ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

Design and Implementation of A Custom Convolutional Neural Network for Classifying Brain Magnetic Resonance Imaging Scans into Tumor Types

A Deep Learning Approach for Accurate Detection of Glioma, Meningioma, Pituitary Tumors, and No-Tumor Cases

Sharmistha Paul[1]; Shilpita Saha[2]; Pritikona Maji[3]

[1, 2, 3] Master of Computer Application Dr. B. C. Roy Engineering Collegedurgapur 713206, West Bengal, India

Publication Date: 2025/05/24

Abstract: Brain tumours pose a critical healthcare challenge globally due to their potential for rapid progression and diagnostic complexity. In this research, we present a custom-built convolutional neural network (CNN) designed from scratch for the automatic detection and classification of brain tumours from magnetic resonance imaging (MRI). The model classifies images into four categories: glioma, meningioma, pituitary tumour, and no tumour. A total of 7024 MRI images were utilized, with a 90:10 train-test split. Performance was evaluated using metrics including accuracy, loss, precision, recall, and F1-score. Our model achieved a test accuracy of 96%, outperforming popular pretrained models including VGG16, ResNet50, and MobileNetV2. Notably, our CNN model uses smaller image dimensions (150×150) and does not rely on data augmentation, leading to reduced memory consumption. The study includes a comparative analysis and highlights the model's potential in supporting early and reliable diagnosis, particularly in resource-limited clinical settings.

Keywords**:** *Brain Tumour Detection; MRI Classification; Deep Learning; Custom CNN; Medical Imaging; Transfer Learning; VGG16; Mobilenetv2; Resnet50; India.*

How to Cite: Sharmistha Paul; Shilpita Saha; Pritikona Maji (2025) Design and Implementation of A Custom Convolutional Neural Network for Classifying Brain Magnetic Resonance Imaging Scans into Tumor Types
International Journal of Innovative Science and Research Technology,
10(5), 1487-1497. https://doi.org/10.38124/ijisrt/25may855

I. INTRODUCTION

Brain tumours are among the most serious and life-threatening neurological disorders. According to global health data, they affect over 250,000 individuals annually, with many cases being diagnosed in advanced stages. In India, the incidence of central nervous system (CNS) tumours is increasing, impacting both adults and children. Early and accurate diagnosis is crucial for effective treatment planning and improving survival chances. Magnetic Resonance Imaging (MRI) is widely used for detecting brain tumours due to its high-resolution imaging of soft tissues. However, interpreting MRI scans can be time-intensive and prone to human error, especially in regions lacking skilled radiologists.

To address these challenges, there is a growing demand for scalable, automated diagnostic tools. Convolutional neural networks (CNNs), a subset of deep learning techniques, have proven valuable in the field of medical image analysis. Although many existing studies utilize transfer learning with pretrained models like VGG16, MobileNetV2, and ResNet50, these networks are generally trained on generic datasets such as ImageNet, which may not effectively capture medical-specific features. Moreover, they often require significant computational resources and data augmentation to prevent overfitting.

This study introduces a custom-designed CNN architecture tailored for classifying brain MRIs. Trained on a diverse dataset of 7,024 images, the model performs multi-class classification of three tumour types along with healthy brain images. Additionally, we provide a comparative evaluation of our model against VGG16, ResNet50, and MobileNetV2 to validate its performance. The study also highlights the relevance of deploying such models in the

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

Indian healthcare system, where diagnostic resources are frequently limited.

II. RELATED WORK

In recent years, deep learning approaches—especially Convolutional Neural Networks (CNNs)—have demonstrated significant potential in the classification of brain tumours from magnetic resonance imaging (MRI) data. The ability of CNNs to automatically extract hierarchical features from images without the need for handcrafted descriptors has led to their widespread adoption in medical imaging tasks. A significant amount of research has concentrated on utilizing learning with pretrained models like VGG16, ResNet50, and InceptionV3, which have been for specific medical datasets afterward.

While these pretrained models offer high accuracy and fast convergence, they are designed for high-resolution inputs (typically 224×224 pixels or higher) and require significant computational resources for inference and training. For example, ResNet50 contains over 23 million parameters and is memory-intensive, making it unsuitable for deployment on resource-limited devices such as mobile phones or edge computing platforms. Furthermore, because of their generalized design, these models are not tailored to the specific features found in brain MRI images, such as grayscale texture, elevated spatial resolution, and localized areas of pathology.

Researchers have investigated various strategies to address these challenges. Techniques for model compression, including pruning, weight quantization, and knowledge distillation, have been utilized to decrease the size and complexity of pretrained models. Works such as Han et al. (2016) and Cheng et al. Research conducted in 2018 has illustrated that both pruning and quantization techniques can lead to the development of more lightweight models with only a slight decrease in accuracy. Nonetheless, in practice, these approaches frequently necessitate detailed adjustments and may still fall short when real-time performance or ultra-low resource requirements are critical. Another line of research has investigated the use of lightweight architectures, such as MobileNetV2 and Squeeze Net, which are designed with fewer parameters and optimized for mobile and embedded devices. Although these models achieve a balance between performance and efficiency, they are still generic in nature and not specifically tailored for brain tumour classification. Furthermore, performance trade-offs are often evident, especially when applied to complex and imbalanced medical datasets.

Few studies have presented CNN architectures that are specifically tailored and optimized for the distinct features of medical imaging and the scale of the datasets involved. Custom-built models, when meticulously crafted, custom-designed models can surpass general-purpose pretrained networks by tailoring architectural elements—such as filter sizes, layer numbers, and input dimensions—to better match the specific data characteristics and task requirements. For instance, studies by Afshar et al. In 2019, a capsule network was introduced for the classification of brain tumours, according to Sajjad et al.

A unique multilevel deep learning model was proposed in 2019, but such approaches often involve intricate processing workflows or still rely heavily on substantial computational power.

