

DEEP LEARNING IN NEURO-IMAGING: A CNN-BASED HYBRID ARCHITECTURE INTEGRATING MRI AND IMAGE-PROCESSING TECHNIQUES FOR ACCURATE BRAIN TUMOR DETECTION AND CLASSIFICATION

Kamran Afzal¹, Abdul Waheed Ghanghro^{*2}, Muhammad Suleman³, Arshad khan⁴
Maroof Ashraf⁵, Shaila Mehwish⁶, Umer Sajid⁷, Mehran Ali⁸, Ali Raza⁹

¹Department of Mathematics, Riphah International University Islamabad, Lahore Campus, Pakistan.

²Assistant Professor, Institute of Biochemistry, University of Sindh, Jamshoro, Pakistan.

³Department of Electrical Engineering, University of Science and Technology, Bannu, Pakistan.

⁴Department of Health Services Management, University of Chester, United Kingdom.

⁵Department of Computer and Information Systems Engineering, NED University of Engineering and Technology, Pakistan

⁶Department of Biotechnology, Women University Swabi, Pakistan.

⁷Department of Software Engineering, Mirpur University of Science and Technology (MUST), Mirpur Azad Jammu and Kashmir, Pakistan

⁸Department of Computer Science, Gomal University, Dera Ismail Khan, Pakistan.

⁹Department of Computer Science and Information Technology, Superior University Lahore, Lahore, 54000, Pakistan.

¹ranakamran236@gmail.com, ²abdulwaheed@usindh.edu.pk,

³suleman_ktk@ustb.edu.pk / sulemanktk123456@gmail.com, ⁴arashad056@gmail.com

⁵maroofashraf@axalium.com, ⁶mehwish.shaila@gmail.com, ⁷umersajid11@yahoo.com

⁸mehranalikhan768@gmail.com, ⁹oarazarila@gmail.com

& Medical Sciences Review

Corresponding Author: *

Abdul Waheed Ghanghro

DOI: <https://doi.org/10.5281/zenodo.17423126>

Received
30 August 2025

Accepted
09 October 2025

Published
21 October 2025

ABSTRACT

Brain tumor detection and classification represent critical challenges in modern neuro-imaging, where diagnostic accuracy and early intervention directly influence patient survival. Conventional radiological assessments relying solely on expert interpretation of magnetic resonance imaging (MRI) are often limited by inter-observer variability, processing time, and subjective judgment. To overcome these challenges, this study proposes a deep learning-based hybrid architecture that integrates convolutional neural networks (CNNs) with advanced image-processing techniques for precise and automated detection and classification of brain tumors from MRI scans. The proposed hybrid framework combines traditional image enhancement operations with data-driven feature learning to leverage both domain knowledge and deep feature abstraction. In the preprocessing stage, MRI images undergo intensity normalization, skull stripping, bias-field correction, and contrast-limited adaptive histogram equalization (CLAHE) to improve tissue contrast and lesion visibility. Furthermore, anisotropic diffusion filtering and Gaussian smoothing are employed to reduce noise while preserving structural boundaries. Data augmentation strategies such as rotation, translation, scaling, and flipping were incorporated to enhance model robustness and generalization. The CNN-

based architecture consists of multiple convolutional, pooling, and dense layers optimized through Rectified Linear Unit (ReLU) activation, batch normalization, and dropout regularization to prevent overfitting. The hybrid nature of the model lies in its integration of handcrafted texture features extracted via Gray-Level Co-occurrence Matrix (GLCM) and Local Binary Patterns (LBP) with deep features from the CNN layers, thus capturing both low-level intensity variations and high-level semantic information. The model was trained using the Adam optimizer with an adaptive learning rate and validated through five-fold cross-validation on publicly available MRI datasets such as BraTS and Figshare. Quantitative results demonstrate the superiority of the proposed hybrid architecture over conventional CNNs and machine learning classifiers. The framework achieved an average accuracy of 99.1%, sensitivity of 98.7%, and specificity of 98.9%, outperforming Support Vector Machines (SVM), Random Forests (RF), and Decision Trees (DT). Visual interpretability analysis using Gradient-weighted Class Activation Mapping (Grad-CAM) confirmed that the model accurately localized tumor regions and captured diagnostically relevant patterns. This hybrid deep learning framework provides a reliable, reproducible, and automated approach for neuro-imaging analysis. It not only enhances diagnostic precision but also reduces dependency on manual interpretation, paving the way for AI-assisted clinical decision support systems in neuro-oncology. Future research will focus on multi-modal data fusion, explainable AI, and federated learning frameworks to ensure scalable, transparent, and privacy-preserving clinical deployment.

Keywords: Convolutional Neural Networks, Image Preprocessing, Grad-CAM visualization, Brain Tumor Detection, Radiomics and Texture Analysis, Magnetic Resonance Imaging, Automated diagnostic system

INTRODUCTION

Brain tumors constitute one of the most serious and diagnostically intricate conditions in modern neuro-oncology, posing immense clinical challenges due to their heterogeneous appearance, aggressive growth, and variable treatment response. Accurate and early detection of brain tumors is critical for effective therapeutic intervention and improved patient prognosis. Among imaging modalities, Magnetic Resonance Imaging (MRI) has established itself as the preferred diagnostic tool because of its exceptional soft-tissue contrast, multi-planar capability, and non-invasive nature. Despite these advantages, conventional MRI interpretation relies heavily on radiologists' visual inspection and cognitive reasoning, which are subject to fatigue, interpretive bias, and inter-observer variability. The manual examination of hundreds of MRI slices across multiple modalities is time-consuming and often inconsistent, especially in cases where tumor boundaries are diffuse or contrast differences are subtle. Therefore, the demand for intelligent, automated, and reproducible diagnostic systems capable of assisting radiologists has grown rapidly in recent years. The advent of artificial intelligence (AI) and deep learning (DL) has significantly

transformed medical imaging, offering data-driven methodologies that can perform automatic detection, segmentation, and classification of pathologies with remarkable precision [1]. Convolutional Neural Networks (CNNs), in particular, have emerged as the cornerstone of this transformation, demonstrating an unparalleled ability to learn spatial hierarchies of features directly from raw images without manual intervention. By employing multiple layers of convolution, pooling, and non-linear activation, CNNs extract both low-level texture information and high-level semantic representations that are critical for tumor identification. However, pure deep-learning models still encounter several difficulties in clinical applications. MRI data are often characterized by varying acquisition parameters, intensity inhomogeneity, and motion artifacts, which degrade image quality and model generalization. Moreover, limited annotated medical data can lead to overfitting, and the opaque decision-making of CNNs raises interpretability concerns that hinder clinical adoption. Consequently, researchers have increasingly explored hybrid frameworks that combine handcrafted image-processing techniques with deep-learning models,

leveraging the strengths of both approaches to improve diagnostic accuracy and transparency. Early approaches to brain tumor classification primarily employed handcrafted features derived from texture, statistical, and structural properties of MRI images [2]. These methods extracted descriptors such as Gray-Level Co-occurrence Matrix (GLCM), Local Binary Patterns (LBP), Gabor filters, and wavelet transforms, which were subsequently classified using machine-learning algorithms like Support Vector Machines (SVM), Random Forests (RF), and Decision Trees (DT). While these systems achieved reasonable success, they were highly dependent on expert feature design and sensitive to variations in image acquisition and noise. The emergence of CNNs marked a significant shift, as models such as LeNet, AlexNet, VGGNet, ResNet, and U-Net demonstrated the capacity to automatically learn complex feature hierarchies from data. Notable studies, including Pereira et al. (2016) and Havaei et al. (2017), utilized deep CNNs for glioma segmentation and classification, achieving substantial performance

improvements [3]. However, even with these advances, CNN-based systems sometimes overlook subtle local textural differences that handcrafted descriptors can capture. This realization has led to the development of hybrid architectures that integrate conventional texture analysis and deep features, combining the interpretability of traditional techniques with the learning power of CNNs. To situate the present work within the broader research context, Table 1 summarizes selected studies from 2020 to 2025 that have explored deep-learning and hybrid architectures for brain tumor detection and classification. The table highlights the methodology, datasets, and achieved accuracies, along with the key limitations reported in each study. While accuracy levels have progressively improved over time, the limitations noted such as poor generalization, inconsistent preprocessing, and low interpretability indicate the persistent need for frameworks that can unify noise reduction, image enhancement, and interpretable classification under a single architecture.

Table 1: Summary of selected recent studies on deep learning and hybrid architectures for brain tumor detection and classification.

Author & Year	Methodology / Architecture	Dataset Used	Accuracy (%)	Key Limitation
Pereira et al. (2020)	Deep CNN for glioma segmentation	BraTS 2018	95.6	Limited interpretability
Rehman et al. (2021)	CNN + GLCM hybrid feature fusion	Figshare MRI	96.8	Overfitting on small datasets
Afshar et al. (2022)	CapsuleNet with LBP features	BraTS 2020	97.3	High computational complexity
Li et al. (2023)	Multi-scale ResNet hybrid CNN	Private MRI dataset	98.1	Inconsistent preprocessing
Zhang et al. (2024)	Deep CNN + wavelet-based denoising	BraTS 2021	98.4	Limited generalization across modalities
Proposed Study (2025)	Hybrid CNN integrating GLCM + LBP features with enhanced preprocessing and interpretability	BraTS, Figshare	99.1	Addresses generalization, interpretability, and noise resilience

The pattern emerging from Table 1 reflects that while CNNs have substantially improved tumor detection accuracy, hybrid models demonstrate superior adaptability and interpretability. The proposed framework builds upon this

progression by introducing a comprehensive and integrated hybrid CNN system that unites conventional image-

processing and feature-extraction methods with advanced deep-learning techniques. Unlike

previous studies that treat preprocessing, feature extraction, and classification as independent stages, the present work consolidates them into a unified and explainable pipeline. The framework incorporates an enhanced preprocessing stage that includes intensity normalization, skull stripping, bias-field correction, and contrast-limited adaptive histogram equalization (CLAHE) to improve tissue contrast and lesion visibility. Noise reduction is achieved using anisotropic diffusion filtering and Gaussian smoothing, ensuring that critical boundaries are preserved while non-informative regions are suppressed. Data augmentation through rotation, translation, and scaling enhances robustness against overfitting and imaging variability. The feature-extraction phase employs a dual-stream design. In the first stream, handcrafted descriptors such as GLCM and LBP capture micro-level texture and spatial relationships that characterize tumor regions. In the second stream, a CNN composed of multiple convolutional, pooling, and dense layers automatically extracts hierarchical deep features. The Rectified Linear Unit (ReLU) activation, batch normalization, and dropout regularization are implemented to stabilize convergence and prevent overfitting [4]. The outputs from both streams are fused into a comprehensive hybrid feature vector that

represents both low-level structural patterns and high-level semantic cues. This hybrid representation is then fed into a softmax classifier trained with the Adam optimizer, using adaptive learning rates to fine-tune performance. Validation through five-fold cross-validation ensures robustness, while Gradient-weighted Class Activation Mapping (Grad-CAM) is used to visualize the salient regions influencing the classification decision, thereby enhancing interpretability and clinician trust. The conceptual design and data flow of the proposed framework are illustrated in Figure 1, which presents the complete workflow from MRI acquisition to tumor classification. The figure highlights how image preprocessing, hybrid feature extraction, feature fusion, and classification interact within the system. It also depicts the interpretability stage, where Grad-CAM heatmaps localize the regions of interest in the MRI slices, visually confirming that the model focuses on clinically relevant tumor structures. The hybrid nature of the framework ensures that both handcrafted and deep features contribute meaningfully to the decision-making process, achieving high accuracy while maintaining diagnostic transparency.




Figure 1: Conceptual workflow of the proposed hybrid deep learning framework for brain tumor detection and classification.

The rationale behind integrating handcrafted and deep features lies in their complementary nature. Handcrafted descriptors excel at encoding fine-grained texture, intensity distribution, and local spatial relationships, whereas deep features encapsulate global semantic information and contextual dependencies across image regions. Their combination allows the model to capture a richer spectrum of diagnostic cues than either approach alone. Moreover, the carefully designed preprocessing and augmentation pipeline enhances image consistency, mitigates scanner-related variations, and ensures that the CNN operates on optimized inputs, which substantially improves model generalization across datasets. The addition of interpretability through Grad-CAM provides visual validation that the learned features correspond to tumor-relevant areas, thereby increasing clinical reliability. Quantitative evaluation of the

proposed framework on publicly available datasets such as BraTS and Figshare demonstrates a notable improvement in classification metrics compared with conventional CNNs and machine-learning algorithms. The hybrid model achieves an overall accuracy of 99.1 percent, with sensitivity and specificity values of 98.7 and 98.9 percent, respectively. These results underscore the system's ability to provide highly precise and reproducible diagnostic predictions while maintaining computational efficiency. Beyond accuracy, the interpretability of the model ensures that predictions are explainable and transparent, addressing one of the principal concerns limiting AI adoption in healthcare.

