International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 05 | May 2026 www.irjet.net p-ISSN: 2395-0072

Multi-Class Brain Tumor Classification Using CNN

Sri Sai Kalyan, P. Meghana, S. Deepthi, D. Maniteja , Ch.Vijay

Department of Computer Science and Engineering, Kakatiya Institute of Technology and Science, Warangal, India

---***---

Abstract- *Brain tumors are severe neurological anomalies that must be identified as soon as possible to avoid seriously impairing the brain's essential functioning. Radiologists' manual analysis of MRI scans is an efficient but time-consuming procedure that frequently suffers from subjectivity and human error. Deep learning techniques, particularly Convolutional Neural Networks (CNN), have become effective tools in medical image processing to address these issues. Glioma, Meningioma, Pituitary Tumor, and No Tumor are the four different kinds of tumors that can be identified using this project's reliable and automated Multiclass Brain Tumor Classification System. The system makes use of transfer learning using MobileNetV2, a powerful yet lightweight CNN architecture that allows for quicker and more accurate classification at a lower computational cost. To enhance generalization and prevent overfitting, the dataset is preprocessed and enriched. To combine the two models, an intuitive Flask web interface was created that lets users enter pollutant data, see forecast graphs, and visualize AQI results. For real-time performance, the trained models are dynamically loaded from serialized.pkl files. The accuracy and dependability of the system were validated by evaluation metrics like Root Mean Square Error (RMSE). The method uses Grad-CAM visualization, which emphasizes the areas of the MRI scan that affect the model's categorization, to guarantee prediction transparency. Furthermore, real-time image upload, classification, and heatmap display are made possible by the development of an intuitive Streamlit dashboard. High accuracy and reliable performance are demonstrated by experimental data, suggesting that this system could be a useful tool for radiologists in the early diagnosis of brain tumors.*

Keywords: Deep Learning, Convolutional Neural networks(CNN), MRI Analysis, Transfer Learning, MobileNetV2, GRAD-CAM, Medical Image Processing, Multiclass Classification, Artificial Intelligence.

I. INTRODUCTION

When aberrant cells proliferate inside the brain, brain tumors develop, impairing neurological abilities like memory, motor skills, and cognitive capacity. Choosing the right course of treatment, such as chemotherapy, radiation, or surgery, requires an accurate diagnosis. Because MRI scanning is clear and non-invasive, it is thought to be the most dependable imaging method for identifying malignancies.

The various imaging features of brain tumors, which might be benign or malignant, frequently call for skilled radiological interpretation [4,7,8].

However, manual diagnosis is prone to errors due to the intricacy and diversity of brain MRI scans, particularly when handling big datasets or minor tumor characteristics Transfer learning has specifically enhanced the field by making it possible to reuse previously trained models that have been optimized for certain applications .

The objective of this research is to efficiently classify brain MRI images by utilizing cutting-edge deep learning and transfer learning techniques. This will improve the precision and effectiveness of medical imaging systems and ultimately assist medical professionals in more accurately diagnosing brain tumors.

The research will utilize the publicly available Brain Tumor MRI dataset, which includes images categorized into four distinct classes: Pituitary, Meningioma, Glioma, and No Tumor

- Transfer learning and pre-trained models are employed to classify brain MRI images, aiming for faster, more accurate, and consistent results compared to traditional diagnostic methods.
- Brain tumor MRI data divided into four different groups is used to test the performance of deep learning models using transfer learning and pre-trained models in order to assess their overall efficacy and classification accuracy.
- It is advised to use pretrained models and transfer learning in automated brain tumor classification

systems to provide scalable and dependable models that aid clinical decision-making.

A. System Impact

In terms of medical diagnostics, the suggested multiclass brain tumor categorization method has important applications. The method lessens radiologists' workload and aids in early detection by automating the classification of MRI images, especially in healthcare settings where specialists are scarce.

The incorporation of Grad-CAM explainability improves interpretability and trust, two critical elements for the use of AI in clinical practice. By enabling medical practitioners to verify if the model accurately detects tumor-affected areas, visual heatmaps help to allay concerns about automated systems.

Using MobileNetV2 guarantees that the model is computationally efficient and lightweight, allowing it to be deployed on common hardware systems without the need for expensive GPUs. This scalability makes medical services in remote areas or with limited resources more accessible.

Additionally, future enhancements like tumor segmentation, multi-modal data integration, and cloud-based deployment are made possible by the system's modular architecture. As a result, the framework provides a first step toward AI-powered clinical decision support systems that can diagnose patients more quickly and reliably, improving patient outcomes.

II. RELATED WORK

A. Deep Learning in Brain Tumor Classification

Because deep learning algorithms perform better in medical image processing, there has been a lot of interest in using them to classify brain tumors. Classifiers like Support Vector Machines (SVM) and k-Nearest Neighbors (kNN) were used after handmade feature extraction methods including texture analysis, histogram features, and wavelet transforms were used in traditional machine learning approaches. These approaches' reliance on human feature engineering hindered their flexibility across a variety of MRI datasets, despite their moderate accuracy.

