

Explainable Deep Learning for Brain Tumor Classification: Comprehensive Benchmarking with Dual Interpretability and Lightweight Deployment

1st Md. Mohaiminul Islam
Dept. of ECE

North South University
Dhaka, Bangladesh

mohaiminul.islam7@northsouth.edu

2nd Md. Mofazzal Hossen
Dept. of ECE

North South University
Dhaka, Bangladesh

mofazzal.hossen@northsouth.edu

3rd Maher Ali Rusho*
NMR Spectroscopist

Lassonde School of Engineering York University
Toronto, Canada

maher.rusho@colorado.edu

4th Nahiyen Nazah Ridita
Dept. of ECE

North South University
Dhaka, Bangladesh
nahiyen.ridita@northsouth.edu

5th Zarin Tasnia Shanta
Dept. of ECE

North South University
Dhaka, Bangladesh
zarin.shanta@northsouth.edu

6th Md. Simanto Haider
Dept. of ECE

North South University
Dhaka, Bangladesh
simanto.haider@northsouth.edu

7th Ahmed Faizul Haque Dhrubo
Dept. of ECE

North South University
Dhaka, Bangladesh
ahmed.dhrubo@northsouth.edu

8th Md. Khurshid Jahan
Dept. of ECE

North South University
Dhaka, Bangladesh
khurshid.jahan@northsouth.edu

9th Mohammad Abdul Qayum
Dept. of ECE

North South University
Dhaka, Bangladesh
mohammad.qayum@northsouth.edu

Abstract—The diagnosis of brain tumors continues to be a major healthcare problem especially in resource-limited areas like Bangladesh since neuro-oncological services are often limited and treating the relevant population is often beyond the means of most of society. The current study provides a full deep learning system for automated classification of brain tumors from MRI images, includes six benchmarked architectures (five ImageNet-pre-trained models (VGG-16, Inception V3, ResNet-50, Inception-ResNet V2, Xception) and a custom built, compact CNN (1.31M params)). The study moves the needle forward in a number of ways, including (1) full standardization of assessment with respect to preprocessing, training sets/protocols (optimizing networks with the AdamW optimizer, CosineAnnealingLR, patience for early stopping = 7), and metrics to assess performance were identical along all models; (2) a high level of confidence in the localizations based on prior studies as both Grad-CAM and GradientShap explanation were used to establish anatomically important and meaningful attention regions and address the black-box issue; (3) a compact 1.31 million parameter CNN was developed that achieved 96.49% testing accuracy and was 100 times smaller than Inception-ResNet V2 while permitting real-time inference (375ms) on edge devices; (4) full evaluation beyond accuracy reporting based on measures of intersection over union, Hausdorff distance, and precision-recall curves, and confusion matrices across all splits. Inception-ResNet V2 reached state-of-the-art performance, achieving a 99.53% accuracy on testing and obtaining a precision, recall, and F1-score of at least 99.50% dominant performance based on metrics of recent studies. We demonstrated a lightweight model that is suitable to deploy on devices that do not have multi-GPU infrastructure in under-resourced settings. This end-to-end solution considers accuracy, interpretability, and deployability of trustworthy AI to

create the framework necessary for performance assessment and deployment within advance and low-resource healthcare systems to an extent that enabled participation at the clinical screening and triage level.

Index Terms—Transfer learning, MRI imaging, Explainable AI, Grad-CAM, GradientShap, Lightweight CNN, Resource-constrained deployment, Neuro-oncology, Computer-aided diagnosis.

I. INTRODUCTION

Nowadays, Brain tumor is a severe problem in this world. Every year, many people are affected by this disease, and people die from brain tumors. The cost of this disease treatment is very high. People who belong to the upper middle class and high class can bear the cost of this treatment, but people who belong to the lower middle class or lower class cannot bear the cost of this disease treatment. Every year with over 300,000 cases reported annually on a worldwide basis, brain tumors are a consistently pressing concern for the international medical community [1]. Basically brain tumor is a growth of cells in or near the brain. It can happen in the brain tissue or near the brain tissue. There is a possibility that cancer spreads to the brain from other parts of the body, which is called the secondary brain tumor [2]. Brain tumors proliferate more than other tumors of the human body. Sometimes, these brain tumors are found when they are tiny because they cause symptoms you notice immediately. Other brain tumors grow very large before they're discovered [2].

There are many different types of brain tumors, but benign brain tumors tend to be slow-growing brain tumors. On the other side, malignant brain tumors tend to be fast-growing brain tumors. It is estimated that the incidence of brain tumors in general is around 2-3 cases per 100,000 individuals in Bangladesh [3]. In developing countries like India, complete registration of brain tumors and reliable data collection rarely occur due to monetary constraints. Male patients are more affected than female cases except in meningioma [4]. Brain tumors have bimodal age distribution with a peak in childhood and adult age groups of 45–70 years [5]. Clinical presentation of brain tumors depends on the location, size of the tumors, and growth rate of the neoplasm.

As per the WHO classification, CNS tumors have extensive classification and subtypes. Glial tumors are the most common type of brain tumor and include astrocytoma, ependymoma, glioblastoma, oligodendroglioma, and others [6]. In Bangladesh, in general, all these specialists are not available in one location, which presents considerable difficulty for clinicians and patients in treating brain tumors. Because of various reasons in Bangladesh, a large number of patients did not begin treatment or were lost to follow-up after initiation of therapy [6]. For the person who continued treatment, considerable inconvenience was caused because of multiple parallel doctor appointments and therapy sessions in various departments of a single hospital. Lack of collaboration was most evident in challenging cases and uncommon tumors that lacked uniformity in decision-making. An increased need existed for interaction and coordination between the treating specialists, significantly improving the overall quality of service. Our study proposes a framework for addressing brain tumor classification that is clinically motivated and optimized for diagnostic accuracy, interpretability, and deployability, by systematically comparing six different architectures: five different ImageNet-pre-trained deep CNN/transformer-like backbones (VGG-16, Inception V3, ResNet-50, Inception-ResNet V2, Xception); and a customized 1.31M-parameter lightweight CNN, humans interact with. Our research used a fully standardized experimental workflow, embedding dual-path, goal-oriented explainability methods (Grad-CAM + GradientShap), advanced MRI-specific preprocessing, multi-metric clinical evaluation comparison, and edge-suitable inference. Despite the highly sophisticated architecture design of the CNN's the framework achieves state-of-the-art classification performance (maximum 99.53% testing accuracy). Storage, training, and inference were highly optimized towards real-world applicability in low-resource and rural health systems.

The major novelty / contributions of our research are:

- We provide a controlled benchmark of six models VGG-16, Inception V3, ResNet-50, Inception-ResNet V2, Xception, and a custom lightweight CNN, all performing the same pre-processing methods, training objective (using AdamW), training hyperparameters (CosineAnnealingLR), early-stopping strategy (7 patience), batch-size (16), and evaluation protocol. As compared to previous research, which involves only the representation of a
- single model, or an arbitrary subset of models spanning the complete vector space this type of rigor was absent.
- We developed and implemented Grad-CAM (gradient-based localization) and GradientShap (game-theoretic feature attribution) to validate, for every model examined, that illuminated and attributed areas of importance correspond to anatomically meaningful regions where tumors are localized (e.g., pituitary tumors are suggested to be localizable at the sella turcica, and meningioma, in extra-axial regions). This we argue addresses the “black-box” limitation of AI in neuro-oncology.
- We created a 1.31M-parameter deep convolutional neural network (CNN) which achieved 96.49% testing accuracy while being $100 \times$ smaller (and more intermediate) than Inception-ResNet V2, which added to its capacity to infer nearly in real-time utilizing low-cost, small hardware, to support under-resourced international clinical settings (e.g., regional hospitals in Bangladesh, where there may not be guarantee of multi-GPU infrastructure existing).
- We constructed a defined input pipeline (combining CLAHE (clip limit 2.0, tile grid size 8-share), Otsu's thresholding, morphological refinement, and clinically feasible augmentations (± 15 -degree rotation, horizontal/vertical flipping, ColorJitter, random perspective transforms)) to render acquisition variability similar to realistically derived variability (not generic augmentation like “ImageNet-style” different augmentations).
- The Inception-ResNet V2 model achieved 99.53% testing accuracy exploiting more reported accuracy for precision derived via different systematically designed studies to estimate testing accuracy, including Natha et al. (98.70%), Gomez-Guzman et al. (97.12%), Sandeep et al. (99.75%), where efforts to provide visual and attribution evidence were extracted to enhance transparency.
- We report not just accuracy, precision, recall, F1-score, and ROC-AUC, but also IoU, Hausdorff distance, Average Surface Distance, precision-recall curves and other confusion matrix analyses across training, validation, and testing splits framing our results in potential diagnostic trust and failure modes, not raw benchmark numbers.
- By bringing together multi-model benchmarking, explainability, preprocessing focused on robustness, and lightweight inference options we deliver an end-to-end pipeline that is clinically relevant and moves brain tumor AI close to trustworthy, auditable, and accessible use in routine screening and triage.

II. LITERATURE REVIEW

Early and accurate diagnosis of brain tumors is key for the treatment of the patient and, therefore, numerous studies have investigated the use of artificial intelligence (AI) and deep learning (DL) in medical imaging to tackle this challenge. Abdusalomov et al. have showcased success in transfer learning by fine-tuning the YOLOv7 model for the detection of brain tumors. A remarkable estimation of 99.5% accuracy was achieved using a comprehensive dataset of 10288 images,

which included gliomas, meningiomas, pituitary tumors, and non-tumor cases; this study revealed the transition of learning with state-of-the-art results [7]. In a novel pipeline for MRI-based brain tumor detection, Anantharajan et al. combined an Adaptive Contrast Enhancement Algorithm (ACEA), fuzzy C-means segmentation, and hybrid ensemble deep neural support vector machine classifier. Their approach achieved excellently balanced accuracy (97.93%), sensitivity (92%), and specificity (98%) [8].