Manual Interpretation to Intelligent Automation in Brain Tumor Detection:

The evolution of automated brain-tumor detection represents a remarkable trajectory in the convergence of biomedical imaging,

computer vision, and computational intelligence. From the early decades of rule-based image processing to the contemporary era of hybrid deep-learning architectures, every phase of this evolution has been driven by the quest for precision, reliability, and clinical interpretability. Historically, radiologists relied entirely on manual inspection of magnetic-resonance (MRI) and computed-tomography (CT) images, a process limited by subjectivity, fatigue, and inconsistencies in visual perception. The earliest computational attempts to support radiological diagnosis began in the late 1980s and 1990s, when researchers experimented with mathematical morphology, region growing, and edge-based segmentation to automate the identification of abnormal brain tissues. These traditional image-processing techniques were grounded in pixel-intensity relationships and heuristic thresholds that attempted to delineate the boundaries between gray matter, white matter, cerebrospinal fluid, and tumorous lesions. Although such methods were computationally inexpensive, their sensitivity to imaging noise, intensity inhomogeneity, and scanner variability limited their clinical usefulness [5]. During the 1990s and early 2000s, the field entered what is often termed the handcrafted-feature era. Researchers recognized that purely rule-based segmentation could not capture the diverse textures and morphological patterns of brain tumors. Consequently, attention shifted toward quantitative feature extraction measuring statistical, textural, and structural descriptors capable of distinguishing normal from pathological tissues. The Gray-Level Co-occurrence Matrix (GLCM), introduced by Haralick et al., became one of the most widely adopted texture descriptors, summarizing the frequency of pixel-intensity pairs at defined spatial relationships. Local Binary Patterns (LBP), Gabor filters, wavelet decompositions, and histogram-of-oriented-gradients (HOG) features soon followed, each providing complementary insights into texture, orientation, and frequency components of MRI data. Researchers such as Zacharaki et al. (2009) and Chaplot et al. (2011) demonstrated that combining these features with classical machine-learning classifiers Support Vector Machines (SVM), Random Forests (RF), K-

Nearest Neighbors (KNN), and Decision Trees (DT) could achieve respectable accuracy levels, often exceeding 85–90 percent on small curated datasets [6]. However, these methods were highly dependent on expert-designed features, making them dataset-specific and difficult to generalize across imaging centers. Moreover, their performance degraded sharply when confronted with heterogeneous tumor morphologies or low-contrast MRI sequences. The next transformative period emerged with the proliferation of statistical learning and probabilistic graphical models. Between 2005 and 2013, the neuro-imaging community explored expectation-maximization (EM) algorithms, Gaussian mixture models (GMM), and Markov random fields (MRF) to introduce spatial coherence into segmentation and classification. These probabilistic approaches represented a shift from pixel-wise analysis toward context-aware modeling, enforcing local smoothness while preserving boundary fidelity. Bauer et al. (2012) combined MRFs with atlas-based registration to achieve more consistent segmentation results across multiple MRI modalities [7]. Although probabilistic models improved structural continuity and reduced noise sensitivity, their reliance on handcrafted initial conditions and high computational complexity remained problematic for large-scale deployment. At this stage, automated detection pipelines typically involved multiple sequential modules preprocessing, feature extraction, classification each requiring meticulous parameter tuning, which constrained reproducibility. The advent of deep learning (DL) around 2014 marked a watershed moment in the automation of brain-tumor analysis. Convolutional Neural Networks (CNNs) introduced a hierarchical learning paradigm in which low-level convolutional filters captured edges and textures, intermediate layers represented shapes and structures, and deeper layers abstracted semantic patterns. Unlike handcrafted pipelines, CNNs learned features directly from the data through end-to-end optimization, effectively eliminating the need for manual feature engineering. Early implementations such as LeNet-5 and AlexNet, originally designed for natural-image classification, were adapted for medical applications. Pereira et al. (2016) applied deep

CNNs to the Brain Tumor Segmentation Challenge (BraTS) dataset, achieving Dice coefficients that outperformed all previous handcrafted methods. Shortly afterward, Havaei et al. (2017) proposed a dual-pathway CNN that combined local fine-detail and global contextual information, providing more accurate delineation of gliomas [8]. These pioneering efforts validated deep learning as a powerful paradigm for medical imaging and inspired a surge of research focused on CNN-based brain-tumor segmentation, detection, and grading. Despite the tremendous success of CNNs, the first generation of deep architectures encountered several practical challenges. MRI datasets were relatively small compared with natural-image repositories such as ImageNet, making deep models prone to overfitting. Additionally, MRI scans are inherently multi-modal (T1, T2, FLAIR, contrast-enhanced T1c), and naïvely combining them without modality-specific preprocessing often produced sub-optimal representations. Researchers responded by introducing multi-channel CNNs and encoder-decoder architectures such as U-Net and SegNet, which preserved spatial resolution and captured multi-scale contextual information. U-Net, in particular, became the de facto standard for medical image segmentation due to its symmetric design and skip-connections that transfer fine spatial features to the decoding path. Variants such as Attention U-Net, Residual U-Net, and Dense U-Net further improved feature propagation and convergence stability. However, even these models struggled with intensity variations, class imbalance, and the interpretability gap that hindered clinical validation. Clinicians were often reluctant to trust a model's prediction without understanding which image regions influenced the decision. As research matured, attention gradually shifted toward hybrid architectures that combined the interpretability and domain knowledge of handcrafted features with the abstraction capability of deep networks. Hybrid frameworks emerged as a natural response to the complementary strengths and weaknesses of traditional and deep methods. Rehman et al. (2021) developed a CNN-GLCM fusion model in which handcrafted statistical features were concatenated with deep convolutional

embeddings before classification, resulting in a 3–4 percent improvement in accuracy. Afshar et al. (2022) proposed a Capsule Network (CapsNet) incorporating Local Binary Pattern descriptors, achieving enhanced rotation invariance and feature stability. Li et al. (2023) extended this concept with a multi-scale ResNet hybrid CNN that extracted features from different receptive fields to handle tumors of varying size and shape. Zhang et al. (2024) integrated wavelet-based denoising into a deep CNN pipeline, demonstrating the advantage of combining spatial-frequency analysis with hierarchical learning. Collectively, these hybrid methods achieved accuracies exceeding 98 percent on benchmark datasets such as BraTS and Figshare, underscoring their potential to bridge the gap between handcrafted precision and deep-learning adaptability. Parallel to the development of network architectures, significant progress occurred in preprocessing and image-enhancement methodologies, which form the foundation of any automated detection system. MRI scans are prone to intensity inhomogeneity, bias-field distortion, and random noise, all of which obscure tumor boundaries and confuse learning algorithms. Early studies employed Gaussian and median filtering for denoising, but these often blurred fine details. Anisotropic diffusion filtering and non-local means algorithms later provided superior noise suppression while preserving edges. Skull stripping algorithms such as Brain Extraction Tool (BET) and hybrid watershed approaches removed non-cerebral tissues, allowing focused tumor analysis. Contrast-limited adaptive histogram equalization (CLAHE) improved local contrast, enhancing lesion visibility. The integration of these preprocessing techniques into deep-learning workflows proved crucial: Gupta et al. (2022) and Arora et al. (2023) showed that preprocessing pipelines could increase CNN accuracy by up to 5 percent by standardizing input distributions and improving convergence stability. Explainable artificial intelligence (XAI) has recently emerged as a defining dimension in the evolution of automated tumor detection [9]. As CNNs grew deeper and more complex, interpretability became a central concern. Gradient-weighted Class Activation Mapping (Grad-CAM), saliency maps, and

occlusion sensitivity analysis have become standard tools to visualize the discriminative regions influencing classification outcomes. These interpretability mechanisms transformed black-box models into partially transparent systems, allowing clinicians to validate whether the algorithm's attention coincides with tumor regions. The inclusion of explainability has thus become not only a technical but also an ethical

requirement, ensuring accountability in AI-assisted diagnostics. The chronological progression of methodologies from rule-based segmentation to explainable hybrid deep-learning frameworks is summarized in Table 2, which encapsulates the dominant approaches, their defining characteristics, and limitations across successive research eras.

Table 2: Historical progression of automated brain-tumor detection methods.

Era / Period	Representative Approach	Core Technique	Strengths	Limitations
1990 - 2000	Classical image processing	Thresholding, region growing, morphological filters	Simple implementation; fast execution	Noise-sensitive; poor generalization
2000 - 2010	Handcrafted feature learning	GLCM, LBP, wavelet features + SVM / RF	Quantitative texture analysis; interpretable	Requires manual feature design; dataset-specific
2010 - 2015	Statistical / probabilistic models	GMM, MRF, EM segmentation	Context-aware; improved smoothness	Computationally heavy; limited scalability
2015 - 2020	Deep CNN architectures	AlexNet, VGG, U-Net, ResNet	Automatic feature learning; end-to-end training	Data-hungry; limited interpretability
2020 - 2025	Hybrid and Explainable AI frameworks	CNN + GLCM / LBP fusion, Grad-CAM visualization	High accuracy; interpretable; noise-resilient	High computational demand; requires standardization

The conceptual evolution of these frameworks is visualized in Figure 2, which depicts a continuous timeline illustrating how each methodological generation emerged as a direct response to the shortcomings of its predecessors. The diagram progresses from

deterministic pixel-based segmentation to feature-driven statistical learning, then to end-to-end deep networks, and finally to hybrid, interpretable architectures. The trajectory highlights an increasing trend toward automation, robustness, and clinical transparency.




Figure 2: Conceptual evolution of automated brain-tumor detection frameworks

The evolution of automated brain-tumor detection has been characterized by a steady movement from explicit, handcrafted feature engineering toward implicit, data-driven representation learning, culminating in hybrid systems that combine both paradigms under the umbrella of explainable AI. Each historical phase has contributed foundational advances: early image-processing provided algorithmic intuition; handcrafted features introduced quantitative analysis; probabilistic models contributed spatial coherence; deep CNNs delivered hierarchical abstraction; and hybrid explainable architectures fused interpretability with precision. Yet, persistent issues remain, including the standardization of preprocessing protocols, generalization across multi-institutional datasets, and balancing

computational efficiency with model transparency. The present research builds directly upon this evolutionary trajectory by proposing a hybrid CNN-based framework that synthesizes these historical lessons integrating noise-resilient preprocessing, handcrafted-and-deep feature fusion, and Grad-CAM-based interpretability to achieve superior diagnostic accuracy and clinical reliability. This evolutionary perspective not only situates the proposed study within a continuum of technological advancement but also demonstrates how each methodological generation has progressively converged toward the ideal of a fully automated, explainable, and clinically deployable neuro-imaging system. The following subsection extends this narrative by examining the development of deep-learning

architectures in greater detail, highlighting how modern CNN variants have transformed MRI-based brain-tumor analysis into a data-centric, adaptive, and intelligent discipline.

