Brain Tumor Detection Using Convolutional Neural Networks

Steven Strange
Kent State University
Kent, Ohio
sstrange@kent.edu

Suwei Gao
Kent State University
Kent, Ohio
sgao16@kent.edu

Lalith Kumar Done
Kent State University
Kent, Ohio
ldone@kent.edu

***Abstract* - One of the most serious medical issues is brain tumors, which must be diagnosed promptly and accurately to enhance patient outcomes. Automated, AI-driven solutions are necessary because manual brain MRI scan analysis is frequently laborious and prone to human error. This research uses the potent machine learning framework TensorFlow to create a Convolutional Neural Network (CNN) that can identify brain tumors from MRI images. The model trains on labeled datasets to attain high tumor identification accuracy after preprocessing MRI images to improve quality. Optimizing computational performance, managing imbalanced datasets, and data augmentation were among the main obstacles. By providing a quicker and more accurate tool to help radiologists diagnose brain tumors, the developed system exemplifies the potential of AI in healthcare.**

Keywords – Artificial Intelligence, Convolutional Neural Networks (CNN), TensorFlow

I. Motivation

Magnetic Resonance Imaging (MRI) is a key tool in diagnosing and managing brain tumors, but interpreting these scans has its challenges. Even experienced radiologists can occasionally make mistakes due to factors such as the large volume of work they handle daily. A study published in *Radiology* highlights that subtle abnormalities or high pressure situations can lead to errors, making it clear that tools designed to enhance accuracy and support decision making [1].

MRI's importance in brain tumor detection is difficult to overstate. It provides detailed images of brain structures, enabling early identification and understanding of tumors. Unlike many other imaging methods, MRI is particularly useful at highlighting soft tissue differences, making it an important resource for neurologists and oncologists. According to the American Health Imaging Network, timely MRIs significantly improve patient outcomes by detecting tumors at treatable stages, offering non-invasive, personalized care that can make a large difference in treatment success [2].

Despite its benefits, the rising demand for MRIs has placed increasing pressure on radiologists. Managing more scans while maintaining high diagnostic standards is increasingly challenging. Neural networks offers a practical solution to this growing problem. Research highlights how neural networks can streamline routine tasks, such as identifying normal scans or flagging abnormalities, allowing radiologists to focus on more complex cases [3]. By quickly highlighting high risk patients, AI can significantly improve the speed and efficiency of diagnosis and treatment planning, ultimately benefiting patient care.

This project explores the potential of Convolutional Neural Networks (CNNs) in analyzing medical images, especially for brain tumor detection. CNNs are incredibly effective at processing and understanding complex image data, often surpassing traditional methods in both speed and accuracy. They can detect patterns and anomalies in MRI scans that might be challenging for humans to identify consistently [4]. In addition to this, their ability to process large datasets quickly ensures patients in need of care are able to receive it as fast as possible.

II. Project Description

The system is designed to use deep learning techniques for automatically classifying brain tumors. In this research, the dataset has images of MRI scans and comes under four classes of the categorization: Glioma, Meningioma, Pituitary Tumor, and No Tumor. It consists of a training dataset and a test dataset. This data set has images labeled, therefore best for a supervised learning task. This dataset is chosen because of its diversity and the comprehensive representation of different tumor types, which enables the model to generalize well and classify MRI images into the specified categories. The dataset will help the model learn from labeled examples for each class, improving the distinguishing features of each tumor type, thus improving its predictive accuracy during inference.

This project's core is a Convolutional Neural Network (CNN). CNNs are specific deep learning architectures that are designed to deal with visual information. Since they can automatically extract spatial and hierarchical features from images, they perform exceptionally well in tasks like image recognition and classification. In this project, MRI images are analyzed by CNNs to find patterns and categorize them into some preestablished groups. It requires that the network have a series of layers for extracting both local and global features from MRI scans; it needs convolutional layers to extract the features, while the pooling layers are necessary to reduce the dimensions. CNN is suitable for this classification because of the reliable way it guarantees for the identification of the pattern and relationships visually. We use them when working with image data.

TensorFlow is an open-source and powerful machine learning library developed to provide all the necessary elements in building, training, and deploying deep learning models. It will be used in this project to implement the CNN architecture, defining the layers of the network, and managing the training process. It offers an extensive set of functions covering data preprocessing, optimization of the training process, and the evaluation of model performance. TensorFlow is chosen because of its flexibility, scalability, and a wide range of built-in functionalities that ease the implementation of complex deep learning tasks. We use TensorFlow at every stage of the project, starting from the preprocessing of the dataset and defining the CNN model to training the network and testing its accuracy on test data. It ensures efficiency and scalability in the workflow through its ecosystem, something quite essential to implement in this project.