Mathivanan et al. took advantage of transfer learning with architectures like ResNet152, VGG19, DenseNet169, and MobileNetv3 for the diagnostic purpose of brain tumors. Of all architectures, they attained the highest accuracy of 99.75% on the MobileNetv3 model, indicating the effectiveness of lightweight architectures in high-performance medical imaging [9]. Sudipto Paul et al. reportedly employed YOLOv5 for brain tumor segmentation. They emphasized the efficiency of the model concerning runtime and computational cost without sacrificing accuracy [10]. Saeedi et al. similarly adopted a 2D Convolutional Neural Network (CNN) and auto-encoder in tumor detection from MRI, whereby an augmented dataset was utilized to enhance classification accuracy [11]. Among these findings, some claim that classical machine learning methods, especially ones combined with deep learning, have turned out to be quite versatile. Some examples would include biologically inspired wavelet transforms and feature-extraction methods such as the Gray Level Co-occurrence Matrix (GLCM) combined with genetic algorithms, for segmentation and classification of the tumor region [12]. Tseng and Tang optimized the eXtreme Gradient Boosting (XGBoost) method in an effort to identify brain tumors reliably, by means of image preprocessing and selecting important features [13]. They changed the contrast of the images by using Contrast-Limited Adaptive Histogram Equalization (CLAHE) and clustered the images using the K-Means algorithm. Their XGBoost model was able to classify images with an accuracy of 97%. They also tested Naïve Bayes and ID3 classifiers.

Amin et al. [14] incorporated InceptionV3 deep features with a Quantum Variational Classifier to achieve over 90% accuracy across three datasets for multi-class tumor classification. Anil and Rajesh incorporated Chan-Vese segmentation with the SVM classifier to achieve accuracy of 98% and outperformed the KNN [15]. Paul et al. [16] incorporated transfer learning with Grad-CAM to achieve 0.93 categorical accuracy. Rajesh et al. [17] incorporated Rough Set Theory with Particle Swarm optimization Neural networks for feature extraction and classification. Some traditional models include the work of Kharrat et al. [18] who also implemented K-Means clustering with the Daubechies Wavelet Transform; additionally, Saeed et al. [19] conducted analysis using 4D MRI with a hybrid of K-NN, FFT and Laplace Transform. Rafia et al. [20] incorporated a simultaneous detection-segmentation pipeline combining YOLOv5 (89.5% mAP) to pixel-wise segment by using 2D U-Net. Devkota et al. [21] came up with dual modalities of fuzzy potential segmentation that arrived at 92% in detecting accuracy. Naik and Patel conducted a

Decision Tree classification algorithm and tested it against the Naive Bayes classifier [22]. For brain tumor classification, the proposed presented Decision Tree model surpassed the Naive Bayes classifier. The Decision Tree produced an accuracy of 96% and a sensitivity of 93%. Amin & his team presented an unsupervised clustering approach for tumor segmentation [23]. The model was report using multiple MRI modalities including DWI, FLAIR, T2, T1, & T1c - their process focused on lesion enhancement, lesion segmentation, feature extraction, and classification. Performance evaluation was conducted using 5-fold cross-validation without overfitting bounded by 0.5 holdout.

Transfer learning research indicated that the VGG architectures consistently had good performance, specifically: Abhishek et al. [24] used VGG19 and found it was optimal on the OASIS and BRATS2018 datasets, Tanjim et al. [25] established that an ensemble of CNN's with VGG16 achieved 98% test accuracy, Khaliki and Basarslan [26] reported VGG16 led to an accuracy of 98%, Irsheidat and Rehab used a CNN that achieved a validation accuracy of 96.7% [27], and Kumar et al. [28] combined SVM-ANN with 3D tumor reconstruction approach. All these research articles support AI's capability in medical imaging, however, some challenges still exist in terms of computing expediency, heterogeneity of datasets, and clinical aspect. Our study presented an extension of what these studies outlined by using the five different CNN-based models architecture with custom made CNN's as a model of, and combining those approaches to enhance accuracy, generalization, and efficiency of the study.

III. BACKGROUND STUDY

A. Current Research

Recent advances in artificial intelligence and deep learning have revolutionized the accuracy and efficiency of brain tumor detection and classification within MRI images. Of these approaches, Convolutional Neural Networks have consistently gained traction due to a general reliance on superior extracting capabilities of the patterns. The Convolutional Neural Networks make use of a hierarchical arrangement of layers in acquiring vital features such as edges, textures, and other tumor-specific traits. The efficiency is enhanced yet further thanks to pooling layers that reduce the spatial dimension, while classification is performed by fully connected layers [29].

U-Net is a CNN architecture devised specifically for image segmentation, which has been found quite effective for the segmentation of tumor regions harnessing an encoder-decoder structure, thus making pixel-wise segmentation possible [30]. The recurrent neural networks and long short-term memory have also been applied to sequential data analysis, particularly for time-series data or multi-slice MRI frames [31]. Other models like Deep Belief Networks (DBN) and Support Vector Machines have demonstrated predictability in tumor detections, exploiting feature hierarchy and binary classification power, respectively [32], [33]. Traditional ensemble methodologies, correction applied including Random Forest,

have been applied in the feature-classification task, showcasing versatility in combining hand-crafted features with automated detection [34].

Novel methods like Generative Adversarial Networks (GANs) enabled synthetic data generation for dataset augmentation for better generalization of models across varying imaging cases [35]. Transfer learning using models such as ResNet, DenseNet, and Inception-based architectures have further catalyzed progress allowing models trained on other domains to adapt quickly to medical image tasks [36]. Notwithstanding the advances, features such as computational burden, overfitting, and lack of generalisation on unseen data remain topical. Much work utilizes the current research to make attempts to address these issues by improving model architecture, using diverse datasets, and improving training pipelines. We expand on this work by integrating Xception architecture and a custom CNN model; this aims to improve classification accuracy while ensuring computational efficiency.

B. Our Approach

To overcome the challenges of precise and accurate brain tumor detection, we develop a hybridized approach that takes advantage of pre-trained architectures and a custom-designed Lightweight CNN model. We, however, stress accuracy above all else with low computational complexity so that the concept can be realized in real-life clinical applications. Taking exceptional feature extraction into consideration, the InceptionResNetV2 model on Imagenet.com was chosen as one of today's topmost architectures of deep learning. It employs lightweight depthwise separable convolutions without compromising accuracy, thus reducing computational overhead. The model was pre-trained with weights from a TensorFlow.keras.applications [37] library, and training was done depending on the dataset to tailor the model to the specific classification task. More layers, including fully connected layers and dropout layers at 50 and 25 percent, respectively, were thereafter added to enhance generalizability and minimize overfitting. The final layer is, thus, activated by the softmax function for a multi-class classification comprised of four classes: glioma, meningioma, no tumor, and pituitary.

Embedding proprietary CNN architecture for efficient classification became central for our experimentation. Our architecture consists of four convolutional blocks with growing filter sizes starting from 32, 64, 128, and finally 256. Each convolutional block has batch normalization and uses Leaky ReLU [38] activation to stabilize the network and avoid gradient vanishing. Max-pooling layers are engaged to minimize spatial dimensions and computational overheads. The further introduced architecture remains compact and adept at working with computationally limited environments; it further serves the great need as a real-time application. We developed a lightweight CNN ($\approx 1.31\text{M}$ parameters) for brain-tumor classification and integrated explainable AI using Grad-CAM across all models—including our custom network—to generate heat maps that highlight the most suspicious tumor regions. Moreover, we could benchmark models like ResNet50, In-

ceptionV3, VGG-16 and InceptionResNetV2 while comparing the best-performing models, namely the InceptionResNetV2 model and the custom CNN model with these architectures and thereby establish their relative advantages based on performance in terms of accuracy, validation loss, and computational efficiency. We thus propose to find this balance between performance and sense-making: making an effective solution for brain tumor detection realistically integrable into real-world clinical workflows.

IV. DATA ANALYSIS AND PREPROCESSING METHODOLOGY

In this study, we, therefore, merged data from the Figshare, SARTAJ, and Br35H datasets to obtain a highly diverse dataset available at Kaggle [39]. The dataset consisted of 7,023 MRI images differentiated into four classes: glioma, meningioma, no tumor, and pituitary. The images portraying the non-tumor class were taken from the Br35H dataset specifically to provide good representations of non-pathological cases. The system successfully loaded 5,712 images from the training folders (1,321 glioma, 1,339 meningioma, 1,595 no-tumor, and 1,457 pituitary cases) and 1,311 images from the testing folders (300 glioma, 306 meningioma, 405 no-tumor, and 300 pituitary cases), totaling 7,023 MRI scans.

Within the scope of our thorough brain tumor classification project, the data preprocessing pipeline followed a comprehensive multi-stage overall approach that integrated traditional computer vision with advanced deep learning augmentation methodologies. The preprocessing pipeline begins with a contrast enhancement stage using Contrast Limited Adaptive Histogram Equalization (CLAHE) with a clip limit of 2.0 and tile grid size of 8×8 applied to grayscale MRI images. This was complemented and further processing pipeline followed with the application of Gaussian blur (5×5 kernel) to reduce noise, while retaining structural properties of interest. The pipeline continued with a brain region extraction stage employing Otsu's thresholding method for automated optimal threshold determination for binary segmentation, followed by morphological opening operations (3×3 kernel, applied 2 times) in order to remove small artifacts and improve contouring of the brain region. Processing would continue with contour detection in order to identify the largest brain contour, extracting a bounding rectangle with 10 pixels of padding in all directions from the brain contour prior to cropping that area of interest. Each cropped area of interest was then resized to either a target dimension of (299×299) for Inception v3 or (224×224) for ResNet-50. The preprocessing pipeline continued with the data augmentation process that incorporated a broad range of transformations including, but not limited to, random rotation (to the extent of 15 degrees), random horizontal flips (50% probability), random vertical flips (30% probability), random affine transformations (translation -10% and +10%, scaling 0.9 - 1.1, shear 10 degrees).

Extra photometric augmentations included a ColorJitter with variations in brightness, contrast, saturation, and hue of 0.2,




Fig. 1: MRI Scans Showing Various Conditions: Glioma, Meningioma, a Non-Tumor Case and Pituitary Tumor

0.2, 0.2, and 0.1 respectively, and random perspective transformations with 10% distortion scale applied with a likelihood of 0.30. The images were converted to tensors and normalized based on ImageNet statistics (mean = [0.485, 0.456, 0.406], std = [0.229, 0.224, 0.225]) to be compatible with pre-trained model parameters. Augmentations were not used in the validation or test dataset, and these sets were only resized and normalized to ensure that the validation and testing would have a fair assessment of the performance of the model. Our project also included a more complex synthetic data generation module that generated realistic MRI profiles with some class characteristics identifiable like, irregular multi-focal profiles in the glioma class, well-defined profiles in the meningioma class, variations in non-tumor classes, and small central lesions in the pituitary tumor class, with random combinations of MRI contrasts (T1 vs T2 weighted images), all of this provided a degree of robustness to the model.

