing to improved discriminative performance.

Tumor/Tissue type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	ROC-AUC
Glioma	95.1	95.8	95	95.4	0.976
Meningioma	94.2	94.1	93.8	93.95	0.968
Pituitary	97.1	97.2	96.9	97.05	0.987
Normal (No Tumor)	98.72	98.5	98.7	98.6	0.997
Average/Overall	96.28	96.4	96.1	96.2	0.982

TABLE 2: Performance metrics.
ROC-AUC, Receiver Operating Characteristic Area Under the Curve

The quantitative performance measures presented in Table 2 are directly correlated with our qualitative and architectural evaluations (Figures 4, 5, and 6). A robust measure of the operational capacity of the multi-head attention weight distributions (Figure 5) is the macro F1 score, with 96.20%. The structural topologies are also preserved by the patch-level graphs (Figure 4), with an F1 score of 99.28%. This local sensitivity is supported visually in the attention heatmaps (Figure 6), with particularly good detection in different anatomical structures like Pituitary lesions.




FIGURE 7: Normalized confusion matrix of the PL-GAT framework.

PL-GAT, Patch-Level Graph Attention Network

To explicitly measure the class-wise predictive reliability, and to see the boundaries of confusion of features, a normalized confusion matrix was computed, as shown in Figure 7. The proposed PL-GAT model presents a strong diagonally accurate result, with maximum accuracy of 98.70% and 96.90% in Normal tissue cohort and Pituitary lesions cohort, respectively. The highest classification difficulty was between the Glioma and Meningioma classes with 3.90% of real Meningioma scans being classified as Gliomas. This minor error trend is indicative of common radiographically similar textures in the vicinity of complex dural borders. Yet there is a fairly high true-positive density throughout the diagonal, confirming the structural capacity of the model to segregate discrete pathologies with low false-positive density crossovers.

How to cite this article:




FIGURE 8: Class-wise receiver operating characteristic (ROC) curves and individual area under the curve (AUC) scores for multi-class brain pathology validation.

The ability of our framework to detect the class was assessed using class-specific ROC curves on the independent patient test set as shown in Figure 8. The proposed multi-head PL-GAT model shows good discriminative thresholds for all pathological tissue variations with AUC values of 0.976 for Glioma, 0.968 for Meningioma, and 0.987 for Pituitary tumor. Notably, the model performs best in slices without tumor (Normal) with AUC of 0.997. This steep curve progression towards the top-left coordinate mathematically validates the network's ability to have a high True Positive Rate and also low false-positives clinical assignments.

In this context, the benchmarking results shown in Table 3 clearly highlight the clear performance gains of going from a classical grid-based convolutional structure to a non-Euclidean graph topological one. The baseline models, including the widely used deep convolutional layers (ResNet-50 and DenseNet-121), are inherently limited by the fixed local receptive fields that encode local patterns within a local neighborhood of the image, which is not effective in modeling long-range regional interactions between non-adjacent structural units of the brain. The proposed PL-GAT framework is able to

How to cite this article:

address these limitations by dividing the 256×256 pixel input image into a standardized arrangement of 32×32 pixel patches to capture micro-level radiographic textures and then mapping these patches as interconnected nodes in a localized graph network. The model has been deployed in a multi-head graph attention routing engine over these nodes to deliver dynamic message-passing, active evaluation of edge transitions and correlation of long-range tissue heterogeneity, resulting in a significant performance boost over the best baseline (DenseNet-121), reaching an Overall Accuracy of 96.28%, Macro F1-Score of 96.20%, and highly robust Mean ROC-AUC of 0.982 for computer-aided multi-class diagnostics.

Model/Architecture	Overall accuracy (%)	F1-Score (%)	Mean ROC-AUC
ResNet-50	93.45	93.1	0.942
DenseNet-121	94.2	93.95	0.955
Proposed PL-GAT	96.28	96.20	0.982

TABLE 3: Quantitative performance comparison.
ROC-AUC, Receiver Operating Characteristic Area Under the Curve

Conclusions

In this paper, a PL-GAT approach to automated multi-class MRI brain tumor classification is introduced. The model uses fine-grained local textures and global anatomical context by breaking down MRI slices into localized patches and representing them as graph nodes. This method overcomes the drawback of traditional CNNs of losing long-range spatial dependency. The graph attention mechanism dynamically adjusts the relationship between different patches with weights, emphasizing the relationship of the patch that is important for diagnosis and filtering out the background relationship with an unimportant patch. Experimental results show that the hybrid CNN-GAT architecture is more superior than the conventional CNN model. Additionally, the strategy based on attention provides more interpretability because it analyzes the importance of regions, giving clinically meaningful support for decisions in the field of neuro-oncology workflows.