III. Project Design

This project design is well-structured to ensure a deep learning model for the detection of brain tumors is systematically implemented, from the stage of data preparation down to model evaluation. All the steps are very significant to ensure that accurate and reliable results are achieved.

A. Data preparation

The preprocessing of data in any machine learning project, including image classification tasks, is a must. Examples are given for MRI images with their respective classes, including Glioma, Meningioma, Pituitary Tumor, and No Tumor. Data division is made into two subsets, including a training dataset of 5,712 images and a testing one consisting of 1,311 images. The preparation of such data is needed in a way to get the proper and valid results about its classification.

B. Data Preprocessing

The pre-processing of the dataset involves some crucial steps to train a machine learning model. All MRI images in this dataset are resized to a uniform size of 128 x 128 pixels. The resizing makes all images with dimensions, which will be quite helpful for efficient computation without errors due to different sizes. The pixel intensity values have been normalized between [0 1]. Normalization is a must to bring all features to a similar scale for the model to learn efficiently and also to speed up convergence during training.

C. Data Augmentation

Data augmentation is used on the brain MRI dataset to improve its diversity, hence increasing the model's performance. Considering the classes of this dataset are images categorized as Glioma, Meningioma, Pituitary Tumor, and No Tumor, there are augmentations of rotations, flipping, zooming, and shifting on the images. With the increased variability within classes due to augmentation, overfitting decreases, thereby enhancing the generalizability of the model on previously unseen MRI images and resulting in better classification performance for tumors.

D. Model Building

Two models were used in this project to classify brain tumors: a Custom Convolutional Neural Network (CNN) that was created from scratch and a Transfer Learning strategy that made use of the VGG16 model that was pretrained on the ImageNet dataset. Each model was designed to categorize MRI images into four groups: meningioma, glioma, pituitary tumor, and no tumor.

1) Convolutional Neural Networks (CNN)

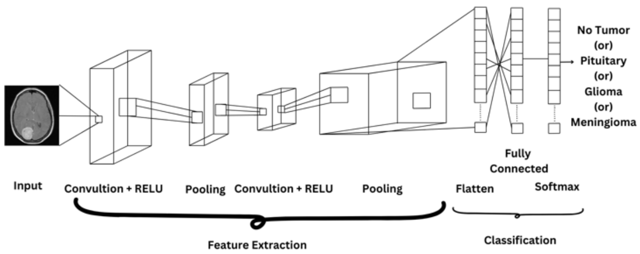

Figure 1: Convolutional Neural Network (CNN)

Images are resized with three RGB color channels and 150x150 pixel dimensions made up the model's input. Features like edges, textures, and patterns were extracted by the architecture's multiple convolutional layers with ReLU activation. Max-pooling layers were then used to downsample feature maps, lowering spatial dimensions without sacrificing important information. To enhance convergence and stabilize the training process, batch normalization was used. Dropout layers, which randomly deactivated neurons during training to prevent overfitting, improved the model's ability to generalize. In order to learn intricate patterns, the extracted features were flattened into a 1D vector and then sent through fully connected dense layers. The final softmax layer generated class probabilities. The Adam optimizer with a learning rate of 0.001 and the categorical crossentropy loss function were used to train the model. While early stopping and learning rate reduction were used to avoid overfitting, data augmentation techniques like scaling, flipping, and rotation were used to enhance generalization.

2) Transfer Learning

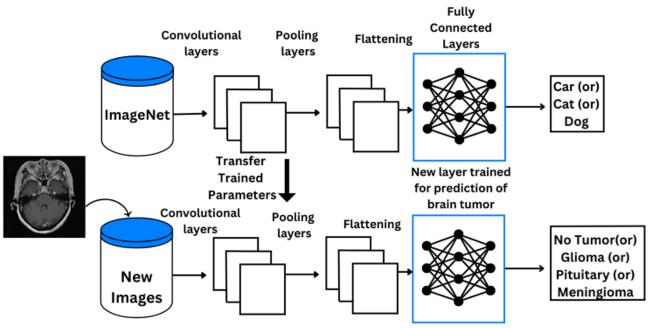

Figure 2: Illustration of transfer learning

E. Model Training

The VGG16 transfer learning model and the custom CNN underwent separate training procedures. The augmented training dataset was used to train both models, with 20% of the data reserved for validation. Scaling, flipping, and rotation are examples of data augmentation techniques that were used to increase the diversity of the training data and boost the generalization abilities of the models.