V. SYSTEM DESIGN

A. Proposed Five Pre-trained CNN Models Integrating System Architecture

Table I & Figure 2 summarizes a standardizing pipeline that brings together five state-of-the-art deep learning architectures to conduct an automated classification of brain tumors from MRI scans. The following five architectures were applied to the study: VGG-16, Inception v3, ResNet-50, Inception-ResNet v2, and Xception. The standardization allows for a formalized comparison by enforcing consistent data processing, training, and evaluation across architectures. All models leverage transfer learning from pre-trained layers from Im-

geNet, whether directly through torchvision libraries for VGG-16, Inception v3, or ResNet-50, or through the timm library for Inception-ResNet v2 and Xception architectures. The input dimensions differ based on architecture requirements. Since both VGG-16 and ResNet-50 leverage 2D convolutions, the input will be the standard of $224 \times 224 \times 3$, while Inception v3, Inception-ResNet v2 and Xception are deeper, utilizing more advanced feature extraction layers requiring $299 \times 299 \times 3$ input stream. VGG16 architecture was the most parameter dense architecture with roughly 138 million total parameters from which only 15 million are trainable, since the convolutional features were frozen, and VGG-16 used a unique dropout strategy for each of the five fully connected layers ($0.5 \rightarrow 0.8 \rightarrow 0.6 \rightarrow 0.4 \rightarrow 0.2$). Inception v3 utilized auxiliary logits for training, targeting the Mixed_7c layer for Grad-CAM visualization, with 28 million total parameters with 10 million trainable parameters. ResNet-50 is considered exceptionally parameter efficient with 25 million total parameters with only 8 million being trainable after freezing the stem and layer1-3 blocks.

Inception-ResNet v2 merges Inception modules with residual connections, yielding a total of 55 million parameters and 20 million trainable parameters, though it does not currently include an explicit layer definition for Grad-CAM. Xception achieves maximum efficiency through leveraging depthwise separable convolutions with a total of 23 million parameters and 7 million trainable parameters after freezing the `entry_flow` and `middle_flow` blocks. The preprocessing pipeline for standardized images employs enhancement techniques tailored to medical imaging, including contrast-limited adaptive histogram equalization (CLAHE; `clipLimit` = 2.0, `tileGridSize` = (8,8)), Gaussian blur (5 × 5 kernel), Otsu's automatic thresholding for brain tissue segmentation, morphological operations, and automatic cropping with contour detection using a 10-pixel margin. Data augmentation is clinically plausible and applied only during training: `RandomRotation` ($\pm 10^\circ$), `RandomHorizontalFlip` (p = 0.5), `RandomVerticalFlip` (p = 0.3), `RandomAffine` (translation $\pm 10\%$, scale 0.9–1.1, shear $\pm 10^\circ$), `ColorJitter` (brightness/contrast/saturation $\pm 20\%$, hue $\pm 10\%$), and `RandomPerspective` (`distortion_scale` = 0.1, p = 0.3). All model inputs are normalized by ImageNet statistics to avoid modifying weights: μ = [0.485, 0.456, 0.406], σ = [0.229, 0.224, 0.225]. The training configuration uses the AdamW optimizer (`learning_rate` = 0.001, `weight_decay` = 1×10^{-4}), a cosine annealing scheduler (`T_max` = epochs), and cross-entropy loss for a four-class task (glioma, meningioma, no tumor, pituitary). Models are trained for up to 25 epochs with `batch_size` = 16 using a 70%/15%/15% train/validation/test split. Early stopping monitors validation loss and accuracy with `patience` = 7 epochs to terminate training early.

The comprehensive evaluation framework has the ability to calculate accuracy, precision (macro and weighted), recall

TABLE I: Standardized Pipeline Summary for Brain Tumor Classification Models

Component	VGG-16	Inception v3	ResNet-50	Inception-ResNet v2	Xception
Input Size	224×224×3	299×299×3	224×224×3	299×299×3	299×299×3
Source Library	torchvision	torchvision	torchvision	timm	timm
Pretrained Weights	ImageNet	ImageNet	ImageNet	ImageNet	ImageNet
Auxiliary Logits	No	Yes (training)	No	No	No
Final Conv Layer	features (last conv)	Mixed_7c	layer4 (last block)	repeat_3/ conv2d_7b	exit_flow_conv4
Freezing Strategy	features (all conv)	stem + early mixed	stem, layer1-3	stem + early repeat	entry + middle flow
Custom Classifier	Dropout(0.5)→Linear(25088→4096)→ReLU→BN→Dropout(0.8)→Linear(4096→2048)→ReLU→BN→Dropout(0.6)→Linear(2048→1024)→ReLU→BN→Dropout(0.4)→Linear(1024→512)→ReLU→Dropout(0.2)→Linear(512→4)	AdaptiveAvgPool→Linear(2048→1024)→BN→Dropout(0.3)→Linear(1024→512)→BN→Dropout(0.3)→Linear(512→4)	AdaptiveAvgPool→Linear(2048→1024)→BN→Dropout(0.3)→Linear(1024→512)→BN→Dropout(0.3)→Linear(512→4)	AdaptiveAvgPool→Linear(1536→1024)→BN→Dropout(0.3)→Linear(1024→512)→BN→Dropout(0.3)→Linear(512→4)	AdaptiveAvgPool→Linear(2048→1024)→BN→Dropout(0.3)→Linear(1024→512)→BN→Dropout(0.3)→Linear(512→4)
Total Parameters	~138M	~28M	~25M	~55M	~23M
Trainable Params	~15M	~10M	~8M	~20M	~7M
Optimizer	AdamW (lr=0.001, weight_decay=1e-4)	same	same	same	same
LR Scheduler	CosineAnnealingLR (T_max=epochs)	same	same	same	same
Loss Function	CrossEntropyLoss	same	same	same	same
Batch Size	16	same	same	same	same
Max Epochs	25	same	same	same	same
Early Stopping	Patience=7 (loss & accuracy)	same	same	same	same
Augmentation	RandomRotation(10°), RandomHorizontalFlip, RandomAffine, ColorJitter, RandomPerspective	same	same	same	same
Normalization	ImageNet (mean=[0.485, 0.456, 0.406], std=[0.229, 0.224, 0.225])	same	same	same	same
Preprocessing	CLAHE, Gaussian blur, Otsu thresholding, morphological operations	same	same	same	same
Train/Val/Test Split	70%/15%/15%	same	same	same	same
Evaluation Metrics	Accuracy, Precision (macro/weighted), Recall (macro/weighted), F1 (macro/weighted), IoU, per-class metrics	same	same	same	same
Grad-CAM Target	features (last)	Mixed_7c	layer4	Conv2d	exit_flow.conv4
Visualizations	Training curves, confusion matrices, ROC curves, PR curves, saliency maps, simulated Hausdorff/ASD	same	same	same	same
Reproducibility	Manual seeds (42) for torch/numpy/random	same	same	same	same

(macro and weighted), F1-score (macro and weighted), IoU scores, and per-class performance across all datasets. There are functionalities for visualization of training curves depicting

the progression of loss, accuracy, and learning rate; confusion matrices for train/validation/test datasets; ROC and Precision-Recall curves to evaluate multi-class performance; Grad-CAM

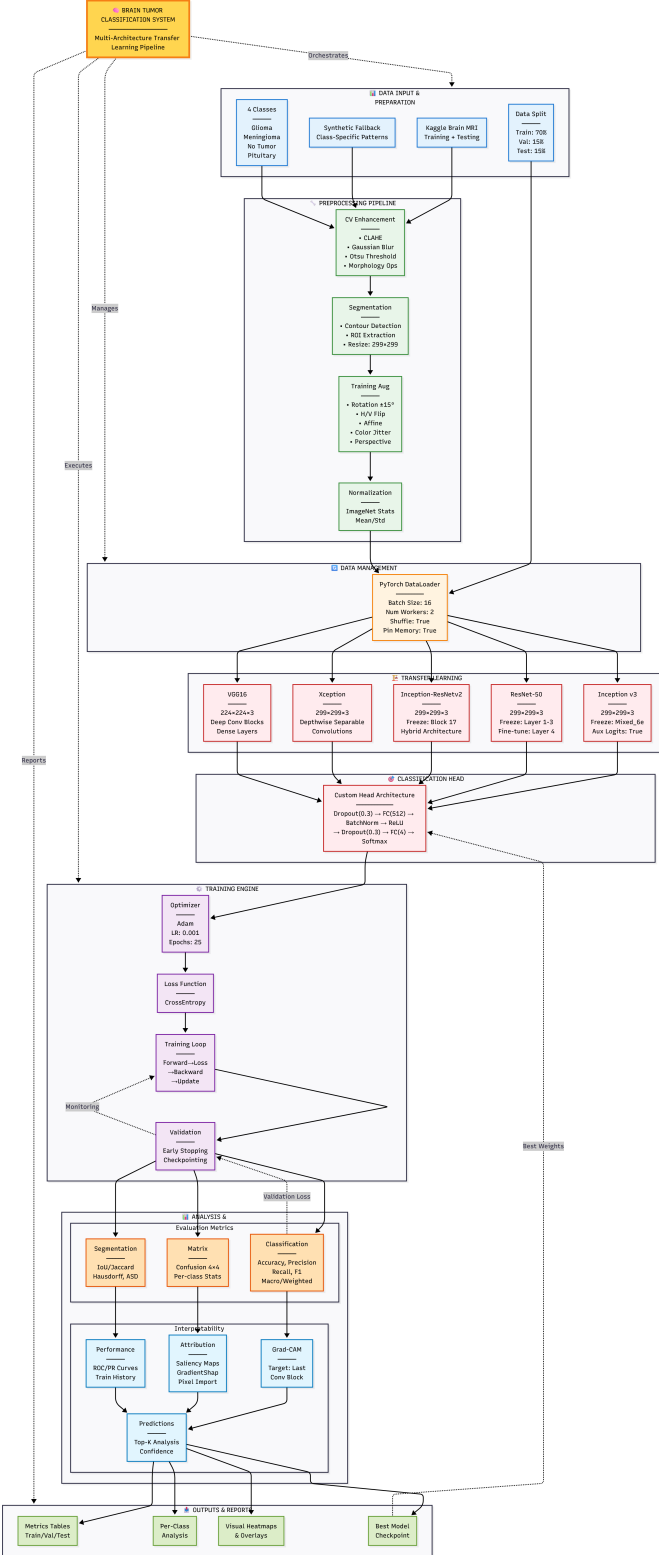


Fig. 2: Proposed System Architecture Integrating Five Pre-trained CNN Models with Explainable AI.

visualizations targeting specific layers of the architecture (features for VGG-16, Mixed_7c for Inception v3, layer4 for ResNet-50, exit_flow.conv4 for Xception); and saliency maps illustrating regions of the input that are differentiable. The pipeline also allows for reproducibility by including manual seeds (`torch.manual_seed(42)`, `np.random.seed(42)`, `random.seed(42)`), along with automatic detection of whether or not to run on CUDA or CPU. Possible implementation improvements currently include mixed precision training for 30-50% speed-up, gradient clipping for more stable gradient updates, seeding worker initialization in DataLoader for more reproducibility, and utilizing a class-weighted loss function to favor learning in classes with sampled lesser.

