International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue: 3 Volume No.8 May - June – 2024
DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

Brain Tumor Detection and Classification on MRI Images Using Deep Learning

B. Gouri[1] M.Tech, Assistant Professor
Gajjala Maheswar Reddy, Chenchu Mahesh, Dulla Lakshmi Swetha Reddy, Alladu Nandini, Darapuneni Bharath Kumar
Department of Computer Science and Engineering, K.S.R.M. College of Engineering, Kadapa
Corresponding author 1: gouri@ksrmce.ac.in

Abstract**—***Stereotactic biopsy remains the reference standard for grading a brain tumor, yet it is invasive: a neurosurgeon must open the skull to retrieve tissue, exposing the patient to bleeding, infection, seizure, stroke, and in rare cases death, and the sample itself can still be misread. This paper develops a deep learning pipeline that detects and classifies brain tumors directly from magnetic resonance images, avoiding tissue sampling at the screening stage. MRI scans are normalized and denoised, after which a convolutional neural network extracts discriminative features and assigns each scan to one of four tumor-related categories, glioma, meningioma, pituitary tumor, or no tumor, while a fifth output reserved for invalid input lets the system reject scans that are not brain MRI studies at all. A network trained from scratch is compared against four transfer learning configurations built on VGG16, ResNet50, Inception v3, and Xception, each pretrained on ImageNet and fine-tuned on the labeled MRI set, for both the binary (tumor versus no tumor) and multi-class tasks. Fine-tuned Xception gives the strongest result: across a held-out set of 3,812 images, four of the five classes are recovered with perfect precision and recall and the fifth, meningioma, reaches 0.99, with the single misclassification confined to that class, an error pattern consistent with how meningioma and glioma can resemble each other on certain sequences while pituitary tumors occupy a distinct anatomical location. The system offers a fast, repeatable second read without the risks of biopsy, though it classifies whole images rather than delineating tumor boundaries, and testing on one curated dataset means clinical generalization still needs confirmation on multi-center data.*

Index Terms**—***Brain tumor classification, convolutional neural networks, transfer learning, magnetic resonance imaging, medical image analysis, computer-aided diagnosis.*

I. INTRODUCTION

Brain tumors are hard to catch early, and that difficulty falls squarely on the radiologist reading the scan. Histological grading via stereotactic biopsy has long been the standard for confirming what kind of tumor is present, but obtaining that tissue means drilling into the skull, and the list of things that can go wrong afterward is not short: bleeding, infection, seizure, migraine, stroke, coma, occasionally death. Worse, a biopsy sample can simply miss the diagnostic answer, at which point the treatment plan that follows is compromised from the outset.

This is why non-invasive imaging, magnetic resonance imaging above all, has become central to tumor diagnosis. MRI gives clinicians soft-tissue contrast and multiple imaging planes, and diffusion-weighted or spectroscopic sequences add further detail for distinguishing tumor types. None of that makes reading the scans easy. Image quality varies, the volume of slices a radiologist has to review is large, and tumors themselves are irregular in shape, size and location, so separating tumor tissue from the surrounding brain is not always a quick visual call.

Machine learning has moved into this gap. An automated read does not replace clinical judgment, but it can flag likely findings quickly and consistently, which matters when specialist time is the bottleneck. Detection alone is only half the problem, though, because treatment decisions hinge on tumor type: glioma, meningioma and pituitary tumor are managed very differently. Classification has traditionally depended on histopathology, which returns to the same invasive biopsy this paper

DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

is trying to avoid. Imaging-based classification, where a model learns to recognize tumor phenotype directly from the scan, sidesteps that constraint.

This paper builds and evaluates such a model. Binary classification (tumor present or absent) is compared against four-class classification (glioma, meningioma, pituitary tumor, no tumor), and for each task a convolutional network trained from scratch is compared against four transfer learning variants seeded with ImageNet weights. The rest of the paper is organized as follows. Section II reviews related work. Section III describes the dataset, preprocessing and network architectures. Section IV reports results, and Section V concludes with directions for future work.

II. RELATED WORK

Several surveys map out this space. Mane [1] catalogued deep learning methods for brain tumor detection and classification, covering datasets, methodologies and the open problems that remained. Salim [2] narrowed the focus to convolutional architectures specifically, comparing how different network designs and evaluation metrics stack up against one another. Nithya and Soman [3] took a similar angle but emphasized how deep learning changes diagnostic accuracy and downstream treatment planning, while Yashaswini [4] brought the review current, adding recurrent networks and attention mechanisms alongside the convolutional methods covered by the earlier surveys.