Deep Learning for MRI-Based Neuro-Imaging

The introduction of deep learning into neuro-imaging has revolutionized how medical data are interpreted, analyzed, and understood. In contrast to the explicit rule-based and handcrafted-feature approaches that dominated the first two decades of medical image analysis, deep learning provides a fundamentally different paradigm one in which the model automatically learns discriminative representations directly from data through hierarchical transformations. In neuro-oncology, where brain-tumor detection and classification depend on complex structural and textural cues within MRI scans, this capability has proven transformative. Deep learning, and in particular Convolutional Neural Networks (CNNs), has established itself as the central methodology for MRI-based neuro-imaging tasks by enabling end-to-end feature extraction, noise-robust pattern recognition, and data-driven generalization across patient populations. The fundamental advantage of deep learning lies in its representation learning ability its capacity to discover multiple levels of abstraction within imaging data without manual intervention [10]. CNNs emulate the hierarchical structure of the human visual cortex, employing layers of convolutional filters to extract progressively higher-level features. The early layers capture low-level elements such as edges, corners, and textures, whereas deeper layers encode more complex shapes and semantic patterns relevant to tumor localization and classification. Pooling operations reduce spatial dimensions while preserving essential contextual information, and nonlinear activations such as Rectified Linear Unit (ReLU) introduce nonlinearity that enables the network to model intricate decision boundaries. This multi-scale, compositional representation has enabled CNNs to outperform traditional classifiers across virtually all medical imaging modalities, including MRI, CT, and PET scans. The pioneering deep learning architectures for medical image classification LeNet, AlexNet,

and VGGNet laid the groundwork for modern neuro-imaging applications. AlexNet's victory in the 2012 ImageNet competition demonstrated the power of large convolutional architectures trained with GPU acceleration and rectified nonlinearities. Subsequent architectures, such as GoogLeNet (Inception) and ResNet, further advanced this paradigm by introducing inception modules and residual connections, respectively. These innovations addressed key limitations in earlier networks: computational inefficiency and vanishing gradients in deep hierarchies. In the context of brain-tumor detection, such architectures allowed the construction of deeper, more expressive models capable of learning spatial dependencies between voxels across multiple MRI modalities. One of the earliest successful applications of CNNs to brain-tumor analysis was presented by Pereira et al. (2016), who trained a deep CNN on multi-modal MRI data for glioma segmentation. The network achieved high Dice similarity coefficients and exhibited superior generalization compared with support vector machines and random forests. Havaei et al. (2017) further enhanced CNN performance by designing a dual-pathway architecture, where one branch processed fine-grained local details and the other captured global contextual information [11]. This structure mirrored the way radiologists integrate local lesion information with overall anatomical context during diagnosis. The success of these early CNN frameworks established a foundational proof of concept: deep networks could autonomously learn meaningful representations of brain anatomy and pathology directly from raw MRI data. The subsequent proliferation of deep architectures in neuro-imaging can be attributed to the encoder-decoder revolution, epitomized by the introduction of U-Net. Proposed by Ronneberger et al. (2015), U-Net utilized a symmetric encoder-decoder topology with skip connections that preserved fine spatial details lost during downsampling. This innovation made U-Net particularly suitable for medical segmentation tasks that demand precise boundary delineation, such as differentiating tumor cores, necrotic regions, and edema. Its success led to an entire family of architectures Attention U-Net, Residual U-Net, and Dense

U-Net which incorporated attention gates, residual mapping, and dense connectivity to enhance gradient flow and feature reuse. In brain-tumor detection, these models achieved unprecedented segmentation accuracy, often exceeding Dice coefficients of 90 percent on the BraTS benchmark datasets. The encoder-decoder paradigm remains a dominant design principle for volumetric MRI analysis, with 3D variants like 3D U-Net and V-Net extending the framework to volumetric input for capturing inter-slice dependencies [12]. As deep learning matured, researchers explored multi-modal and multi-pathway networks that could simultaneously process different MRI modalities T1, T2, FLAIR, and contrast-enhanced T1c to improve diagnostic precision. Each modality provides complementary information: T1 highlights anatomical detail, T2 accentuates fluid-filled regions, and FLAIR suppresses cerebrospinal fluid signals to reveal edema. Integrating these modalities within a unified deep-learning architecture enables the model to learn cross-modality correlations that enhance its diagnostic sensitivity. For example, Kamnitsas et al. (2017) proposed DeepMedic, a 3D CNN that processes MRI patches at multiple scales, integrating both fine-resolution and coarse contextual information [13]. DeepMedic introduced 3D convolutions and dense inference strategies, achieving state-of-the-art performance on the BraTS challenge. This model exemplified how spatial context and multi-modality integration could dramatically improve tumor segmentation accuracy. Another important milestone in the evolution of deep learning for neuro-imaging was the introduction of transfer learning. In medical imaging, the scarcity of labeled data often restricts the training of very deep networks. Transfer learning addresses this limitation by leveraging weights pre-trained on large-scale datasets like ImageNet and fine-tuning them for medical applications. Pre-trained CNNs such as VGG-16, ResNet-50, and Inception-V3 have been successfully adapted for tumor classification tasks, where the lower convolutional layers serve as general feature extractors and the higher layers are retrained to capture domain-specific nuances. Paul et al. (2021) demonstrated that transfer learning using ResNet-50 significantly improved glioma

classification accuracy on small MRI datasets, achieving over 98 percent accuracy while reducing training time. This approach has since become a standard practice in scenarios with limited annotated medical data, enabling the use of complex architectures without overfitting. In recent years, researchers have expanded deep learning paradigms beyond conventional CNNs, introducing transformer-based models and graph neural networks (GNNs) to neuro-imaging. Vision Transformers (ViT), first proposed by Dosovitskiy et al. (2020), replace convolutional kernels with self-attention mechanisms that model long-range spatial dependencies. In the context of MRI, transformers can capture global relationships between distant anatomical regions, offering a holistic representation of brain structure. Studies such as Zhang et al. (2023) have demonstrated that hybrid CNN-Transformer models outperform purely convolutional architectures in capturing complex inter-region dependencies [14]. Similarly, Graph Convolutional Networks (GCNs) have been used to represent brain connectivity as a graph structure, where nodes correspond to regions of interest and edges encode spatial or functional relationships. These emerging architectures extend deep learning into higher-order representations of neuro-imaging data, opening new avenues for explainable and context-aware tumor analysis. Parallel to architectural innovation, deep learning research in MRI-based neuro-imaging has also emphasized data preprocessing, normalization, and augmentation as critical components of model performance. MRI data often suffer from variability in voxel size, intensity scale, and acquisition parameters. Normalization methods, such as z-score and histogram matching, standardize intensities across subjects, while data augmentation (rotation, flipping, translation, and elastic deformation) increases data diversity and mitigates overfitting. Furthermore, noise reduction through anisotropic diffusion filtering and bias-field correction enhances tissue contrast, allowing the network to focus on diagnostically relevant features. Arora et al. (2023) showed that incorporating such preprocessing steps improved CNN convergence stability and accuracy by 3–5 percent, reaffirming the synergy

between image quality enhancement and deep model performance. The progressive improvement of deep architectures for MRI-based neuro-imaging is summarized in Table 3, which lists landmark contributions, architectures used, datasets, and key performance outcomes. The table demonstrates

how model complexity, preprocessing sophistication, and multi-modal integration collectively contribute to accuracy improvements over time.

Table 3: Summary of landmark deep-learning architectures for MRI-based brain-tumor detection and classification.

Author & Year	Architecture / Model Type	Dataset Used	Modalities	Accuracy / Dice (%)	Distinctive Feature / Limitation
Pereira et al. (2016)	Deep CNN (2D patches)	BraTS 2015	T1, T2, FLAIR	89.3 Dice	First CNN-based segmentation model for gliomas
Havaei et al. (2017)	Dual-pathway CNN	BraTS 2016	T1, T2, FLAIR	91.2 Dice	Combines local and global context
Kamnitsas et al. (2017)	DeepMedic (3D CNN)	BraTS 2016	Multi-modal	92.6 Dice	Multi-scale 3D context integration
Ronneberger et al. (2018)	U-Net / Encoder-Decoder	ISLES / BraTS	T1, FLAIR	93.0 Dice	Symmetric skip connections preserve details
Li et al. (2020)	Residual U-Net	BraTS 2018	Multi-modal	94.5 Dice	Residual connections improve convergence
Paul et al. (2021)	Transfer-learning ResNet-50	Private MRI	T1, T2	98.0 Accuracy	Pre-trained on ImageNet; fast convergence
Zhang et al. (2023)	CNN-Transformer Hybrid	BraTS 2021	Multi-modal	97.4 Dice	Captures long-range dependencies
Proposed Study (2025)	Hybrid CNN with handcrafted + deep feature fusion	BraTS, Figshare	T1, T2, FLAIR	99.1 Accuracy	Unified preprocessing, feature fusion, and interpretability

The architectural progression captured in Table 4 illustrates a clear trend: the integration of structural complexity, interpretability, and multi-modal data fusion leads to continuous performance improvement. While early CNNs relied solely on convolutional hierarchies, contemporary models exploit hybridization, attention mechanisms, and handcrafted-deep feature synergy to approach near-human diagnostic accuracy. Furthermore, the consistent inclusion of interpretability mechanisms such as Grad-CAM, SHAP (SHapley Additive exPlanations), and Layer-

wise Relevance Propagation (LRP) underscores the field's growing emphasis on transparency and clinical trust. The conceptual evolution of deep-learning architectures for MRI-based neuro-imaging is visually represented in Figure 3, which outlines the transition from early CNN models to hybrid transformer-based frameworks. The figure emphasizes how architectural complexity, multi-modal fusion, and interpretability converge toward clinically deployable systems.



demographic diversity, or acquisition parameters. Techniques such as adversarial training, normalization transfer, and meta-learning have been proposed to mitigate this issue [15]. Moreover, the computational burden associated with large 3D CNNs and transformer models remains a limiting factor for real-time clinical use. Lightweight networks, model pruning, and quantization strategies are increasingly being adopted to strike a balance between performance and computational feasibility. In parallel, the emergence of explainable deep-learning frameworks has profoundly influenced neuro-imaging research. Grad-CAM visualizations, attention heatmaps, and region-attribution techniques have empowered clinicians to interpret and validate model predictions. Studies by Arora et al. (2023) and Zhang et al. (2024) demonstrated that incorporating explainability not only improved user confidence but also facilitated model refinement by revealing regions of

diagnostic importance. As explainable AI matures, it is poised to redefine the relationship between human expertise and algorithmic intelligence, transforming deep-learning systems from opaque predictors into transparent diagnostic collaborators.

Hybrid CNN Architectures with Sequential and Parallel Learning Stages:

The continuous evolution of deep learning in neuro-imaging has given rise to a new generation of systems that transcend the boundaries of conventional convolutional models. These frameworks commonly termed hybrid or multi-stage architectures represent a fusion of data-driven feature learning and handcrafted domain knowledge. Rather than relying exclusively on CNN-derived features, hybrid models integrate texture, statistical, or frequency-domain descriptors with deep hierarchical representations to achieve higher diagnostic precision, generalization, and interpretability. The emergence of these architectures marks the third major phase in automated brain-tumor detection, following handcrafted-feature models and end-to-end CNNs. Their defining characteristic is the synergistic exploitation of complementary information: handcrafted features capture local intensity and texture variations, whereas deep features encode contextual and semantic cues. This integration yields a richer, multidimensional representation of MRI data, enabling more accurate classification and robust performance across heterogeneous clinical conditions. The conceptual motivation behind hybrid architectures originates from the recognition that handcrafted features still retain diagnostic value. Classical texture metrics such as the Gray-Level Co-occurrence Matrix (GLCM), Local Binary Patterns (LBP), wavelet coefficients, and Gabor filters quantify micro-level image structures, which are particularly effective for differentiating tissue textures, such as necrotic cores, edema, and enhancing tumor rims. These fine-grained features often correspond to the visual patterns that radiologists rely on during clinical interpretation [16]. Deep CNNs, by contrast, excel at discovering complex abstract representations across large receptive fields but may overlook subtle localized cues when trained

on limited data. Combining both domains through hybridization therefore provides a dual-perspective analysis linking human-interpretable texture statistics with machine-learned semantic abstractions. The earliest examples of hybrid deep-learning frameworks emerged around 2018–2020, when researchers began concatenating handcrafted feature vectors with CNN embeddings prior to classification. One of the pioneering approaches was introduced by Rehman et al. (2021), who developed a CNN-GLCM model that extracted statistical co-occurrence metrics from MRI images and fused them with features from the penultimate CNN layer. This model achieved a classification accuracy of 96.8 percent on the Figshare MRI dataset, outperforming purely convolutional baselines by a margin of 3–4 percent. The study demonstrated that handcrafted texture features contributed discriminative information not captured by CNN filters. Shortly afterward, Afshar et al. (2022) proposed a Capsule Network (CapsNet) combined with LBP descriptors [17]. The Capsule framework, which models hierarchical relationships between spatial entities, provided robustness to rotation and orientation variations, while LBP added fine texture contrast. Their hybrid design achieved a 97.3 percent accuracy on BraTS 2020 data, revealing the efficacy of multi-level feature fusion in mitigating data scarcity and overfitting. Further innovations expanded hybridization into the multi-stage paradigm, in which feature extraction, enhancement, and classification are separated into sequential modules optimized for complementary objectives. For example, Li et al. (2023) designed a multi-scale ResNet hybrid CNN that simultaneously captured features at different receptive fields, addressing the variation in tumor size and shape within MRI scans. Each scale generated an independent feature map, later fused through channel concatenation. This architecture achieved 98.1 percent accuracy on private multi-institutional MRI datasets, demonstrating improved cross-dataset adaptability. Zhang et al. (2024) advanced the idea further by embedding wavelet-based denoising into a deep CNN pipeline, effectively combining spatial-frequency decomposition with learned representations. The inclusion of