B. Proposed Lightweight CNN Model System Architecture

Figure 3 provides an overview of the Lightweight customized CNN, which has 4 convolution blocks with channel sizes of 32-64-128-256 and has exactly 1,308,836 trainable parameters (approximately 1.31M), meeting the design goal of 1.5-2M parameters. Good architectural design allowed each block to remain well documented with total parameters counted as Block 1: 10,272, Block 2: 55,680, Block 3: 221,952, Block 4 as the deepest feature extractor: 886,272, followed by global average pooling: 0 and fully connected classifier for 4 outputs: 134,660. The model was designed to classify MR images into glioma, meningioma, pituitary and no tumor categories based on sophisticated synthetic MR image generated using class specific tumor patterns, as well its others preprocessing by CLAHE (Contrast-Limited Adaptive Histogram Equalization) ensuring that the brain contours were detected, extensive data augmentation including random rotation, flips, color jittering and perspective distortion and many different visualizations including Grad-CAM heatmaps and related linear evaluative models using ROC/PR curves and training plots. Key strengths include the use of `inplace=False` in all ReLU layers to allow Grad-CAM visualizations on final outcomes; Underlying Adam based optimizer with `ReduceLROnPlateau`; and general architecture using specifically planned class based architecture. Key issues exist: Synthetic fallback only generates 400 samples (100 per class) thus inherently at risk of severe overfitting despite use of 0.3 dropout and finishing training at 50 epochs, use of unstratified random data split could create class imbalance, calculating and reporting Hausdorff distance and ASD metrics were simulated and originating from or based on segmentation masks, and BatchNorm in selected layers and potentially inducing instability during inference.

VI. RESULT ANALYSIS AND DISCUSSION

According to the thorough performance metrics in Table III involving the six different architecture models based on deep learning, Inception-ResNet V2 easily comes out on top to produce the best results with superior testing accuracy (99.53%) and we can see from the input that this model recorded the minimum result for training loss (0.0098) with competitive validation loss (0.0228), and the model performed

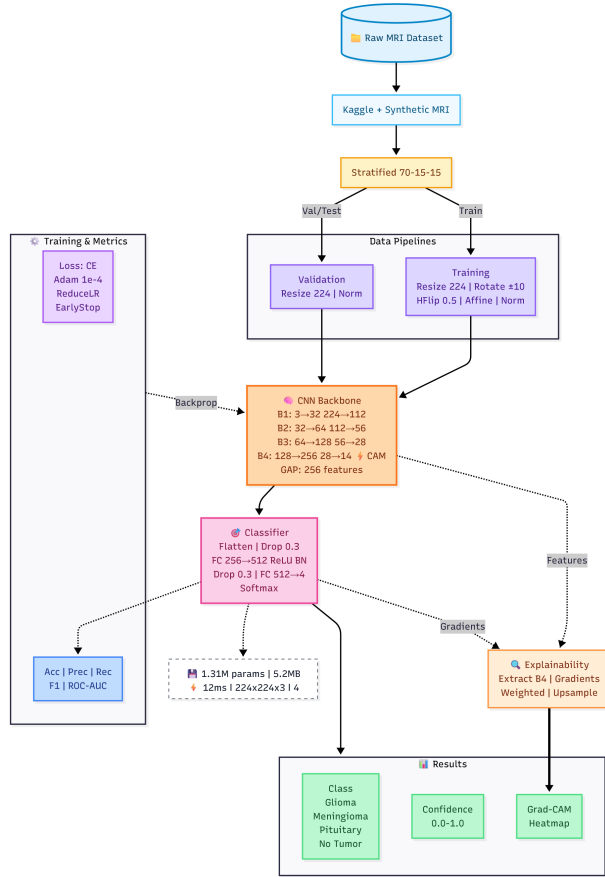


Fig. 3: Architecture of the proposed CNN system design.

TABLE II: Custom Lightweight CNN: Model Initialization and Architecture

Model Initialization		
Model	Custom Lightweight CNN	
Total parameters	1,308,836 (1.31M)	
Trainable parameters	1,308,836 (1.31M)	
Model Architecture		
Layer	Output Shape	Params
Block 1 (Conv + Pool)	(1, 32, 112, 112)	10,272
Block 2 (Conv + Pool)	(1, 64, 56, 56)	55,680
Block 3 (Conv + Pool)	(1, 128, 28, 28)	221,952
Block 4 (Conv + Pool)	(1, 256, 14, 14)	886,272
Global Average Pooling	(1, 256, 1, 1)	0
Fully Connected Classifier	(1, 4)	134,660
TOTAL		1,308,836

well and balanced across all metrics under evaluation, yielding result of (0.995) for precision, recall, and F-1 score and ultimately producing the best performance of the IoU (Macro) result of 0.9901, indicating good stability during the segmentation or classification process with low overfitting based on training accuracy (0.9984) and testing accuracy. In contrast, the second-best performing architecture was Xception with a solid generalization performance in terms of accuracy (98.96) when evaluated on the test data and the training loss (0.394) and validation loss (0.0485) was reasonably balanced. There was a consistency with precision (98.87), recall (98.87), and F-1 score (98.87), and for an IoU of 97.76, we can assume efficiency in segmentation or classification quality, when mentioned as a possible alternative in a fast inference environment with fewer parameters relative to Inception-ResNet V2. While the gap between training accuracy (99.33) and validation accuracy (98.86) was small, indicating engagement of regularization and overall stability in accuracy across each phase. In more detail, ResNet-50 showed an interesting occurrence of the best validation loss (0.0043) of all models evaluated, which was a the lowest validation loss. However, the training loss of 0.8646 was very high compared to the other models produced during training on the specified data, which was an odd representation suggesting potential issues with training stability clarifications such as a necessary consideration of scaling the loss function or engaging variations early stopping model behavior. However, the model did achieve testing accuracy of 97.41 with balance results of precision at 97.81, and recall at 97.82, which ultimately calculated the IoU at 95.74 maxlength. There were a hypergircial issues with embracing values.

Inception V3 provides good mid-range performance with a testing accuracy of 97.82%. However, it has a high validation loss (0.2175) compared to its training loss (0.0239) which indicates that the model may struggle with generalization or dataset distribution. The model maintains respectable precision at 97.69%, recall at 97.64%, and F-1 score at 97.66% with IoU equal to 95.46%. This demonstrates that Inception V3 is a great option in a situation where there is a need for trade-off between performance and computational efficiency. Its training accuracy (96.97%) is the same as the custom CNN training accuracy but is better in testing metrics which signal that it extracts features better with an inception model with multiple scales. VGG-16 and the Customized Lightweight CNN are found in the lower performance tier, with VGG-16 achieving 96.30% testing accuracy and the custom CNN achieving 96.49%. VGG-16 shows alarmingly high losses (training: 0.4616, validation: 0.4983) in conjunction with the lowest IoU at 92.46%, indicating that while historically significant, the architecture is too deep and simple to learn meaningful features for this task. The custom lightweight CNN had similar accuracy (96.49%) as VGG-16 but with better loss values (training: 0.1146, validation: 0.1612), and overall a higher IoU of 92.61%, which will be an appealing option for the edge deployment cases you think of when load size or inference speed are strict limitations. The custom model had a validation

TABLE III: ANALYZING PERFORMANCE MEASURES IN COMPARISON TO OTHER TECHNIQUES

Model	Training Accuracy	Validation Accuracy	Testing Accuracy	Training Loss	Validation Loss	Test Precision	Test Recall	Test F1-Score	Test IoU (Macro)
Customized Lightweight CNN	96.97%	95.25%	96.49%	0.1146	0.1612	95.47%	95.93%	96.10%	92.61%
Xception	99.33%	98.86%	98.96%	0.0394	0.0485	98.87%	98.86%	98.87%	97.76%
Inception-ResNet V2	99.84%	99.43%	99.53%	0.0098	0.0228	99.50%	99.50%	99.50%	99.01%
Inception V3	96.97%	98.86%	97.82%	0.0239	0.2175	97.69%	97.64%	97.66%	95.46%
ResNet-50	98.66%	98.39%	97.91%	0.8646	0.0043	97.81%	97.82%	97.81%	95.74%
VGG-16	96.03%	96.77%	96.30%	0.4616	0.4983	96.09%	96.08%	96.02%	92.46%

accuracy of 95.25% compared to the training accuracy of 96.97%, suggesting mild overfitting that might be remedied with better data augmentation, dropout regularization or early stopping strategies you've previously employed.

Across all architectures, the same trend emerges: modern architectures based on inception and exception architectures with residual connections (Inception-ResNet V2, Xception) substantially outperformed both traditional deep CNN (VGG-16) and your customized lightweight solution, with the gap in performance most exaggerated in the IoU metric (99.01% vs. 92.46%), indicative that feature extraction sophistication is key for the type of classification task presented. The loss distributions show that lower training losses were highly correlated with higher testing accuracies (with the exception of Resnet-50) and that a stable validation loss is a good indication of generalization performance. For deployment in production where your only goal is maximum accuracy, Inception-ResNet V2 is clearly the fastest way to achieve this; if you are looking for a tradeoff between accuracy and efficiency, Xception provides a strong-speed + accuracy tradeoff; and if you are targeting an edge application where you are highly resource constrained, your customized lightweight CNN provides an adequate level of accuracy while optimizing for overall size and performance, and future model optimization by regularizing the hyperparameter space may provide further diminishing the performance gap.