Additional Information

Author Contributions

All authors have reviewed the final version to be published and agreed to be accountable for all aspects of the work.

Concept and design: Maheswari KP, Kalaiselvi T

Acquisition, analysis, or interpretation of data: Maheswari KP, Kalaiselvi T

Drafting of the manuscript: Maheswari KP, Kalaiselvi T

Critical review of the manuscript for important intellectual content: Maheswari KP, Kalaiselvi T

Supervision: Kalaiselvi T

Disclosures

Human subjects: All authors have confirmed that this study did not involve human participants or tissue. **Animal subjects:** All authors have confirmed that this study did not involve animal subjects or tissue. **Conflicts of interest:** In compliance with the ICMJE uniform disclosure form, all authors declare the following: **Payment/services info:** All authors have declared that no financial support was received from any organization for the submitted work. **Financial relationships:** All authors have declared that they have no financial relationships at present or within the previous three years with any organizations that might have an interest in the submitted work. **Other relationships:** All authors have declared that there are no other relationships or activities that could appear to have influenced the submitted work.

Data Availability Statements

The dataset and code used in the experiment are available in the below link:
https://drive.google.com/drive/folders/1iU8hdu1aCg_rkpecFmSh8BgdPfvQo7OU?usp=drive_link

Acknowledgements

This research paper was presented in the International Conference on Applied Intelligence and Innovative Computational Solutions (IC-AIICS 2026) during February 9, 2026 to February 11, 2026, organized by the Department of Computer Science and Computer Applications, The Gandhigram Rural Institute (deemed to be university), Dindigul.

References

1. Zacharaki EI, Wang S, Chawla S, Dong SY, Wolf R, Melhem ER, Davatzikos C: [Classification of brain tumor type and grade using MRI texture and shape in a machine learning scheme](#). Magnetic Resonance in Medicine. 2009, 62:1609-1618. [10.1002/mrm.22147](#)
2. Chaplot S, Patnaik LM, Jagannathan NR: [Classification of magnetic resonance brain images using wavelets as input to support vector machine and neural network](#). Biomedical Signal Processing and Control. 2006, 1:86-92. [10.1016/j.bspc.2006.05.002](#)
3. El-Dahshan ESA, Hosny T, Salem ABM: [Hybrid intelligent techniques for MRI brain images classification](#). Digital Signal Processing. 2010, 20:433-441. [10.1016/j.dsp.2009.07.002](#)
4. Krizhevsky A, Sutskever I, Hinton GE: [ImageNet classification with deep convolutional neural networks](#). Communications of the ACM. 2017, 60:84-90. [10.1145/3065386](#)
5. Pereira S, Pinto A, Alves V, Silva CA: [Brain tumor segmentation using convolutional neural networks in MRI images](#). IEEE Transactions on Medical Imaging. 2016, 35:1240-1251. [10.1109/tmi.2016.2538465](#)
6. Payer C, Pienn M, Bálint Z, et al.: [Automated integer programming based separation of arteries and veins from thoracic CT images](#). Medical Image Analysis. 2016, 34:109-122. [10.1016/j.media.2016.05.002](#)
7. Kamnitsas K, Ledig C, Newcombe VFJ, et al.: [Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation](#). Medical Image Analysis. 2016, 36:61-78. [10.1016/j.media.2016.10.004](#)
8. He K, Zhang X, Ren S, Sun J: [Deep residual learning for image recognition](#). IEEE Conference on Computer Vision and Pattern Recognition. 2016, 770-778. [10.1109/CVPR.2016.90](#)
9. Huang G, Liu Z, Van Der Maaten L, Weinberger KQ: [Densely connected CNNs](#). IEEE Conference on Computer Vision and Pattern Recognition. 2017, 4700-4708. [10.1109/CVPR.2017.243](#)
10. Bengio Y, Courville A, Vincent P: [Representation learning: a review and new perspectives](#). IEEE Transactions on Pattern Analysis and Machine Intelligence. 2013, 35:1798-1828. [10.1109/tpami.2013.50](#)
11. Wang X, Girshick R, Gupta A, He K: [Non-local neural networks](#). IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2018, 7794-7803. [10.1109/CVPR.2018.00813](#)
12. Nie D, Wang L, Gao Y, Shen D: [Fully convolutional networks for multi-modality isointense infant brain image segmentation](#). IEEE International Symposium on Biomedical Imaging. 2016, 2016:1342-1345. [10.1109/isbi.2016.7493515](#)
13. Ullah F, Salam A, Abrar M, Amin F: [Brain tumor segmentation using a patch-based convolutional neural network: a big data analysis approach](#). Mathematics. 2023, 11:1635-1635. [10.3390/math11071635](#)