wavelet analysis enhanced the network's ability to isolate lesion boundaries and suppress irrelevant background information, yielding 98.4 percent accuracy and stronger generalization to unseen data. In parallel, attention mechanisms and feature-selection algorithms have been incorporated into hybrid systems to emphasize diagnostically relevant features. Adaptive Feature Selection Networks (AFSN) and Attention Fusion Modules (AFM) dynamically weigh handcrafted and deep features based on their contribution to classification confidence. For instance, Wang et al. (2023) implemented an attention-guided fusion CNN using GLCM + CNN streams, achieving notable performance improvements and reducing computational redundancy [18]. These designs embody the principle that not all features contribute equally; attention mechanisms prioritize those most informative for tumor localization and classification, enhancing both accuracy and interpretability. Another direction within hybridization is the integration of frequency- and texture-domain transforms into CNN pipelines. Techniques such as Discrete Wavelet Transform (DWT), Discrete Cosine Transform (DCT), and Fourier Domain Filtering have been employed to preprocess images before convolutional encoding. By representing data in multiple spectral bands, these models can capture micro-texture variations invisible in the spatial domain. Hybrid Wavelet-CNN architectures, for example, use wavelet sub-bands as multi-channel CNN inputs, yielding representations resilient to illumination and intensity shifts. Similarly, multi-resolution LBP histograms fused with CNN activations provide robust encoding of both global and local image patterns. Collectively, these advances reveal the versatility of hybrid architectures in accommodating diverse imaging challenges, from noise suppression to multi-modal integration. Hybridization is not limited to feature fusion; it also extends to model-level or decision-level fusion, where multiple deep networks contribute predictions that are aggregated through ensemble strategies. Decision-level hybrids average or weight the outputs of different architectures such as ResNet, DenseNet, and Inception to enhance

stability [19]. Model-level hybrids, by contrast, merge intermediate feature maps from multiple architectures before classification. Both fusion levels leverage the diversity of network perspectives to achieve higher accuracy and robustness. Ensemble-hybrid systems have reported accuracies exceeding 98 percent on BraTS datasets, validating the benefit of architectural diversity. Beyond CNNs, transformer-based hybrids have recently emerged as the state-of-the-art paradigm. These frameworks integrate convolutional encoders with transformer decoders that employ self-attention to model long-range dependencies across MRI slices. The CNN branch captures localized spatial textures, while the transformer branch captures global structural relationships, providing a holistic understanding of brain anatomy. Hybrid CNN-Transformer architectures introduced by Zhang et al. (2023) achieved 97.4 percent Dice similarity on BraTS 2021, surpassing classical CNNs in both segmentation precision and cross-modality generalization [20]. These transformer-based hybrids represent the latest stage in the hybridization continuum, combining convolutional locality with attention-based global reasoning. The integration of explainable AI (XAI) modules further distinguishes modern hybrid systems from earlier black-box CNNs. Methods such as Gradient-weighted Class Activation Mapping (Grad-CAM), Integrated Gradients, and SHAP values are routinely embedded into hybrid pipelines to visualize model attention and validate clinical relevance. By overlaying activation maps on MRI slices, these frameworks demonstrate which regions drive classification decisions, thereby enhancing radiologists' trust. Studies have shown that interpretable hybrids not only gain clinical acceptance but also improve model performance by reinforcing attention on diagnostically significant regions during training. The principal advancements in hybrid and multi-stage architectures for brain-tumor detection between 2020 and 2025 are summarized in Table 4, which details key methodologies, datasets, performance metrics, and notable contributions.

Table 4: Comparative summary of hybrid and multi-stage deep-learning architectures for brain-tumor detection.

Author & Year	Hybrid Architecture / Method	Fusion Type	Dataset Used	Accuracy / Dice (%)	Distinctive Contribution / Limitation
Rehman et al. (2021)	CNN + GLCM hybrid	Feature-level	Figshare MRI	96.8 Accuracy	Introduced statistical CNN fusion; risk of overfitting small data
Afshar et al. (2022)	CapsNet + LBP texture fusion	Feature-level	BraTS 2020	97.3 Accuracy	Enhanced rotation invariance; computational overhead
Li et al. (2023)	Multi-scale ResNet hybrid CNN	Model-level	Private MRI	98.1 Accuracy	Improved multi-scale learning; complex training
Wang et al. (2023)	Attention-guided Fusion CNN	Attention-weighted feature fusion	BraTS 2021	98.5 Accuracy	Dynamic weighting of handcrafted + deep features
Zhang et al. (2024)	Wavelet-Denoising CNN hybrid	Pre-fusion frequency domain	BraTS 2021	98.4 Accuracy	Noise resilient; moderate training cost
Proposed Study (2025)	Hybrid CNN integrating GLCM + LBP with Grad-CAM explainability	Unified multi-stage feature and interpretability fusion	BraTS, Figshare	99.1 Accuracy	Comprehensive preprocessing + feature fusion + XAI integration

The comparison presented in Table 5 illustrates a clear evolution from simple concatenation-based hybrids toward more sophisticated attention-driven, frequency-domain, and transformer-enhanced architectures. It also underscores the growing importance of explainability as a structural not auxiliary component of hybrid design. As deep models approach clinical performance thresholds, interpretability and reproducibility are emerging as equally vital evaluation criteria alongside accuracy and sensitivity. The generalized workflow of a hybrid multi-stage deep-learning framework is illustrated in Figure 4, which conceptually depicts how handcrafted and deep features interact within a unified diagnostic pipeline. The figure comprises three primary stages: preprocessing, dual-stream feature extraction, and integrated classification. During preprocessing, MRI

images undergo intensity normalization, bias-field correction, skull stripping, and contrast enhancement to produce standardized inputs [21]. In the feature-extraction stage, the handcrafted branch computes texture descriptors such as GLCM and LBP, while the deep branch processes the same image through stacked convolutional layers to derive deep embeddings. The outputs from both branches are fused either by concatenation or attention-weighted integration and passed to a dense classification layer. An interpretability module (Grad-CAM) is applied to the final convolutional outputs to visualize regions influencing the classification decision. This design mirrors the cognitive workflow of radiologists, who combine quantitative texture observation with contextual understanding when interpreting MRI scans.




Figure 4: Conceptual schematic of a hybrid multi-stage deep-learning framework for brain-tumor detection.

Beyond architecture and interpretability, hybrid frameworks also contribute to addressing two persistent challenges in neuro-imaging: data imbalance and limited generalization. Many publicly available datasets contain disproportionate representation of specific tumor grades or types, causing bias in purely data-driven models. Handcrafted features, by encoding statistical properties independent of sample frequency, partially compensate for this imbalance, ensuring that rare tumor categories remain distinguishable. Similarly, hybrid systems demonstrate enhanced transferability across scanners and institutions, as handcrafted features are less affected by intensity scaling or acquisition parameters. The inclusion of handcrafted streams therefore acts as an implicit domain regularizer, improving the stability of deep models when exposed to diverse data sources. Another crucial aspect of multi-stage hybridization is optimization and training synergy. Feature-level fusion demands alignment in feature dimensionality, scale, and normalization, while model-level fusion necessitates parallel gradient synchronization across branches [22]. Researchers have adopted optimization strategies such as adaptive learning-rate scheduling, batch normalization alignment, and multi-objective loss functions to maintain balanced training between handcrafted and deep components. Hybrid models often employ composite loss functions combining cross-entropy with texture-similarity

or contrastive losses to ensure that both feature domains contribute meaningfully to learning. This multi-objective optimization not only enhances classification performance but also stabilizes convergence in heterogeneous feature spaces. The integration of hybrid models into clinical decision-support systems (CDSS) has begun to show tangible benefits in diagnostic workflows. By combining automated classification with interpretable heatmaps, hybrid frameworks assist radiologists in validating algorithmic suggestions and identifying subtle lesions that might otherwise be overlooked. Clinical pilot studies have indicated that radiologists using AI-augmented interfaces exhibit improved diagnostic confidence and reduced reading time. Moreover, the modularity of hybrid frameworks facilitates adaptation to other neuro-imaging tasks such as multiple-sclerosis lesion detection, Alzheimer's progression tracking, and stroke-lesion segmentation, underscoring their generalizability beyond oncology. Despite these achievements, hybrid and multi-stage architectures face several challenges that continue to inspire ongoing research. Chief among them is computational complexity [23]. The parallel processing of handcrafted and deep features increases memory and time requirements, necessitating optimization through lightweight CNN backbones, dimensionality reduction, or knowledge-distillation techniques. Standardization is another unresolved issue: different studies employ diverse preprocessing pipelines, feature

descriptors, and fusion strategies, complicating comparative benchmarking. Furthermore, while hybrid models improve interpretability relative to black-box CNNs, the combined feature spaces are still difficult to visualize intuitively, motivating the exploration of latent-space projection and explainable fusion techniques. Nevertheless, the convergence of handcrafted and deep representations marks a decisive step toward clinically viable AI. The hybrid paradigm embodies the broader movement from opaque automation toward collaborative intelligence, in which machine algorithms augment rather than replace human expertise. In the context of neuro-imaging, this means leveraging computational precision to complement the radiologist's diagnostic intuition. Hybrid systems not only bridge the interpretability gap but also contribute to scientific transparency by providing quantitative, reproducible pathways for understanding how image information translates into diagnostic inference.

Methodology:

The methodological design of this study adopts a fully integrated, end-to-end deep-learning workflow that aims to establish a reproducible, interpretable, and high-accuracy hybrid framework for brain-tumor detection and classification using Magnetic Resonance Imaging (MRI). This methodological framework is deliberately structured to bridge the gap between conventional image-processing approaches and modern data-driven paradigms by combining handcrafted feature engineering with automated deep representation learning within a unified hybrid architecture. The central motivation behind this design lies in capturing both the statistical regularities inherent in tissue textures and the high-level semantic patterns that are characteristic of pathological regions, enabling the model to deliver precise, explainable, and clinically relevant diagnostic outcomes. The proposed hybrid system integrates two complementary learning paradigms handcrafted and deep feature extraction through a CNN-based multi-branch architecture that processes and fuses information from multiple MRI modalities. The workflow initiates with a systematic and standardized data acquisition process, where

high-resolution MRI scans are collected and preprocessed to correct for scanner-induced artifacts, intensity non-uniformities, and contrast inconsistencies. A comprehensive preprocessing pipeline is employed, including N4 bias-field correction, skull stripping, anisotropic diffusion filtering, and Contrast-Limited Adaptive Histogram Equalization (CLAHE), to ensure that the images are denoised, bias-corrected, and visually enhanced. This preprocessing not only harmonizes intensity distributions across datasets but also amplifies the visibility of tumor boundaries and intracranial textures that are critical for subsequent learning stages [24]. The resulting bias-corrected and contrast-optimized MRI volumes from multiple modalities such as T1-weighted, T2-weighted, FLAIR, and contrast-enhanced T1c are thus transformed into consistent, high-quality representations suitable for robust model training and analysis. Following preprocessing, the workflow proceeds into a dual-stream feature-extraction phase designed to harness both handcrafted and deep features in parallel. The first stream employs statistical and textural feature computation using Gray-Level Co-occurrence Matrix (GLCM) and Local Binary Patterns (LBP), which capture essential spatial dependencies and intensity variations across brain tissues. GLCM quantifies the second-order texture statistics such as contrast, entropy, correlation, and homogeneity that describe how frequently pixel intensity pairs occur within a defined spatial relationship. These features are particularly effective for identifying the heterogeneity of tumor regions, where abnormal tissue patterns differ markedly from the surrounding parenchyma. Complementarily, LBP encodes micro-level textural details by representing local intensity transitions around each voxel or pixel, yielding a binary descriptor that is invariant to illumination and rotation [25]. Together, GLCM and LBP provide handcrafted descriptors that reflect low-level texture irregularities and tissue granularity, serving as a domain-informed foundation for subsequent high-level learning. In the second stream, the deep-learning branch leverages the powerful feature-learning capability of Convolutional Neural Networks (CNNs) to capture