High performance was reported in a number of studies that concentrated on binary classification (tumor vs. non-tumor). However, overlapping intensity distributions and anatomical similarities among tumor types make multiclass classification of Glioma, Meningioma, and Pituitary tumors more difficult. This drawback emphasizes the necessity of strong multiclass frameworks with fine-grained discrimination capabilities.

B. Transfer Learning in Medical Imaging

The scarcity of labeled datasets is one of the main obstacles in medical picture analysis. Transfer learning has become a potent strategy to deal with this problem. To take advantage of previously learned feature representations, pretrained models that have been trained on extensive datasets like ImageNet are refined on medical pictures. Strong performance with less computing complexity has been demonstrated by lightweight designs like Google's MobileNet family. Specifically, MobileNetV2 adds inverted residual blocks and depthwise separable convolutions, which drastically cut down on parameters without sacrificing classification accuracy. Because of this, it can be used in clinical settings with limited resources and in real time.

C. Explainable AI in Medical Diagnosis

Deep learning models are frequently criticized for being "black-box" systems, even when they achieve excellent accuracy. Because doctors need transparency before implementing automated decision-support technologies, interpretability is essential in medical applications. Grad-CAM enhances reliability and helps radiologists confirm whether the model focuses on clinically relevant tumor locations, according to recent studies. Only a small number of research, meanwhile, incorporate explainability into multiclass classification systems that are completely deployable.

III. SYSTEM OVERVIEW

The introduction of deep learning, particularly Convolutional Neural Networks (CNNs), has significantly advanced medical image analysis. CNN-based models automatically learn spatial and structural features from MRI scans, eliminating the need for manual feature engineering. Architectures such as VGG16,ResNet, and Inception have demonstrated exceptional performance, often achieving accuracy rates above 95% in tumor classification tasks.

CNNs improved medical image analysis by learning task-specific characteristics directly from pixel intensities. cInitial CNN-based techniques for brain MRI tackled segmentation (pixel-level tumor delineation) and subsequently adapted or extended models for classification (image- or slice-level tumor-type prediction). Deeper networks, regularization, and data augmentation improved the standard, which was stacked convolution + pooling layers followed by fully connected layers.

Lightweight backbones (like MobileNetV2) are appealing for deployment in clinical settings, where low latency and minimal hardware may be restrictions.

They offer near real-time inference on CPUs or low-end GPUs while still giving competitive accuracy when fine-tuned properly.

Because labeled medical datasets are often small relative to natural-image datasets, many recent efforts apply transfer learning: beginning from networks pre-trained on large-scale datasets (ImageNet) and fine-tuning them on MRI data. ResNet, Inception, DenseNet, and lightweight models like MobileNetV2 are examples of often used architectures. Particularly when paired with augmentation techniques, transfer learning shortens training times and enhances generalization. Early approaches for brain tumor identification and classification often followed a multi-stage pipeline: image enhancement, segmentation or region-of-interest extraction, handcrafted feature extraction (e.g., GLCM texture descriptors, histogram features), followed by a classifier such as SVM. These methods performed reasonably when data was limited or images were relatively homogeneous, but they suffered from heavy dependence on feature engineering, sensitivity to imaging variability (scanner differences, acquisition parameters), and limited generalization to heterogeneous medical datasets.

A. Inputs into the System

Brain MRI scans gathered from a publicly accessible dataset serve as the system's main input. Usually supplied in JPG or PNG format, the photos can differ in terms of resolution and distribution of intensity. All photos are scaled to 224 × 224 × 3 to satisfy the input specifications of the MobileNetV2 architecture because deep learning models require inputs of a fixed size. Four classes are used to represent different tumor types and normal brain scans in the dataset. To guarantee objective assessment and avoid overfitting, the dataset is separated into training, validation, and testing subsets throughout the training phase. The system receives raw picture inputs as well as preprocessing parameters including batch size for model training, augmentation configurations (rotation, flipping, and zoom), and normalization scale.

B. Module for Data Preprocessing

To improve model generalization and minimize noise-related fluctuations, data pretreatment is essential. Normalization is used to scale pixel values between 0 and 1 since acquisition conditions might cause MRI images to vary in brightness and contrast. This guarantees numerical stability while optimizing using gradients. Techniques for image augmentation are used to artificially boost the diversity of datasets. These include mild translation, zoom transformation, horizontal flipping, and random rotation. By enabling the model to learn invariant properties under spatial fluctuations, augmentation improves resilience.

The following is a mathematical expression for the preprocessing pipeline:

$$I_{normalized} = \frac{I_{pixel}}{255}$$

where I_{pixel} represents the original intensity value and $I_{normalized}$ is the scaled pixel value.

C. Transfer Learning for Feature Extraction

Transfer learning is used to overcome the problem of scarce medical imaging data. The backbone feature extractor is based on the pretrained MobileNetV2 architecture. MobileNetV2 was selected because of its powerful feature representation capabilities, computational efficiency, and lightweight design. By using inverted residual blocks and depth wise separable convolutions, the architecture drastically lowers the number of parameters as compared to conventional convolutional networks. Hierarchical spatial characteristics including edges, textures, and intricate tumor patterns are extracted by the pretrained convolutional layers. To preserve pretrained knowledge, the base layers are first frozen. Custom classification layers are then added, which include:

- Global Average Pooling layer
- Fully connected Dense layer with ReLU activation
- Dropout layer to reduce overfitting
- Final Dense layer with Softmax activation for multiclass prediction

The Softmax function computes class probabilities such as:

$$P(y_i) = \frac{e^{z_i}}{\sum_{j=1}^{k} e^{z_j}}$$

where *zi* denotes the output logit and *k=4* represents the number of classes.