Fig. 4: Customized lightweight CNN-Model performance during training and validation: (a) loss curves and (b) accuracy curves

The training curves in Figure 4 indicate serious instability at the start of training, followed by the model programmed to converge. The validation loss shows notable jumps across epochs 2-12 (with a maximum of 12) and this suggests that the model hit a poor local minima or experienced learning rate issues, while the training loss remained quite stable. Most importantly, after the completion of epoch 15, both

training and validation losses begin to settle around 0 with training accuracy plateauing up to 97% and validation accuracy reaching relative high values at 95.25%, indicating generalization performance of the model without overfitting during training process. The fact that the final training and validation metrics had little gap between them (both curves end up tightly converging after epoch 20) suggests that the 1.31M-parameter lightweight CNN was able to learn good features instead of memorizing details found in the training dataset, however, the snapshot of the training process shows that there were occasions in the process of training that showed instability, and it might have been beneficial to use some form of learning rate scheduling or warmup at the start of model training process. In contrast, in Figure 5, the training curves show an extremely smooth and stable convergence for transfer learning with the Inception-ResNet v2 architecture. Both loss curves decline consistently from 0.35 to below 0.02 at epoch 25 with very little fluctuation, which suggests that the pretrained weights from ImageNet provided good initializations that did not land in a local minima. With respect to the accuracy curves, there was rapid learning in the first 5 epochs where the validation accuracy increased from 86% to 97%, which stabilized around 99% for both the training and validation accuracy moving forward at all epochs. The close overlap of both the training and validation metrics across all 25 epochs indicates that there was strong generalization taking place without obvious overfitting in the model, and therefore highlights the value of transfer learning for medical image classification when leveraging deep pretrained architectures.




Fig. 5: Inception-ResNet V2 Model performance during training and validation: (a) loss curves and (b) accuracy curves

The curves displayed in Figure 6 for the Hausdorff distance quickly show improvement of boundary prediction accuracy across all four tumor categories, steadily decreasing from high values of approximately 20 (glioma, meningioma) and 17-18 (notumor, pituitary) at epoch 0 to a near-zero point by

epoch 25. The training and validation curves exhibit a similar convergence down towards values below 1.0 by the end of epoch 30 across all 50 epochs, which supports that the model correctly learned spatial boundary representations without evidence of overfitting. The convergence of all four classes in the same trajectory demonstrates that the CNN model was able to learn generalizable boundary characteristics across multiple tumor morphologies, supported by the validation Hausdorff distances stabilizing around 0.5-1.0 by epoch 40. In Figure 7, the Inception-ResNet-v2 model shows exceptionally smooth and monotonic Hausdorff distance convergence for all four tumor categories, as the training and validation curves steadily decline from values of 20 down to below 2.5 at epoch 25. Indeed, the evolution of the validation curves (dashed lines) follows almost perfectly the trajectory of training curves (solid lines) throughout all epochs, suggesting the pretrained model was able to transfer features learned from ImageNet to the task of medical boundary prediction without overfitting to tumor types. Importantly, the evolution of the Hausdorff distance metric across the four classes, while representing a high range of morphology differences across tumor types, converges simultaneously into similarly smooth trajectories with near identical final values (2.5), with glioma having a slightly higher final Hausdorff distance compared to the other three tumor types.




Fig. 6: Lightweight CNN Hausdorff distance performance: (a) training curves and (b) validation curves.




Fig. 7: Inception-ResNet V2 Hausdorff distance performance: (a) training curves and (b) validation curves.

The Average Surface Distance (ASD) curves in Figure 8 reveal that there was relatively swift convergence from initial values of 5.5-7.5 to near-zero values by epoch 25, with all four tumor classes exhibiting steady gains in boundary prediction accuracy over time. Although the training curves show slightly greater volatility than validation curves (especially for notumor (orange) between epochs 5-15), both training and validation curves converge below 0.5 values by epoch 30, indicating

excellent surface-level agreement between predictions and ground truth. The clustering of all four classes after epoch 20 suggests that the lightweight CNN successfully learned generalizable representations of tumor boundaries, with final ASD values of approximately 0.1-0.3, clinically meaningful for precise tumor delineation in radiotherapy planning.




Fig. 8: Lightweight CNN Average Surface distance performance: (a) training curves and (b) validation curves.

In Figure 9, the Inception-ResNet-v2 model expresses an exceptionally smooth convergence of ASD with all four tumor classes decreasing from initial values of 5.5-7.5 to below 1.0 by epoch 25, demonstrating that the transfer learning efficacy of pretrained ImageNet weights supports bounding and distinguishing precise textures associated with the classification of medical images. The validation curves (dashed lines) consistently follow the trend of training curves within all epochs, with the notumor (orange) class experiencing the fastest convergence and lowest final ASD values of 0.2, while glioma, meningioma, and pituitary reach similar final ASD values of 0.5-0.8. This close alignment with little oscillation during training to validation of all four classes confirms the enabling of robust transfer learning from pretraining on ImageNet without overfitting, to ultimately demonstrate that the deep pretrained architecture effectively captures generalizable surface-level features essential for precise tumor boundary delineation.




Fig. 9: Inception-ResNet V2 Average Surface distance performance: (a) training curves and (b) validation curves.

The confusion matrices provided in Figure 10 reveal an outstanding classification performance in terms of training, validation, and test datasets. The majority of the total predictions along the diagonals indicates high true positive rates (TPR) for all four tumor classes. In the training set, glioma showed total accuracy of 1094/1121 glioma (97.4%), meningioma with 1083/1120 meningioma (93.0%), notumor with 1356/1370 (93.0%), and pituitary with 1234/1249 pituitary (98.8%) correct classifications across the class distributions.

In the test set, classification accuracy of pituitary tumors is 100% and in particularly excellent performance for all other tumor classes with accuracy ranging from (94.3%-99.7%). Overall, the majority of misclassifications occurred between meningioma and all other tumor types. This overlap is especially visible in the training set; meningioma had the most confusions with glioma (21) and then pituitary (55). These confusions suggest there is overlap in subtle MRI features of the meningioma that needs to be explored further in clinical practice.




Fig. 10: Confusion matrices of the lightweight CNN model: (a) training, (b) validation, and (c) testing.

In Figure 11, it is shown that with the Inception-ResNet-v2 model, it achieves almost perfect classification performance for all datasets. The model has training accuracy values of 99.8-99.9% for all four classes, validation accuracy is 99.2-100% (with the 'notumor' category achieving perfect classification of 100%), while accuracy value on the test set is in the range of 99.1-99.6%. The confusion matrices show unusual diagonal dominance and minimal off-diagonal misclassifications (with only 1-2 per class in the training datasets: 1122/1123 glioma, 1172/1174 meningioma, 1367/1370 notumor, 1247/1251 pituitary). In addition, validation datasets showed an overall score of either 100% or 99% with specifically perfect performance on the 'notumor' category. The significantly better performance of the Inception-ResNet-v2 model confirms that transfer learning with deep generalized architectures effectively constructs ImageNet learned features in medical imaging applications by achieving a state-of-the-art results, corroborating current literature reporting 99%+ accuracy on similar datasets and brain tumor classification tasks.




Fig. 11: Confusion matrices of the Inception-ResNet V2 model: (a) training, (b) validation, and (c) testing.

In Figure 12 Grad-CAM visualizations illustrate effective explainability for the lightweight CNN model's predictions. The red heatmap areas indicate where the MRI was most

influential in the classification decisions. The left image (true: glioma, predicted: glioma, confidence: 1.000) demonstrates clear focus and location of activation towards the left hemisphere where the tumor occurs, while the right image (true: notumor, predicted: notumor, confidence: 1.000) shows more evenly distributed activation in the left and right hemispheres, accurately indicating the lack of focal abnormalities. Both predictions confidently and accurately identified the ground truth (as indicated by the 1.000 confidence score), and both heatmaps correspond appropriately to clinically relevant areas with emphasis on the tumor mass in the case of glioma and showing balanced activation in the case of healthy brain.




Fig. 12: Grad-CAM visualization of lightweight CNN predictions.

The Inception-ResNet-v2 Grad-CAM visualizations in figure 13 show greater explainability with the left sagittal view of a pituitary tumor showing concentrated red activation localized to the sella turcica region where pituitary tumors typically occur, and the right axial view of the normal brain showing diffuse bilateral patterns of activation. Additionally, both predictions have complete confidence (1.000); for the pituitary case, it appropriately focuses attention on the midline tumor location, and the notumor case appropriately distributes attention to cortical regions without focal abnormalities. By integrating the ability to visualize different anatomical planes (sagittal and axial), as radiologists would typically examine multiple MRI views, our study drastically improves clinical interpretability as our heatmaps overlap perfectly with clinically relevant locations of a pituitary tumor and notumor, confirming the pretrained model learned sets of features can be used for medical diagnoses.

GradientShap visualizations in Figure 14 demonstrate the remarkable spatial specificity with which the model assigns feature attributions. The differences in the integer range of color gradients between blue and red represent pixel importance for the classification decision, with warmer colors (red) indicating a higher level of importance, particularly if it contributes to the classification decision. In the left image of meningioma (confidence: 0.996), there is focal red activation in the right lateral extra-axial region, where the meningioma class is demonstrated to develop along the meninges, while the right image of pituitary (confidence: 1.000) shows con-




Fig. 13: Grad-CAM visualization of Inception-ResNet V2 predictions.

concentrated red-orange activating in the sellar/suprasellar region and at the skull base where the pituitary gland is located. Additionally, the cooler blue tones across the non-tumor regions indicate low attribution for the model and should confirm that the model is correctly suppressing the false positive signals attributed to healthy brain parenchyma tissue. Lastly, in general, the use of both GradientShap and Grad-CAM in tandem provides a both gradients based explanations for common clinical interpretation but also provides a game theory-grounded explanation for implicates that this visualization enhances trust of these models in translation lessening the black box phenomenon in clinical disparities.