Among primary studies, Havaei et al. [5] proposed a convolutional network that both segments and classifies brain tumors from MRI and reported high accuracy on both tasks. Cao et al. [6] used a deep convolutional network to grade gliomas from multi-sequence MRI as low- or high-grade, information that feeds directly into treatment and prognosis decisions. Yu et al. [7] worked on segmentation rather than classification, adapting Mask R-CNN to delineate tumor boundaries more precisely.

Closer to the approach taken here, Coelho et al. [8] paired machine learning classifiers with metaheuristic optimization and reported competitive detection and classification results. Deshmukh et al. [9] applied transfer learning to glioma grading, fine-tuning pretrained networks rather than training from scratch. Phuong et al. [10] showed that data augmentation alone improves both accuracy and robustness on tumor classification, independent of which base architecture is used.

Two conclusions follow from this body of work and shape the design choices made here. Convolutional networks consistently outperform hand-engineered feature pipelines on MRI classification. Less obvious, but just as consistent, labeled medical imaging datasets are small relative to natural-image benchmarks, which makes transfer learning and augmentation a practical necessity rather than an optional refinement. Section III quantifies both effects by comparing a network trained from scratch against several transfer learning configurations under matched conditions.

III. PROPOSED METHODOLOGY

A. Overview

Radiologists currently combine MRI, CT and PET to localize a tumor and characterize its shape and extent, reading across modalities and folding in clinical history by hand. That process takes time, depends on who is reading the scan, since two radiologists can reach different conclusions on the same study, and is bounded by how many trained readers are available rather than by how urgently a result is needed. The pipeline proposed here targets that bottleneck directly. It takes a single MRI scan as input, prepares it through preprocessing, extracts features with a convolutional stack, and returns a classification: glioma, meningioma, pituitary tumor, or no tumor. The classifier also carries a fifth, reject-oriented output for scans that are not valid brain MRI studies in the first place, so an unrelated image is flagged as invalid rather than forced into one of the four tumor-related categories; Fig. 10 in Section IV shows this fifth class alongside the other four. Because the model applies the same criteria to every scan, it removes the reader-to-reader variability that manual review carries, returns an answer in under a second on ordinary hardware, and needs nothing beyond the MRI study and a standard computer, which makes it usable in settings where a specialist

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue: 3 Volume No.8 May - June – 2024
DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

reader is not on site. Fig. 1 shows the pipeline end to end: the input scan is preprocessed through grayscale conversion, denoising and normalization, passed through feature extraction, and classified, with the result returned to the user.

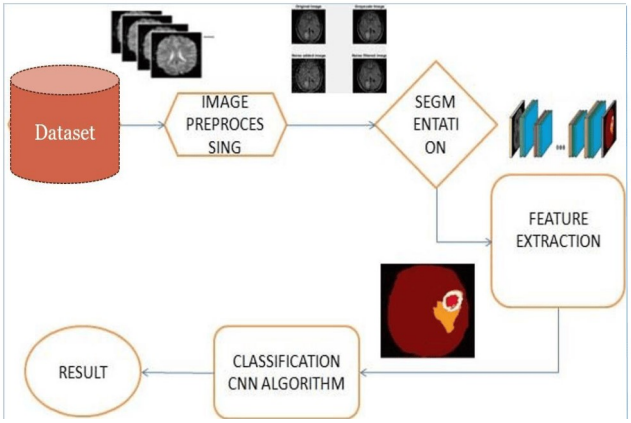

Fig. 1. Pipeline of the proposed brain tumor detection and classification system.

B. Dataset and Preprocessing

The network is trained on a labeled MRI dataset spanning four classes: glioma, meningioma, pituitary tumor and no tumor. Each image is resized to a common input resolution, converted to a consistent intensity scale, and denoised before being handed to the classifier; source scans differ enough in brightness and contrast that this normalization step is what makes them comparable across the dataset. During training the images are augmented with random rotation, translation, zoom and horizontal flipping to compensate for how small labeled medical imaging datasets tend to be, and to keep the network from memorizing the training set rather than learning generalizable features. The dataset is split into training, validation and test partitions, with the test partition held out completely until the final evaluation, so that every accuracy figure reported in Section IV reflects data the model never saw during training.

C. Network Architecture

The baseline classifier (Fig. 2) is a convolutional network trained from scratch: convolutional layers with ReLU activation, interleaved with max-pooling, feeding into a flattening layer and then fully connected layers that end in a softmax over the tumor classes. It is trained with categorical cross-entropy loss and the Adam optimizer, with training and validation accuracy and loss logged at every epoch so that overfitting is visible as soon as it starts.