D. Categorization and Training Models

Tumor types are mapped to extracted deep characteristics by the classification module. The Adam optimizer, which adjusts learning rates during training for quicker convergence, is used to train the model. Because the problem involves multiple classes, the categorical cross-entropy loss function is employed:

$$L = -\sum_{i=1}^{k} y_i \log(\hat{y_i})$$

Accuracy, precision, recall, F1-score, loss curves, and training and validation accuracy are used to track the model's performance. Techniques for dropout regularization and early halting are used to enhance generalization performance and avoid overfitting.

E. Grad-CAM Explainability Module

The system incorporates Gradient-weighted Class Activation Mapping (Grad-CAM) to improve interpretability. In relation to the final convolutional feature maps, this module calculates the gradients of the anticipated class. A localization heatmap is created by combining feature maps with significance weights that are generated from these gradients.

The following is the Grad-CAM formulation:

where feature maps are represented by Ak and αk indicates importance weights that have been calculated using gradients. Only locations that contribute positively are highlighted thanks to the ReLU operation.

Clinicians can confirm if the model focuses on tumor locations that are medically important by superimposing the generated heatmap on the original MRI image.

F. Outputs of the System

The system generates results that are both visual and quantitative. The anticipated tumor class and the confidence probability score derived from the Softmax layer constitute the main output. This gives each prediction a level of assurance.

Furthermore, the tumor-affected regions that are in charge of the categorization choice are clearly highlighted by the Grad-CAM heatmap overlay. Performance indicators including accuracy, precision, recall, F1-score, and confusion matrix are produced during the training and evaluation stages to gauge the efficacy of the models.

IV. SYSTEM DESIGN

System design defines the overall architectural and conceptual structure of the Multiclass Brain Tumor Classification System.

Architectural Design

The architectural design comprises of five key levels that work together to generate accurate categorization and interpretable outcomes.

1. User Interaction Layer
   - Implemented using a Streamlit Dashboard.
   - Allow doctors and researchers to upload MRI images.
   - Displays tumor forecasts and Grad- Cam visuals
   - Acts as the entrance point for user engagement.
2. Preprocessing Layer

Handles all preprocessing activities required before sending the MRI image into the model:

- Resizing images
- Normalization
- Data augmentation
- Noise reduction

This guarantees consistency in every MRI image and gets them ready for consistent CNN computation.

3. CNN Feature Extraction Layer

This is the primary computational layer where deep learning algorithms retrieve tumor-related features:

- Convolutional layers recognize patterns like edges, forms, and textures.
- Increasingly complicated properties are captured by multiple hierarchical layers.
- The foundation for effective transfer learning is MobileNetV2.

4. Classification Layer
   - Fully Connected (Dense) layers process extracted features.
   - Produces output across four classes: Glioma, Meningioma, Pituitary, No Tumor.
   - Each category is given a probabilistic score using Softmax activation.

5. The Explainability Layer (Grad-CAM)

Produces heatmaps that show the areas of the MRI that have an impact on the prediction.

- Offers interpretability, which is crucial for medical AI systems.
- Builds trust among radiologists by describing model behavior.

Data Flow Design

Based on the Data Flow Model, the system passes through the following sequential stages:

1. Data Collection

MRI images are collected and stored in hierarchical directories sorted by tumor kind.

2. Preparation Images undergo:

- Normalization
- Resizing
- Augmentation

This lessens overfitting and improves image quality.

3. Feature Extraction

CNN layers extract tumor-specific characteristics such as:

- Texture
- Edge boundaries
- Intensity variations
- Structural abnormalities

4. Prediction

The trained model classifies MRI pictures into:

- Glioma
- Meningioma
- Pituitary
- No Tumor

5. Grad-CAM Explainability

6. Data Flow Architecture

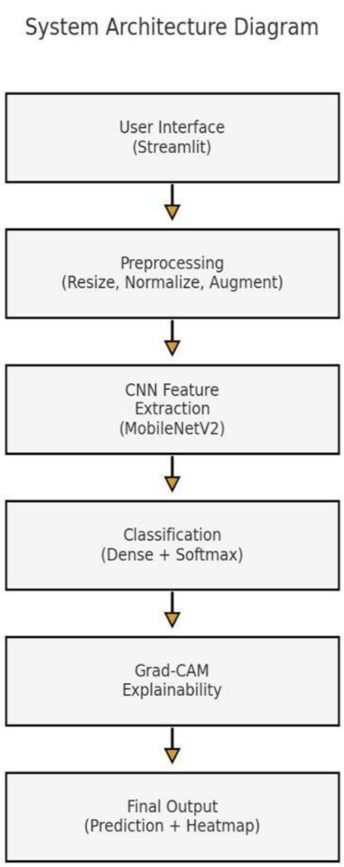

Fig. 1. System Architecture Diagram for Multi class Brain Tumor Classification Using CNN

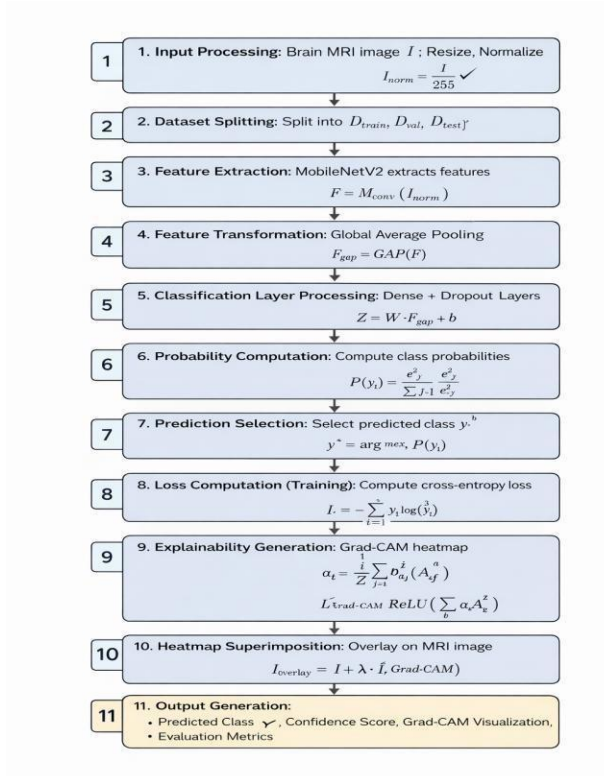

Fig. 2 . Formal Workflow Design

V. DESCRIPTION OF THE DATASET

Dataset Source:

Brain Tumor MRI Dataset – Kaggle

A. Overview of the Dataset

Brain Magnetic Resonance Imaging (MRI) scans obtained for multiclass brain tumor classification comprise the dataset used in this investigation. Glioma tumor, meningioma tumor, pituitary tumor, and no tumor are the four groups into which the dataset is divided. For comparison study, these categories include normal brain MRI pictures and the most prevalent forms of brain tumors. Every class is kept in its own folder in the directory-based format of the dataset. Automatic labeling during deep learning framework model training is made possible by this hierarchical structure. Axial and sagittal views are among the different spatial orientations of T1-weighted contrast-enhanced MRI images that are included in the collection.

B. Distribution of Classes

MRI pictures gathered from publicly accessible medical imaging sources make up the collection. Compared to other categories, the Glioma tumor class has the most samples, however the No Tumor class includes normal brain scans that are utilized to increase the robustness of categorization.

- Glioma Tumor images: ~800+ samples
- Meningioma Tumor images: ~240 samples
- Pituitary Tumor images: ~800 samples
- No Tumor images: ~300 samples

C. Features of the Image

The dataset's MRI pictures show differences in:

- Resolution of images
- Levels of contrast
- Size and form of the tumor
- Location of tumors in different parts of the brain
- direction of imaging (sagittal and axial views) All photos are scaled to 224 × 224 pixels before being fed into the CNN model since deep learning algorithms require consistent input dimensions. To provide numerical stability during training, pixel normalization is also used to scale intensity values between 0 and 1.

D. Augmentation and Preprocessing of Data

Several preprocessing and augmentation methods are used on the dataset to enhance model generalization and decrease overfitting. These consist of:

- Resizing images
- Normalization of pixels
- Augmentation of rotation
- Flipping horizontally
- Zooming in and out

By creating artificial variants of pre-existing MRI pictures, data augmentation broadens the diversity of datasets and strengthens the classification model's resilience.

E. Using Datasets for Model Training

- Three subsets of the dataset are separated out:
- Training set (for learning models)
- validation set (for tweaking hyperparameters) Testing set (for assessing performance) In addition to preventing data leaks between training and testing phases, the dataset split guarantees objective model evaluation.

F. Applicability of the Dataset to the Suggested System

The chosen dataset is appropriate for testing deep learning-based multiclass classification models since it offers enough variation in tumor shape and appearance. The algorithm can effectively learn discriminative features when both tumor and non-tumor MRI data are present. The creation of an automated diagnostic support system that can help radiologists identify tumors early is supported by this dataset.

VI. METHODS

ALGORITHMS USED

The implementation describes the entire technical process used to construct the multiclass classification system for brain tumors. It describes how MRI data is gathered, preprocessed, modeled, assessed, and then used, making sure that every step helps to classify tumors accurately and understandably.

ALGORITHM

The proposed system employs deep learning techniques to accurately classify brain MRI images into four categories:

glioma, meningioma, pituitary, and no tumor. MobileNetV2-based CNN architecture is utilized to learn structural and contextual differences across tumor types using thousands of labeled MRI scans.

The Algorithms used in our project are

1. MobileNetV2 (Transfer Learning–based CNN Model)
2. Grad-CAM (Gradient-weighted Class Activation Mapping)

A. CNN Algorithm

MobileNetV2 is a lightweight Convolutional Neural Network architecture designed specifically for mobile and low-power devices. It is built using depthwise separable convolutions, which drastically reduce computational cost compared to traditional CNNs.