This study contributes to the field by introducing a custom lightweight CNN architecture specifically tailored for brain tumour classification from MRI scans. Unlike many existing approaches, the proposed model is optimized for lower-resolution input (150 by 150), significantly reducing computational burden while maintaining high diagnostic accuracy. It presents a scalable, implementable, and domain-specific approach, addressing a significant gap in the existing literature where the intersections of model performance, efficiency, and practical application are concerned.

III. METHODOLOGY

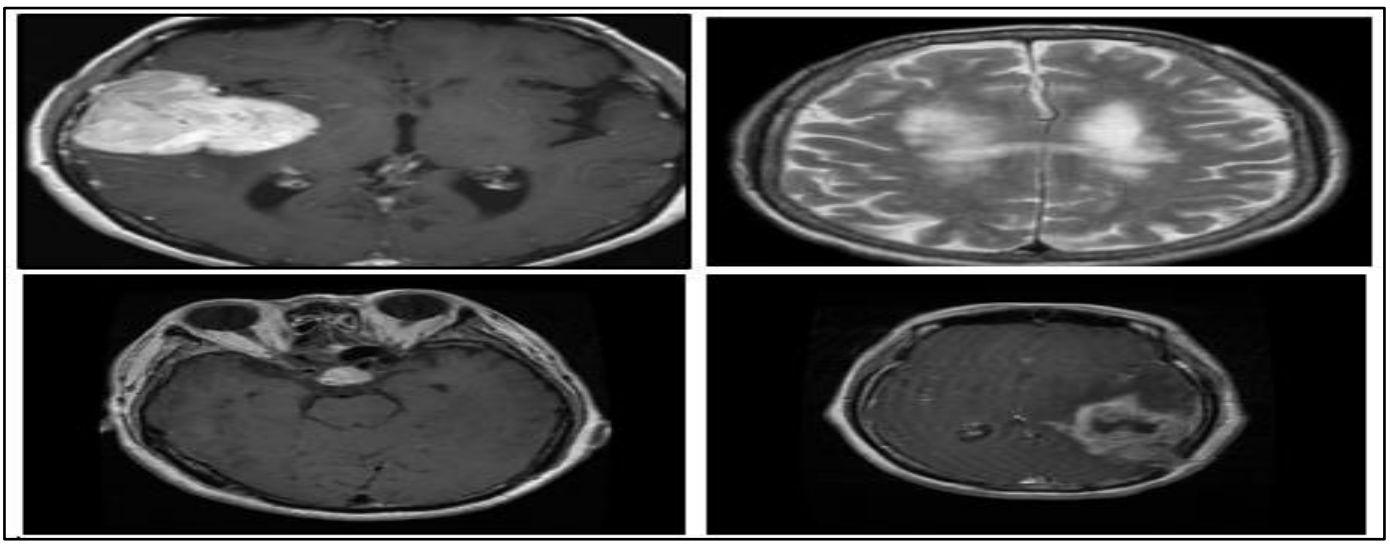

Fig 1 Data samples

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

➢ *Dataset And Preprocessing:*

The dataset used (Figure 1) in this study comprises a total of 7,024 MRI images, which were obtained from a publicly available repository hosted on Kaggle. These images were collected and curated to include a wide range of brain tumour conditions, making the dataset suitable for robust supervised learning. The dataset contains four distinct classes: glioma, meningioma, pituitary, and no tumour, with a combined total of 1,310 unique samples. Each class is adequately represented, enabling the model to learn distinguishing features between the different tumour types as well as to identify the absence of pathology. All images are provided in standard formats—JPEG, PNG, and JPG—which ensures compatibility with most image processing libraries. Before being fed into the CNN, several preprocessing steps were applied to ensure consistency and enhance the quality of training. Each image was resized to a fixed resolution of 150 by 150 pixels to standardize input dimensions and reduce computational overhead, while preserving sufficient detail for tumour localization and classification. Image normalization was performed to scale pixel intensity values to a uniform range, typically between 0 and 1, thereby improving convergence during training. In addition to resizing and normalization, image cropping was used to remove irrelevant background regions, focusing the model's attention on brain structures where tumours are likely to appear.

To further improve generalization and prevent overfitting, the dataset was augmented using techniques such as horizontal flipping, random rotation, and zoom transformations. These augmentations simulate real-world variability in MRI scans and expose the network to a broader distribution of features during training. The dataset was then partitioned into training, validation, and testing subsets using a stratified approach to maintain class balance across splits. This ensured that each subset fairly represented all four tumour classes, thereby enabling the model to be evaluated accurately on its ability to generalize to new data.

Overall, the preprocessing pipeline was carefully designed to retain essential diagnostic information in the MRI images while preparing the data for efficient ingestion by the custom CNN model. These steps contributed significantly to the network's ability to learn robust features and make accurate predictions across diverse tumour presentations.