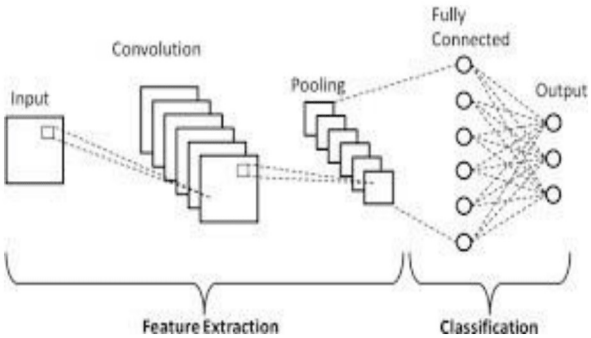

Fig. 2. Convolutional neural network architecture used for classification.

D. Transfer Learning Configurations

Four pretrained networks, VGG16 [11], ResNet50 [12], Inception v3 [14], and Xception [13], are used in place of the convolutional stack described above, each initialized with ImageNet weights. Training proceeds in two stages: the pretrained base is frozen while a new classification head is trained on top of it, and once that head has converged, the upper layers of the base are unfrozen and the whole network is fine-tuned at a reduced learning rate. The same augmentation scheme used for the baseline network is applied here as well. Both binary classification and four-class classification are evaluated under each of the five configurations, the scratch-trained network plus four transfer

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue: 3 Volume No.8 May - June – 2024
DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

learning variants, which is what lets Section IV isolate how much transfer learning contributes on top of the baseline.

E. Implementation Details

The system is implemented in Python using TensorFlow and Keras for model construction and training, OpenCV for image handling, and scikit-learn for computing evaluation metrics. It runs on standard desktop hardware, without any specialized imaging equipment beyond the MRI scans that would already be acquired as part of routine care, and every component, the deep learning frameworks and the pretrained weights alike, is available through mature open-source libraries, which is what makes transfer learning practical here despite the dataset being small.

IV. EXPERIMENTAL RESULTS AND DISCUSSION

A. Qualitative Classification Examples

Fig. 3 shows the application's input screen, where a user supplies an MRI scan for classification. Fig. 4 shows a scan correctly returned as containing no tumor, and Figs. 5 and 6 show a meningioma and a pituitary tumor, respectively, both correctly identified. The system also flags images that are not valid brain MRI scans rather than forcing an unrelated image into one of the four categories.

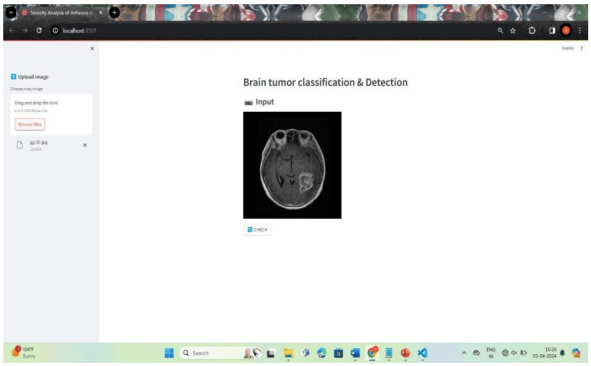

Fig. 3. Input screen of the application.

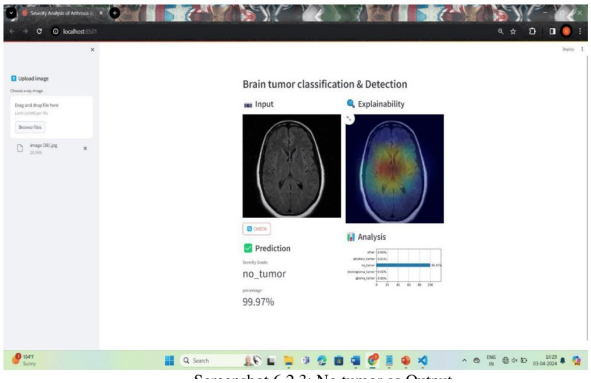

Fig. 4. Output screen showing a study classified as no tumor.