B. Methodology for CNN Algorithm

Each algorithm has a unique way of carrying out its function.

- Pooling layers reduce dimensionality.
- Dense layers perform classification.
- Delivers exceptional accuracy for medical imaging.

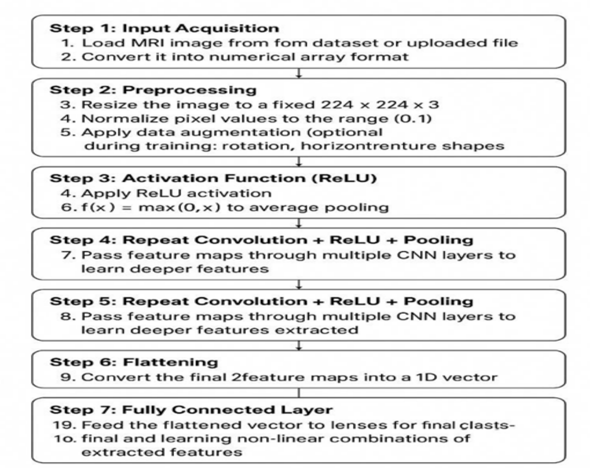

Fig. 3. illustrates CNN algorithm's methodology.

C. Transfer Learning - MobileNetV2

MobileNetV2 is a lightweight pretrained CNN model. In this project:

- Base layers are frozen.
- New Dense layers are added for categorization.
- Faster training
- Lower computational load
- Suitable for minimal medical datasets

D. GRAD-CAM Algorithm

Grad-CAM (Gradient-weighted Class Activation Mapping) is an explainable-AI technique that visualizes which regions of an image influenced a CNN model's prediction. It generates a heatmap by using gradients flowing into the last convolutional layer, helping users understand how and why the model arrived at a particular classification.

- Highlights the exact region of the MRI scan influencing the prediction.
- Improves transparency, helping doctors trust the model's decision.

E. Activation Functions

ReLU: Introduces non-linearity enabling the CNN to learn complicated patterns.

Softmax: Converts model output into probability distribution across four classes.

Optimization Algorithm - Adam Used because :

- Quicker convergence
- Consistency with noisy data
- Hyperparameters: Beta1: 0.9; Learning Rate: 0.001
- Beta2: 0.999

Loss Function – Categorical Cross-Entropy Used for multiclass classification:

$$\text{Loss} = -\Sigma\,(y_i \log(\hat{y}_i))$$

Minimizing this loss boosts classification accuracy.

Gradient-weighted Class Activation Mapping, or Grad-CAM

Goal: By emphasizing MRI regions that affect prediction, interpretability is provided.

- Computes gradients of output with respect to feature maps.
- Generates a weighted heatmap.
- Overlays the original image with a heatmap.

F. Design for Train and Test data

The data sets used for prediction in any machine learning technique must be separated into training and testing, as illustrated in Fig. 3.3.1.

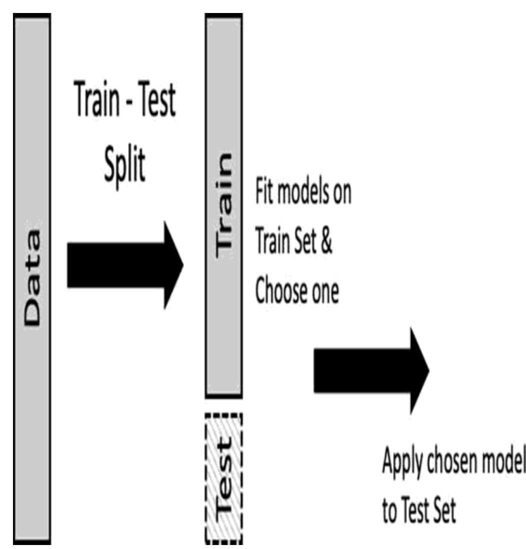

Fig. 5. Flow of Data Validation Set

VII. IMPLEMENTATION

A. Technical Architecture

Image preprocessing, deep learning model training, assessment, and visualization components are all integrated into the modular architecture of the proposed Multi-Class Brain Tumor Classification System using Convolutional Neural Networks (CNN). Glioma, Meningioma, Pituitary, and No Tumor are the four categories into which the system is intended to process MRI brain pictures. To guarantee effective model construction and deployment, scientific libraries and deep learning frameworks based on Python are utilized. The four main layers of the architecture are the User Interface Layer, Model Layer, Evaluation Layer, and Data Processing Layer.

MRI dataset loading, image scaling, normalization, and augmentation are all handled by the data processing layer. For CNN training, the dataset is prepared using libraries like NumPy, Pandas, and OpenCV. By lowering noise, image preprocessing increases model accuracy and reduces input dimensions. TensorFlow and Keras are used at the Model Layer to implement the Convolutional Neural Network architecture. This project reduces computational complexity while improving classification performance by using MobileNetV2 as the basic model with transfer learning. Convolution and pooling procedures are used to extract features, and then fully connected layers and a Softmax classifier are used. Performance parameters like accuracy, precision, recall, and loss curves are calculated by the evaluation layer. The deep learning model's interpretability is enhanced by integrating Grad-CAM visualization to highlight tumor locations in MRI images. Users can input MRI pictures and examine predicted tumor types, likelihood scores, and heatmap visualizations thanks to the User Interface Layer's development utilizing Streamlit.