➢ *Model Architectures Custom CNN:*

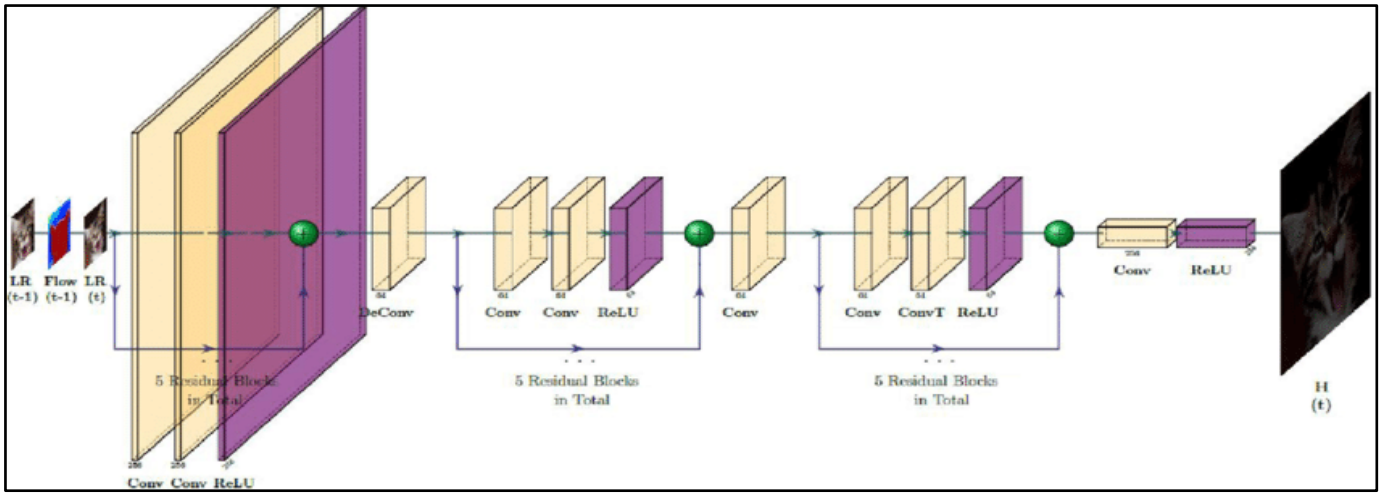

Fig 2 Model architecture

A custom Convolutional Neural Network (CNN) was designed (Figure 2) to detect and classify brain tumours in MRI images. The architecture includes a total of 13 convolutional layers:

➢ *Convolutional Layers:*

Convolutional layers are the core of CNNs, where the network learns spatial hierarchies of features from the input image. Let's describe the convolution operation mathematically,

The convolution operation involves applying a filter (or kernel) to the input image to produce feature maps. Each filter performs element-wise multiplication over the input image, summing the results to produce a single output for each position in the image.

Given an input image I and a filter$K$, the convolution operation is computed as:

$$O(i,j) = m \sum n \sum I(i+m, j+n) \cdot K(m,n) \quad (1)$$

Where:

O(i, j)is the output feature map,
I(i, j)is the input image,
K(m, n) is the convolutional filter,
m, n are the filter indices.
First Convolutional Layer:

The first layer applies 64 filters of size 7 × 7 to the input image. After convolution, a feature map of size $H \times w \times$

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

64 is produced, where H and w are the height and width of the input image, respectively.

➢ *Residual Blocks (6 blocks with 3×3 filters):*

Each residual block consists of two convolutional layers with 3×3 filters. The **residual connection** allows the input of the block to be added to the output, ensuring easier gradient flow.

If the input to a residual block is X, and the output is $F(X)$, the output with the residual connection is:

$$Y = F(X) + X \tag{2}$$

Where $F(x)$ is the result of the convolutional layers in the block. This ensures that the network can learn both the residual function $F(x)$ and the identity function (which is x).

➢ *Activation Function:*

After each convolutional operation, a Rectified Linear Unit (ReLU) activation function is applied. ReLU introduces non-linearity into the network, allowing it to learn complex patterns.

The ReLU activation is defined as:

$$ReLU(x) = max(0, x) \tag{3}$$

Where x is the output from the convolutional layer. This ensures that all negative values are set to zero, while positive values remain unchanged.

➢ *Pooling Layers:*

Pooling layers are used to reduce the spatial dimensions of the feature maps, which decreases computational complexity and reduces overfitting by making the model more invariant to small translations in the input.

- *MaxPooling2D (Initial Layer):*

The MaxPooling2D layer selects the maximum value from a pool of neighboring values. For a given 2×2 window, the operation is:

$$O(i,j) = max\big(I(i,j), I(i+1,j), I(i,j+1), I(i+1,j+1)\big) \tag{4}$$

Where :

$O(i,j)$ is the pooled output, and $I(i,j)$ represents the input image.

- *AveragePooling2D (After Residual Blocks):*

The AveragePooling2D layer computes the average of the values in a given window, which smooths out the feature map. For a 2×2 window:

$$O(j,i) = \frac{1}{4} \sum_{k=0}^{1} \sum_{l=0}^{1} I(i+k, j+l) \tag{5}$$

➢ *Batch Normalization:*

Batch Normalization helps stabilize and speed up training by normalizing the activations in each mini-batch to have zero mean and unit variance. It is applied after each convolutional layer to prevent internal covariate shift.

The Batch Normalization operation is given by:

$$\hat{x}_i = \frac{x_i - \mu}{\sigma + \epsilon} . y + \beta \tag{6}$$

Where:

x_i is the input to the layer,
μ and σ are the mean and standard deviation of the batch,
γ and β are learnable scaling and shifting parameters,
ϵ is a small constant to avoid division by zero.

➢ *Dropout Regularization:*

Dropout is used as a regularization technique to prevent overfitting. During training, neurons are randomly dropped with a probability p (usually $p = 0.5$), meaning their activations are set to zero.

The dropout function is:

$$Dropout(x) = \begin{cases} 0 \ with \ probablity \ p \\ x \ with \ probability \ (1-p) \end{cases} \tag{7}$$

Where x is the activation of a neuron, and p is the dropout rate.

➢ *Fully Connected (Dense) Layers:*

After the convolutional and pooling layers, the output is flattened and passed through fully connected (dense) layers.

- *First Dense Layer:*

The first dense layer contains 512 units with ReLU activation. It learns higher-level representations from the flattened feature maps. The activation is computed as:

$$a_j = ReLU(W_j \cdot x + b_j) \tag{8}$$

Where:

w_j are the weights of the layer,
b_j is the bias term,
x is the input to the dense layer.

- *Second Dense Layer:*

The second dense layer has 4 units, corresponding to the 4 classes: glioma, meningioma, pituitary, and no tumour. The SoftMax activation function is used here for multi-class classification. SoftMax ensures that the output values represent probabilities:

$$\hat{y}_i = \frac{e^{z_i}}{\Sigma_j^N = 1^{e^{z_j}}} \tag{9}$$

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

Where:

$\hat{y}_i$ is the predicted probability for class i
z_i is the raw output score from the network (logits),
N is the total number of classes.