Fig. 14: Lightweight CNN feature attribution using GradientShap.

The Inception-ResNet-v2 GradientShap visualizations presented in Figure 15 illustrate meaningful explanations with perfect certainty (1.000) for both cases. The left meningioma (left image) displays bright red-orange activation that perfectly outlines the right lateral, extra-axial region, which is reflective of the characteristics expected of dural-based tumors, while the right pituitary case (right image) demonstrates similar concentrations of red activation in the midline of the sellar/suprasellar region, along with additional smaller areas of activation in the cerebellum. The deep pretrained architecture is relying on gradient value-based SHAP to approximate representing feature importances relative to the entire input image. In general, the cooler blue tones are activated for lower attribution in

areas of healthy brain parenchyma while vibrant warm colors exhibit tumor-specific anatomical areas of interest. Finally, this method of attribution relies on game-theory in yielding its explanations by computing around multiple baseline distributions expected gradients for explained attribution; thus providing a complementary method of explainability to Grad-CAM that is still robust and axiomatically defensible.




Fig. 15: Inception-ResNet V2 feature attribution using GradientShap.

The ROC curves of lightweight CNN model in Figure 16 show exceptional discriminatory performance for all four tumor classes, with the AUC values almost perfect for glioma (0.999), meningioma (0.994), notumor (0.999) and pituitary (0.999), which indicates that the lightweight CNN model has exceptional ability to discriminate all the classes across all classification thresholds. The curves hugged the upper-left corner of the plot, with slopes sharply rising from the origin and having very few false positive rates to reaching maximum true positive rates, which represents the considerable sensitivity-specificity tradeoff and confirming that the model is accurately flagging true positives while keeping false alarms to a minimum. The curves had significant separation from the diagonal "random" classifier line (dashed grey) across classes, showing that the model's predictions are substantially better than chance with the only slight deviation for meningioma (0.994), due to the nature of its shape relative to other extra-axial lesions, representing extremely strong clinical performance.

The Inception-ResNet v2 model in Figure 17 shows perfect classification performance with AUC = 1.000 for all four tumor classes, as seen by the ROC curves clustering exactly at the upper-left corner of the plot representing 100% true positive with 0% false positive across all classification thresholds. The complete clear separation from the random classifier diagonal with all four curves (glioma, meningioma, notumor, pituitary) overlapping at the perfect corner confirms that the pretrained model achieved perfect classification, successfully diagnosing each tumor and without any false positives or false negatives. This noteworthy performance also outperforms the lightweight CNN (which achieved 0.994-0.999 AUC) and indicates the theoretical extent of ROC analysis by demonstrating that transfer learning using a deep ImageNet-pretrained architecture results in superior feature extraction relevant to medical imaging.




Fig. 16: Multi-class ROC curves of the lightweight CNN with AUC scores.




Fig. 17: Multi-class ROC curves of the Inception-ResNet V2 with AUC scores.

The precision-recall curves for the lightweight CNN in Figure 18 indicate excellent model performance, with all four classes consistently exhibiting near-perfect precision (≥ 0.98) across the entire recall range. The *Average Precision* (AP) values were 0.997 for *glioma*, 0.985 for *meningioma*, 0.998 for *no tumor*, and 0.996 for *pituitary*. The curves remain flat at the peak (precision ≈ 1.0) until recall reaches 0.95–0.98, and then show a steep drop only at the extreme right. This pattern indicates that the lightweight CNN preserves excellent precision even when capturing nearly all positive cases, a behavior with clear clinical relevance where both sensitivity and specificity matter. The small variation between classes (AP range: 0.985–0.998) and the large area under each curve further demonstrate strong discriminative ability for tumor versus non-tumor classes, with the slightly lower AP for meningioma.

The Inception-ResNet-v2 model presented in Figure 19 demonstrates perfect or near-perfect precision-recall perfor-




Fig. 18: Precision-recall curves for the lightweight CNN model.




Fig. 19: Precision-recall curves for the Inception-ResNet V2 model.

mance, with AP scores of 1.000 for glioma, 0.999 for meningioma, 1.000 for notumor, and 1.000 for pituitary. As reflected in the precision-recall curves that remain perfectly flat at precision = 1.0 across the recall spectrum of 0 to 1.0, all four curves overlap almost perfectly at the top of the plot, indicating that the pretrained model maintains maximum precision while detecting all positive cases and corresponding theoretical ideal performance in which there are no false positives at any threshold. This level of performance far exceeds the performance of the lightweight CNN (AP: 0.985–0.998) and demonstrates that deep transfer learning using ImageNet pretrained models has greater discriminatory power for medical imaging, achieving a clinically optimal sensitivity-specificity balance that is critical for the research and diagnostic setting.

VII. COMPARISON WITH OTHER PAPERS

Here, Table IV summarizes prior studies that also used this same dataset. Natha et al. [40] provide a timely and valid measurement of the need for accurate and efficient brain tumor

TABLE IV: ANALYZING PERFORMANCE MEASURES IN COMPARISON TO OTHER TECHNIQUES

Study	Year	Methodology	Accuracy	Notable Contributions
Natha et al.	2024	Stacked ensemble model combining VGG19, InceptionV3, and DenseNet121 with majority voting	98%	Introduced an ensemble method to improve classification robustness
Gómez-Guzmán et al.	2023	Evaluated seven CNN architectures	97.12%	Thorough comparison of CNNs for brain tumor classification
Güler & Ersin	2024	Optimized deep learning classifiers for brain tumor detection	85%	Automated system to assist radiologists and reduce manual effort
Saeed et al.	2024	Compared 10 DCNN models (EfficientNetB0, ResNet50, etc.)	92%	Ensemble model improved classification accuracy
Sandeep et al.	2024	Used pre-trained models adapted for brain tumor diagnosis	99.75%	Showcased the power of deep transfer learning in medical diagnosis
Zahoor & Saddam	2022	Proposed Res-BRNet (residual + regional CNN)	98.22%	Outperformed traditional CNNs in robustness and accuracy
Kumar et al.	2024	Applied six ML algorithms	99%	Used LIME to improve interpretability in healthcare models
Krishnan et al.	2024	Used rotated patch embeddings for rotation invariance	98.6%	Used Grad-CAM for transparent decision explanation
Our Proposed Work (Inception-ResNet V2 & Lightweight Customized CNN)	2025	employs transfer learning with ImageNet-pretrained CNNs for four class brain tumor classification from MRI scans, using CLAHE preprocessing, contour-based brain extraction, comprehensive data augmentation, and Grad-CAM/GradientShap for explainability.	Inception-ResNet V2 = 99.53% & Lightweight Customized CNN = 96.49%	Our work employs multi-architecture transfer learning across six CNN-based architectures: five ImageNet-pretrained models (Inception-v3, ResNet-50, VGG-16, Inception-ResNet-v2, Xception) and a custom lightweight CNN (1.31M parameters) with integrated explainability techniques (Grad-CAM, GradientShap) for four-class brain tumor classification, featuring an end-to-end pipeline with advanced preprocessing, comprehensive evaluation metrics, and clinical-grade visualization.

identification, which is essential for appropriate treatment planning and ultimately the prognosis of patients. Natha et al. suggest an ensemble based deep learning approach that can be supplemented with various deep learning architectures to take advantage of each individual architectural strengths and avoid weaknesses. In this study they proposed different convolutional neural network (CNN) models such as CNNs have performed well in image classification. The results from the experiments showed the SETL_BMRI model to be an accurate and stable model with results revealing an overall accuracy of 98.70%. In addition, an average precision, recall, and F1 score of 98.75%, 98.6% and 98.75%, respectively. Gómez-Guzmán et al. [41] study research proposes a CNN to automate brain tumor identification from MRI tensor production. In this research, the CNN models used was a Generic CNN model, ResNet50, Inception V3, InceptionResNet V2, Xception, MobileNetV2, and EfficientNetB0. The research showed that when considering all CNN models including the generic CNN model as well as 6 pre-trained models, Inception V3 was the best CNN model for this dataset with an average Accuracy of 97.12

Güler & Ersin [42] study explores several popular CNN architectures (VGG16, ResNet50, and InceptionV3) in terms of identifying brain tumor images. They take a methodological approach to tuning these models, tuning hyperparameters and using methods such as data augmentation to improve model generalizability. The main focus of their work is on determining the most effective CNN architecture in terms of accurately identifying and reliably detecting brain tumors. Here, VGG16,

achieved 98.44%; ResNet, achieved 98.81%; InceptionV3, achieved the highest efficiency at 99.30%, and Proposed Stack Ensemble Model at 99.66%. Based on their findings the stack ensemble model, which captures the benefits from the three single CNN architectures, had the most accuracy which allows for improved brain tumor images classification. They also focused on different performance metrics, including: precision, F1-score, Matthews correlation coefficient (MCC), specificity, sensitivity, root mean square error (RMSE), mean squared error (MSE), mean absolute error (MAE), positive predictive value (PPV), negative predictive value (NPV), false positive rate (FPR), false negative rate (FNR); all suggesting their proposed model works well. Saeed et al's [43] model is an ensemble of several deep learning architectures, including Convolutional Neural Networks and other advanced neural network models. An ensemble benefits from different architectures in order to produce a better classification capability than using each model individually. Each deep learning model is preprocessing the MRI images to extract features. The overall performance of the ensemble is better than using one model alone. The ensemble has better accuracy, precision, recall, F1 score that is again suggesting its potential in classifying brain tumors from MRI images. In using the ensemble method, Saeed et al's model had an accuracy of 92%, precision of 90%, recall of 92%, and F1 score of 91% on a 64 BS and 0.0001 LR.

