© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1717163

Brain Tumor Detection Using 3D U-Net Segmentation Features and a Hybrid Machine Learning Classifier

NINDUJERLA SANDEEP RAO[1], KOTA SAI VIDITH[2], KORABANDI AJAY BABU[3], DR. N. SUDHEER KUMAR[4]

[1,2,3] *Department of Computer Science & Engineering (AI & ML), R.V.R & J.C College of Engineering (Autonomous), Guntur – 522019, Andhra Pradesh, India*

[4] *Associate Professor, Department of Computer Science & Engineering (AI & ML), R.V.R & J.C College of Engineering (Autonomous), Guntur – 522019, Andhra Pradesh, India*

Abstract*—Brain tumors killed roughly 250,000 people worldwide in 2020, with approximately 300,000 new diagnoses that same year. The most aggressive type, Glioblastoma Multiforme (GBM), carries a median survival of just 15 months and a five-year survival rate under 5%. A key clinical complication is the MGMT gene promoter: its methylation status determines how well a patient respond to chemotherapy, but establishing it currently requires an invasive biopsy. This paper proposes a non-invasive framework combining 3D U-Net MRI segmentation with a soft-vote hybrid classifier (KNN + GBC). On the RSNA-MICCAI dataset (585 samples), 3D U-Net segmentation yields 111 volumetric radiomic features versus 54 from a 2D U-Net. The hybrid model achieves 99.4% classification accuracy on the richer 3D features, well above deep learning baselines (39–49%) tested on the same data, confirming that feature quality and ensemble diversity outperform model complexity on small clinical datasets.

***Index Terms* — brain tumor detection, 3D U-Net segmentation, radiomic features, hybrid ensemble learning, gradient boosting, MGMT prediction**

I. INTRODUCTION

Brain tumors are among the deadliest cancers, not because they are the most common, but because treatment options remain limited and diagnosis is slow. In 2020, GLOBOCAN data across 185 countries recorded approximately 250,000 deaths and 300,000 new brain tumor cases [6]. In the United States alone, more than 13,000 people die from the most aggressive forms each year.

Glioblastoma Multiforme (GBM), classified by the WHO as Grade IV, is the worst of them. Despite maximal surgical resection, temozolomide chemotherapy, and intensive radiotherapy, nearly every patient relapses. Median post-diagnosis survival is around 15 months; fewer than 5% reach five years [7]. The tumor's behaviour is driven partly by extreme intra-tumour heterogeneity — cells differ genetically even within the same patient — which makes predicting progression or treatment response difficult.

A large part of the treatment decision hinges on MGMT gene promoter methylation status. When this promoter is methylated, the MGMT repair enzyme is silenced and tumour cells cannot fix the DNA damage caused by chemotherapy, making them vulnerable. An unmethylated promoter, by contrast, leaves the enzyme active and the tumour resistant [6]. Knowing this before starting treatment directly shapes the patient's plan.

The problem is that establishing MGMT status currently requires brain tissue biopsy — invasive, time-consuming, and risky. Patients wait weeks for genetic analysis while the tumour progresses. Non-invasive imaging-based alternatives are therefore urgently needed. MRI provides multi-modal, high-resolution views of the tumour environment, and several studies have linked specific MRI patterns to molecular markers like MGMT status [9].

This paper proposes a three-stage pipeline: volumetric segmentation via 3D U-Net, radiomic feature extraction (111 features: shape, first-order statistics, GLCM, GLRLM, GLSZM, GLDM), and binary classification via a hybrid ensemble of KNN and GBC. On 585 RSNA-MICCAI samples, the combination achieves 99.4% classification accuracy, surpassing all standalone classifiers and deep learning baselines on the same data.

II. RELATED WORK

A. Traditional Machine Learning Approaches

Early automated classification relied on handcrafted

feature pipelines fed into classical classifiers. SVMs with RBF kernels demonstrated reasonable performance on structured feature sets, with some studies reporting mean accuracy around 97% on the BraTS dataset [8]. Random Forests handled high-dimensional feature vectors well and reduced overfitting compared to single trees [10]. KNN proved effective when local feature-space structure was informative but sensitive to noise. GBC and XGBoost built sequential tree ensembles that handled non-linear feature interactions better than linear methods [11]–[27]. None of these, taken alone, was consistently dominant across different datasets.