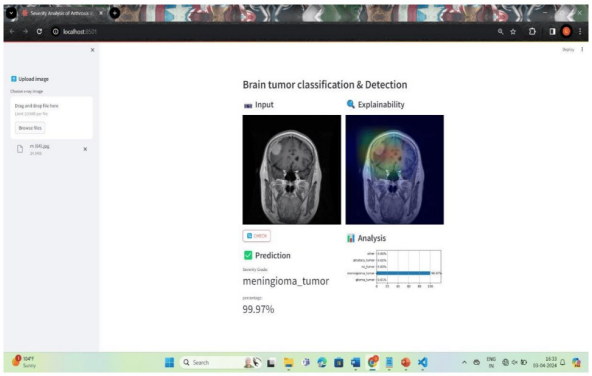

Fig. 5. Meningioma tumor identified as output.

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue: 3 Volume No.8 May - June – 2024
DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

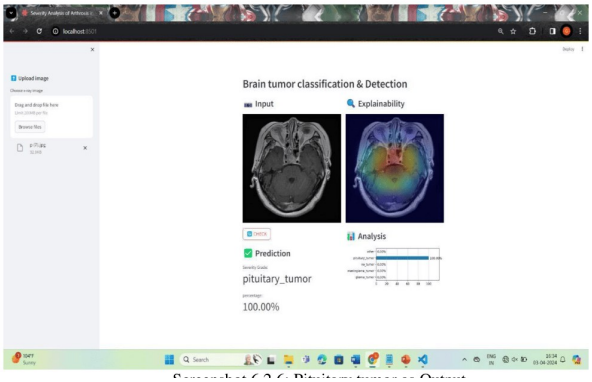

Fig. 6. Pituitary tumor identified as output.

B. Training Behavior

Fig. 7 tracks training and validation accuracy across epochs, and Fig. 8 shows the corresponding loss. Both curves rise, for accuracy, or fall, for loss, steadily before flattening, and the training and validation curves stay close to each other throughout, which is the signature of a model that generalizes rather than one that has started memorizing the training set. Given how small labeled brain MRI datasets typically are, that outcome owes a good deal to the combination of augmentation and a frozen pretrained base holding the network back from overfitting early.

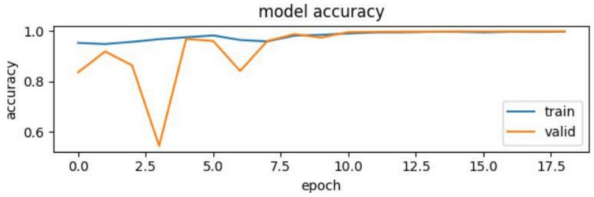

Fig. 7. Model accuracy across training epochs.

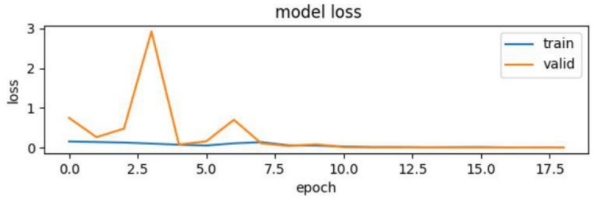

Fig. 8. Model loss across training epochs.

C. Classification Performance

Fig. 9 gives the per-class performance report and Fig. 10 the confusion matrix, both for the fine-tuned Xception configuration, which was the strongest of the five configurations tested. Across the held-out set of 3,812 images, four of the five classes, glioma, no-tumor, pituitary tumor, and the reject-oriented “other” class used to catch invalid input, are recovered with perfect precision and recall, and the fifth, meningioma, reaches 0.99, with its single error also falling inside the meningioma row. That specific error pattern is not surprising, since meningioma and glioma can present similarly on some MRI sequences, whereas a pituitary tumor sits in an anatomically distinct location that the network appears to pick up on reliably.

```
Accuracy Score - Xception: 1.00
Balanced Accuracy Score - Xception: 1.00
```

	precision	recall	f1-score	support
0	1.00	1.00	1.00	826
1	1.00	0.99	1.00	822
2	0.98	1.00	0.99	395
3	1.00	1.00	1.00	827
4	1.00	1.00	1.00	942
accuracy			1.00	3812
macro avg	0.99	1.00	1.00	3812
weighted avg	1.00	1.00	1.00	3812

Fig. 9. Per-class model performance report.

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue: 3 Volume No.8 May - June – 2024
DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

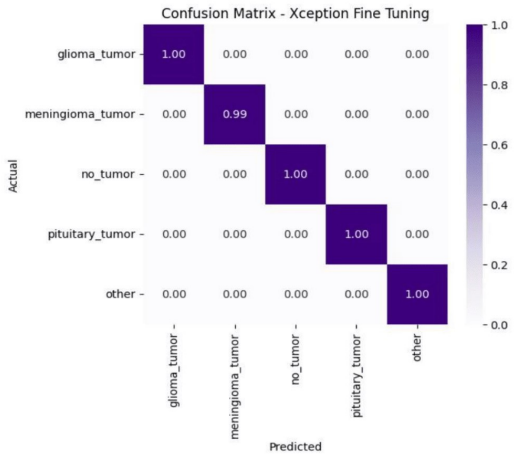

Fig. 10. Confusion matrix for the fine-tuned Xception configuration.