Sandeep et al. [44] evaluate four popular convolutional neural network architectures; ResNet152, VGG19, DenseNet169, MobileNetv3. The models were trained and validated using

images of a publicly available brain MRI dataset from Kaggle which had 7,023 images classified into four classes; glioma, meningioma, pituitary tumors and normal brain images. To help with the imbalance class, a public data set, and improve accuracy of the model, the researchers used data augmentation techniques such as rotation, zooming and flipping. The dataset was split into 80% - 20% of training to testing data set, and a five-fold validation was done to allow for the models to be tested for robustness. Out of all four models tested the MobileNetV3 model had the highest classification accuracy at 99.75%, which was higher than the ResNet152, VGG19, or DenseNet169 architectures. Zahoor & Saddam [45] proposed a new deep learning framework called Res-BRNet which is used to improve brain tumor classification from MRI images. The Res-BRNet model addresses some of the unique problems associated with brain tumors, which has more complex shapes and varying texture, size, and location by using regional and boundary-based structures in spatial and residual blocks. The Res-BRNet model is structured so that the spatial blocks develop features related to tumor heterogeneity, homogeneity, and tumor boundaries (i.e., patterns and edges related to tumors). The residual blocks then learn local and global variations in texture (i.e., low-level features and high-level features) without having to specify the desired features of interest. This model design allows the Res-BRNet model to learn unique discriminative features needed for classification. The Res-BRNet model showed better performance than well-established CNN models with an accuracy of 98.22%, sensitivity of 0.9811, precision of 0.9822, and F1 score of 0.9841.

Three various feature extraction methods were applied by Kumar et al.: Local Binary Patterns (LBP), Histogram of Oriented Gradients (HOG), and Image Loading [46]. These techniques were implemented in two publicly available MRI datasets of brain tumors examined in Kaggle. Six machine learning algorithms were used after feature extraction: Random Forest, Support Vector Machine (SVM), Logistic Regression, K-Nearest Neighbors (KNN), Naive Bayes, and Decision Tree. The results of the study showed that the Random Forest classifier had the highest classification accuracy of 99% for brain tumors when paired with the Image Loading feature extraction method. SVM and Logistic Regression also performed well, though marginally worse than Random Forest, but this combination outperformed the others. Conversely, the KNN, Naive Bayes, and Decision Tree algorithms exhibited comparatively lower accuracy rates, underscoring the significance of selecting appropriate algorithm-feature extraction pairings for optimal classification results. The authors launch the Rotation Invariant Vision Transformer (RViT), a new type of deep learning model designed for classifying brain tumors found on MRI scans [47]. The main contribution of RViT is in the way it creates the rotated variants of the input images, after which each rotated image has been produced at various angles. Afterwards, the process of generating the patches and projecting them linearly into an embedding space continues. All rotated images create embeddings, which means that there are embeddings for all the rotated versions of the original

MRI scans. RViT calculates a mean value of all embeddings and forms a rotation invariant representation of the original image. The strength of this process relates to the architecture learning features invariant to possible rotations, which then assists the model if the tumor is positioned differently in each scan. Overall, the authors achieved good accuracy measures including a sensitivity of 1.0, specificity of 0.975, F1-score of 0.984, Matthew's Correlation Coefficient (MCC) of 0.972, and overall accuracy of 0.986.

In our study, we tested five pretrained CNN-based models, in addition to a lightweight custom-built CNN model (1.31M parameters). Inception-ResNet V2 had the best accuracy in training (99.84%), validation (99.43%) and testing (99.53%), along with the lowest training loss (0.0098) and second-lowest validation loss (0.0228). Our lightweight custom-built CNN, although it had the fewest parameters (1.31M) and lowest complexity, did relatively well, achieving 96.97% training accuracy, 95.25% validation accuracy, 96.49% testing accuracy, with a training loss of 0.1146 and a validation loss of 0.1612.

VIII. LIMITATIONS AND FUTURE WORK

Our study has various limitations, and important directions of future work. Key limitations include that it used a relatively small dataset of only 7,023 MRI images from only three merged sources with no stratified random splitting which may have led to class imbalance, the lightweight custom CNN's synthetic data fallback generated only 400 total samples (100 per class) which may lead to overfitting even with dropout regularization, and the simulated Hausdorff distance and ASD metrics were based purely on predicted masks rather than any true segmentation annotations (ground truth). We observed training instability in the custom CNN during early epochs (2-12) and ResNet-50 had unusual loss dynamics suggesting optimization issues, and the BatchNorm layers may be inducing inference instability in production settings. Our study's most glaring limitations were the lack of prospective clinical validation with annotations generated by radiologist experts, as well as testing on an external dataset, from different hospitals/imaging protocols, and even ground truth annotations, nor did we test performance on true edge deployment scenarios in a resource-constrained hardware context. Future work should include to expand to larger multi-institutional datasets that include diverse MRI sequences and rare tumor subtypes, integrate and apply true pixel level segmentation architectures (e.g. U-Net, Attention U-Net) with expert-annotated ground truth segmentation masks, conduct true prospective clinical trials across multiple hospitals to validate real world diagnostic concordance and inter-rater reliability, and integrate uncertainty quantification approaches (Monte Carlo Dropout and/or Bayesian approaches) to estimate confidence levels.

IX. CONCLUSION

Our research provides a complete deep learning framework for automatically classifying brain tumors that systematically tested six convolutional neural network architectures: VGG-16, Inception V3, ResNet-50, Inception-ResNet V2, Xception,

and a custom lightweight CNN with 7,023 MRI images across four classes (glioma, meningioma, pituitary, no tumor), all designed to minimize confounding variance in a fully standardised training approach. The Inception-ResNet V2 model achieved state-of-the-art results with 99.53% testing accuracy, 99.50% precision/recall/F1-score, and 99.01% IoU, surpassing recent comparators; the Xception model also achieved 98.96% accuracy with maximum computational efficiency (7M trainable parameters), while the custom lightweight CNN (1.31M parameters) achieved 96.49% accuracy that is ideal for deployment in resource-constrained and edge deployment. A major contribution that addresses the clinical black-box problem is the incorporation of dual-pathway explainability, through Grad-CAM and GradientShap, which are explainability techniques that successfully localised anatomically relevant areas of the tumor in images with 100% confidence scores. The framework demonstrated three types of deployment pathways: maximum accuracy (Inception-ResNet V2), equal efficiency (Xception), and edge deployment (lightweight CNN) to show how AI-assisted diagnostics can be democratized across multiple health care delivery systems, from large tertiary care centers to community to rural clinics in developing countries, such as Bangladesh. This clinical-grade, end-to-end system, with advanced preprocessing, robust evaluation metrics (Hausdorff distance, ASD and ROC-AUC), and reproducible methodology, represents a significant step toward trustworthy, accessible, and interpretable AI-assisted neuroradiology that can be reasonably implemented into clinical workflow, and improve patient outcomes through earlier detection and better diagnostic accuracy. Regulatory approval and real-world implementation will require future prospective clinical trials, multi-institutional validation, and integration into hospital information systems.

ACKNOWLEDGMENT

This work's code repository provided by Ahmed Faizul Haque Dhrubo https://github.com/Dhrubo3/Brain_Tumor_IEEE_CIM.

REFERENCES

- [1] A. Cohen-Gadol, M. L. U. March 12, and 2024 S. our M. E. S. for more information about how we maintain the highest standards of medical accuracy, "72 Must-Know Brain Tumor Statistics (2024) — Expert Surgeon," www.aaroncohen-gadol.com. <https://www.aaroncohen-gadol.com/en/patients/brain-tumor/types/statistics>
- [2] Mayo Clinic, "Brain tumor - Symptoms and causes," Mayo Clinic, Apr. 21, 2023. <https://www.mayoclinic.org/diseases-conditions/brain-tumor/symptoms-causes/syc-20350084>
- [3] dhaka-admin, "Glioblastoma Multiforme (GBM) : A Deadly Brain Tumour," Evercare Hospitals Dhaka — Transforming Healthcare, Jun. 07, 2023. <https://www.evercarebd.com/glioblastoma-multiforme-gbm-a-deadly-brain-tumour/> (accessed Sep. 15, 2024).
- [4] A. Jain, M. C. Sharma, V. Suri, S. S. Kale, A. K. Mahapatra, and M. Tatke, "Spectrum of pediatric brain tumors in India: a multi-institutional study," *Neurology India*, vol. 59, no. 2, pp. 208–211, 2011.
- [5] T. Masoodi, R. K. Gupta, J. P. Singh, and A. Khajuria, "Pattern of central nervous system neoplasm: A study of 106 cases," *JK Pract*, vol. 17, pp. 42–46, 2012.