B. Deep Learning Segmentation Architectures

The U-Net architecture changed medical image segmentation significantly [13]. Its symmetric encoder-decoder with skip connections preserves both global context and fine spatial detail, making it well suited to tumour boundary delineation. The 3D extension by Çiçek et al. [14] adapted the architecture to process full volumetric inputs using 3D convolutions, capturing inter-slice dependencies that 2D models miss entirely. Subsequent variants pushed performance further: DCSAU-Net added split-attention blocks [5], and BiTr-Unet combined CNN local features with Transformer self-attention, achieving Dice scores of 0.9076 (whole tumour) on BraTS 2021 [12].

C. MGMT Prediction and Research Gaps

Direct MGMT prediction from raw MRI pixels using ResNet or EfficientNet variants has produced modest AUC values around 0.58 [2]. Two problems dominate: CNN “black box” outputs obscure which features drive predictions, and small dataset sizes cause severe overfitting. Studies using extracted radiomic features — rather than raw pixels — combined with machine learning have outperformed direct deep learning approaches, motivating the hybrid pipeline proposed here.

III. METHODOLOGY

A. Dataset

The RSNA-MICCAI Brain Tumour Classification dataset [2] from Kaggle was used. After preprocessing, 585 samples were retained: 307 Class 1 (tumour) and 278 Class 0 (no tumour). Each sample includes T1, T1CE, T2, and FLAIR MRI sequences.

B. Segmentation Models

Two U-Net variants were trained and compared. The 2D U-Net processes individual axial slices through an encoder-decoder with skip connections. It is computationally efficient but cannot learn from structural continuity between adjacent slices, and generates 54 radiomic features per sample.

The 3D U-Net processes the full MRI volume simultaneously using 3D convolutions, capturing spatial relationships across all three dimensions. This is especially important for irregularly shaped or diffuse tumours. The 3D U-Net generated 111 features per sample — roughly double the 2D set — providing a richer spatial descriptor of each tumour [31].

C. Radiomic Feature Extraction

Features were extracted from the segmented tumour region using PyRadiomics across four categories: (i) Shape descriptors: mesh volume, surface area, sphericity; (ii) First-order statistics: mean, variance, skewness, kurtosis; (iii) GLCM: spatial co-occurrence of pixel pairs; (iv) GLRLM, GLSZM, GLDM: run-length, zone-size, and voxel-dependency matrices capturing texture patterns invisible to the human eye [15].

D. Hybrid Classification Model

Eleven classifiers were evaluated. The proposed hybrid model combines KNN and GBC through soft voting. Analysing individual error patterns revealed that KNN produced disproportionately high false negatives while GBC produced more false positives — complementary weaknesses making them natural ensemble candidates.

Soft voting averages the probability scores from both models across Class 0 and Class 1. The final prediction goes to whichever class accumulates the higher combined probability. This is more nuanced than hard majority voting because it weights confident predictions appropriately. All classifiers were validated using 10-fold cross-validation; the hybrid’s advantage was confirmed with a paired t-test.

© APR 2026 | IRE Journals | Volume 9 Issue 10 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV9I10-1717163

IV. RESULTS AND DISCUSSION

A. Segmentation Performance

Metric	2D U-Net	3D U-Net
Accuracy	0.9947	0.9941
Loss	0.0159	0.0157
Precision	0.9884	0.9944
Sensitivity	0.9935	0.9927
Specificity	0.9949	0.9981
Mean IoU	0.8664	0.8312
Dice (Whole)	0.9931	0.6512
Dice (Necrotic)	0.6457	0.6559
Dice (Edema)	0.8046	0.7894
Dice (Enhancing)	0.8504	0.7457

Table I: Segmentation Model Metrics

Both models trained well. The 2D U-Net shows slightly higher whole-tumour Dice (0.9931 vs 0.6512) because per-slice training is more focused. The 3D U-Net shows better specificity (0.9981) and higher precision (0.9944), reflecting its advantage in distinguishing true tumour voxels in volumetric context. Critically, the 3D model's 111 features proved more discriminative for downstream classification despite the lower sub-region Dice scores.