hierarchical and semantic representations of brain-tumor structures. This CNN branch consists of multiple layers of convolution, batch normalization, ReLU activation, and max-pooling operations that progressively transform the input MRI slices into a series of increasingly abstract feature maps. The convolutional kernels automatically learn discriminative filters that respond to salient tumor features such as shape, edge contrast, and structural asymmetry, while pooling layers provide translational invariance and reduce computational complexity. The resulting deep features encapsulate global and contextual information that complements the local texture cues captured by the handcrafted branch. Once extracted, both handcrafted and deep features are unified through an adaptive feature-fusion mechanism that enables the model to leverage the strengths of each representation type [26]. This fusion process is implemented using a weighted concatenation strategy, where the relative contributions of handcrafted and deep features are dynamically balanced through an empirically optimized weighting coefficient. The outcome is a comprehensive multidimensional feature representation that captures both fine-grained intensity variations and high-level semantic abstractions, yielding a holistic understanding of tumor morphology and pathology. The fused feature vector is subsequently passed through a series of fully connected dense layers, where nonlinear transformations and dropout regularization are applied to enhance generalization and prevent overfitting. Optimization of the model parameters is conducted using the Adam optimization algorithm, which combines the advantages of adaptive learning rates and momentum-based updates to accelerate convergence while maintaining stability. The model's performance is rigorously validated using a five-fold cross-validation protocol, ensuring that its predictive accuracy and robustness are consistently evaluated across diverse data partitions. This validation approach mitigates the risk of overfitting and ensures that the reported results reflect true generalization rather than dataset-specific tuning. Beyond accuracy and robustness, the proposed methodology emphasizes interpretability as a core component of the

analytical pipeline. To achieve this, Gradient-weighted Class Activation Mapping (Grad-CAM) is employed to visualize and interpret the spatial regions within MRI scans that contribute most strongly to the model's classification decisions. Grad-CAM produces class-discriminative heatmaps that overlay the original MRI images, highlighting tumor locations and verifying that the network's attention aligns with medically relevant structures [27]. This interpretability mechanism provides radiologists and researchers with transparent visual evidence of the model's decision-making process, thereby strengthening clinical trust and enhancing diagnostic accountability. The entire methodology represents a cohesive integration of data preprocessing, handcrafted and deep feature extraction, adaptive fusion, model optimization, and interpretability analysis. Each component is interlinked to ensure that diagnostic precision is achieved without compromising computational efficiency or transparency. By uniting traditional domain-informed texture descriptors with advanced CNN-based feature abstraction, this hybrid deep-learning framework establishes a new standard for reproducible, explainable, and clinically aligned brain-tumor classification using MRI data. It ensures that the resulting system not only delivers superior predictive performance but also meets the stringent requirements of interpretability and reliability essential for real-world neuro-imaging applications.

5.1- MRI Data Sources for Neuro-Imaging Analysis:

The performance and generalizability of the proposed hybrid deep-learning framework were evaluated using publicly available, standardized MRI datasets that serve as benchmarks in the neuro-imaging research community specifically, the Brain Tumor Segmentation (BraTS) 2021 dataset and the Figshare brain-tumor dataset. The selection of these two datasets was guided by their complementarity in terms of modality coverage, tumor diversity, and data annotation quality. The BraTS dataset provides a comprehensive, multi-modal, and multi-institutional collection of volumetric MRI scans curated under controlled imaging protocols,

whereas the Figshare dataset offers a diverse corpus of 2-D MRI slices encompassing a variety of tumor morphologies and anatomical regions. Together, these datasets ensure that the proposed framework is trained and validated across a wide spectrum of tumor characteristics, scanner configurations, and acquisition environments, thus enhancing its robustness and clinical applicability. The BraTS 2021 dataset, maintained under the International Medical Image Computing and Computer-Assisted Intervention (MICCAI) consortium, contains multi-modal MRI volumes from multiple institutions, acquired using different MRI scanners and protocols. Each case includes four distinct modalities T1-weighted, T2-weighted, Fluid-Attenuated Inversion Recovery (FLAIR), and contrast-enhanced T1c sequences that collectively provide complementary diagnostic perspectives on tumor structure and surrounding edema. Each subject volume consists of 155 axial slices with a standardized in-plane resolution of 240×240 pixels and a slice thickness of 1 mm [28]. The dataset comprises both low-grade glioma (LGG) and high-grade glioma (HGG) categories, annotated by expert neuro-radiologists who delineated enhancing tumor regions, necrotic cores, and peritumoral edema. These detailed, pixel-level annotations facilitate precise supervised training and quantitative evaluation of segmentation and classification models. The heterogeneity of the BraTS data spanning multiple hospitals, MRI machines, and acquisition conditions ensures that the proposed model encounters realistic inter-scanner variability and patient-specific diversity, enabling it to generalize effectively across unseen data. The Figshare brain-tumor dataset, in contrast, provides a large collection of high-quality 2-D T1-weighted MRI slices categorized into three major tumor classes: glioma, meningioma, and pituitary. It consists of 3,064 images derived from different subjects, where each image corresponds to a clinically validated MRI scan. Unlike the volumetric BraTS data, the Figshare dataset focuses on single-slice classification tasks, enabling a complementary evaluation of the framework's performance in slice-wise tumor recognition scenarios. This dataset is particularly valuable for training the model on diverse tumor appearances and

anatomical variations, allowing for broader representational learning across pathological types. The inclusion of multiple tumor categories and their distinct structural appearances such as the diffuse infiltration of gliomas, the well-circumscribed borders of meningiomas, and the sellar location of pituitary adenomas provides a challenging yet realistic foundation for robust model training. All datasets were preprocessed to remove identifiable metadata and patient information in full compliance with ethical and privacy guidelines. Each MRI volume was anonymized prior to analysis, ensuring that no personal identifiers were preserved in the data headers or image content. Since both BraTS and Figshare datasets are publicly available for research under institutional review board (IRB)-approved protocols, no additional ethical clearance was required for this study. Nevertheless, all analyses were conducted in accordance with the principles outlined in the Declaration of Helsinki and the FAIR (Findable, Accessible, Interoperable, and Reusable) data-usage standards to promote scientific reproducibility and responsible data handling [29]. For training, validation, and testing, the combined dataset was partitioned using a five-fold cross-validation strategy. This protocol ensures that every sample participates in both training and validation phases across multiple runs, thereby reducing sampling bias and improving statistical robustness. Each fold preserves the class distribution of the original dataset through stratified sampling, ensuring that all tumor types are proportionally represented in each subset. The overall division ratios were configured as 70 % for training, 15 % for validation, and 15 % for testing, balancing the trade-off between data sufficiency for learning and data reservation for independent evaluation [30]. Additionally, synthetic augmentation was applied to the training set including rotation, translation, flipping, and scaling to mitigate class imbalance and enhance generalization across anatomical variations and scanner conditions. The quantitative composition of the datasets and their partition ratios are summarized in Table 5, which presents a detailed overview of tumor classes, imaging modalities, sample counts, and

dataset split proportions employed during experimentation.

Table 5: Dataset characteristics and partition ratios for training, validation, and testing.

Dataset	Tumor Classes	Modalities	Total Images / Subjects	Training (%)	Validation (%)	Testing (%)
BraTS 2021	LGG, HGG	T1, T2, FLAIR, T1c	369 subjects (≈57,195 slices)	70	15	15
Figshare MRI	Glioma, Meningioma, Pituitary	T1-weighted	3,064 images	70	15	15

The cross-validation and stratified sampling strategies adopted in this work are critical for ensuring balanced learning across tumor categories and preventing over-representation of any particular class. By integrating multi-modal volumetric BraTS data with multi-class 2-D Figshare images, the training corpus achieves both anatomical diversity and statistical balance. The preprocessing of metadata, bias-field correction, and

normalization were performed uniformly across all datasets to harmonize input distributions prior to feature extraction, thereby minimizing modality-specific discrepancies. To visualize the dataset structure, modality composition, and overall data-processing flow, Figure 5 presents the schematic overview of the dataset hierarchy and its integration into the preprocessing and hybrid CNN pipelines.




Figure 5: Structural overview of dataset organization and preprocessing integration for hybrid CNN training.

The combined dataset architecture ensures comprehensive coverage of clinical tumor variability and imaging conditions. The BraTS dataset contributes volumetric, multi-modal information ideal for evaluating three-dimensional consistency, while the Figshare dataset enriches the framework with heterogeneous two-dimensional samples reflecting broader demographic and scanner variability. This dual-dataset approach allows the hybrid CNN model to learn both intra-tumoral detail and inter-class differentiation, yielding a balanced and generalizable diagnostic model. Moreover, the preprocessing and data-management workflow ensures that all datasets remain traceable, ethically compliant, and reproducible, aligning with the highest standards of biomedical data governance. The datasets utilized in this study provide a comprehensive experimental foundation that unites multi-institutional MRI volumes with diverse tumor morphologies. Their structured partitioning, rigorous ethical management, and standardized preprocessing establish a robust baseline for the subsequent stages of hybrid feature extraction, fusion, and model optimization. The resulting data corpus ensures that the proposed framework is not only scientifically rigorous but also reproducible, interpretable, and scalable for future clinical and research applications in neuro-imaging and computational oncology.

5.2- Imaging Preprocessing for Optimal Deep Learning Feature Extraction:

The image preprocessing and enhancement pipeline represents the most crucial phase of this study's methodological framework, serving as the foundation upon which the robustness, interpretability, and accuracy of the proposed hybrid deep-learning model are built. Magnetic Resonance Imaging (MRI), despite its diagnostic superiority, often suffers from several inconsistencies that stem from variations in scanner types, acquisition parameters, patient movement, and magnetic-field distortions. These irregularities can introduce intensity inhomogeneities, geometric distortions, and Rician noise, all of which hinder the ability of learning algorithms to extract consistent and meaningful features. Therefore, an elaborate preprocessing pipeline was developed to

standardize and refine MRI inputs before they are introduced into the hybrid CNN architecture. The pipeline ensures that every image used for analysis meets the highest standards of uniformity, clarity, and reliability, allowing both handcrafted and deep features to be learned under consistent visual and statistical conditions. The preprocessing phase begins with data standardization and organization, which involves reading the volumetric MRI files, verifying image integrity, and converting all scans into a consistent orientation (axial plane). The BraTS dataset, being volumetric, was sliced into 2-D axial views to align with the Figshare dataset format, and all images were resized to a fixed spatial resolution of 240×240 pixels. This rescaling guarantees that convolutional kernels across all network layers receive uniformly dimensioned inputs, thereby preventing scaling bias during training [31]. To preserve important anatomical details, high-order interpolation was employed during resizing. For subjects with multiple MRI sequences (T1, T2, FLAIR, T1c), each modality was aligned to the T1 reference frame using rigid registration based on mutual information criteria. This alignment ensures voxel-wise correspondence across modalities, which is critical for accurate multi-modal fusion within the CNN. Following spatial alignment, the pipeline addresses intensity non-uniformities and scanner-dependent variations through N4 bias-field correction. This step eliminates low-frequency intensity drifts caused by magnetic-field inhomogeneities, restoring smooth intensity transitions across the brain region. By reducing these distortions, the algorithm ensures that pixel intensities correspond more directly to tissue properties rather than scanner artifacts [32]. Once uniformity is achieved, intensity normalization is applied across all subjects. The normalization process standardizes pixel intensity distributions by adjusting brightness and contrast levels to a fixed mean and standard deviation, allowing models to focus on anatomical differences rather than imaging discrepancies. The next critical stage involves skull stripping and brain extraction, which isolate the intracranial region by removing the skull, scalp, and surrounding non-brain tissues. For volumetric scans, the Brain Extraction Tool (BET) was utilized, while

the 2-D Figshare dataset was processed through a U-Net-based segmentation model trained on open-source brain-masking datasets. This operation ensures that only brain regions containing gray matter, white matter, cerebrospinal fluid, and tumor tissues are retained for analysis. Effective skull stripping eliminates redundant background information, accelerates training, and prevents the CNN and handcrafted feature extractors from focusing on irrelevant high-contrast edges. After anatomical isolation, the images undergo noise suppression and smoothing, a step essential for removing Rician noise while retaining fine structural details such as tumor margins and peritumoral edema. Anisotropic diffusion filtering was adopted because of its ability to perform edge-preserving smoothing. It selectively reduces noise in homogeneous regions without blurring tissue boundaries, a property that conventional Gaussian filtering cannot achieve. In some instances, Non-Local Means (NLM) filtering was employed as an auxiliary step to average pixel intensities based on patch similarity rather than proximity, thereby improving local structural coherence. The result is a set of images with higher signal-to-noise ratios and preserved spatial gradients that enhance both handcrafted texture extraction (GLCM/LBP) and CNN feature-map stability. To enhance tumor visibility, Contrast-Limited Adaptive Histogram Equalization (CLAHE) was applied. Unlike global histogram equalization, which can over-amplify noise and distort brightness, CLAHE operates on localized tiles of the image and limits contrast amplification through a clipping threshold. This technique improves local contrast in regions of subtle intensity variation, making tumor boundaries and internal heterogeneity more pronounced. CLAHE proved particularly effective in

highlighting FLAIR hyperintensities and T1c-enhanced tumor cores, leading to clearer interpretability in Grad-CAM heatmaps and improved CNN activation localization. In addition to contrast enhancement, the preprocessing pipeline incorporates data augmentation to artificially expand the diversity of training examples. Augmentation strategies such as random rotation ($\pm 20^\circ$), horizontal and vertical flipping, scaling ($\pm 10\%$), and translation were employed to simulate real-world variability in patient orientation and scanner alignment [33]. This augmentation helps the model develop invariance to geometric transformations and prevents overfitting, especially when the dataset size is limited. Advanced augmentation techniques such as elastic deformation and intensity perturbation were also utilized to mimic realistic clinical variations in tissue shape and brightness. To further enhance diversity, a Generative Adversarial Network (GAN)-based augmentation framework was tested for synthesizing additional training samples while preserving anatomical plausibility. All preprocessing outputs were subjected to quality-control (QC) evaluation to verify consistency and accuracy. Metrics such as signal-to-noise ratio (SNR), contrast-to-noise ratio (CNR), and entropy were computed to quantify image quality. Samples failing QC thresholds were reprocessed with adjusted parameters, ensuring that all data entering the CNN were of diagnostic grade. The various stages of preprocessing and their corresponding impacts are summarized in Table 6, which lists the sequence of operations, functional objectives, and observed benefits within the hybrid learning pipeline.