With a learning rate of 0.001, the Adam optimizer was used to train the model from scratch for the custom CNN. The loss metric for multi-class classification was the categorical crossentropy loss function. A learning rate reduction callback was used to dynamically lower the learning rate in the event that the validation performance plateaued, and early stopping was used to end training if the validation loss stopped getting better. By using these techniques, overfitting was avoided and the training process was optimized.

Only the custom layers were trained for the VGG16-based transfer learning model; the pretrained base layers were initially frozen. To ensure stable fine-tuning, the model was compiled using the Adam optimizer and a lower learning rate of 0.0001. The model was able to improve its feature extraction on the brain tumor dataset by unfreezing some of the base layers after the custom layers had been trained. The model's accuracy in classifying MRI images was further enhanced by this two-step fine-tuning procedure.

F. Model Evaluation

Metrics like accuracy, loss, and confusion matrices were used to assess the VGG16 and custom CNN transfer learning models on the testing dataset. MRI images were successfully classified by the custom CNN, but because it lacked pretrained knowledge, it needed more training epochs and processing power. By using pretrained features and fine-tuning, on the other hand, the VGG16 model showed better accuracy and generalization with less training time. The outcomes demonstrated the benefits of transfer learning, as VGG16 outperformed the custom CNN in terms of robustness and accuracy. The effectiveness of transfer learning for medical image classification tasks was highlighted by this comparison.

IV. Problems Faced

V. Implimentation

This project implemented a brain tumor detection neural network using two deep learning methods, a custom CNN and a transfer learning model with a pre-trained VGG16 architecture. The process included data preparation, model training, and testing.

The dataset used for this project was sourced from Kaggle: Brain Tumor MRI Dataset [5]. It consisted of 7,043 labeled MRI images divided into four categories: Glioma, Meningioma, Pituitary Tumor, and No Tumor. The dataset was organized into two subsets: 5,712 images for training and 1,331 images for testing. The testing set was further split into validation and test subsets to ensure a balanced evaluation.

All images were resized to 224×224 pixels to standardize input dimensions. Data augmentation was applied to the training set to increase data diversity and improve model generalization. Augmentation techniques included horizontal flipping, brightness adjustments, and small positional shifts.

The custom CNN model is designed to classify MRI images into four categories: Glioma, Meningioma, Pituitary Tumor, and No Tumor. The architecture starts with a convolutional layer containing 32 filters of size 3×3, using the ReLU activation function to extract key visual features such as edges and textures. This is followed by a second convolutional layer with 64 filters of size 3×3, also using ReLU activation, to identify more complex patterns in the images. After each convolutional layer, a max pooling layer with a size of 2×2 reduces the spatial dimensions of the feature maps, helping the model focus on the most important features while reducing computation. To prevent overfitting, dropout regularization is applied after each pooling layer, randomly turning off 20% of the neurons during training.

Once the features are extracted, the model flattens the two-dimensional feature maps into a one-dimensional vector to prepare the data for the fully connected layers. The first dense layer has 128 neurons and uses ReLU activation to learn high-level patterns from the extracted features. Batch normalization is used here to make training more stable and faster. A second dense layer with 64 neurons and ReLU activation further processes these features. Dropout regularization with a rate of 0.1 is applied after this layer to further reduce overfitting. Finally, the model uses a softmax output layer with four neurons to predict the probabilities of each image belonging to one of the four categories.

In order to further improve performance, we used transfer learning with the pre-rained VGG16 model. VGG16 is a well-known deep convolutional neural network trained on the ImageNet dataset, which contains millions of labeled images. [9] It is effective at extracting features from images due to its deep, layered architecture. For this project, we kept the convolutional layers of VGG16 as they are. These layers act as a feature extractor for the MRI images.

The flattened output from the VGG16 layers is passed into a dense layer with 128 neurons and ReLU activation, which learns new relationships from the features. Batch normalization is used to make training more stable and efficient, and dropout regularization with a rate of 0.1 is applied to reduce overfitting. A second dense layer with 64 neurons and ReLU activation further processes the features, with another dropout layer added to prevent overfitting. Finally, a softmax output layer with four neurons is used to classify the images into the four tumor categories.

VI. Results

After Training for 20 epochs our custom CNN model was able to achieve an accuracy score around 80%. In the same number of epochs our model trained using VGG16 as a starting point was able to an accuracy score around 95%.

