Brain Tumor Diagnosis Using Convolutional Neural Network on Magnetic Resonance Imaging Data

Aryan Mukherjee
University High School, Irvine, CA

Abstract

Brain tumors, characterized by abnormal cell growth, are commonly detected by magnetic resonance imaging (MRI). The primary tumor types– meningioma, glioma, and pituitary tumor– exhibit distinct patterns on MRI scans. This work trains a convolutional neural network (CNN) to perform multinomial classification of brain tumors, using more than 7,000 MRI scans from the publicly available Kaggle dataset.

Keywords: Brain tumor, glioma, meningioma, pituitary tumor, magnetic resonance imaging, convolutional neural network, ReLU activation function

1 Introduction

1.1 Background

According to the National Cancer Institute, over 18,000 people in the United States are projected to die from brain or other nervous system cancers in 2025. From 2015 to 2021, the five-year relative survival rate for these cancers was just 33% [1].

A brain tumor is an abnormal mass of cells in or around the brain. If malignant, such tumors can severely impair brain function and cognitive ability.

Historical records indicate that brain tumors were first described in ancient Egypt before being more rigorously documented during the Renaissance era. However, it was not until the late 19th century that tumors were first operated on successfully, and more advanced diagnosis techniques such as the X-ray and MRI imaging were not widely accepted until the early and late 20th century, respectively. In the mid-1900s, chemotherapy, now a cornerstone of cancer treatment, was introduced, and soon applied to brain tumor management [2].

Brain tumors are broadly classified into two categories: primary and metastatic (secondary) [3]. Although secondary tumors are more common, this research will focus on primary brain tumors due to the greater availability and granularity of data.

Primary brain tumors originate within the brain tissue and can be further subdivided into several major types:

- **Meningioma**, accounting for over 40% of all brain tumors, arises from the meninges—the protective layers surrounding the brain and spinal cord. While typically non-malignant, meningiomas are persistent and difficult to eliminate fully. They are also classified based on their anatomical location within the cranial region [4].

- **Gliomas** are a more aggressive and heterogeneous group of tumors arising from glial cells that support neurons. These tumors can range from low to high grade and are often more invasive and difficult to treat surgically than meningiomas [5].

- **Pituitary tumors** develop in the pituitary gland at the base of the brain. Though generally benign, they can disrupt hormonal balance and lead to serious health consequences if left untreated [6].

In 2023, approximately 94,390 Americans were diagnosed with a primary brain tumor, with nearly 19,000 deaths recorded. In 2024, gliomas accounted for 80% of malignant brain tumors, meningiomas represented over 25% of benign cases, and pituitary tumors comprised 16% of all primary brain tumors [7].

1.2 Literature Review

Extensive literature has been covered on convolutional neural networks, several of which feature tumor-related datasets. Techniques from such literature have been adapted and modified for this study, and many notable examples of such reports are worth highlighting.

For instance, Tiwari et al. [8] used a CNN with four convolution layers, one fully connected layer, and one classification layer to classify 224 x 224 pixel square MRI images from a dataset of size 3264 and four classes: no tumor, meningioma, pituitary, and glioma. Their model runs with 30 epochs and a batch size of 32, uses a categorical cross entropy loss function, and achieves an accuracy in training of 99%.

Kibriya et al. [9] proposed a 13-layer CNN, classifying 3064 MRI images known to correspond with a tumorous patient further into their specific grouping: meningioma, pituitary, or glioma. Their model achieved 97.2% accuracy on the same database, but validation was also performed on a separate database, verifying model robustness.

Jaspin et al. [10] employed a multi-class CNN designed for feature extraction on the Brain Tumor Image Segmentation (BRATS) dataset from 2015, coupled with data from the Figshare data repository. Their model demonstrated low complexity and low output from the loss function, retaining an average accuracy of 97.5% between experiments performed.

Mukkapati et al. [11] performed a 5-class classification of MRIs from four datasets: the reference image dataset for assessing therapeutic response, REMBRANDT, TCGA-LGG, and a set of TI-weighted contrast enhanced pictures. The 5 classes were meningioma, pituitary, glioma, no tumor, and metastatic (secondary) tumor, and performance of the 25-layer CNN was evaluated at 92.98% accuracy.

Singh et al. [12], on several publicly available datasets, compared CNN models for various purposes: binary classification, multinomial classification, and severity classification, achieving 99.74%, 96.26%, and 99.18% accuracies respectively. They employed Hybrid Particle Swarm Grey Wolf Optimization techniques rather than a grid search technique like the previous paper used.