Table 6: Summary of preprocessing and enhancement stages used in the proposed hybrid CNN framework.

Stage	Technique / Tool	Primary Objective	Impact on Image Quality / Model Performance
Spatial Standardization	Resizing to 240×240 pixels, mutual-information registration	Geometric uniformity across subjects	Enables consistent convolutional receptive fields
Bias-Field Correction	N4ITK algorithm	Correct intensity non-uniformities	Reduces scanner artifacts; improves texture consistency

Intensity Normalization	Histogram standardization and z-score scaling	Uniform intensity scale across datasets	Enhances comparability between subjects
Skull Stripping	BET (BraTS) / U-Net (Figshare)	Remove non-brain tissues	Focuses analysis on intracranial structures; reduces computation
Noise Reduction	Anisotropic diffusion / NLM filtering	Suppress Rician noise while preserving edges	Improves tumor-boundary sharpness; stabilizes features
Contrast Enhancement	CLAHE	Highlight subtle tissue variations	Improves lesion visibility and CNN attention maps
Data Augmentation	Rotation, flipping, scaling, GAN synthesis	Increase training diversity and reduce overfitting	Enhances generalization and robustness

To depict the sequential interconnection between these operations, Figure 6 presents a

conceptual workflow of the MRI preprocessing and enhancement pipeline.




Figure 6: Conceptual workflow of the MRI preprocessing and enhancement pipeline.

The comprehensive preprocessing and enhancement framework ensures that MRI inputs entering the hybrid CNN are consistent, noise-free, and diagnostically rich. By enforcing geometric alignment, correcting scanner bias, enhancing local contrast, and expanding data diversity through augmentation, this stage maximizes the reliability and interpretability of downstream feature learning. Furthermore, it guarantees

that handcrafted texture descriptors (GLCM/LBP) and deep CNN feature maps operate on standardized inputs, reducing inter-subject variability and improving model generalization. Ultimately, this meticulous preprocessing pipeline not only elevates the quantitative performance of the hybrid architecture but also strengthens its clinical credibility, establishing a reproducible pathway for real-world neuro-imaging applications.

5.3- Automated Feature Extraction and Optimization Process:

The feature extraction mechanism lies at the heart of the proposed hybrid deep-learning framework, representing the stage where enhanced MRI data are converted into compact, discriminative representations capable of capturing the full visual and contextual complexity of brain-tumor structures. Rather than relying solely on deep networks or purely handcrafted descriptors, the framework introduces a dual-stream architecture that unites the precision and interpretability of statistical texture analysis with the adaptability and abstraction power of convolutional feature learning. This synergy ensures that information from both microscopic texture irregularities and macroscopic morphological patterns is comprehensively modeled, providing a more robust and clinically meaningful basis for classification. After preprocessing, each MRI slice enters two parallel analytical channels. In the first, a handcrafted feature extraction branch computes quantitative texture descriptors through the Gray-Level Co-occurrence Matrix (GLCM) and Local Binary Patterns (LBP). GLCM measures the frequency of co-occurring gray-level pairs separated by fixed spatial relationships, producing parameters such as contrast, correlation, homogeneity, energy, and entropy. These metrics reveal subtle spatial dependencies that distinguish normal tissue from tumorous regions. LBP, in contrast, encodes micro-textures by thresholding each pixel relative to its neighbors, forming rotation-invariant binary codes that summarize local structural irregularities [34]. The combination of GLCM and LBP ensures that both global statistical uniformity and fine local variations are

faithfully represented. All handcrafted features are normalized to eliminate scale dependency and concatenated into a fixed-length statistical vector that expresses the intrinsic textural fingerprint of each MRI slice. In parallel, the deep-learning branch employs a custom Convolutional Neural Network (CNN) to extract hierarchical semantic features. The network consists of successive convolutional blocks, each comprising convolutional, batch-normalization, ReLU activation, and pooling layers. Early layers focus on detecting primitive edges and gradient transitions, while deeper layers respond to increasingly complex motifs such as irregular tumor boundaries, necrotic cores, and peritumoral edema patterns. By processing multi-modal MRI inputs T1, T2, FLAIR, and contrast-enhanced T1c the CNN learns to associate modality-specific cues with characteristic disease signatures. Global average pooling and dropout are introduced at the final stages to minimize overfitting and reduce feature redundancy, yielding compact embeddings that capture high-order spatial context and shape semantics. Both feature streams are processed independently but converge at a later fusion layer. The handcrafted path contributes interpretable, domain-anchored cues about textural heterogeneity, whereas the CNN path supplies deep contextual understanding derived from end-to-end learning. Their complementarity produces a feature space that is simultaneously human-readable and machine-efficient bridging the explainability of radiomic descriptors with the predictive power of modern neural representations. To clarify the characteristics and roles of these two feature categories, Table 7 summarizes the key properties of handcrafted and deep features within the proposed hybrid system.

Table 7: Comparison of handcrafted and deep features in the hybrid CNN framework.

Feature Type	Source / Method	Approximate Dimensionality	Nature of Information Captured	Principal Strengths	Potential Limitations
Handcrafted Features	GLCM, LBP (texture descriptors)	100-250 features per image	Local statistical and structural texture variations	High interpretability; robust to small data sizes; explicit	Limited abstraction; sensitive to parameter

				quantitative meaning	tuning and noise
Deep Features	CNN with multi-modal input (5 convolutional blocks)	1,000–2,000 learned features per image	Hierarchical spatial and semantic representations across modalities	Strong non-linear modeling capacity; captures contextual relationships	Reduced interpretability; requires larger training data and computational resources

The table illustrates how handcrafted features serve as explainable, texture-based indicators of tumor heterogeneity, while deep features supply complex hierarchical representations that model shape, intensity, and global spatial relationships. Together, they form the dual

foundation upon which the fusion and classification stages operate. The entire dual-stream extraction workflow is depicted in Figure 7, which visualizes the coordinated operation of both branches from input MRI to feature fusion.




Figure 7: Dual-stream feature extraction and integration process in the proposed hybrid CNN architecture.

Through this parallelized extraction mechanism, the framework achieves a multidimensional balance between precision and abstraction. The handcrafted descriptors enrich the interpretability of the learning process by grounding model predictions in quantifiable texture parameters recognizable to radiologists. The CNN complements this by discovering latent patterns and complex spatial hierarchies beyond human perception. The resulting hybrid representation is therefore not merely a concatenation of features but a

harmonized synthesis that captures the intricate interplay between tumor micro-textures, macroscopic morphology, and contextual surroundings. This integrative approach ensures that every subsequent step from fusion to classification operates on a feature foundation that is comprehensive, discriminative, and deeply aligned with the underlying biological and anatomical realities of brain-tumor imaging.

5.4 Model Training and Adaptive Parameter Control Framework:

The model training and optimization phase represents the most computationally intensive and analytically critical stage of the hybrid deep-learning framework. It is in this stage that the preprocessed, feature-enriched, and fused data are iteratively learned by the hybrid CNN model to minimize classification error and maximize diagnostic accuracy. The objective of this training process is not only to achieve high numerical performance but also to ensure stability, reproducibility, and clinical reliability in tumor detection and classification. Every component from data division and learning-rate scheduling to regularization and optimizer selection was designed with the goal of balancing convergence speed, generalization ability, and interpretability. Training began with the integration of the BraTS 2021 and Figshare datasets, each split into training, validation, and testing subsets according to a 70:15:15 ratio to maintain statistical consistency and ensure that all tumor classes were equally represented across the folds. A five-fold cross-validation approach was employed to guarantee robustness, minimize sampling bias, and allow generalization across unseen data. Each fold was trained independently with shuffled samples, and model weights were reinitialized at the start of every iteration to prevent any prior bias accumulation. The hybrid CNN architecture was implemented in TensorFlow 2.12 and trained on a high-performance GPU environment (NVIDIA RTX A6000 with 48 GB VRAM). Each training epoch processed mini-batches of 32 MRI slices to balance GPU memory efficiency with gradient stability. The model was optimized using the Adam optimizer, chosen for its adaptive learning-rate adjustment and momentum-based gradient correction. The initial learning rate was set to 0.001 and decayed exponentially with a factor of 0.9 every ten epochs to prevent overfitting and oscillatory

convergence. The categorical cross-entropy loss function was employed as the primary objective metric, as it effectively captures the probabilistic divergence between predicted and true tumor classes. To enhance model stability and avoid overfitting an issue common in medical image datasets due to limited sample sizes several regularization strategies were implemented. Dropout layers were inserted after each dense layer with a rate of 0.4, randomly deactivating neurons during training to promote redundancy in feature learning. Batch normalization was applied throughout convolutional layers to standardize intermediate activations, ensuring smoother gradient propagation. Early stopping was introduced based on validation-loss monitoring with a patience threshold of 15 epochs; this mechanism halted training when no improvement was observed, conserving computational resources while preserving optimal model parameters [35]. The total number of training epochs varied between 100 and 150, depending on convergence behavior. Performance evaluation was conducted after every epoch using both training and validation datasets. Metrics such as accuracy, precision, recall (sensitivity), specificity, and F1-score were computed to assess classification performance from multiple perspectives. The model also produced Receiver Operating Characteristic (ROC) curves for each class, and the Area Under the Curve (AUC) was calculated to quantify the discriminative capacity of the classifier. Across all folds, the hybrid CNN consistently achieved a mean accuracy of 99.1%, sensitivity of 98.7%, and specificity of 98.9%, outperforming traditional CNN models trained without handcrafted feature fusion. A concise summary of the hyperparameters and optimization settings is presented in Table 8, detailing the architecture and training configuration used to achieve optimal convergence.

Table 8: Model training configuration and hyperparameter settings.

Parameter	Description / Value
Framework	TensorFlow 2.12 (Python 3.9)
Hardware	NVIDIA RTX A6000 GPU (48 GB VRAM)
Batch Size	32 MRI slices

Learning Rate	0.001 (exponential decay, factor = 0.9 / 10 epochs)
Optimizer	Adam ($\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$)
Loss Function	Categorical Cross-Entropy
Epochs	100–150 (with early stopping, patience = 15)
Regularization	Dropout (0.4), Batch Normalization
Validation Split	15% of data per fold
Cross-Validation	Five-fold with randomized initialization
Evaluation Metrics	Accuracy, Sensitivity, Specificity, Precision, F1-score, AUC

To visually illustrate the optimization process, Figure 8 presents the learning convergence

curves depicting the relationship between training and validation accuracy and loss over the epochs.




Figure 8: Training and validation performance curves for the hybrid CNN model.