Figure 3: Accuracy and loss up to 20 epochs (custom CNN)

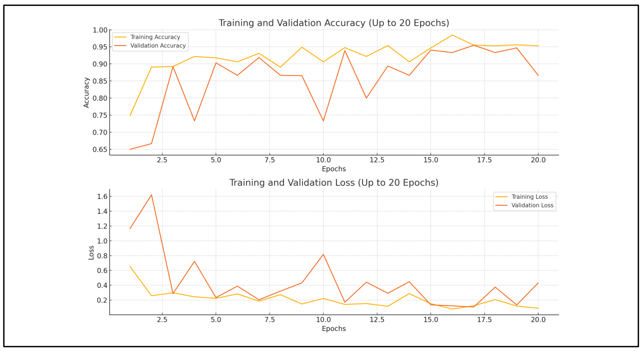

Figure 4: Accuracy and loss up to 20 epochs (Transfer Learning)

These results demonstrate the effectiveness of transfer learning for image classification. Not only were better results achieved, but higher accuracy was seen in fewer epochs.

When classifying images our model trained using transfer learning often gave more confident results while our custom CNN model would have more uncertainty even when it arrived at the correct classification.

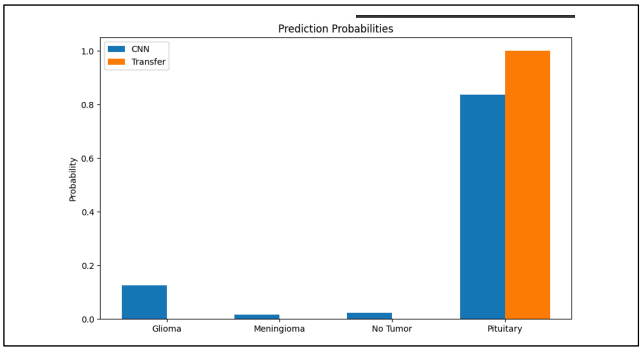

Figure 4: Example output (correct answer Pituitary)

VII. References

[1] S. Van der Gijp, M. R. Ravesloot, R. H. van Wijnhoven, et al., “Radiological error in daily practice: fatigue and cognitive overload as root causes,” *European Journal of Radiology*, vol. 144, 2023. [Online]. Available:https://www.ejradiology.com/article/S0720-048X%2823%2900346-7/fulltext.

[2] “Can an MRI detect a brain tumor?” *American Health Imaging Network*, 2023. [Online]. Available: https://neurolaunch.com/can-an-mri-detect-a-brain-tumor.

[3] Lamontagne, “AI shows potential to reduce radiologists’ chest x-ray workload in outpatient settings,” *Radiology Business*, 2023. [Online]. Available: https://radiologybusiness.com/topics/artificial-intelligence/ai-shows-potential-reduce-radiologists-chest-x-ray-workload-outpatient-setting.

[4] D. Gupta, M. R. Shah, et al., “Deep learning-based CNNs for medical imaging: Application to brain tumor detection,” *BMC Medical Imaging*, vol. 24, no. 1, 2024. [Online]. Available: https://bmcmedimaging.biomedcentral.com/articles/10.1186/s12880-024-01292-7.

[5] Sharma, Manav & Sharma, Pramanshu & Mittal, Ritik & Gupta, Kamakshi. (2021). Brain Tumour Detection Using Machine Learning. Journal of Electronics and Informatics. 3. 10.36548/jei.2021.4.005.

[6] S. Asif, W. Yi, Q. U. Ain, J. Hou, T. Yi and J. Si, "Improving Effectiveness of Different Deep Transfer Learning-Based Models for Detecting Brain Tumors From MR Images," in IEEE Access, vol. 10, pp. 34716-34730, 2022, doi: 10.1109/ACCESS.2022.3153306.

[7] Z. Jia and D. Chen, "Brain Tumor Identification and Classification of MRI images using deep learning techniques," in IEEE Access, doi: 10.1109/ACCESS.2020.3016319.

[8] M. Nickparvar, "Brain Tumor MRI Dataset," Kaggle, 2021. DOI: 10.34740/KAGGLE/DSV/2645886. [Online]. Available: https://www.kaggle.com/datasets/masoudnickparvar/brain-tumor-mri-dataset. [Accessed: Dec. 7, 2024].

[9] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” arXiv, 2014. [Online]. Available: https://arxiv.org/abs/1409.1556.