How to cite this article:

KP M, T K (June 30, 2026) MRI Brain Tumor Classification using Patch-Level Graph Attention Networks. Cureus J Comput Sci 3 : es44389-026-00144-5. DOI <https://doi.org/10.7759/s44389-026-00144-5>

14. Akkus Z, Galimzianova A, Hoogi A, Rubin DL, Erickson BJ: [Deep learning for brain MRI segmentation: state of the art and future directions](#). Journal of Imaging Informatics in Medicine. 2017, 30:449-459. [10.1007/s10278-017-9983-4](#)
15. Gai D, Zhang J, et al.: [GL-Segnet: global-local representation learning net for medical image segmentation](#). Frontiers in Neuroscience. 2023, 17:1-16. [10.3389/fnins.2023.1153356](#)
16. Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G: [The graph neural network model](#). IEEE Transactions on Neural Networks. 2009, 20:61-80. [10.1109/tnn.2008.2005605](#)
17. Parisot S, Ktena SI, Ferrante E, Lee M, Guerrero R, Glocker B, Rueckert D: [Disease prediction using graph convolutional networks: application to autism spectrum disorder and Alzheimer's disease](#). Medical Image Analysis. 2018, 48:117-130. [10.1016/j.media.2018.06.001](#)
18. Bronstein MM, Bruna J, LeCun Y, Szlam A, Vandergheynst P: [Geometric deep learning: going beyond Euclidean data](#). IEEE Signal Processing Magazine. 2017, 34:18-42. [10.1109/msp.2017.2693418](#)
19. Veličković P, Cucurull G, Casanova A, Romero A, Liò P, Bengio Y: [Graph attention networks](#). International Conference on Learning Representations 2018 Conference Track Proceedings. 2018, [10.17863/CAM.48429](#)
20. Zhang F, Cetin KS, Hoffmann N, Rath Y, Golby AJ, O'Donnell LJ: [Deep white matter analysis \(DeepWMA\): fast and consistent tractography segmentation](#). Medical Image Analysis. 2020, 65:101761. [10.1016/j.media.2020.101761](#)
21. Li X, Zhou Y, Dvornek N, et al.: [BrainGNN: interpretable brain graph neural network for fMRI analysis](#). Medical Image Analysis. 2021, 74:102233. [10.1016/j.media.2021.102233](#)
22. Liu M, Zhang Z, Dunson DB: [Graph auto-encoding brain networks with applications to analyzing large-scale brain imaging datasets](#). NeuroImage. 2021, 245:1-15. [10.1016/j.neuroimage.2021.118750](#)
23. Çinar A, Yildirim M: [Detection of tumors on brain MRI images using the hybrid convolutional neural network architecture](#). Medical Hypotheses. 2020, 139:109684. [10.1016/j.mehy.2020.109684](#)
24. Holzinger A, Langs G, Denk H, Zatloukal K, Müller H: [Causability and explainability of artificial intelligence in medicine](#). wiley interdisciplinary reviews. data mining and knowledge discovery. 2019, 9:1-13. [10.1002/widm.1312](#)
25. Montavon G, Samek W, Müller KR: [Methods for interpreting and understanding deep neural networks](#). Digital Signal Processing. 2018, 73:1-15. [10.1016/j.dsp.2017.10.011](#)
26. Samek W, Montavon G, Vedaldi A, Hansen LK, Müller KR: [Explainable AI: Interpreting, Explaining and Visualizing Deep Learning](#). Springer, Cham; 2021. [10.1007/978-3-030-28954-6](#)

How to cite this article: