IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Hybrid Deep Learning Architectures for MRI-Based Brain Tumor Diagnosis

Aarti Verma[1] Prof. Ashish Tiwari[2]

Research Scholar[1], Assistant Professor[2]

[12]Department of Computer Science and Engineering, Swami Vivekanand College of Engineering Indore, MP

Abstract

Hybrid deep learning architectures have emerged as a powerful solution for addressing the complexity of MRI-based brain tumor diagnosis, a domain characterized by heterogeneous tumor morphology, multimodal imaging protocols, and the need for accurate, early detection. Traditional convolutional neural networks (CNNs), although highly effective at learning spatial features, often struggle to capture long-range dependencies, multi-sequence relationships, and contextual variations inherent in brain tumor MRI data. By integrating complementary modules such as recurrent neural networks, transformers, generative adversarial networks, and classical machine-learning classifiers, hybrid models leverage the strengths of each component to enhance feature representation, robustness, and generalization across scanners and institutions. Advances in multimodal MRI—including T1, T2, FLAIR, and contrast-enhanced sequences—further motivate hybrid architectures by providing diverse inputs that benefit from sophisticated fusion strategies. Preprocessing pipelines involving skull stripping, normalization, and denoising underpin the reliability of these systems, while hybrid CNN–GAN and CNN–Transformer approaches demonstrate improvements in tumor detection, classification, and segmentation, particularly for complex tasks such as delineating glioma subregions or distinguishing tumor grades. Benchmarking on datasets such as BraTS, TCGA, and REMBRANDT reveals consistent performance gains for hybrid models, although issues of domain shift, annotation scarcity, and limited cross-dataset validation remain significant barriers. Despite impressive accuracy metrics—including enhanced Dice scores, F1-scores, and reduced false positives—clinical integration is hindered by computational cost, interpretability challenges, and regulatory constraints. Overall, hybrid deep learning architectures represent a promising path toward more reliable and clinically useful MRI-based brain tumor diagnosis, while highlighting the need for improved datasets, standardized evaluation protocols, and greater emphasis on robustness and transparency.

Keywords: Deep Learning, MRI, Brain Tumor, Convolutional Neural Networks, T1, T2, FLAIR and GAN

1. Introduction

Brain tumours continue to pose significant challenges in neurology and oncology, due to their diverse histopathological types, variable growth patterns, and the critical importance of early detection. Magnetic resonance imaging (MRI) has long been the primary imaging modality for diagnosing and monitoring brain tumours, given its high soft-tissue contrast, ability to acquire multiple sequences and planes, and its non-ionising nature. Yet, despite these advantages, interpretation of brain tumour MRIs remains labour-intensive and subject to inter-observer variability. Advances in artificial intelligence (AI), especially deep learning (DL), have begun to support radiologists by automating tasks such as tumour segmentation, classification and grading, thereby promising more consistent and efficient workflows. For example, a recent review highlighted how DL models are enabling automated tumour segmentation and even non-invasive prediction of molecular biomarkers from MRI alone. (Dorfner et al., 2025)

Nonetheless, the complexity of brain tumour MRI data—characterised by heterogeneous tumour morphology, variable contrast across MRI protocols and multi-modality acquisition—means that single-network DL models can struggle to capture all relevant spatial, contextual and modal features. To overcome these limitations, emerging hybrid deep learning architectures—which combine different neural network types (e.g., convolutional neural networks, recurrent or transformer modules) or fuse deep learning with classical machine-learning classifiers—have shown considerable promise. A systematic review of hybrid machine learning and deep learning models demonstrated that combining CNNs with traditional classifiers improved classification accuracy and reduced false positives compared to standalone networks.

2. Background on Brain Tumors and MRI Importance

Brain tumors represent a diverse group of intracranial neoplasms that vary widely in biological behavior, aggressiveness, and prognosis, making accurate diagnosis essential for guiding treatment and improving survival. As these tumors often infiltrate surrounding tissues or exhibit heterogeneous internal composition, they can be difficult to assess through clinical symptoms or conventional diagnostic methods alone. Magnetic Resonance Imaging (MRI) has therefore become the cornerstone of brain tumor evaluation because of its superior soft-tissue contrast, non-ionizing acquisition process, and ability to visualize anatomical and pathological changes with high precision. MRI provides detailed structural information, enabling clinicians to identify key features such as tumor boundaries, peritumoral edema, necrosis, cystic components, and mass effect. For example, a comprehensive review emphasizes that MRI remains the primary modality for diagnosing and monitoring primary adult brain tumors owing to its unmatched imaging versatility. (Martucci et al., 2023)

Beyond structural assessment, advanced MRI techniques—including diffusion-weighted imaging, perfusion imaging, susceptibility mapping, and MR spectroscopy—provide valuable functional and metabolic insights that help distinguish tumor types, assess tumor grade, and monitor post-treatment changes. These modalities aid in differentiating true tumor progression from treatment-related effects such as pseudoprogression or radiation necrosis, which can mimic tumor recurrence. However, interpretation of MRI remains challenging due to heterogeneous imaging protocols, subtle lesion appearances, and the need for volumetric analysis across multiple sequences. As a result, radiologists often face significant cognitive load and variability in interpretation. A pictorial review highlights that while conventional MRI is powerful, its limitations in distinguishing tumor from non-tumor pathology underscore the need for more advanced analytic approaches. (Sawlani et al., 2020)

2.1. Motivation for AI and Hybrid Deep Learning Approaches

Although MRI provides extensive structural and functional information, the complexity of multi-sequence brain tumor imaging poses significant challenges for manual interpretation and traditional computational methods. Tumor heterogeneity, differences in MRI acquisition parameters, and the high dimensionality of 3D volumetric data make it difficult for radiologists and classical machine-learning algorithms to consistently extract meaningful patterns. These challenges have fueled the adoption of artificial intelligence (AI), particularly deep learning (DL), which can automatically learn hierarchical feature representations from raw imaging data. However, single-architecture DL models—such as conventional convolutional neural networks—may still struggle to capture long-range dependencies, multimodal relationships, or subtle global context within MRI volumes. To address these limitations, researchers have increasingly developed hybrid deep learning models that combine complementary network types, merge DL with classical machine-learning classifiers, or integrate multimodal MRI features into unified frameworks. Such hybrid architectures leverage the strengths of each component, improving model robustness, enhancing feature expressiveness, and enabling better generalization across institutions and MRI protocols. A recent systematic review found that hybrid AI approaches significantly improve diagnostic accuracy, demonstrating their potential to outperform traditional DL pipelines while offering more reliable performance in real-world clinical environments. (Satushe et al., 2025)

3. Overview of MRI Modalities Used in Brain Tumor Diagnosis

Magnetic resonance imaging (MRI) plays a central role in brain tumour evaluation thanks to its ability to provide high resolution structural images across multiple contrast sequences, thus enabling visualization of tumour morphology, edema, necrosis and infiltration. Moreover, advanced functional and parametric MRI techniques—including diffusion, perfusion and spectroscopy—offer insight into tumour cellularity, vascularity and metabolism, thereby enhancing diagnostic specificity and helping guide biopsy and treatment

decisions (Sawlani et al., 2020). The adoption of multi-modal MRI protocols means clinicians and researchers now routinely integrate several sequence types to capture complementary tumour characteristics and improve the robustness of image-analysis pipelines (Zhang et al., 2021).
Despite these strengths, MRI interpretation for brain tumours remains demanding because of heterogeneous tumour appearances, variable protocol parameters across centres, and the challenge of reliably isolating infiltrative margins or differentiating tumour tissue from treatment-related changes. These factors underscore the need to understand the key MRI modalities used in brain tumour diagnosis, their relative strengths and limitations, and the added value of combining them into multi-modal workflows for enhanced diagnostic accuracy and deep-learning model performance.

3.1. Common MRI Sequences: T1, T2, FLAIR, and Contrast-Enhanced

Here is an introduction and then the key sequences described point-wise:
In brain tumour imaging, four widely used MRI sequences are T1-weighted, T2-weighted, FLAIR (Fluid-Attenuated Inversion Recovery) and contrast-enhanced T1 (often with gadolinium). These sequences each provide distinct tissue contrast and together form the backbone of many neuro-oncology imaging protocols.

- **T1-weighted (T1):** Provides good anatomical detail and delineation of tissues; cerebrospinal fluid (CSF) appears dark, white matter is relatively brighter, and many tumours are iso- or hypointense compared to grey matter. (Martucci et al., 2023)
- **T2-weighted (T2):** Highlights fluid content and edema by showing CSF and other fluid as bright; useful for visualising peritumoral edema and infiltrative zones which often appear hyperintense. (Stall et al., 2010)
- **FLAIR:** Suppresses CSF signal to better visualise lesions adjacent to ventricles or cortex and improves detection of tumour-associated edema or infiltration by reducing fluid background brightness. (Stall et al., 2010)
- **Contrast-Enhanced T1 (T1-CE):** After administration of gadolinium-based contrast agent, enhancing lesions (indicating breakdown of the blood-brain barrier or tumour neovascularity) appear bright; critical for defining active tumour margins, ring enhancement, and for treatment monitoring.

4. Strengths and Limitations of MRI for Tumor Visualization

MRI offers several important advantages for brain tumour imaging. Among them: excellent soft-tissue contrast enabling differentiation of tumour from normal brain structures; the ability to acquire volumetric images in multiple planes; non-ionising radiation making repeated scans safe; and the availability of multiple sequences to assess different tissue properties. However, MRI also has limitations. One key limitation is that conventional structural MRI sequences sometimes lack specificity — infiltrative tumour margins may appear indistinct from surrounding edema or post-treatment changes, leading to challenges in accurate delineation. Differences in MRI acquisition parameters, scanner hardware, and protocol heterogeneity across centres can also reduce reproducibility of imaging biomarkers and hamper generalisation of analytic models. A review noted that while MRI is central in neuro-oncology, its ability to discriminate between tumour tissue and non-tumour lesions remains limited without advanced sequences. (Sawlani et al., 2020)

4.1. Role of Multimodal MRI in Improving Diagnostic Accuracy

- Combining multiple MRI sequences (T1, T2, FLAIR, T1-CE) enables complementary contrast information—structural detail from T1, fluid/edema from T2/FLAIR and neovascular enhancement from T1-CE—thus improving tumour margin delineation and reducing mis-classification of peritumoral tissue (Buchner et al., 2023).
- Integration of functional MRI sequences (such as perfusion or diffusion) with structural imaging allows assessment of tumour micro-environment (e.g., high rCBV or low ADC) which helps in grading gliomas and distinguishing high-grade from low-grade lesions.
- Multimodal data fusion supports deep-learning models by providing diverse inputs that improve feature learning and generalisation, especially in heterogeneous datasets collected from different scanners or patient populations (Zhang et l., 2021).
- Multi-sequence MRI enables better monitoring of treatment response or recurrence, as changes in enhancement pattern (T1-CE) combined with evolving edema or diffusion changes (T2/FLAIR, DWI) provide richer biological context than single-modality imaging (Sawlani et al., 2020).

- Use of multimodal MRI protocols standardises imaging across centres and supports benchmarking and segmentation tasks (for example via the BRATS dataset), promoting reproducibility and comparability of research and clinical results (Buchner et al., 2023).

4.2. Deep Learning in Medical Imaging: A Background

Deep learning (DL) has rapidly transformed the field of medical imaging by enabling automated feature extraction and high-performance image analysis tasks that were previously extremely labour-intensive. Traditional medical-image analysis often relied on handcrafted features (shape, texture, intensity) followed by classifiers such as support vector machines or random forests; however these approaches struggled with the complex, high-dimensional nature of modalities such as MRI, CT, and PET. Deep neural networks, by contrast, learn hierarchical representations directly from raw image data, significantly improving performance in tasks like segmentation, detection and classification of pathology, including in neuroimaging. For example, one broad survey reported that for radiological imaging, deep learning models achieved notably better diagnostic accuracy compared with classical methods. (Aggarwal et al., 2021)

Nevertheless, DL in medical imaging also brings new challenges—such as the need for large labelled datasets, domain shift across scanners and protocols, interpretability, and robustness to real-world clinical variability. A comprehensive review highlighted that while DL algorithms have achieved impressive success, they must be carefully adapted to the imaging domain given issues like small sample sizes, volumetric data, and multi-modal inputs in neuroimaging. (Zhou et al., 2021) These observations underscore the importance of understanding both the promise and the limitations of DL, especially when applied to complex tasks such as brain tumour diagnosis from MRI.

5. Conventional Machine Learning vs. Deep Learning

In conventional machine-learning (ML) workflows for medical imaging, analysts first design features manually—such as region of interest (ROI) shape descriptors, histogram and texture features—and then feed these into classifiers like SVMs, logistic regression or decision trees. These pipelines depend heavily on domain expertise for feature engineering and often struggle to scale or generalize when faced with heterogeneous imaging data or novel tumour presentations. By contrast, deep learning eliminates (or greatly reduces) the need for manual feature engineering: convolutional layers automatically learn relevant spatial filters, and deeper layers extract progressively more abstract semantic features from the image data.

Another distinction is how the models handle data volume and complexity. Deep learning models excel when large sets of annotated images exist and when variability is high—because they can learn from examples rather than relying on pre-defined feature sets. For brain tumour MRI analysis, DL approaches have shown markedly improved accuracy over traditional ML methods, particularly when large multi-modal datasets are available. However, deep learning is also more computationally intensive, requires more data, and may be more prone to overfitting or domain-shift issues compared with simpler ML models. Overall, while both ML and DL have roles in medical imaging, the evolution toward deep learning marks a significant shift in capability and practice.

5.1. Key DL Architectures: CNNs, RNNs, GANs, Transformers

Convolutional Neural Networks (CNNs): Designed specifically to process grid-like data (images), CNNs apply convolution and pooling operations to learn spatial hierarchies of features, making them especially effective for segmentation and classification tasks in medical imaging. (Zhang et al., 2023)

Recurrent Neural Networks (RNNs): Primarily used for sequential or time-series data, RNNs (and their gated variants such as LSTM/GRU) can model inter-slice dependencies in volumetric MRI or temporal changes in longitudinal imaging studies.

Generative Adversarial Networks (GANs): Comprise a generator and discriminator network that compete, and are used in medical imaging to synthesise realistic annotated images for augmentation, or to perform modality translation (e.g., generating contrast-enhanced images from non-contrast MRI).

Transformers / Vision Transformers (ViTs): Based on self-attention mechanisms, transformer architectures capture global dependencies and long-range context in images; when adapted to medical imaging, they can improve performance especially when combined with CNN backbones to handle both local and global features. (Huang et al., 2022)

6. Hybrid Deep Learning Architectures: Concepts and Taxonomy

Hybrid deep learning architectures combine two or more different modelling techniques or network types to leverage their complementary strengths, address shortcomings of single-architecture systems, and enhance overall performance in complex tasks such as MRI-based brain tumour diagnosis. In essence, these hybrid systems aim to merge spatial feature extraction, temporal or contextual modelling, modality-fusion or ensemble strategies into a unified pipeline. The rationale is that while conventional single-type deep networks (e.g., CNNs) excel at capturing local spatial features, they may struggle with long-range dependencies, temporal relationships, or multi-modal integration; hybrid designs mitigate this by embedding additional network modules (e.g., RNNs, transformers), combining deep-learned features with classical machine-learning classifiers, or utilising generative models for augmentation or translation. In medical image analysis, hybrid architectures have shown enhanced robustness, better generalisation across centres, and improved performance on heterogeneous datasets. For instance, reviews of DL in medical imaging highlight that hybrid approaches—such as combining CNNs with RNNs or transformers—are increasingly used to handle volumetric data, multi-slice dependencies and multi-sequence MRI inputs. (Li et al., 2023) Moreover, studies of hybrid vision transformer models report that integrating convolutional and attention-based modules yields superior segmentation or classification outcomes compared to standalone architectures. (Pu et al., 2024) As the MRI-based brain tumour diagnosis task demands integration of multi-modal inputs, contextual/temporal information and high spatial resolution, hybrid deep learning systems provide a compelling framework. A clear taxonomy helps us categorise such systems, understand their design trade-offs and assess their suitability for different tasks (detection, classification, segmentation) in neuro-oncological imaging.

6.1. Types of Hybrid Systems (CNN–RNN, CNN–GAN, CNN–Transformer, Ensembles)

Hybrid deep learning systems in the brain tumour MRI domain can be broadly grouped into the following categories:

- **CNN–RNN hybrids:** In this configuration, a convolutional neural network (CNN) serves as the feature extractor from MRI slices or volumes, capturing spatial patterns, while a recurrent neural network (RNN) (e.g., LSTM or GRU) models sequential or inter-slice dependencies or longitudinal temporal changes. For example, in a longitudinal study of tumour growth, the CNN extracts features from each time-point scan while the RNN integrates temporal progression, providing improved grading or progression prediction compared with a CNN alone. The sequential modelling of RNN augments spatial extraction of the CNN and helps capture dependencies across slices or time-points.
- **CNN–GAN hybrids**: A generative adversarial network (GAN) is typically used to augment the training data (especially when annotated tumour MRI scans are scarce) or to perform domain adaptation (e.g., synthesising contrast-enhanced images from non-contrast MRI). The CNN then performs classification or segmentation on the augmented data, benefiting from improved robustness and variety of input. This hybrid approach addresses data scarcity and enhances generalisation of downstream CNN tasks.
- **CNN–Transformer hybrids:** Here, the CNN backbone extracts local spatial features from MRI and one or more transformer modules apply self-attention to model global context, long-range dependencies or multi-slice correlation across volumes. For instance, a hybrid architecture might use CNN layers to build feature maps per slice and then a transformer encoder to capture inter-slice relationships or integrate multi-modal channels. Recent studies show that CNN-Transformer hybrids outperform pure CNN or pure transformer models in segmentation and classification of medical images. (Pu et al., 2024)
- **Ensemble hybrids / Feature-fusion systems**: In ensemble hybrids, multiple networks—possibly of different architectures—are trained separately (for example, a CNN, a CNN-RNN, and a CNN-Transformer) and their outputs are combined by weighted voting or meta-classifier. Alternatively, feature fusion approaches extract features from several networks or modalities (e.g., CNN features + radiomic features + handcrafted features) and fuse them before classification. These hybrids aim to improve robustness and mitigate weaknesses of any single model. The choice among these hybrid paradigms depends on task complexity (detection vs segmentation), dataset size, modality heterogeneity, computational resources and the requirement for interpretability or generalisation.

7. MRI Preprocessing and Feature Engineering Techniques

Effective MRI-based analysis for brain tumours requires rigorous preprocessing and feature engineering to ensure that downstream deep learning models receive clean, consistent, and meaningful input. Preprocessing plays a critical role in reducing scanner- and protocol-induced variability, removing irrelevant tissue or

artefacts, and aligning the data across subjects and modalities. Without these steps, models may learn spurious features, suffer from reduced generalisability, or perform poorly when confronted with data from different centres. Once preprocessing is complete, feature engineering comes into play: this includes augmentation, synthetic generation of images to bolster training sets, and fusion or reduction of features (hand-crafted or learnt) to improve model efficiency and performance. A recent review of MRI brain imaging highlighted that preprocessing (including registration, skull extraction, intensity normalisation and resizing) plus feature engineering (e.g., radiomic feature extraction and selection) are widely used but also noted considerable variability in pipelines across studies. (Ottoni et al., 2025)

In the context of hybrid deep learning architectures, feature engineering gains additional importance: hybrid models often combine deep-learned representations with classical features, or require multi-modal fusion of data (e.g., T1, T2, FLAIR) before triggering the hybrid network pipeline. As such, preprocessing must not only standardise the data but facilitate multisequence registration, modality alignment, and feature fusion in a way that downstream models can meaningfully exploit complementary information. Moreover, since hybrid architectures may include generative modules (e.g., GANs) or attention-based modules (e.g., transformers), the quality of preprocessing (skull-stripping, denoising, normalisation) directly impacts the reliability of those more advanced components. For example, one study evaluating tumour segmentation found that the benefit of preprocessing steps such as skull stripping or intensity normalisation was surprisingly limited in some contexts, emphasising that preprocessing should be carefully designed rather than applied as a one-size-fits-all step. (Kondrateva et al., 2024)

7.1. Preprocessing Methods: Skull Stripping, Normalization, Denoising

Preprocessing of MRI data for brain tumour analysis generally begins with skull stripping (brain extraction) to remove non-brain tissues (e.g., skull, scalp, eyes, neck) from the volume, allowing the model to focus solely on intracranial tissues and reducing irrelevant background noise. Skull stripping improves segmentation and classification accuracy and reduces computational burden. Several algorithms exist (morphology-based, thresholding, atlas-based, and deep-learning based) for this task. (Pei et al., 2022) After skull stripping, intensity normalization is applied: this may include bias-field correction, histogram equalisation, z-score normalisation, or min-max scaling, to account for scanner/protocol differences and ensure consistent intensity distributions across the dataset. Normalisation is essential before combining data from different centres or modalities. Denoising is another key step: MRI scans often contain noise (thermal, physiological, artefacts) and inhomogeneity (due to magnetic field). Denoising algorithms (e.g., non-local means, anisotropic diffusion, deep-learning denoising) help improve image quality, enhance signal-to-noise ratio, and support better downstream feature extraction. Good preprocessing ensures that the subsequent feature engineering and model training steps operate on stable and relevant data. (Kalavathi & Prasath, 2015)

8. Hybrid Models for Brain Tumor Detection, Classification, and Segmentation

Hybrid deep learning models have emerged as a powerful approach for brain tumor detection because they combine the spatial learning strengths of convolutional neural networks (CNNs) with additional modules that capture temporal, contextual, or higher-level representations from MRI data. Tumor detection models must not only identify abnormal regions but also differentiate them from normal anatomical variations or post-treatment changes. Hybrid architectures often use CNNs for low-level spatial feature extraction and integrate recurrent neural networks (RNNs) or transformers to model cross-slice dependencies in 3D MRI volumes. This multi-component structure enhances the model's ability to detect subtle lesions that may be missed by single-architecture systems. The inclusion of generative adversarial networks (GANs) for synthetic MRI generation is another common hybrid approach that addresses data scarcity by augmenting datasets with realistic tumor-bearing images. For example, GAN-augmented CNN models have been shown to improve tumor-presence classification performance, particularly when training data is limited or imbalanced.

For tumor classification, hybrid models integrate multiple architectures to better capture the heterogeneity of tumor morphology, texture, and contrast patterns across MRI sequences. Classification tasks are complex because tumors such as gliomas present with wide variability in location, intensity, and infiltration patterns. Traditional CNN networks struggle with long-range dependencies and multimodal integration, but hybrid systems mitigate this by combining CNN-extracted features with attention modules or classical machine-learning classifiers to strengthen decision-making. Ensemble hybrids—where outputs from several deep networks are combined through weighted averaging or a meta-learner—have demonstrated improved tumor subtype and grade classification performance, especially on heterogeneous datasets. These models benefit from fusing multi-sequence MRI inputs, such as T1-weighted, T2-weighted, FLAIR, and contrast-enhanced

imaging. A recent review of hybrid and ensemble models in neuro-oncological imaging noted substantial improvements in robustness, with classification accuracy surpassing that of individual networks alone. (Li et al., 2023) Hybrid architectures that combine CNN feature extractors with SVM or random forest classifiers also show strong performance, indicating that blending deep and classical methods can enhance interpretability and generalization across imaging centers. (Ottoni et al., 2025)

8.1. Classification and Grading Using Combined Models

Brain tumor classification and grading benefit substantially from hybrid deep learning architectures because these models integrate complementary feature-extraction and decision-making mechanisms to handle variability in tumor appearance across MRI modalities. Grading tasks, such as distinguishing low-grade from high-grade gliomas, require capturing fine-grained features related to tumor infiltration, necrosis, and contrast enhancement — patterns that may not be fully captured by CNNs alone. Hybrid classification models typically begin with deep CNNs or 3D CNNs to extract spatial and textural features from multi-sequence MRI inputs. These features are then fused with attention mechanisms, RNN modules, transformers, or classical machine-learning classifiers to enhance discriminative power. In particular, CNN–Transformer combinations have demonstrated strong performance, as transformers capture long-range spatial dependencies and global contextual cues that help distinguish between subtle grade differences. Ensemble-based hybrids further improve classification reliability by aggregating outputs from diverse models to reduce overfitting and variance. Classical ML components, such as SVMs or XGBoost classifiers, are often used at the final decision stage to improve interpretability and generalization across different MRI scanners. Evidence from recent research indicates that hybrid architectures consistently outperform single-architecture systems for glioma grading tasks, demonstrating higher accuracy and better robustness to dataset heterogeneity. (Aamir et al., 2025)

9. Datasets and Benchmarking Protocols for Hybrid DL Models

High-quality publicly-available MRI datasets play a foundational role in developing and benchmarking hybrid deep-learning (DL) models for brain tumour diagnosis. These datasets provide both imaging and ground-truth annotation (segmentation masks, tumour grades, sometimes genomics) and enable reproducible evaluation across research groups. The widespread adoption of such datasets has helped standardise tasks like segmentation, classification and detection of brain tumours using MRI. Nonetheless, benchmarking protocols—such as how data is split into training/validation/test sets, how metrics are computed, and how models generalise across centres—remain variable, thereby complicating direct comparison of hybrid architectures. A recent review observed that although many studies report high performance on internal splits of standard datasets, fewer undertake rigorous cross-dataset evaluation or multi-centre validation, which limits clinical translation. (Dorfner et al., 2025)

In the context of hybrid deep-learning models, proper benchmarking is even more critical: because such architectures can combine network types, fuse multi-modal inputs or ensemble outputs, consistent evaluation of their added value requires standardised protocol, transparent reporting of dataset splits, metrics and external validity. Use of challenge datasets like BraTS (Brain Tumor Segmentation) helps in this regard, but many studies still rely on private data or ad-hoc splits. To build confidence in hybrid models (e.g., CNN-Transformer, CNN-GAN hybrids), it is essential to follow best practices in data splitting, reporting performance on unseen centres, and assessing model robustness to domain shift.

9.1. Overview of Popular Datasets: BRATS, TCGA, REMBRANDT

Here is a summary of key public datasets widely used in MRI-based brain tumour research:

- **BraTS (Brain Tumor Segmentation Challenge dataset):** Initiated in 2012, BraTS provides multi-institutional, multi-modal MRI volumes (T1, T1-CE, T2, FLAIR) of glioma patients with manual sub-region labels (enhancing tumour, tumour core, oedema) and increasingly includes post-treatment and non-glioma tumour cases. It has become a benchmark for segmentation and hybrid model evaluation. (Menze et al., 2014)
- **TCGA (The Cancer Genome Atlas) glioma collections:** These datasets (e.g., TCGA-GBM, TCGA-LGG) link brain tumour MRIs with genomic and clinical data. They enable radiogenomic studies and classification/grading tasks beyond segmentation, making them suitable for hybrid DL models that fuse imaging with other modalities. (Mohsen et al., 2025)
- **REMBRANDT (Repository for Molecular Brain Neoplasia Data):** REMBRANDT includes MRI scans, pathological and genomic data from glioma patients, and offers longitudinal clinical outcome

information. It supports studies on tumour progression, survival prediction and hybrid modelling combining imaging and non-imaging features. (Mohsen et al., 2025)

9.2. Data Splitting and Benchmarking Standards

Proper dataset splitting and benchmarking are vital for fair evaluation of hybrid DL models. Standard practice includes dividing data into independent training, validation and test sets, with the test set held out during model development. Cross-validation or k-fold approaches are often used for internal validation, but they must avoid data leakage (e.g., same patient appearing in multiple splits). Furthermore, benchmarking standards require consistent metric definitions (e.g., Dice coefficient for segmentation, accuracy/F1 for classification), clearly reported hyper-parameters, and ideally external test sets from independent institutions. Without these 'gold-standard' protocols, reported gains for hybrid models may simply reflect over-fit or dataset-specific bias rather than real improvement. (Dorfner et al., 2025)

9.3. Cross-Dataset Evaluation and Generalization Issues

One of the primary challenges in MRI-based brain tumour modelling is generalisation across scanners, protocols, institutions and patient populations. Hybrid DL models may perform well on the dataset they are trained on but degrade on unseen data due to domain shift—differences in acquisition, resolution, contrast or patient demographics. Studies have found that cross-dataset evaluation (i.e., training on one dataset and testing on another) is much less common, yet essential to assess model robustness and translational readiness. For example, one recent work used cross-dataset testing of brain tumour MRI classification and showed a substantial drop in performance when models were applied to unseen data sources. (Tian et al., 2025) Hybrid architectures need this type of evaluation more than ever because their increased complexity could exacerbate overfitting or reliance on dataset-specific artefacts rather than generalisable features.

10. Performance Analysis of Hybrid Architectures Models

Evaluating hybrid deep learning architectures for MRI-based brain tumour tasks requires detailed analysis of how they perform under realistic conditions, both in internal validations and external/generalization settings. Hybrid systems—such as combinations of CNNs with transformers, GANs, or classical machine-learning classifiers—have shown markedly improved performance in many studies. For example, a systematic review found that hybrid models demonstrated superior diagnostic power compared with single-method approaches, showing higher resistance to false positives and class imbalance. Moreover, one recent study of MRI-based brain tumour classification reported very high accuracy (e.g., ~99.77%) using a hybrid CNN + XGBoost model with explainable AI components. These findings suggest that hybrid architectures can push performance metrics to new levels.

However, high performance on internal datasets does not always translate into clinical readiness: issues such as over-fitting, domain shift (scanner/protocol variability), unbalanced classes, and lack of interpretability pose challenges. Although many studies report high accuracy or Dice scores, fewer present robustness testing, external validation, or interpretability analyses. A review on DL and machine learning for brain tumour MRI emphasised that performance evaluation must consider model robustness, reliability, interpretability, and generalisation beyond the dataset on which it was trained. (Mohsen et al., 2025) Therefore, a comprehensive performance analysis of hybrid models must include not only accuracy metrics but also comparisons to non-hybrid baselines, evaluation of generalisation, and assessment of interpretability and reliability in clinical workflows.

10.1. Evaluation Metrics: Accuracy, F1-Score, Dice, IoU Models

In classification tasks (e.g., tumour vs non-tumour, tumour subtype), common performance metrics include accuracy, precision, recall, and the F1-score (the harmonic mean of precision and recall). These metrics are intuitive and widely used when classes are balanced, but may mask poor performance in imbalanced scenarios — hence the F1-score is especially relevant when false negatives carry high risk. For segmentation tasks (e.g., tumour region delineation), volumetric overlap metrics such as the Dice similarity coefficient (DSC) and Intersection over Union (IoU) are standard. Dice measures overlap between predicted and ground-truth masks, while IoU quantifies the ratio of intersection to union of the prediction and reference. For example, a review

of DL models in brain tumour MRI summarised that Dice scores often serve as the key benchmark for segmentation, while classification studies routinely report accuracy and F1-score. (Mohsen et al., 2025) Additionally, some studies report Hausdorff distance, volumetric similarity error, and sensitivity/specificity to provide fuller performance context. Using a consistent set of metrics across studies allows meaningful comparison of hybrid vs non-hybrid models, but such consistency is often lacking in the literature.

10.2. Comparative Analysis of Hybrid vs. Non-Hybrid Models

Several studies have directly compared hybrid deep-learning architectures with non-hybrid (single-network) models and found meaningful performance gains for hybrids. For example, the systematic review reported that combining CNNs with traditional ML classifiers or attention mechanisms tended to improve classification accuracy and reduce false positive rates compared with standalone CNNs. These improvements can be attributed to the complementary strengths of the hybrid components: the CNN captures spatial features, while the additional component (e.g., RNN, transformer, classical classifier) handles global context, temporal/sequential information or decision-level fusion. Nevertheless, some studies caution that the performance gains may come with increased complexity, risk of over-fitting, and greater computational cost. It is essential that hybrid models be benchmarked against strong baselines under the same data splits, and that gains are statistically significant and clinically meaningful.

11. Challenges, Limitations, and Clinical Integration Barriers

Despite rapid progress in hybrid deep learning architectures for MRI-based brain tumour analysis, several challenges hinder their adoption in real-world clinical workflows. Hybrid models often require large annotated datasets, yet high-quality medical imaging labels remain scarce, making it difficult for these systems to achieve consistent performance across institutions. Additionally, the heterogeneity of MRI acquisition protocols leads to domain-shift issues that degrade model generalisation when applied to new scanners or populations. Computational complexity presents another barrier, as hybrid models—especially those combining CNNs with transformers or GANs—demand substantial GPU resources, limiting use in resource-constrained settings. Clinical interpretability also remains limited; many hybrid architectures rely on opaque intermediate representations that make it difficult for radiologists to trust automated recommendations. Regulatory challenges, including requirements for explainability, reproducibility and safety evaluation, further slow down deployment in hospitals. A recent review emphasised that despite promising accuracy, deep learning systems in neuro-oncology face major translational barriers, particularly around data scarcity, workflow integration, and the need for rigorous external validation. (Dorfner et al., 2025)

11.1. Data Scarcity, Imbalance, and Annotation Challenges

A core limitation in developing robust hybrid deep learning models is the scarcity of high-quality, expertly annotated MRI datasets. Medical imaging data is often restricted due to privacy regulations, uneven data-sharing practices and institutional silos, resulting in datasets that are insufficiently large for training hybrid architectures with millions of parameters. Even widely used datasets like BraTS, while invaluable, remain relatively small when compared to the scale required for training advanced transformer-based or GAN-based hybrids. In addition, many MRI datasets lack uniform coverage of tumour types, grades or sequences, resulting in modality-specific gaps. This imbalance causes hybrid models to overfit to majority classes or sequences, performing poorly on rare tumours or underrepresented patterns. (Mohsen et al., 2025)

Annotation quality poses another major challenge: accurate tumour segmentation requires considerable time and expertise from neuroradiologists, making manual labelling costly and often inconsistent across institutions. Inter-observer variability further complicates the creation of reliable ground truth, particularly for infiltrative tumour margins where experts may disagree. Hybrid models, which often rely on multi-stage or multi-modal pipelines, are particularly sensitive to noisy labels because errors propagate through multiple components. Researchers have attempted partial solutions—such as weak supervision, semi-supervised learning, or crowdsourcing—but inconsistencies remain. Additionally, many hybrid models require voxel-level labels for segmentation and region-level labels for classification, yet most clinical datasets only contain diagnostic summaries rather than detailed annotations, limiting their utility for training. (Aamir et al., 2025)

11.2. Computational Complexity and Resource Requirements

Hybrid deep learning architectures often combine multiple computationally intensive components—such as CNN backbones, attention-based transformer modules, recurrent layers, or GAN sub-networks—resulting in large memory footprints and long training times. This complexity poses a significant barrier to clinical integration, especially for hospitals with limited computational infrastructure. Inference time can also be problematic: high-resolution 3D MRI volumes require substantial processing, and hybrid models may incorporate multiple passes, ensemble steps, or multi-modal fusion operations that increase latency. Such constraints limit real-time or near-real-time clinical use, particularly in emergency or intraoperative settings where rapid decision-support is critical. Moreover, resource-heavy training requirements impede reproducibility, as many research groups cannot replicate published results without access to high-end GPU clusters. A study evaluating advanced transformer-based architectures highlighted that hybrid systems offered improved accuracy but substantially higher computational cost compared to CNN-only baselines, raising concerns about scalability and deployment feasibility. (Pu et al., 2024)

11.3. Ethical, Legal, and Clinical Adoption Issues

Ethical and legal considerations pose significant obstacles to the clinical deployment of hybrid deep learning models. Patient privacy is a major concern, as MRI datasets often contain identifiable information, and large-scale data sharing is required to train generalisable models. Even anonymised data carries risks of re-identification when combined with external datasets. Issues of algorithmic bias also arise: if training data is skewed toward particular demographic groups or tumour types, hybrid architectures may perform poorly for underrepresented populations, exacerbating healthcare disparities. Regulatory frameworks are still developing, and hybrid systems—which involve complex chains of components—face additional scrutiny regarding traceability, validation, and safety. In many jurisdictions, AI models must provide explainability, yet hybrid architectures often rely on obscure intermediate representations, making regulatory compliance challenging. (Dorfner et al., 2025)

Clinical adoption requires trust, transparency, and seamless integration into existing radiology workflows. Even high-performing hybrid models may be rejected by clinicians if they provide unclear reasoning, unpredictable behaviour, or poor compatibility with reporting systems. Moreover, medico-legal liability remains unclear: if an AI-assisted diagnosis is incorrect, responsibility between radiologists, hospitals, and software developers becomes difficult to assign. A recent review of AI in medical imaging highlighted that reliable interpretability, continuous model monitoring, and well-designed human-AI collaboration strategies are essential prerequisites for widespread use. (Aamir et al., 2025)