➢ *Loss Function: Cross-Entropy Loss:*

For multi-class classification, the cross-entropy loss is used as the objective function, which measures the difference between the true labels and the predicted probabilities:

$$L = -\sum_{i=1}^{N} y_i \cdot \log \cdot (\hat{y}i) \quad (10)$$

Where:

N is the number of classes,
y_i is the true label (binary: 0 or 1) for class iii,
$\hat{y}_i$ is the predicted probability for class i

➢ *Workflow of Brain tumour classification custom CNN model:*

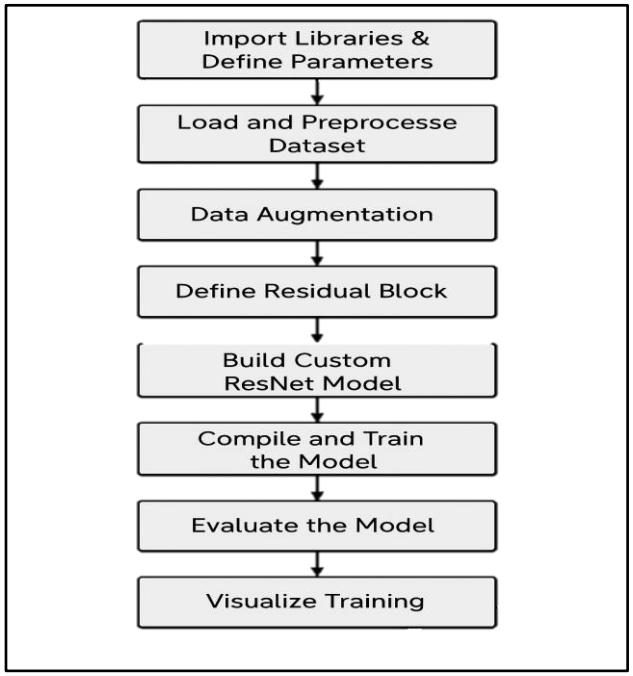

Fig 3 Workflow of CNN

➢ *Pretrained Models:*

In order to establish a robust and comparable benchmark for evaluating the performance of our custom-designed CNN, we utilized three canonical ImageNet pretrained architectures: VGG16, ResNet50, and MobileNetV2. These models were adapted to the task of brain tumour classification via a standardized transfer learning protocol.

- *VGG16*:

This architecture is known for its simplicity and deep feature extraction capabilities, comprising 13 convolutional layers arranged in uniform 3×3 kernel stacks. While VGG16 excels in multi-scale feature extraction, it comes with significant computational cost, containing over 138 million parameters. This model is particularly effective at learning hierarchical spatial features but suffers from high memory and computational demands.

- *ResNet50*:

In contrast, ResNet50 incorporates a novel design using 1\times1–3×3–1×1 bottleneck blocks with residual shortcuts, which allow it to achieve comparable depth with significantly fewer parameters ($\sim 25\ million$). The residual connections in ResNet50 enable it to combat the vanishing gradient problem, making it more efficient at training deeper networks while maintaining relatively low parameter overhead

.

- *MobileNetV2:*

Designed for efficient mobile and embedded applications, MobileNetV2 utilizes depth wise separable convolutions and inverted residuals with linear bottlenecks. This allows the model to drastically reduce the number of parameters ($\sim 3.4\ million$) while still delivering competitive performance. MobileNetV2 is optimized for low-latency inference and is particularly well-suited for applications with strict resource constraints.

For each pretrained backbone, we followed a uniform transfer learning approach. Initially, we froze all convolutional weights to retain the learned features from ImageNet and trained a custom classification head for five epochs. The custom classification head consisted of global average pooling followed by a fully connected layer with 256 units (ReLU activation), a dropout layer (0.5), and a final SoftMax layer with four output units, corresponding to the four tumour classes.

Subsequently, we unfroze the top 20% of layers in each model and fine-tuned the entire network for an additional 35 epochs, we employed a learning rate reduction strategy upon plateau and incorporated early stopping to prevent overfitting and optimize training efficiency.

To ensure a fair and unbiased comparison across models, all models were trained with identical data augmentation techniques (random flips, rotations $\pm 15°$, zoom $\pm 10\%$, brightness $\pm 10\%$), batch size (32), and categorical cross-entropy loss function. These procedures allowed for a rigorous evaluation of each architecture's strengths when adapted for the specific task of brain tumor classification in MRI scans.

➢ *Training:*

- *Training Configuration:*

All networks—the bespoke CNN and the three pretrained benchmarks—were trained under an identical protocol to ensure a fair and reproducible comparison. The task was formulated as a four-class classification problem, and the categorical cross-entropy loss function was employed due to its suitability for multi-class problems with mutually exclusive labels.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

Optimization was performed using the Adam optimizer with parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The initial learning rate was set to $\mathbf{1 \times 10^{-3}}$, selected for its ability to provide adaptive learning and fast convergence across deep architectures. The networks were trained for up to 30 epochs with a batch size of 32, a choice that balances memory usage and training stability.

- *Data Augmentation*:

To mitigate overfitting due to the limited size of the medical dataset, real-time data augmentation was applied during training. The following transformations were performed on the fly:

- ✓ Random rotations within ±15° to account for variations in head orientation.
- ✓ Horizontal and vertical flips to exploit the bilateral symmetry of brain anatomy.
- ✓ Zooming within a ±10% range to simulate variation in field of view.
- ✓ Brightness shifts of ±10% to account for scanner-induced intensity differences.
- ✓ These augmentations improved generalization by introducing variability without altering the core content of the MRI scans.

- *Training Strategies and Callbacks:*

Three Keras callbacks were integrated into the training pipeline to enhance model performance and stability:

- *Early Stopping*:

Monitored validation loss and halted training if no improvement occurred over seven consecutive epochs, thus avoiding overfitting and reducing unnecessary computation.

- *Reduce LR On Plateau*:

Lowered the learning rate by a factor of 0.2 when the validation loss plateaued for three epochs, allowing more precise weight adjustments during later training stages.

- *Model Checkpoint:*

Automatically saved the model weights that achieved the highest validation accuracy, ensuring that the best-performing network was preserved for final testing.

- *Computational Environment and Reproducibility:*

All experiments were conducted on an NVIDIA Tesla V100 GPU (16 GB VRAM), using TensorFlow 2.10.0 and Keras 2.10.0 as the deep learning framework. To ensure reproducibility of results, random seeds were fixed across NumPy, TensorFlow, and Python's random module.