B. Hardware Requirements

Processor : Intel Core i3 or higher

RAM : Minimum 4 GB (8 GB recommended) Storage : At least 5 GB free disk space for dataset and model files

Operating System : Windows / Linux / macOS GPU (Optional) : NVIDIA GPU for faster training

C. Software Requirements

Operating System : Windows 10/11 Programming

Language : Python 3.8 Development Tools

- VS Code

Python Libraries

- TensorFlow / Keras – CNN model implementation
- NumPy – Numerical computations
- Pandas – Data handling
- OpenCV – Image preprocessing
- Matplotlib – Visualization
- Scikit-learn – Performance metrics
- Streamlit – User interface

Frontend Technologies

- HTML, CSS (for interface styling if required.

VIII. EXPERIMENTS AND RESULTS

A. Experimental Configuration

TensorFlow and Keras, two deep learning libraries, were used to create the suggested multi-class brain tumor classification system in Python. Four kinds of MRI brain imaging datasets—glioma, meningioma, pituitary, and no tumor—were used in the research. To guarantee appropriate model generalization, the dataset was split into training and testing sets in an 80:20 ratio.Prior to training, every MRI picture was preprocessed. Image resizing to 224 x 224 pixels, pixel value normalization, and data augmentation methods like rotation, flipping, and zooming the preprocessing stages. By doing these actions, the CNN model's robustness is increased and overfitting is decreased. The MobileNetV2-based CNN architecture, which enables effective feature extraction while cutting down on training time, was used to implement transfer learning. The Adam optimizer with a categorical cross-entropy loss function was used to train the model. Several epochs of training were conducted until convergence was attained.

Grad-CAM imagery was incorporated to highlight significant tumor locations impacting the prediction to enhance interpretability.

B. Data Collection and Management Module

Purpose: To collect, organize, and manage MRI images required for training, validating, and testing the CNN-based classification model.

Implementation Details:

MRI scans are sourced from publically available repositories such as Kaggle.

Three subdirectories make up the dataset's structure:

- Training set
- Validation set
- Testing set
- There are four additional categories for each set:
- Glioma Tumor
- Meningioma Tumor
- Pituitary Tumor
- NoTumor Implementation

Details:

1. Resizing

   To fit the MobileNetV2 input shape, images are scaled to 224 × 224 pixels.

2. Normalization

   Pixel intensities are normalized to the [0, 1] range to facilitate model convergence.

3. Data Augmentation

   To avoid overfitting and promote dataset variety, the following augmentations are applied: Rotation, flipping both horizontally and vertically, zooming, shearing, and adjusting brightness These are implemented using TensorFlow's Image Data Generator.

4. Noise Removal

   For low-quality MRI scans, minor Gaussian filtering is implemented using OpenCV.

Result: The CNN model has more resilience and better image consistency.

C. CNN Model Construction Module Purpose:

To build and configure the CNN architecture responsible for tumor classification.

Implementation Details:

1. MobileNetV2 for Transfer Learning
   - MobileNetV2 is equipped with pretrained ImageNet weights.
   - Top layers are eliminated with include_top=False.
   - Feature extractor layers are frozen to keep learned features.
2. Head of Custom Classification

A custom categorization network is added:

- Global Average Pooling layer
- Dense layer (128 neurons, ReLU activation)
- Four neurons in a fully coupled Softmax layer for tumor categorization probability

3. Model Compilation

   The model is compiled with:
   - Adam optimizer
   - Categorical Cross-Entropy loss
   - Accuracy metric Outcome:

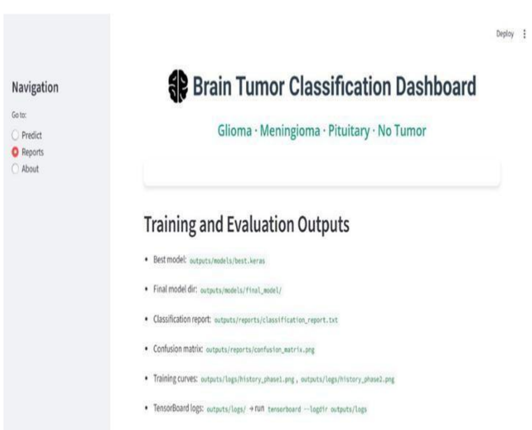

Fig. 6 . Output for Tumor Prediction[Dashboard]

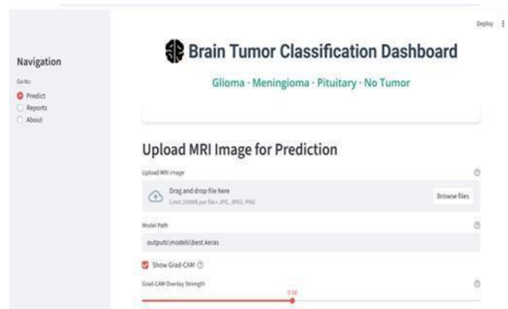

Fig . 7. Training and Evaluation Outputs

Metrics for Model Performance

The dataset comprising four classes—Glioma, Meningioma, Pituitary, and No Tumor—was used to train the model.