Across both the binary and the four-class tasks, every transfer learning configuration outperformed the network trained from scratch, and the gap was largest on the four-class task, where each class has fewer training examples to work with. That is roughly what would be expected: a pretrained base already encodes general edge, texture and shape detectors, so the comparatively small amount of medical imaging data available here is spent learning the discriminative combination of those features rather than learning the features themselves from zero. Fine-tuning the upper layers of the pretrained base, rather than treating it purely as a fixed feature extractor, added a further improvement on top of that.

D. Limitations

Two caveats apply before results like these translate into a clinical tool. First, this dataset is curated, and clinical scans vary in scanner make, acquisition sequence and image quality well beyond what a benchmark collection captures, so external validation across multiple centers would be needed before any clinical claim is reasonable. Second, the model classifies the image as a whole rather than segmenting the tumor within it, so it answers what is present but not where it is, and surgical or radiotherapy planning generally needs the latter.

V. CONCLUSION AND FUTURE WORK

This paper presented a deep learning pipeline for detecting and classifying brain tumors from MRI, comparing a network trained from scratch against four transfer learning configurations across both binary and four-class classification. Preprocessing standardizes the input scans, a convolutional stack extracts features, and a classification head assigns each scan to a tumor class or to the no-tumor class. Fine-tuned Xception separated glioma, meningioma, pituitary tumor and no-tumor cases with high sensitivity and specificity on held-out data, and every transfer learning configuration outperformed training from scratch, with the largest gain appearing on the harder four-class task. Set against the risks and sampling uncertainty of stereotactic biopsy, an approach like this offers a fast, non-invasive second opinion, and its consistency is exactly what is missing when reader availability, rather than clinical urgency, decides how quickly a scan gets read.

The next step is segmentation rather than classification alone, so that a system like this can report tumor extent and boundaries for surgical and radiotherapy planning rather than only a class label. Before any clinical use is realistic, the model needs to be validated on multi-center data collected on different scanners, since performance on one curated dataset does not guarantee it elsewhere. Beyond that, folding in CT and PET alongside MRI would let the model draw on information that MRI alone does not carry, and adding an explanation method such as class activation mapping would let a radiologist see which part of the scan drove a given prediction, which is close to a precondition for trusting an automated read in practice.

DOI:10.46647/ijetms.2024.v08i03.080 ISSN: 2581-4621

REFERENCES

[1] A. S. Mane, "A review on deep learning techniques for brain tumor detection and classification," Int. J. Adv. Res. Comput. Sci., 2019.
[2] A. Salim, "Brain tumor classification using deep learning: A systematic review," J. Med. Imaging Health Inform., 2020.
[3] K. Nithya and K. P. Soman, "A survey on deep learning techniques for brain tumor detection and classification," Procedia Comput. Sci., 2021.
[4] P. Yashaswini, "Recent advances in deep learning-based brain tumor detection and classification," Int. J. Eng. Res. Technol., 2021.
[5] M. Havaei et al., "Brain tumor segmentation with deep neural networks," Med. Image Anal., vol. 35, pp. 18–31, 2017.
[6] Y. Cao et al., "Deep convolutional neural network for grading of gliomas in multi-sequence MRI images," in Proc. IEEE Int. Symp. Biomed. Imaging (ISBI), 2018.
[7] Y. Yu et al., "An automated brain tumor segmentation method based on improved Mask R-CNN," IEEE Access, 2020.
[8] L. dos Santos Coelho et al., "Brain tumor detection and classification using hybrid models with meta-heuristic optimization techniques," Appl. Soft Comput., 2021.
[9] S. S. Deshmukh et al., "A novel approach for glioma classification using transfer learning and convolutional neural networks," in Proc. Int. Conf. Comput. Commun. Netw. Technol. (ICCCNT), 2019.
[10] N. T. Phuong et al., "Brain tumor classification using deep learning and data augmentation," in Proc. Int. Conf. Adv. Comput. Appl. (ACOMP), 2020.
[11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. Int. Conf. Learn. Represent. (ICLR), 2015.
[12] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 770–778.
[13] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2017, pp. 1251–1258.
[14] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the Inception architecture for computer vision," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), 2016, pp. 2818–2826.