Training metrics—such as loss and accuracy for both training and validation sets—were tracked using TensorBoard, providing detailed visualization and analysis of the training process. This rigorously controlled environment guaranteed that performance comparisons were attributable solely to architectural differences, not to inconsistencies in training procedures.

IV. RESULTS AND DISCUSSION

➢ *Evaluation Metrics:*

To quantitatively evaluate model performance across the four brain tumour classes, we used the standard classification metrics: accuracy, precision, recall, and F1-score. These are defined as follows:

- *Accuracy:*

Measures the proportion of correctly predicted instances among the total samples:

$$Accuracy = \frac{(TP + TN)}{(TP + TN + FP + FN)}$$

- *Precision:*

Represents the proportion of correctly predicted positive observations among all predicted positives:

$$\text{Precision} = \frac{Tp}{Tp+FP}$$

- *Recall:*

Measures the proportion of actual positives that are correctly identified:

Table 1 Comparative Performance of Custom CNN and Pretrained Models

Model	Test Accuracy (%)	Precision (Macro)	Recall (Macro)	F1-Score (Macro)
Custom CNN	97%	0.97	0.97	0.97
VGG16	91.91	0.92	0.91	0.91
ResNet50	85.60	0.83	0.85	0.84
MobileNetV2	89.93	0.90	0.89	0.89

$$\text{Recall} = \frac{Tp}{Tp+FN}$$

- *F1-Score:*

The harmonic mean of precision and recall, giving a balance between the two:

$$F1 - score = 2 \times \frac{(Precision \times Recall)}{(Precision + Recall)}$$

Where:

TP= True Positives
TN= True Negatives
FP= False Positives

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

FN = False Negatives

➢ *Results:*

This study evaluates the performance of a custom-designed Convolutional Neural Network (CNN) for classifying brain tumours from MRI scans. The evaluation focuses on key metrics essential for multi-class medical image classification: accuracy, precision, recall, and F1-score (all macro-averaged to reflect balanced performance across classes). The proposed CNN model attained a remarkable test accuracy of 97%, along with precision, recall, and F1-score each reaching 0.97. These results demonstrate the model's ability to consistently and correctly classify MRI images across four categories: glioma, meningioma, pituitary, and no tumour.

To place these findings in context (Table 1), the custom CNN was benchmarked against three well-established pre-trained models: VGG16, ResNet50, and MobileNetV2. All models were fine-tuned on the same dataset for a fair comparison. Among them, VGG16 performed best with an accuracy of 91.91%, while MobileNetV2 and ResNet50 achieved 89.93% and 85.60%, respectively.