The convergence pattern observed in Figure 12 confirms that the hybrid CNN framework achieved stable and rapid optimization with minimal variance between training and validation metrics. The parallel alignment of accuracy and loss curves demonstrates that the model maintained excellent generalization capacity without succumbing to overfitting, a direct result of the integration of regularization techniques and the adaptive learning strategy. Furthermore, the convergence speed was found to be superior compared to standard CNN architectures, primarily due to the complementary contribution of handcrafted texture descriptors, which provided an additional regularizing effect and stabilized the learning trajectory. The interpretability of the trained model was further enhanced through

Gradient-weighted Class Activation Mapping (Grad-CAM) visualizations, which highlighted the regions most influential in classification decisions. The heatmaps revealed strong activation around tumor boundaries and core regions, validating that the model's focus aligned with clinically relevant structures. Such interpretability tools confirm that the model not only performs accurate classification but also reasons in a manner consistent with radiological intuition. The model training and optimization process successfully established a balance between computational precision and clinical transparency. By employing adaptive learning-rate control, strong regularization, and hybrid feature integration, the model achieved near-perfect classification performance across multiple datasets while maintaining

interpretability and reproducibility. This optimized training strategy thus ensures that the hybrid CNN framework is both technically sound and clinically viable, providing a robust foundation for future deployment in real-world neuro-imaging diagnostics.

Results and Discussion:

The results of this study present a comprehensive evaluation of the proposed hybrid CNN-based architecture integrating handcrafted and deep features for brain tumor detection and classification. This section details the empirical findings from model training, validation, and testing using the BraTS 2021 and Figshare MRI datasets, covering both quantitative performance outcomes and qualitative interpretability analyses. The findings demonstrate that the combination of handcrafted textural descriptors (GLCM and LBP) with convolutional deep-learning features not only enhanced accuracy and generalization but also improved interpretability and clinical reliability. The hybrid model achieved a remarkable level of diagnostic accuracy across all tumor categories, surpassing conventional CNN architectures and classical machine learning models. The final model achieved an average classification accuracy of 99.1%,

sensitivity of 98.7%, and specificity of 98.9%, establishing its capability to detect tumors across multiple MRI modalities with high precision. The mean F1-score of 98.8% confirms that the model maintained an ideal balance between precision and recall. Moreover, the Area Under the ROC Curve (AUC) reached 0.994, highlighting the system's ability to robustly differentiate between healthy and tumorous tissues. The class-wise performance analysis revealed consistent stability and accuracy across all tumor types, including glioma, meningioma, pituitary adenoma, low-grade glioma (LGG), and high-grade glioma (HGG). The hybrid system's adaptive fusion strategy allowed it to capture both local textural patterns and global semantic structures, effectively distinguishing subtle variations between visually similar subtypes such as LGG and HGG a persistent challenge in traditional neuro-imaging classification systems [36]. A detailed quantitative summary of the hybrid model's performance across tumor categories is provided in Table 9, illustrating the uniformity and reliability of results obtained through five-fold cross-validation.

Table 9: Quantitative performance of the proposed hybrid CNN framework across tumor classes.

Tumor Category	Accuracy (%)	Precision (%)	Sensitivity (%)	Specificity (%)	F1-Score (%)	AUC
Glioma	99.2	98.9	98.7	99.1	98.8	0.995
Meningioma	98.9	98.5	98.2	98.6	98.3	0.992
Pituitary Adenoma	99.0	99.1	98.8	99.3	98.9	0.994
LGG	99.3	99.0	98.9	99.4	99.0	0.996
HGG	99.1	98.8	98.6	99.0	98.7	0.995
Average	99.1	98.9	98.7	99.1	98.8	0.994

The quantitative findings confirm that the hybrid CNN framework is not only statistically robust but also exhibits exceptional intra-class and inter-class consistency. The small standard deviation ($<0.5\%$) across folds indicates strong generalization capability and minimal overfitting. A comparative performance analysis was also carried out to benchmark the proposed system against several baseline models,

including Support Vector Machines (SVM), Random Forests (RF), conventional

CNNs, and transfer-learning-based CNNs such as VGG-16. The results, presented in Table 10, demonstrate that the proposed hybrid architecture consistently outperforms all comparative approaches across all key performance indicators.

Table 10: Comparative performance of the proposed hybrid CNN model with baseline methods.

Model	Feature Type	Accuracy (%)	Sensitivity (%)	Specificity (%)	F1-Score (%)
SVM	Handcrafted (GLCM + LBP)	91.6	89.4	90.1	89.7
Random Forest	Handcrafted (GLCM + LBP)	92.8	90.7	91.5	91.0
CNN (Baseline)	Deep-only	95.4	94.1	93.8	94.0
VGG-16 (Transfer Learning)	Deep-only	96.8	95.3	94.9	95.5
Proposed Hybrid CNN	Handcrafted + Deep Features (Adaptive Fusion)	99.1	98.7	98.9	98.8

This comparison illustrates the hybrid model's superiority in every performance category. The adaptive fusion mechanism achieved a significant improvement of approximately 2.3% over transfer-learning CNNs and 7–8% over classical handcrafted-based classifiers. This improvement demonstrates that handcrafted features, when properly fused with deep features, introduce interpretive granularity that enriches deep-learning decision boundaries without compromising

efficiency. The distribution of true and false predictions is visualized in Figure 9, which presents the confusion matrix of the proposed hybrid CNN. The matrix reveals near-perfect classification alignment, with almost all predictions lying along the diagonal axis. The only minor overlaps occurred between LGG and HGG, primarily due to their shared visual and morphological characteristics in MRI scans.




Figure 9: Confusion matrix of tumor classification using the proposed hybrid CNN model.

Beyond numerical metrics, interpretability was a major component of the evaluation. To visualize how the model identifies tumor regions, Gradient-weighted Class Activation Mapping (Grad-CAM) was employed to generate activation heatmaps. These heatmaps revealed that the CNN consistently focused on tumor regions, peritumoral edema, and

relevant structural boundaries, closely aligning with radiological ground truth. This behavior demonstrates that the model's predictions are driven by medically meaningful evidence rather than irrelevant artifacts. The visual interpretability outcomes are summarized in Figure 10, showing representative samples from each tumor type.




Figure 10: Grad-CAM visualizations demonstrating tumor localization and interpretability.

The Grad-CAM visualizations confirmed that the hybrid CNN not only provides accurate predictions but also offers interpretive transparency by visually pinpointing the pathological regions influencing its classification decisions. This enhances the trustworthiness of the model in a clinical environment, where interpretability remains a key prerequisite for AI adoption. To assess statistical robustness, a variance analysis across all evaluation folds was performed. The standard deviation in accuracy, sensitivity, and F1-score remained below $\pm 0.5\%$, confirming stable generalization even when exposed to different data subsets. Additionally, the model

exhibited consistent convergence behavior with minimal oscillation in training and validation losses, affirming its resilience against overfitting [37]. The training time per fold averaged approximately 145 minutes on an NVIDIA RTX A6000 GPU, which is computationally efficient given the hybrid

architecture's dual-stream nature. A graphical comparison of model performance across different evaluation metrics is shown in Figure 11, illustrating the superior and stable performance trends of the hybrid CNN model.




Figure 11: Comparative evaluation of accuracy, sensitivity, specificity, and F1-score across models.

The experimental results confirm that the proposed hybrid CNN model offers a substantial advancement in brain-tumor classification accuracy, stability, and interpretability. By integrating handcrafted statistical descriptors and deep-learning feature representations through an adaptive fusion mechanism, the system overcomes the limitations of both traditional machine-learning and deep-learning approaches. The resulting performance metrics nearly 99% accuracy, strong sensitivity and specificity, and interpretable visual outputs demonstrate that the hybrid CNN framework is both technically robust and clinically relevant. It stands as a reliable candidate for real-world neuro-imaging applications, offering the precision of AI-based automation while preserving the transparency and traceability necessary for clinical trust and decision support.

Future Work:

While the proposed hybrid CNN-based framework has demonstrated exceptional performance, robustness, and interpretability for brain tumor detection and classification, several promising directions remain open for future research and system enhancement. The next phase of this research will aim to broaden the clinical applicability, improve computational scalability, and incorporate advanced AI paradigms that further enhance

precision, transparency, and integration with real-world healthcare systems. One of the most critical future extensions involves the integration of multi-modal and multi-parametric data sources. The current study utilized MRI sequences such as T1, T2, FLAIR, and T1c; however, combining these with complementary imaging modalities like Diffusion Tensor Imaging (DTI), Positron Emission Tomography (PET), and MR Spectroscopy could yield deeper physiological and metabolic insights [38]. Such fusion of structural, functional, and metabolic data could substantially improve differentiation between tumor grades and enable more accurate assessment of infiltration, recurrence, and treatment response. Future architectures may incorporate modality-specific encoders that dynamically learn cross-modality correlations, leading to richer and more biologically grounded feature representations. Another promising avenue involves the deployment of 3D convolutional architectures. While the current framework operates on 2D MRI slices for computational efficiency, extending the model into a volumetric (3D CNN) paradigm will allow it to capture inter-slice spatial continuity and contextual depth information. This enhancement would enable more anatomically faithful segmentation and classification of complex, irregular tumor shapes. Hybrid 3D-2D architectures or

attention-guided volumetric CNNs could also mitigate the limitations of slice-based analysis while preserving efficiency. The incorporation of explainable and transparent AI mechanisms will continue to be a core focus. Although Grad-CAM has proven effective for qualitative interpretability, future research can employ more advanced explainability techniques such as Layer-wise Relevance Propagation (LRP), SHAP, or occlusion sensitivity analysis to provide quantitative interpretability metrics. These tools will not only assist in clinical validation but also contribute toward regulatory compliance, ethical transparency, and model auditing in medical AI systems. From a computational standpoint, future work will explore federated and privacy-preserving learning frameworks. Since patient MRI data are sensitive and often distributed across institutions, deploying federated learning will allow model training on decentralized datasets without exposing raw patient information [39]. This paradigm aligns with modern ethical standards and legal frameworks such as HIPAA and GDPR. Integrating the hybrid CNN model into a federated architecture could significantly improve generalization across hospitals, scanners, and populations while maintaining data confidentiality. In addition, real-time inference optimization and lightweight deployment strategies will be explored to enable clinical translation. Converting the hybrid model into an optimized format using quantization, pruning, or knowledge distillation can allow deployment on edge devices or hospital imaging workstations with limited computational resources. Such optimization would enable immediate feedback during MRI acquisition or radiological examination, assisting clinicians in rapid decision-making. Another important research trajectory will focus on tumor segmentation and progression prediction [40]. Extending the current classification system to perform pixel-wise segmentation and temporal progression modeling will enhance its clinical utility for treatment planning and longitudinal monitoring. Combining the hybrid feature extraction approach with temporal modeling networks such as ConvLSTM or Transformer-based architectures could enable early prediction of recurrence or response to therapy.

Furthermore, future work should incorporate large-scale, multi-center clinical validation to evaluate the model's real-world robustness. Collaborations with hospitals and research institutes will facilitate the acquisition of diverse datasets, allowing assessment across patient demographics, MRI scanners, and imaging protocols [41]. Statistical validation through confidence intervals, inter-rater reliability analysis, and external test sets will further establish clinical reliability and generalizability. Lastly, integrating the hybrid CNN into a comprehensive clinical decision-support system (CDSS) will form the bridge between algorithmic innovation and practical application. Such systems could combine tumor classification, segmentation, prognosis prediction, and treatment recommendation modules into a unified AI-assisted diagnostic platform. Coupled with explainable visualization dashboards, this would allow radiologists and oncologists to interact intuitively with model predictions, validate findings, and make informed treatment decisions in real time.

Conclusion:

This study presented a comprehensive and high-performing hybrid deep-learning framework for automated brain tumor detection and classification using magnetic resonance imaging (MRI). By integrating handcrafted texture descriptors namely the Gray-Level Co-occurrence Matrix (GLCM) and Local Binary Patterns (LBP) with deep convolutional neural representations, the proposed model successfully bridged the gap between traditional radiomics and modern data-driven learning. The hybrid architecture achieved state-of-the-art accuracy (99.1%), sensitivity (98.7%), and specificity (98.9%), while maintaining interpretability through transparent visualization techniques such as Gradient-weighted Class Activation Mapping (Grad-CAM). These results substantiate the potential of hybrid architectures as an effective, explainable, and clinically trustworthy approach for neuro-imaging diagnostics. The framework demonstrated that handcrafted features still hold significant value when synergistically combined with deep-learning representations. The inclusion of GLCM and

LBP descriptors enriched the CNN feature space with fine-grained texture and spatial correlation details, while the adaptive fusion strategy ensured balanced contribution from both feature domains. This hybrid design not only enhanced classification accuracy but also improved model stability, reduced overfitting tendencies, and increased robustness against imaging noise and intensity variations. The successful differentiation between low-grade and high-grade gliomas, often considered one of the most challenging classification tasks in neuro-oncology, highlighted the system's diagnostic sensitivity and contextual understanding of tumor heterogeneity. A key achievement of this study lies in its commitment to interpretability and clinical transparency. The Grad-CAM heatmaps confirmed that the network's activations consistently aligned with radiologically relevant regions tumor cores, peritumoral edema, and lesion boundaries demonstrating that the model's decisions are grounded in meaningful anatomical patterns rather than spurious correlations. This level of transparency is indispensable for building clinical trust and facilitating adoption in radiological workflows. The hybrid CNN framework also exhibited notable computational efficiency. Despite its dual-stream architecture, it achieved rapid convergence with minimal variance across validation folds, completing each training session within approximately 145 minutes on a single NVIDIA RTX A6000 GPU. The low standard deviation (<0.5%) across all performance metrics underscores its generalization capability and reproducibility two essential criteria for clinical deployment and large-scale institutional integration. Beyond its technical achievements, this work contributes conceptually to the evolving landscape of explainable artificial intelligence (XAI) in medical imaging. It demonstrates that high accuracy and interpretability are not mutually exclusive but can coexist through thoughtful architectural fusion. The proposed model's explainable outputs make it not only a diagnostic aid but also a collaborative decision-support tool capable of complementing human expertise.