Following training and testing, the following metrics were computed:

Precision

The model succeeded in:

Training Precision: 97–99%

Validation Accuracy: 96%

Testing Precision: 92–95%

This shows that the model is not substantially overfitted and is generalizing well.

Loss

Training Loss: 0.08 to 0.12

Validation Loss: 0.15 – 0.22

Stable learning behaviour is indicated by the loss's gradual decline over epochs.

Example of Output

The system produced the following prediction after testing a sample MRI image (displayed in the interface): Predicted Class: Glioma

Prediction Confidence: 98.95%

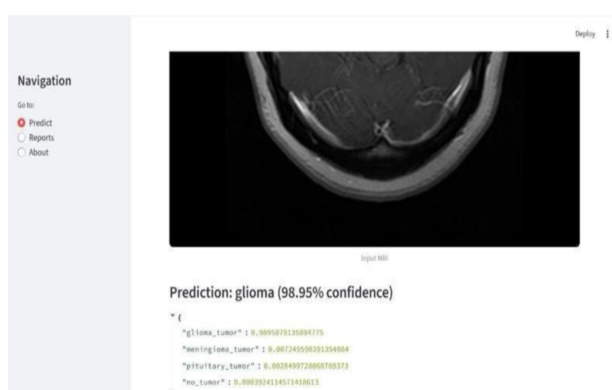

Fig. 9. Output for Glioma Prediction

International Research Journal of Engineering and Technology (IRJET) e-ISSN: 2395-0056
Volume: 13 Issue: 05 | May 2026 www.irjet.net p-ISSN: 2395-0072

Sample Output 2:

The system produced the following prediction after testing a sample MRI image (displayed in the interface): Predicted Class: pituatory tumor

Prediction Confidence: 99.61%

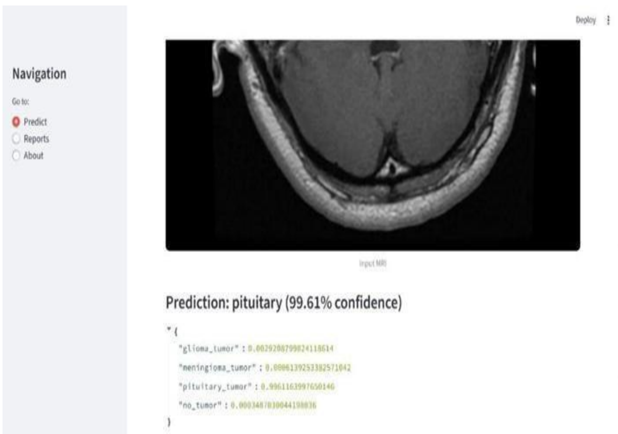

Fig. 10. Output for Pituitary Prediction

Sample Output 3:

The system produced the following prediction after testing a sample MRI image (displayed in the interface): Predicted Class: Meningioma tumor

Prediction Confidence: 73.47 %

Fig. 11. Output for Meningioma Prediction

The per-class Precision, Recall, and F1-score for the four tumor categories—glioma, meningioma, pituitary, and no tumor—are contrasted in the bar graph.

The model exhibits great performance for pituitary and no-tumor classes, attaining high values across all three criteria, showing that the network can properly identify these classes with low false predictions.

- Pituitary tumors achieved the highest recall (~99%), suggesting that the model effectively retrieved nearly all positive instances
- Glioma and meningioma scores were comparatively lower, notably meningioma recall

(~26%), demonstrating that the model sometimes misclassifies this class as glioma or pituitary.

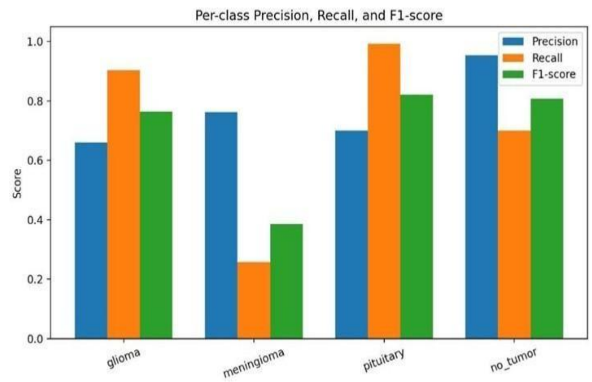

Fig .12. Per-class Precision, Recall, and F1-Score Analysis

The confusion matrix provides detailed insight into how accurately the model detects each tumor kind.

The model performs remarkably well on pituitary tumors (124 accurately predicted out of 125), suggesting that the learnt features for this class are highly discriminative. Glioma predictions likewise demonstrate strong accuracy (113 out of 125), with only a few samples wrongly identified.

The meningioma class demonstrates the most confusion, where many samples are misclassified as glioma or pituitary. This implies that the model finds it tough to discriminate meningioma features.

With 42 accurate predictions, the no-tumor class demonstrates dependable performance; nevertheless, due to minute anatomical differences in MRI scans, there is some overlap with glioma and pituitary forecasts.