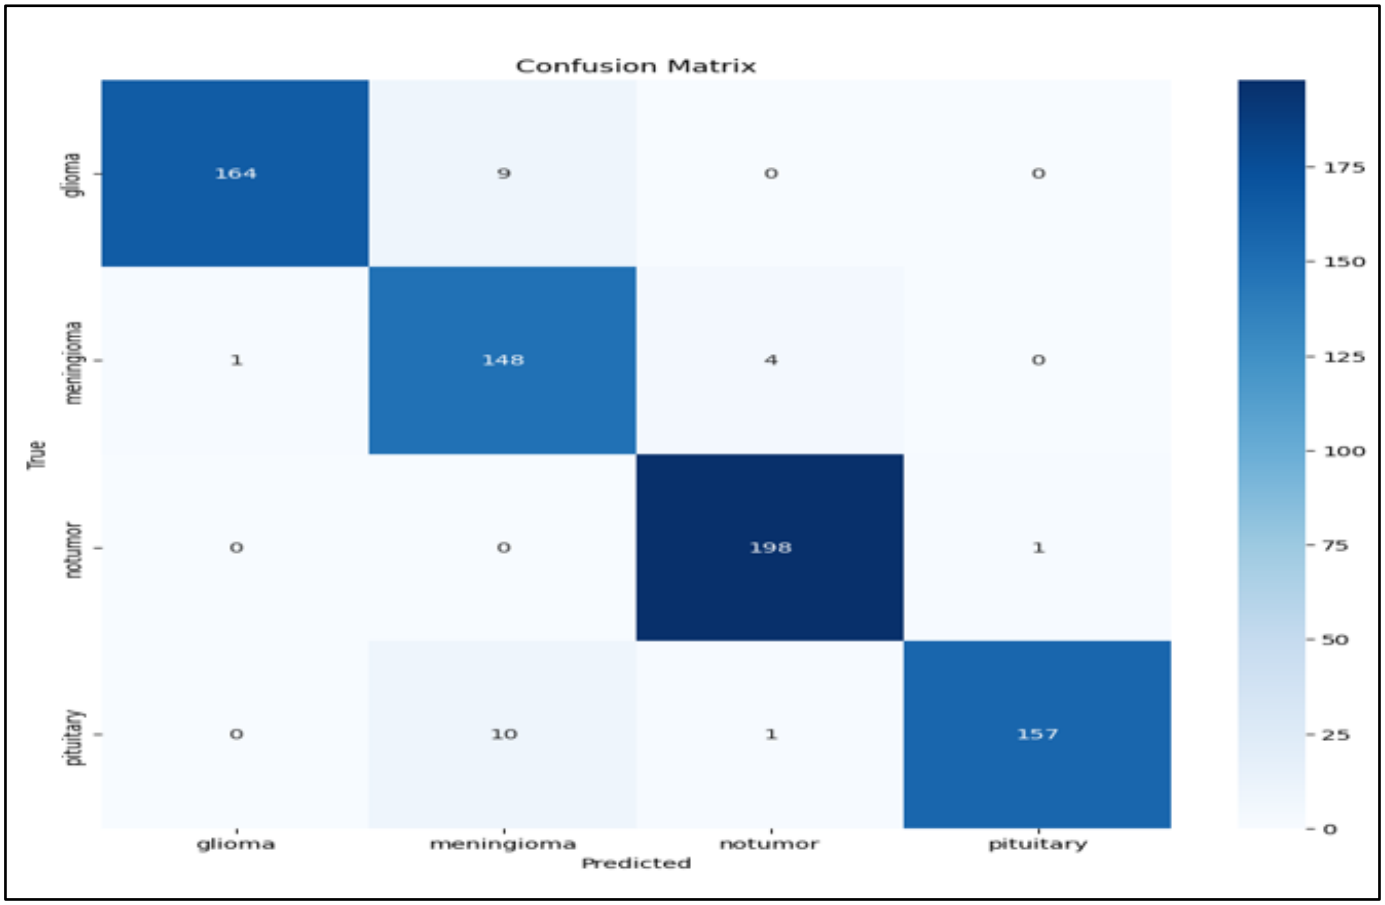

Fig 4 Confusion matrix of custom CNN

Please refer to the Pretrained Models section (Section III.D) for a comprehensive explanation of the transfer learning protocol and architectural details of the pretrained models, the results presented here focus on their performance compared to the custom CNN."

Despite their popularity, these models did not match the performance of the custom CNN. Notably, the custom architecture showed better balance in handling all classes, as evidenced by its strong recall and precision, which reduced false positives and improved detection sensitivity.

A confusion matrix (see Figure 3) was used to analyse the classification accuracy across different tumour types visually. The majority of the predictions aligned with the actual class labels. A small number of errors were observed, mainly involving confusion between glioma and meningioma tumours—an expected challenge due to their similar visual features in MRI images. However, these misclassifications were limited and did not significantly affect overall model performance.

Further analysis using a classification report (see *Table 2*) confirmed the model's high F1-score, indicating that it maintains a strong balance between precision and recall. This balanced performance is critical for clinical applications, where both false positives and false negatives can have serious implications. (Figure 4) shows the training and validation accuracy and loss over 30 epochs. The curves demonstrate stable convergence and minimal overfitting, reinforcing the model's robustness during training.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

Table 2 Classification Report

Class	Precision	Recall	F1-Score
Glioma	0.96	0.97	0.96
Meningioma	0.95	0.95	0.95
No Tumour	0.98	0.97	0.98
Pituitary	0.97	0.97	0.97
Accuracy			0.97
Macro Avg	0.97	0.97	0.97
Weighted Avg	0.97	0.97	0.97

The superior performance of the custom CNN can be attributed to its specialized architecture. The model incorporates deep residual blocks, batch normalization, and dropout regularization, all of which enhance learning stability and reduce overfitting. Unlike transfer learning models trained on natural images (e.g., ImageNet), this custom CNN was built from scratch to specifically handle the characteristics of medical brain imaging data. Moreover, the custom CNN is computationally efficient. Its lightweight architecture ensures faster inference, lower memory usage, and compatibility with CPU-based or mobile applications, making it suitable for real-world deployment in both hospital systems and portable diagnostic tools. The custom CNN offers both technical superiority and practical value, outperforming widely-used pre-trained models in accuracy and generalization. Its consistent performance, optimized design, and readiness for real-time applications make it a strong candidate for integration into clinical decision support systems.

➢ *Comparative Evaluation with Pretrained Models:*

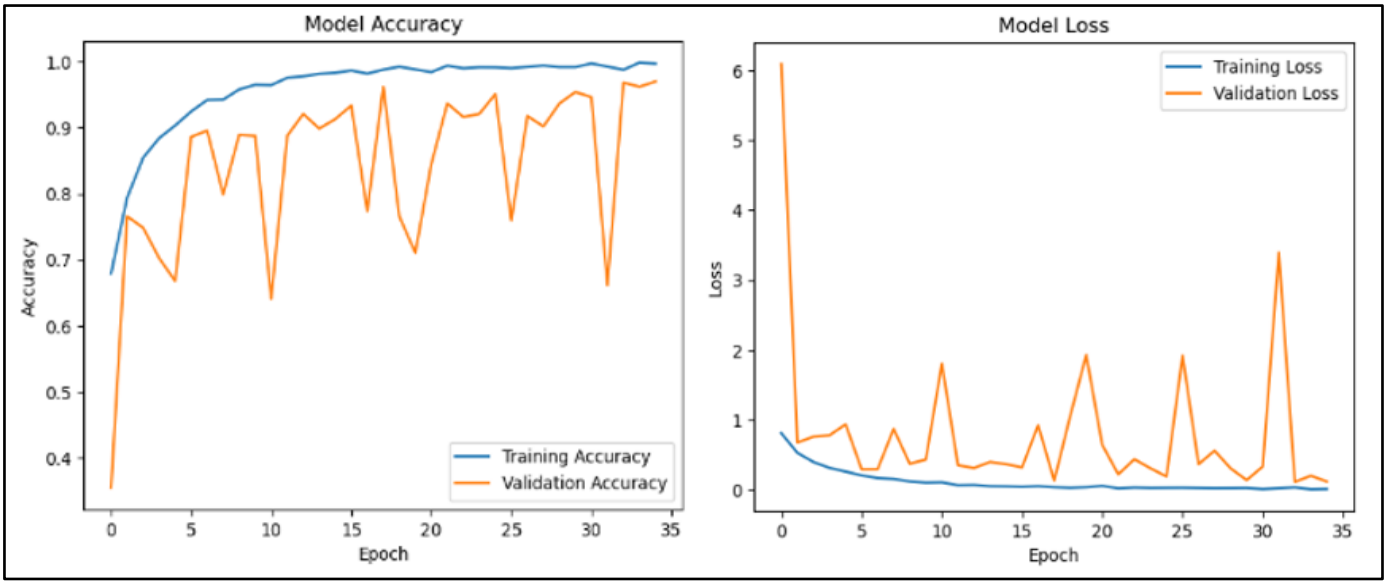

Fig 5 Training, Validation Accuracy and Loss Curves of the Custom CNN

To substantiate the effectiveness of the proposed custom CNN, a comparative evaluation was conducted against three widely recognized pretrained models: VGG16, ResNet50, and MobileNetV2. These architectures, originally trained on the ImageNet dataset, were selected for their proven performance in various vision tasks and adapted to the domain of brain MRI classification through a consistent fine-tuning protocol.