12. Conclusion

Hybrid deep learning architectures represent a decisive evolution in MRI-based brain tumor diagnosis, offering an integrated framework capable of addressing the intrinsic challenges posed by heterogeneous tumor biology and multimodal imaging data. By blending the spatial representational power of CNNs with the contextual modeling capacity of transformers and RNNs, and augmenting learning through GAN-based synthesis or classical machine-learning classifiers, hybrid systems achieve higher diagnostic accuracy, more reliable segmentation, and improved handling of multi-sequence MRI inputs. These gains are evident across benchmark datasets such as BraTS and TCGA, where hybrid architectures consistently outperform single-model baselines on metrics including Dice coefficient, accuracy, and F1-score. Furthermore, hybrid designs inherently support multimodal fusion, enabling the integration of structural, functional, and advanced MRI sequences to capture tumor-specific features such as infiltration patterns, edema spread, and contrast enhancement. This multi-component learning process not only enhances performance but also allows greater flexibility in addressing domain shift, variability in acquisition protocols, and the need for explainability through interpretable components such as attention maps or classical classifier decision weights. Consequently, hybrid architectures hold strong potential for real-world clinical deployment, provided that reliability and interpretability continue to improve.

Despite these advancements, substantial challenges remain before hybrid models can be seamlessly integrated into clinical neuro-oncology workflows. Data scarcity, annotation difficulty, and class imbalance continue to limit the robustness and generalizability of complex hybrid systems, particularly in multi-institutional environments. Computational demands further hinder scalability, as transformer- and GAN-enhanced architectures require extensive resources that are often unavailable in routine clinical settings. Ethical and regulatory issues—including patient privacy, algorithmic bias, and the need for transparency in diagnostic decision-making—pose additional barriers to deployment. To unlock the full potential of hybrid deep learning in brain tumor MRI, future research must prioritize large-scale, standardized, diverse datasets; cross-dataset

validation protocols; efficient model designs; and interpretable hybrid pipelines aligned with clinical needs. With these developments, hybrid architectures can serve as a cornerstone for trustworthy, accurate, and actionable AI-driven neuroimaging tools.

Reference

1. Martucci, H. M., Jain, R., Borges, L. F., & Young, R. J. (2023). Magnetic resonance imaging of primary adult brain tumors. Scientific Data.
2. Sawlani, V., Davies, N., & Dineen, R. (2020). Multiparametric MRI: Practical approach and pictorial review. Clinical Radiology.
3. Satushe, S., et al. (2025). AI in MRI brain tumor diagnosis: A systematic review of machine learning and deep learning applications. Frontiers in Oncology.
4. Buchner, J. A., et al. (2023). Identifying core MRI sequences for reliable automatic brain tumour segmentation across centres. Medical Image Analysis,
5. Zhang, W., et al. (2021). Overview of multi-modal brain tumour MR image segmentation: challenges and methods. PMC.
6. Aggarwal, R., et al. (2021). Diagnostic accuracy of deep learning in medical imaging. NPJ Digital Medicine,
7. Huang, L., et al. (2022). A transformer-based generative adversarial network for brain tumor segmentation with multi-modalities MRI. Frontiers in Neuroscience,
8. Zhang, H., et al. (2023). Applying deep learning to medical imaging: A review. Applied Sciences,
9. Zhou, S. K., Greenspan, H., Davatzikos, C., Duncan, J. S., van Ginneken, B., Madabhushi, A., … Summers, R. M. (2021). A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. Medical Image Analysis,
10. Li, M., et al. (2023). Medical image analysis using deep learning algorithms. Frontiers in Public Health
11. Pu, Q., et al. (2024). Advantages of transformer and its application for medical image segmentation. Biomedical Engineering Online, 23, Article 12.
12. Kalavathi, P., & Prasath, V. B. (2015). Methods on skull stripping of MRI head scan images — A review. Journal of Medical Imaging and Health Informatics, 5(4), 731-742.
13. Kondrateva, E., Druzhinina, P., Dalechina, A., Zolotova, S., Golanov, A., & Shirokikh, B. (2024). Negligible effect of brain MRI data preprocessing for tumour segmentation. Computerized Medical Imaging and Graphics,
14. Ottoni, M., et al. (2025). Machine learning in MRI brain imaging: A review of preprocessing, feature engineering and segmentation. Diagnostics, 15(21), 2692.
15. Pei, L., et al. (2022). A general skull stripping of multiparametric brain MRIs. Scientific Reports, 12, 14983.
16. Aamir, M., et al. (2025). From CNNs to transformers: A review of evolving deep learning architectures for brain tumor classification. Frontiers in Oncology.
17. Dorfner, F. J., et al. (2025). A review of deep learning for brain tumor analysis in MRI. npj Precision Oncology.
18. Menze, B. H., Jakab, A., Bauer, S., Kalpathy-Cramer, J., Farahani, K., Kirby, J., … Weber, M.-A. (2014). The Multimodal Brain Tumor Image Segmentation Benchmark (BRATS). IEEE Transactions on Medical Imaging, 34(10), 1993-2024.
19. Mohsen, S., … (2025). Deep learning and machine learning for brain tumor detection from MRI. Machine Learning Review Journal.
20. Tian, Y., … (2025). Trade-off analysis of classical machine learning and deep learning for cross-domain brain tumor MRI classification. JMIR Artificial Intelligence,