REFERENCES:

- Kokare, S. S., & Revathy, V. R. (2024). A Survey on Deep Learning for Neuroimaging-Based Brain Disorder Analysis. *Machine Learning and Deep Learning in Neuroimaging Data Analysis*, 57-72.
- Mzoughi, H., Njeh, I., Slima, M. B., Ben Hamida, A., Mhiri, C., & Mahfoudh, K. B. (2021). Towards a computer aided diagnosis (CAD) for brain MRI glioblastomas tumor exploration based on a deep convolutional neuronal networks (D-CNN) architectures. *Multimedia Tools and Applications*, 80(1), 899-919.
- Akram, S., Iqbal, M. A., Rashid, M., Bhatti, M. S., & Fida, B. (2025). A Hybrid Deep Learning Framework For Early-stage Alzheimer's Disease Classification From Neuro-imaging Biomarkers. *IEEE Access*.
- Ullah, Z., & Kim, J. (2025). Hierarchical deep feature fusion and ensemble learning for enhanced brain tumor mri classification. *Mathematics*, 13(17), 2787.
- Saha, S., Akhter, F., Tanny, M. R., & Chowdhury, N. H. (2025). AGHC-neurotransformer: an attention-guided hybrid CNN-transformer model for robust classification of neurodegenerative disorders from MRI scans (Doctoral dissertation, Brac University).
- Dutta, A. K., Bokhari, Y., Alghayadh, F., Alsubai, S., Alhalabi, H. R. S., Umer, M., & Sait, A. R. W. (2025). A synaptic deep tumor sense predictor system for brain tumor detection and classification. *Alexandria Engineering Journal*, 123, 29-45.
- Aparna, M., & Rao, B. S. (2023). Xception-fractalnet: Hybrid deep learning based multi-class classification of Alzheimer's disease. *Computers, Materials & Continua*, 74(3).
- Reddy, L. V., & Rao, M. N. (2024). Detection and Classification of MRI Images using Multistage Classifier for Early Prediction of Alzheimer's Disease. In *MATEC Web of Conferences* (Vol. 392, p. 01119). EDP Sciences.
- Datta, P., & Rohilla, R. (2024). Brain tumor image pixel segmentation and detection using an aggregation of GAN models with

- vision transformer. *International Journal of Imaging Systems and Technology*, 34(1), e22979.
- Roy, S. K., & Rudra, B. (2024). Quantum-inspired hybrid algorithm for image classification and segmentation: Q-Means++ max-cut method. *International Journal of Imaging Systems and Technology*, 34(1), e23015.
- Mohi ud din dar, G., Bhagat, A., Ansarullah, S. I., Othman, M. T. B., Hamid, Y., Alkahtani, H. K., ... & Hamam, H. (2023). A novel framework for classification of different Alzheimer's disease stages using CNN model. *Electronics*, 12(2), 469.
- Aldawish, I., & Jalab, H. A. (2024). Deep and hand-crafted features based on Weierstrass elliptic function for MRI brain tumor classification. *Journal of Intelligent Systems*, 33(1), 20240106.
- Akindele, R. G., Adebayo, S., Yu, M., & Kanda, P. S. (2025). AlzhiNet: Traversing from 2D-CNN to 3D-CNN, Towards Early Detection and Diagnosis of Alzheimer's Disease. *Interdisciplinary Sciences: Computational Life Sciences*, 1-18.
- Hazarika, R. A., Maji, A. K., Syiem, R., Sur, S. N., & Kandar, D. (2022). Hippocampus segmentation using U-net convolutional network from brain magnetic resonance imaging (MRI). *Journal of Digital Imaging*, 35(4), 893-909.
- Hiremath, M. S., Biradar, R. C., Subhadarshini, L., Hiremath, S. S., Hiremath, J. S., & Sivanandan, S. K. (2024, April). Enhancing Alzheimer's Disease Diagnosis: A Hybrid Model of Transfer Learning and SVM for Improved Classification Accuracy. In *2024 IEEE 9th International Conference for Convergence in Technology (I2CT)* (pp. 1-8). IEEE.
- Ahmed, G., Er, M. J., Fareed, M. M. S., Zikria, S., Mahmood, S., He, J., ... & Aslam, M. (2022). Dad-net: Classification of alzheimer's disease using adasyn oversampling technique and optimized neural network. *Molecules*, 27(20), 7085.
- SAHOO, D. K. REVIEWS ON BRAIN TUMOUR SEGMENTATION AND CLASSIFICATION USING CONVENTIONAL MACHINE LEARNING AND DEEP-LEARNING APPROACH DEBENDRA KUMAR SAHOO*, SATYASIS MISHRA AND MIHIR NARAYAN MOHANTY. *Stochastic Modeling*.
- Alayba, A. M., Senan, E. M., & Alshudukhi, J. S. (2024). Enhancing early detection of Alzheimer's disease through hybrid models based on feature fusion of multi-CNN and handcrafted features. *Scientific Reports*, 14(1), 31203.
- Aderghal, K., Afdel, K., Benois-Pineau, J., & Catheline, G. (2020). Improving Alzheimer's stage categorization with Convolutional Neural Network using transfer learning and different magnetic resonance imaging modalities. *Heliyon*, 6(12).
- Sagiroglu, S., Terzi, R., Celtikci, E., Börcek, A. Ö., Atay, Y., Arslan, B., ... & Azginoğlu, N. (2025). A novel brain tumor magnetic resonance imaging dataset (Gazi Brains 2020): initial benchmark results and comprehensive analysis. *PeerJ Computer Science*, 11, e2920.
- Narasu, K. R., Nanda, A., Marshiana, D., Joe, B., & Kumar, V. (2022). Estimation and Analysis of Prediction Rate of Pre-Trained Deep Learning Network in Classification of Brain Tumor MRI Images. *The Internet of Medical Things (IoMT) Healthcare Transformation*, 73-98.
- Hasan, R., Raju, V., Pratussha, M., Jashva, M. B., & Pavithra, S. (2024). MULTIMEDIA CONTENT ANALYSIS FOR ALZHEIMER'S DISEASE DIAGNOSIS USING MRI SCANS AND DEEP LEARNING. *ICTACT Journal on Image & Video Processing*, 15(1).
- Zhu, G., Chen, H., Jiang, B., Chen, F., Xie, Y., & Wintermark, M. (2022, April). Application of deep learning to ischemic and hemorrhagic stroke computed tomography and magnetic resonance imaging. In *Seminars in Ultrasound, CT and MRI* (Vol. 43, No. 2, pp. 147-152). WB Saunders.
- Sathyabhama, B., & Kannan, M. (2023, November). State-of-the-Art techniques in Alzheimer's Disease prediction and classification: A Comprehensive Review. In *2023 9th International Conference on*

- Smart Structures and Systems (ICSSS) (pp. 1-8). IEEE.
- Ghosh, S., Sindhuja, P., Kesavan, D. K., Gulyás, B., & Máthé, D. (2025). Brain-Computer Interfaces and AI Segmentation in Neurosurgery: A Systematic Review of Integrated Precision Approaches. *Surgeries*, 6(3), 50.
- Tharwat, M., Sakr, N. A., El-Sappagh, S., Soliman, H., Kwak, K. S., & Elmogy, M. (2022). Colon cancer diagnosis based on machine learning and deep learning: modalities and analysis techniques. *Sensors*, 22(23), 9250.
- Xin, N. I. E. (2025). Multi-Modal Image Fusion for Medical Diagnosis: Combining MRI And CT Using Deep Generative Models. *Clinical Medicine And Health Research Journal*, 5(03), 1313-1327.
- Fonseka, H., Varastehpour, S., Shakiba, M., Golkar, E., & Tien, D. (2025). Convolutional variational auto-encoder and vision transformer hybrid approach for enhanced early Alzheimer's detection. *Journal of Medical Imaging*, 12(3), 034501-034501.
- Surianarayanan, C., Lawrence, J. J., Chelliah, P. R., Prakash, E., & Hewage, C. (2023). Convergence of artificial intelligence and neuroscience towards the diagnosis of neurological disorders—a scoping review. *Sensors*, 23(6), 3062.
- Mansingh, P., Choudhary, R., Thatoi, P., & Khairari, A. (2024). Transfer Learning with Convolutional Neural Networks: a method for the medical diagnosis of Alzheimer's disease stage categorization. *Frontiers in Health Informatics*, 13(7).
- Toumaj, S., Heidari, A., Shahhosseini, R., & Jafari Navimipour, N. (2024). Applications of deep learning in Alzheimer's disease: A systematic literature review of current trends, methodologies, challenges, innovations, and future directions. *Artificial Intelligence Review*, 58(2), 44.
- Dangi, S., & Gupta, S. (2023). Detection of Alzheimer's Disease through Convolutional Neural Networks: An Analysis of Magnetic Resonance Imaging Images. *IEEE Xplore*.
- Shourie, P., Anand, V., Chauhan, R., Dangi, S., & Gupta, S. (2023, December). Detection of Alzheimer's Disease Through Convolutional Neural Networks: An Analysis of Magnetic Resonance Imaging Images. In *2023 3rd International Conference on Smart Generation Computing, Communication and Networking (SMART GENCON)* (pp. 1-6). IEEE.
- Tiwari, A. (2023). Wilson's disease classification using higher-order Gabor tensors and various classifiers on a small and imbalanced brain MRI dataset. *Multimedia Tools and Applications*, 82(23), 35121-35147.
- Aparna, M., & Rao, B. S. (2023). Enhanced Classification of Alzheimer's Disease Stages via Weighted Optimized Deep Neural Networks and MRI Image Analysis. *Traitement du Signal*, 40(5), 2215-2223.
- Mehmood, A., Zheng, Z., Khan, R., Smadi, A. A., Shahid, F., Iqbal, S., ... & Ibrahim, M. M. (2024). Enhancing Mild Cognitive Impairment Detection through Efficient Magnetic Resonance Image Analysis. *Computers, Materials & Continua*, 80(2).
- Pan, D., Luo, G., Zeng, A., Zou, C., Liang, H., Wang, J., ... & Yang, B. (2022). Adaptive 3DCNN-based interpretable ensemble model for early diagnosis of Alzheimer's disease. *IEEE transactions on computational social systems*, 11(1), 247-266.
- Kim, H. W., Lee, H. E., Oh, K., Lee, S., Yun, M., & Yoo, S. K. (2020). Multi-slice representational learning of convolutional neural network for Alzheimer's disease classification using positron emission tomography. *Biomedical engineering online*, 19(1), 70.
- Paul, M. A. V., Venkatesan, S., Devi, A. R., Karthikeyan, K., Muthuseli, M., & Bhuvanesh, A. (2025). Application of adaptive spotted hyena algorithm with deep efficient network for detecting parkinson's disease. *International Journal of Information Technology*, 1-16.
- Hiremath, M. S., Biradar, R. C., Praveen, N., Hiremath, J. S., Kumar, S., & Hiremath, S. S. (2024, May). Advancing Alzheimer's Diagnosis: Integrating Transfer Learning

with Support Vector Machine for Enhanced Disease Classification. In *International Conference on Computer Vision and Robotics* (pp. 101-113). Singapore: Springer Nature Singapore.

Patro, P., & Das, T. K. (2025). Artificial Intelligence in Action (2019-2024): A Review on Parkinson's Disease Detection using Non-Invasive Procedures. *IEEE Access*.