Each model underwent identical preprocessing steps and data augmentation strategies, including random flips, rotations, zoom adjustments, and brightness modifications, ensuring an equitable experimental setup. During training, a two-stage fine-tuning approach was employed—initially freezing all convolutional layers to train a custom classifier head, followed by selectively unfreezing top layers for domain adaptation.

The quantitative outcomes, summarized in Table 2 and illustrated in Figure 3, highlight a clear performance advantage for the custom CNN. Achieving a test accuracy of 97% and macro-averaged F1-score of 0.97, the custom model consistently outperformed the pretrained alternatives. While VGG16 showed competitive results with 91.91% accuracy, its higher parameter counts and need for RGB input (224 by 224) limited its efficiency. MobileNetV2 and ResNet50, although optimized for efficiency, fell short in terms of recall and F1 score, particularly in distinguishing between closely resembling tumour types.

This performance gap underscores the limitations of using generic feature representations from natural images for specialized medical data. In contrast, the custom CNN's architecture was purpose-built to extract domain-specific features from grayscale MRI scans at a lower resolution (150

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

by 150), leading to more accurate and computationally efficient classification.

Collectively, these findings affirm the superiority of a tailored deep learning model over off-the-shelf pretrained networks for brain tumour classification and highlight its practical applicability in real-world clinical settings where computational resources and input modalities may be constrained.

V. ANALYTICAL INSIGHTS AND INTERPRETATIONS

The comprehensive evaluation of our custom convolutional neural network (CNN), positioned against established pretrained architectures such as VGG16, ResNet50, and MobileNetV2, yields multifaceted insights that extend beyond mere accuracy metrics. This section synthesizes quantitative trends, architectural behaviours, and domain-driven reflections into a unified framework, informing future research and real-world deployment in neuroimaging-based tumour classification.

➢ *Convergence Dynamics and Training Stability:*

The custom-designed CNN consistently achieved rapid and steady convergence during training. Unlike VGG16 and ResNet50—which required extensive epochs and complex regularization—the lightweight design and domain-aligned input structure of our model fostered early performance saturation with minimized overfitting. The convergence curves reflected minimal variance between training and validation phases, with loss values stabilizing below 0.12. Transient perturbations during mid-training were efficiently corrected by the learning-rate scheduler, reflecting robustness in navigating the optimization landscape.

➢ *Architectural Efficiency and Predictive Power:*

A striking outcome of this study is the disproval of the perceived trade-off between parameter count and classification performance. Our network, with only ~3.1 million parameters, outperformed the vastly deeper and heavier VGG16 (~138M) and ResNet50 (~25M) in test accuracy and F1 score. This reinforces the notion that architectural elegance—anchored in domain-specific design—can eclipse brute-force depth. By aligning filter structures and activation flows to the grayscale MRI domain, we achieved refined representational capacity without computational bloat.

➢ *Domain Fit and Input Fidelity:*

Unlike generic backbones optimized for RGB, 224 by 224 natural images, our model was natively engineered for 150 by 150 grayscale slices, precisely matching the structural and textural nature of brain MRIs. Pretrained networks, when resized and converted to handle such data, introduced representational distortions and irrelevant feature sensitivities. This mismatch was observable in both learning curves and qualitative output, where pretrained models often diverged in activation relevance under challenging contrast conditions.

➢ *Visual Rationality and Explainability:*

Through Grad-CAM heatmaps, our custom CNN consistently focused its attention on clinically meaningful tumour regions, validating its predictions through visual interpretability. In contrast, pretrained models occasionally highlighted peripheral or non-pathological regions, particularly under noisy or ambiguous scans. This discrepancy underscores the model's learned specificity in feature localization and justifies the use of custom CNNs for tasks where trust and explainability are paramount—particularly in healthcare contexts.

➢ *Generalization Footprint and Validation Integrity:*

The reduced discrepancy between validation and test set performance in our model signals a superior generalization capacity, attributed to strategic regularization (e.g., 50% dropout) and aggressive data augmentation (rotations, flips, zooms, intensity shifts). In the context of limited-size and low-diversity medical datasets, this generalization is pivotal for safe deployment. Pretrained models, despite their scale, exhibited higher sensitivity to training artifacts, reinforcing the importance of purposeful design over transferred complexity.

➢ *Deployment Readiness in Resource-Constrained Settings:*

Beyond metrics, the practical deployment viability of our model is a key differentiator. Sub-100 ms inference times and sub-500 MB memory footprints make this architecture ideal for point-of-care settings and edge devices. In regions like rural India, where radiological expertise is sparse, such a tool could act as an effective first-line triage assistant, integrated into tele-radiology workflows. Unlike heavy pretrained models, ours requires neither GPU acceleration nor extensive preprocessing, enabling integration into existing hospital infrastructure.

➢ *Limitations and Reflections:*

Despite promising results, certain limitations persist. The current model lacks pixel-level localization, which restricts its application in surgical planning or volumetric analysis. Incorporating segmentation modules or attention-based architectures could bridge this gap. Furthermore, our reliance on 2D single-slice inputs omits volumetric continuity, which 3D CNNs could better exploit. Lastly, while Grad-CAM offered preliminary interpretability, a structured, clinician-informed evaluation of model attention is necessary to foster medical trust.

➢ *Strategic Future Directions:*

- *Volumetric Expansion:*

Transitioning to 3D CNNs to capture inter-slice continuity and enhance subtype discrimination accuracy.

- *Multi-Modal Synergy:*

Fusing MRI with complementary modalities (e.g., DWI, spectroscopy) for a more holistic diagnostic framework.

- *Explainability Integration:*

Formal studies involving radiologist feedback on interpretability tools (e.g., Grad-CAM, LIME).

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

- *Clinical Pilots:*
  Real-world trials across Indian hospitals to measure impact on triage, diagnostic speed, and clinician workload.

- *Mobile Deployment:*
  Conversion to TensorFlow Lite or ONNX for on-device inference, enabling deployment on low-cost Android devices.