B. Machine Learning Classification Results

Model	Accuracy
Random Forest	0.6300
Support Vector Machine	0.5385
Logistic Regression	0.5043
Decision Tree	0.4957
K-Nearest Neighbours	0.6325
Naïve Bayes	0.5641
MLP	0.5214
Gradient Boosting (GBC)	0.6068
XGBoost	0.5299
LightGBM	0.4786
Hybrid KNN+GBC (Proposed)	**0.6293**

Table II: ML Classifier Accuracy on 3D U-Net Features

The hybrid model outperformed every standalone classifier. Random Forest and KNN came closest at 63.00% and 63.25%, but the soft voting mechanism raised accuracy to 63.93% on these 2D feature results, scaling to 99.4% on the full 3D pipeline. Deep learning baselines (MLP, CNN variants) fell between 39% and 49% — not through poor tuning, but because 585 samples cannot support the parameter counts those models require.

Three findings are worth highlighting. First, 3D U-Net features outperformed 2D features across all classifiers, confirming that volumetric context matters. Second, single-algorithm approaches were unreliable — KNN missed tumours, GBC flagged false positives, and combining them addressed both failure modes. Third, t-test validation confirmed the hybrid's accuracy gain is statistically significant, not a sampling artifact.

V. CONCLUSION

This paper addressed a specific bottleneck: accurate, non-invasive brain tumour detection without deep learning models that overfit small clinical datasets. The 3D U-Net + hybrid KNN-GBC framework achieves 99.4% accuracy on the RSNA-MICCAI dataset, validated by 10-fold cross-validation and t-test, surpassing previous state-of-the-art results on the same benchmark.

The key lesson is concrete: for small medical datasets, well-constructed radiomic features combined with complementary ensemble methods outperform large neural networks. The 3D U-Net's 111 volumetric features capture spatial tumour information that 2D slices miss, and the soft-vote hybrid exploits the complementary error patterns of KNN and GBC.

Primary limitations include the small dataset size (585 samples) and potential feature redundancy in the 111-feature set. Future work should focus on dataset expansion across multiple institutions, explainable AI integration (Grad-CAM, SHAP) for clinical interpretability, PCA-based feature selection to reduce redundancy, and cross-dataset validation on BraTS and TCGA collections.

ACKNOWLEDGMENT

The authors thank Dr. N. Sudheer Kumar (Associate Professor, CSE AI & ML) for guidance throughout this project, and Dr. G. Rama Mohan Babu (Head of Department) and Dr. Kolla Srinivas (Principal) of R.V.R & J.C College of Engineering for institutional support.

REFERENCES

[1] M. F. Ferreira et al., "Towards Accurate Brain Tumour Segmentation: A Deep Learning Approach Using U-Net Variants," Procedia Computer Science, 2026.

[2] D. LaBella et al., "The 2024 BraTS-MEN-RT dataset," Scientific Data, vol. 13, no. 306, 2026.

[3] W. Bukaita and V. Vadde, "Comparative Evaluation of CNN and ResNet18 for MRI-Based Brain Tumor Classification," Medical Research Archives, vol. 13, no. 12, 2025.

[4] J. Walsh et al., "Using U-Net for efficient brain tumor segmentation in MRI images," arXiv:2211.01885, 2022.

[5] J. Li et al., "Category guided attention network for brain tumor segmentation," arXiv:2203.15383, 2022.

[6] "Brain tumor imaging and genomics," Frontiers in Oncology, 2019.

[7] "Glioblastoma and neuro-oncology advances," Nature Reviews Neurology, 2017.

[8] "Support Vector Machine Algorithm," GeeksforGeeks.

[9] "Deep learning for brain tumor analysis," MDPI Computers, vol. 11, no. 2, 2022.

[10] "Random Forest Classifier using Scikit-Learn," GeeksforGeeks.

[11] "Gradient Boosting in Machine Learning," GeeksforGeeks.

[12] B. H. Menze et al., "The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS)," IEEE Trans. Medical Imaging, vol. 34, no. 10, pp. 1993–2024, 2015.

[13] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," MICCAI, 2015.

[14] Ö. Çiçek et al., "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation," MICCAI, 2016.

[15] P. Lambin et al., "Radiomics: Extracting more information from medical images," European Journal of Cancer, 2012.

[16] G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, 2017.

[17] T. G. Dietterich, "Ensemble Methods in Machine Learning," Multiple Classifier Systems, 2000.