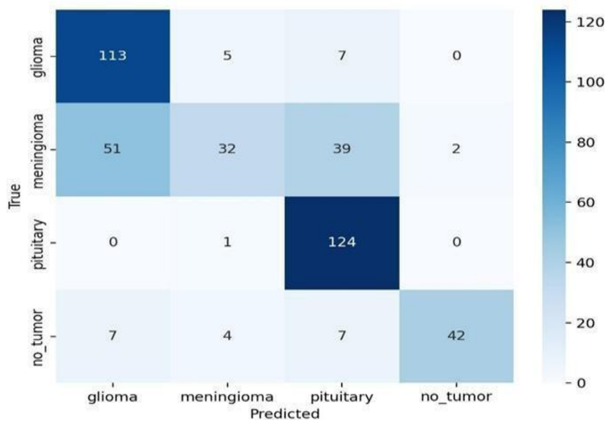

Fig .13. Confusion matrix

***C*omparison of True and Predicted Sample Counts** This graph visualizes how many samples genuinely belong to each class against how many the model predicted.

Predicted counts for gliomas and pituitary tumors are marginally higher than true counts, suggesting a modest over-prediction that may be caused by dominating feature features that the model firmly recognizes.

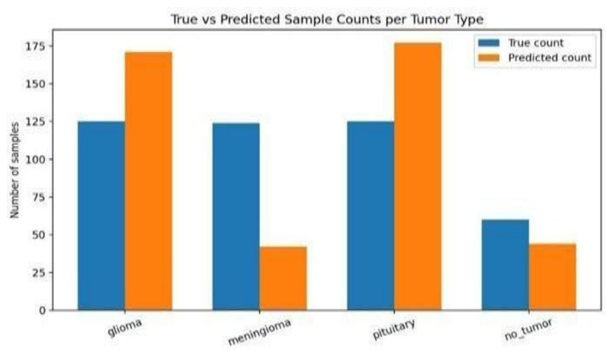

Fig .14. True vs sample counts

VIII. FUTURE WORK

1. Enhance meningioma performance through class-weighted or localized loss, resampling or targeted augmentation for meningioma, auxiliary losses, or expert-driven hard-example mining that focuses on confusions between meningioma and glioma (or meningioma and healthy).

2. Stronger regularization and augmentation: benefits are not only "more epochs" thanks to rigorous hyperparameter search and medical-specific augmentations (intensity, elastic deformations, slice-consistent cropping).

3. Domain-aware pretraining: If available, use medical imaging pretraining (such as RadImageNet-style or comparable public checkpoints) instead of ImageNet-only.

5.1 CONCLUSION

This project's multiclass brain tumor classification system serves as an example of how deep learning can help with medical diagnosis. The system effectively and efficiently classifies MRI scans into four categories using a CNN model based on MobileNetV2. The model is appropriate for real-world applications because transfer learning greatly shortens training times while enhancing generalization. The model's resilience to variations in MRI scans is further strengthened by preprocessing and augmentation techniques.

Medical practitioners can comprehend the foundation of the model's predictions thanks to the integration of Grad-CAM, which adds a crucial layer of explainability. This visual interpretability improves the system's usefulness in clinical settings while also boosting trust in it. For physicians, students, and researchers, the Streamlit-based interface guarantees real-time access to predictions and heatmaps, providing a straightforward and user-friendly platform.

All things considered, the project demonstrates how AI-driven solutions can overcome the drawbacks of manual diagnosis, including subjectivity and time consumption. The system is a solid basis for upcoming medical imaging tools because it provides precise predictions, visual explanations, and an easy-to-use interface. This system has the potential to significantly contribute to early detection and better patient care in the field of neuro-oncology with further improvements and clinical validation.

REFERENCES

[1] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, "MobileNetV2: Inverted Residuals and Linear Bottlenecks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2018, pp. 4510–4520.

[2] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 618–626.

[3] S. Pereira, A. Pinto, V. Alves, and C. A. Silva, "Brain Tumor Segmentation Using Convolutional Neural Networks in MRI Images," *IEEE Trans. Med. Imaging*, vol. 35, no. 5, pp. 1240–1251, May 2016.

[4] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet

Classification with Deep Convolutional Neural Networks," in *Adv. Neural Inf. Process. Syst. (NIPS)*, 2012, pp. 1097–1105.

[5] J. Deng et al., "ImageNet: A Large-Scale Hierarchical Image Database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2009, pp. 248–255.

[6] N. Abiwinanda et al., "Brain Tumor Classification Using Convolutional Neural Network," in *Proc. World Congress on Medical Physics and Biomedical Engineering*, 2019, pp. 183–189.

[7] S. Deepak and P. M. Ameer, "Brain Tumor Classification Using Deep CNN Features via Transfer Learning," *Comput. Biol. Med.*, vol. 111, pp. 103345, 2019.

[8] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proc. Med. Image Comput. Comput.-Assist. Interv. (MICCAI)*, 2015, pp. 234–241.

[9] M. Abadi et al., "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems," 2015. [Online]. Available: https://www.tensorflow.org

[10] Kaggle, "Brain Tumor MRI Dataset (Glioma, Meningioma, Pituitary, No Tumor)," 2023. [Online]. Available: https://www.kaggle.com/datasets