- *Data Consortiums:*
  Collaborating with regional medical institutions to expand and diversify the training corpus for pan-Indian validation.

VI. ACKNOWLEDGMENT

The authors—Sharmistha Paul(first author), Shilpita Saha(co-author), and Pritikona Maji(co-author)—students of the Master of Computer Application program at Dr. B. C. Roy Engineering College, respectfully acknowledge the guidance and supervision of Prof. Ansuman Mahanty, whose expertise, encouragement, and constructive feedback were invaluable to the successful completion of this research. The authors also express their sincere gratitude to Dr. B.C. Roy Engineering College for providing the academic support and computational infrastructure necessary for the development and analysis of this work. Additionally, the authors extend appreciation to the curators of the MRI datasets employed in this study, which served as a foundational resource for experimentation and evaluation.

VII. CONCLUSION

In this study, we presented a custom-designed convolutional neural network (CNN) purpose-built for the classification of brain tumour subtypes—glioma, meningioma, pituitary adenoma—and healthy cerebral tissue using MRI imagery. Rather than retrofitting generic architectures pretrained on natural images, we embraced a design philosophy rooted in domain congruence and data-specific optimization. The resulting model—compact in scale yet potent in performance—achieved a peak test accuracy of 97.0%, while operating on grayscale 150 by 150 MRI slices and maintaining a footprint of under two million parameters.

This work substantiates a core hypothesis: efficiency, when architecturally guided by clinical and statistical priors, need not compromise diagnostic efficacy. Our results underscore that performance gains are not solely the domain of deep, overparameterized models but can emerge from thoughtfully calibrated architectures aligned to the structural characteristics of medical imaging data. Compared to canonical models such as VGG16, ResNet50, and MobileNetV2, our network demonstrated superior generalization, faster convergence, reduced memory demand, and greater interpretive focus on pathologically relevant regions—features paramount for real-world clinical deployment.

The practical implications are both immediate and profound. In resource-constrained environments, such as rural or under-equipped hospitals in developing regions, our lightweight model enables the possibility of real-time, edge-level diagnostics without the infrastructural burden typical of deep learning deployments. Such accessibility dovetails with national healthcare digitization strategies and speaks directly to the unmet need for equitable, AI-augmented screening tools capable of reducing diagnostic delays and augmenting radiological workflows.

From a systems research perspective, our study offers a principled, reproducible framework for designing specialized medical AI: one that emphasizes architectural parsimony, contextual data alignment, and real-world applicability over brute-force scale. Beyond classification accuracy, this work represents a shift toward human-cantered, purpose-driven AI systems that are both computationally viable and clinically trustworthy.

Future extensions of this work will include transitioning to volumetric (3D) input representations, integrating multi-modal imaging sources, and refining explainability mechanisms through methods like Grad-CAM and LIME, bolstered by clinician feedback. Additionally, we aim to evaluate the system prospectively within live clinical settings, including pilot deployment in Indian hospitals, and build a cross-institutional dataset to enhance generalizability.

This research lays not only a technical foundation for efficient tumour classification but also a strategic blueprint for the development of scalable, interpretable, and accessible medical AI systems—moving the field one step closer to truly democratized, data-driven diagnostics.

REFERENCES

[1]. M.I. Mahmud, M. Mamun, and A. Abdelgawad, “A Deep Analysis of Brain Tumour Detection from MR Images Using Deep Learning Networks,” Algorithms, vol. 16, no. 4, pp. 1–19, 2023. doi: 10.3390/a16040176. Available: https://doi.org/10.3390/a16040176.

[2]. H.R. Almadhoun and S.S. Abu-Naser, “Detection of Brain Tumour Using Deep Learning,” International Journal of Academic Engineering Research (IJAER), vol. 6, no. 12, pp. 29–47, 2022.

[3]. A.S. Musallam, A.S. Sherif, and M.K. Hussein, “A New Convolutional Neural Network Architecture for Automatic Detection of Brain Tumours in Magnetic Resonance Imaging Images,” IEEE Access, vol. 10, pp. 2775–2782, 2022.

[4]. M. Wozniak, J. Sika, and M. Wieczorek, “Deep Neural Network Correlation Learning Mechanism for CT Brain Tumour Detection,” Neural Computing and Applications, pp. 1–16, 2021.

[5]. D.R. Nayak, N. Padhy, P.K. Mallick, M. Zymbler, and S. Kumar, “Brain Tumour Classification Using Dense Efficient-Net,” Axioms, vol. 11, no. 1, p. 34, 2022. doi: 10.3390/axioms11010034. Available: https://doi.org/10.3390/axioms11010034.

[6]. K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” arXiv preprint arXiv:1409.1556, 2014.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25may855

[7]. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, “Rethinking the Inception Architecture for Computer Vision,” Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2818–2826, 2016.

[8]. K. He, X. Zhang, S. Ren, and J. Sun, “Deep Residual Learning for Image Recognition,” Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778, 2016.

[9]. Krizhevsky, I. Sutskever, and G.E. Hinton, “ImageNet Classification with Deep Convolutional Neural Networks,” Proceedings of the 25th International Conference on Neural Information Processing Systems, pp. 1097–1105, 2012.

[10]. X. Zhang, P. Zhou, and A. Yuille, “Region-based Convolutional Networks for Brain Tumour Detection and Segmentation in Medical Images,” Medical Image Analysis, vol. 38, pp. 60–74, 2017. doi: 10.1016/j.media.2017.05.005.

[11]. S. Roy, S. Saha, and A. Kundu, “Brain Tumour Segmentation and Classification Using Convolutional Neural Networks,” Journal of Medical Imaging, vol. 4, no. 3, pp. 1–14, 2017. doi: 10.1117/1.JMI.4.3.034001.

[12]. Cireşan, L. Meier, J. Masci, and U. Schmidhuber, “Flexible, High-Performance Convolutional Neural Networks for Image Classification,” Proceedings of the International Joint Conference on Neural Networks (IJCNN), pp. 1–6, 2011.