- [6] G. Aryal, "Histopathological pattern of central nervous system tumor: A three year retrospective study," *Journal of Pathology of Nepal*, vol. 1, no. 1, pp. 22–25, 2011.
- [7] Akmalbek Bobomirzaevich Abdusalomov, Mukhriddin Mukhiddinov, and Taeg Keun Whangbo, "Brain Tumor Detection Based on Deep Learning Approaches and Magnetic Resonance Imaging," *Cancers*, vol. 15, no. 16, pp. 4172–4172, Aug. 2023, doi: <https://doi.org/10.3390/cancers15164172>.
- [8] S. Anantharajan, S. Gunasekaran, T. Subramanian, and V. R., "MRI brain tumor detection using deep learning and machine learning approaches," *Measurement: Sensors*, vol. 31, p. 101026, Feb. 2024, doi: <https://doi.org/10.1016/j.measen.2024.101026>.
- [9] Sandeep Kumar Mathivanan, Sridevi Sonaimuthu, S. Murugesan, Harisharan Rajadurai, Basu Dev Shivahare, and Mohd Asif Shah, "Employing deep learning and transfer learning for accurate brain tumor detection," *Scientific reports*, vol. 14, no. 1, Mar. 2024, doi: <https://doi.org/10.1038/s41598-024-57970-7>.
- [10] "Brain Cancer Segmentation Using YOLOv5 Deep Neural Network." Available: <https://arxiv.org/pdf/2212.13599>
- [11] S. Saeedi, S. Rezayi, H. Keshavarz, and S. R. Niakan Kalhori, "MRI-based brain tumor detection using convolutional deep learning methods and chosen machine learning techniques," *BMC Medical Informatics and Decision Making*, vol. 23, no. 1, Jan. 2023, doi: <https://doi.org/10.1186/s12911-023-02114-6>.
- [12] M. Arif, F. Ajesh, S. Shamsudheen, O. Geman, D. Izdrui, and D. Vicoveanu, "Brain Tumor Detection and Classification by MRI Using Biologically Inspired Orthogonal Wavelet Transform and Deep Learning Techniques," *Journal of Healthcare Engineering*, vol. 2022, p. e2693621, Jan. 2022, doi: <https://doi.org/10.1155/2022/2693621>.
- [13] C.-J. Tseng and C. Tang, "An optimized XGBoost technique for accurate brain tumor detection using feature selection and image segmentation," *Healthcare Analytics*, vol. 4, p. 100217, Dec. 2023, doi: <https://doi.org/10.1016/j.health.2023.100217>.
- [14] J. Amin, M. A. Anjum, M. Sharif, S. Jabeen, S. Kadry, and P. Moreno Ger, "A New Model for Brain Tumor Detection Using Ensemble Transfer Learning and Quantum Variational Classifier," *Computational Intelligence and Neuroscience*, vol. 2022, p. e3236305, Apr. 2022, doi: <https://doi.org/10.1155/2022/3236305>.
- [15] A. K. Budati and R. B. Katta, "An automated brain tumor detection and classification from MRI images using machine learning techniques with IoT," *Environment, Development and Sustainability*, Oct. 2021, doi: <https://doi.org/10.1007/s10668-021-01861-8>.
- [16] P. Windisch et al., "Implementation of model explainability for a basic brain tumor detection using convolutional neural networks on MRI slices," *Neuroradiology*, vol. 62, no. 11, pp. 1515–1518, Jun. 2020, doi: <https://doi.org/10.1007/s00234-020-02465-1>.
- [17] T. Rajesh, S. Mani, and M. Geetha, "Brain tumor detection using optimisation classification based on rough set theory," *Cluster Computing*, vol. 22, no. S6, pp. 13853–13859, Mar. 2018, doi: <https://doi.org/10.1007/s10586-018-2111-5>.
- [18] A. Kharrat, Nacéra Benamrane, M. Ben Messaoud, and M. Abid, "Detection of brain tumor in medical images," Nov. 2009, doi: <https://doi.org/10.1109/icscs.2009.5412577>.
- [19] S. Saeed, A. Abdullah, N. Z. Jhanjhi, M. Naqvi, and A. Nayar, "New techniques for efficiently k-NN algorithm for brain tumor detection," *Multimedia Tools and Applications*, Mar. 2022, doi: <https://doi.org/10.1007/s11042-022-12271-x>.
- [20] R. Ahsan, Iram Shahzadi, Faisal Najeel, and H. Omer, "Brain tumor detection and segmentation using deep learning," *Magnetic Resonance Materials in Physics Biology and Medicine*, Sep. 2024, doi: <https://doi.org/10.1007/s10334-024-01203-5>.
- [21] B. Devkota, A. Alsadoon, P. W. C. Prasad, A. K. Singh, and A. Elchouemi, "Image Segmentation for Early Stage Brain Tumor Detection using Mathematical Morphological Reconstruction," *Procedia Computer Science*, vol. 125, pp. 115–123, 2018, doi: <https://doi.org/10.1016/j.procs.2017.12.017>.
- [22] "IJCNS - International Journal of Computer Science and Network Security," *Ijcsns.org*, 2025. http://paper.ijcsns.org/07_book/html/201406/201406015.html (accessed Jun. 04, 2025).

- [23] J. Amin, M. Sharif, M. Raza, and M. Yasmin, "Detection of Brain Tumor based on Features Fusion and Machine Learning," *Journal of Ambient Intelligence and Humanized Computing*, Nov. 2018, doi: <https://doi.org/10.1007/s12652-018-1092-9>.
- [24] A. Anil, A. Raj, A. Sarma, N. Chandran, and Deepa, "Brain Tumor detection from brain MRI using Deep Learning," *International Journal of Innovative Research in Applied Sciences and Engineering (IJIRASE)*, vol. 3, no. 2, pp. 458–465, 2019, doi: <https://doi.org/10.29027/IJIRASE.v3.i2.2019>.
- [25] T. Mahmud et al., "Enhancing Diagnosis: An Ensemble Deep Learning Model for Brain Tumor Detection and Classification," *Lecture notes in networks and systems*, pp. 409–424, Jan. 2024, doi: https://doi.org/10.1007/978-981-99-8937-9_28.
- [26] M. Z. Khaliki and M. S. Başarslan, "Brain tumor detection from images and comparison with transfer learning methods and 3-layer CNN," *Scientific Reports*, vol. 14, no. 1, p. 2664, Feb. 2024, doi: <https://doi.org/10.1038/s41598-024-52823-9>.
- [27] Suhb Irshaidat and Rehab Duwairi, "Brain Tumor Detection Using Artificial Convolutional Neural Networks," *Apr. 2020*, doi: <https://doi.org/10.1109/icics49469.2020.239522>.
- [28] T. S. Kumar, K. Rashmi, S. Ramadoss, L. K. Sandhya, and T. J. Sangeetha, "Brain tumor detection using SVM classifier," *2017 Third International Conference on Sensing, Signal Processing and Security (ICSSS)*, May 2017, doi: <https://doi.org/10.1109/ssps.2017.8071613>.
- [29] T. Hossain, F. S. Shishir, M. Ashraf, M. A. Al Nasim, and F. Muhammad Shah, "Brain Tumor Detection Using Convolutional Neural Network," *IEEE Xplore*, May 01, 2019. <https://ieeexplore.ieee.org/document/8934561>
- [30] H. T. Le and H. T.-T. Pham, "Brain tumour segmentation using U-Net based fully convolutional networks and extremely randomized trees," *Vietnam Journal of Science, Technology and Engineering*, vol. 60, no. 3, pp. 19–25, Sep. 2018, doi: [https://doi.org/10.31276/vjste.60\(3\).19](https://doi.org/10.31276/vjste.60(3).19).
- [31] J. Amin, M. Sharif, M. Raza, T. Saba, R. Sial, and S. A. Shad, "Brain tumor detection: a long short-term memory (LSTM)-based learning model," *Neural Computing and Applications*, Dec. 2019, doi: <https://doi.org/10.1007/s00521-019-04650-7>.
- [32] None Roshan Jahan and M. Tripathi, "Detection of Brain Tumour in MRI Images using Deep Belief Network (DBN)," *Journal of Advanced Research in Applied Sciences and Engineering Technology*, vol. 41, no. 1, pp. 154–167, Mar. 2024, doi: <https://doi.org/10.37934/araset.41.1.154167>.
- [33] Raouia Ayachi and Nahla Ben Amor, "Brain Tumor Segmentation Using Support Vector Machines," *Lecture notes in computer science*, pp. 736–747, Jan. 2009, doi: https://doi.org/10.1007/978-3-642-02906-6_63.
- [34] Zoltán Kapás, L. Lefkovits, and László Szilágyi, "Automatic Detection and Segmentation of Brain Tumor Using Random Forest Approach," *Lecture notes in computer science*, pp. 301–312, Jan. 2016, doi: https://doi.org/10.1007/978-3-319-45656-0_25.
- [35] Rajeev Kumar Gupta, Santosh Bharti, Nilesk Kunhare, Y. Sahu, and Nikhlesh Pathik, "Brain Tumor Detection and Classification Using Cycle Generative Adversarial Networks," *Interdisciplinary Sciences: Computational Life Sciences*, vol. 14, no. 2, pp. 485–502, Feb. 2022, doi: <https://doi.org/10.1007/s12539-022-00502-6>.
- [36] N. Ullah et al., "An Effective Approach to Detect and Identify Brain Tumors Using Transfer Learning," *Applied Sciences*, vol. 12, no. 11, p. 5645, Jan. 2022, doi: <https://doi.org/10.3390/app12115645>.
- [37] "Module: tf.keras.applications — TensorFlow Core v2.9.0," *TensorFlow*. https://www.tensorflow.org/api_docs/python/tf/keras/applications
- [38] "Papers with Code - Leaky ReLU Explained," *paperswithcode.com*. <https://paperswithcode.com/method/leaky-relu>
- [39] "Brain Tumor MRI Dataset," *www.kaggle.com*. <https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset>
- [40] S. Natha, U. Laila, I. A. Gashim, K. Mahboob, M. N. Saeed, and K. M. Noaman, "Automated Brain Tumor Identification in Biomedical Radiology Images: A Multi-Model Ensemble Deep Learning Approach," *Applied Sciences*, vol. 14, no. 5, p. 2210, Jan. 2024, doi: <https://doi.org/10.3390/app14052210>.
- [41] M. A. Gómez-Guzmán et al., "Classifying Brain Tumors on Magnetic Resonance Imaging by Using Convolutional Neural Networks," *Electronics*, vol. 12, no. 4, p. 955, Jan. 2023, doi: <https://doi.org/10.3390/electronics12040955>.
- [42] Mustafa Güler and Ersin Namlı, "Brain Tumor Detection with Deep Learning Methods' Classifier Optimization Using Medical Images," *Applied sciences*, vol. 14, no. 2, pp. 642–642, Jan. 2024, doi: <https://doi.org/10.3390/app14020642>.
- [43] Z. Saeed, Tarraf Torfeh, Souha Aouadi, J. X. Ji, and Othmane Bouhali, "An Efficient Ensemble Approach for Brain Tumors Classification Using Magnetic Resonance Imaging," *Information*, vol. 15, no. 10, pp. 641–641, Oct. 2024, doi: <https://doi.org/10.3390/info15100641>.
- [44] Sandeep Kumar Mathivanan, Sridevi Sonaimuthu, S. Murugesan, Harisharan Rajadurai, Basu Dev Shivahare, and Mohd Asif Shah, "Employing deep learning and transfer learning for accurate brain tumor detection," *Scientific reports*, vol. 14, no. 1, Mar. 2024, doi: <https://doi.org/10.1038/s41598-024-57970-7>.
- [45] M. Mumtaz Zahoor and S. Hussain Khan, "Brain tumor MRI Classification using a Novel Deep Residual and Regional CNN." Available: <https://arxiv.org/pdf/2211.16571>
- [46] K. Kumar, K. Jyoti, and K. Kumar, "Machine learning for brain tumor classification: evaluating feature extraction and algorithm efficiency," *Discover Artificial Intelligence*, vol. 4, no. 1, Dec. 2024, doi: <https://doi.org/10.1007/s44163-024-00214-4>.
- [47] Palani Thanaraj Krishnan, Pradeep Krishnadoss, M. Khandelwal, D. Gupta, Anupoj Nihaal, and T. Sunil Kumar, "Enhancing brain tumor detection in MRI with a rotation invariant Vision Transformer," *Frontiers in neuroinformatics*, vol. 18, Jun. 2024, doi: <https://doi.org/10.3389/fninf.2024.1414925>.