Gautam et al. [13] achieved 93% accuracy with a 6-layer CNN on the BRATS 2020 dataset for 3-class classification, where all 2475 brain scans in the study hosted a tumor of one of three types: meningioma, glioma, or pituitary adenoma. Here, principal component analysis was utilized to optimize model features, as well as for reduction of each image's 256 x 256 square pixel dimensions to 128 x 128 square pixels.

Kumar et al. [14] compiled a 25-layer CNN model that achieved a maximum accuracy of 86.23% when used with an Adam optimizer for hyperparameter tuning. Their model was trained from two public datasets aggregating on 3580 total MRI images. However, its classification task was one of the most precise in the studied literature, forecasting not only tumor type but grade (severity).

Lastly, Ozkaraca et al. [15] delivered a hybrid CNN architecture from the VGG16Net model and the dense-layer structure of DenseNet, both of which implement convolutional neural networks themselves. Through k-fold cross-validation, they identified a consistent success rate of 95-97% for the modified CNN, far surpassing a basic CNN architecture and transfer learning techniques.

1.3 Scope and Structure of the Study

This paper describes the first in a series of three research studies focused on predictive modeling for brain tumor data. Here, we explore convolutional neural networks (CNNs), starting with their theoretical foundations, followed by practical applications and a discussion of results. The second paper will address brain tumor diagnostics using biomedical data, while the third will examine survival prediction using approaches such as the Cox model, parametric regressions, and random survival forests. The latter two papers will be published in subsequent issues of this journal.

2 Multinomial Classification with Convolutional Neural Networks

2.1 Theoretical Framework

A convolutional neural network (CNN) is a type of neural network designed to extract patterns from data structured as matrices, most commonly images, and classify those data based on the identified patterns. In this study, we employ a CNN for multiclass classification of various brain tumor types using MRI images. While CNNs are predominantly used for image analysis, they are also applicable to other data types such as video and audio.

Here we outline the theoretical steps for fitting a convolutional neural network. A more detailed explanation can be found in [16].

To build a CNN, we first prepare and label a training dataset. Each MRI scan is assigned a numeric label corresponding to its class: 0 for no tumor, 1 for glioma, 2 for meningioma, and 3 for pituitary tumor. A model is then initialized with random parameters and trained using gradient descent, an optimization technique that iteratively updates the model to minimize prediction error. This process continues until the CNN can accurately classify new, unseen images.

Specifically, if layer L_n has m nodes and layer L_{n-1} has k nodes, with n representing the number of layers in the CNN and $m = 4$ for our case, then we can represent:

$$L_{n,j} = W_{j,1}L_{n-1,1} + W_{j,2}L_{n-1,2} + \cdots + W_{j,k}L_{n-1,k} + B_j$$

where W and B denote weights and biases that, when aggregated with the connected node from a previous layer, provide the value for the jth node in the nth layer. We can use this to compute the initialized cost function for our network, namely,

$$C_0 = -\sum_{j=1}^{m} y_j \log(L_{n,j})$$

where y_j denotes the desired value at the jth output node. For backpropagation to transpire, we must minimize this cost function using the gradient with respect to each weight and bias, propagating backward in the neural network and repeating this across all layers for our entire training set. Specifically, by the chain rule, we have that

$$\frac{\partial C_0}{\partial L_{n-1,i}} = \sum_{j=1}^{m} \frac{\partial R^{-1}(L_{n,j})}{\partial L_{n-1,i}} \cdot \frac{\partial L_{n,j}}{\partial R^{-1}(L_{n,j})} \cdot \frac{\partial C_0}{\partial L_{n,j}}$$

where R is the ReLU function discussed later in this section, and $i \in [1, k]$. This result enables us to modify the matrix of weights and biases, feeding through the network based on the signs and relative magnitudes of each component in the gradient. Of course, this just shows an example for the cost from a single training image, and by repeating this fine-tuning process and averaging the cost gradient across the training set, we can successfully attain the local minimum.

With this refined model, post-validation, we can analyze the accuracy on a test set of images. Specifically, for our testing set, and more generally, for all images run through a complex CNN, each image must first be converted into matrix form. An image is split into three matrices for red, green, and blue coloring, and each pixel in the image is assigned three unique 8-bit values from 0 to 255 corresponding to the relative intensities of the red, green, and blue components comprising the matrix entries.

Next, the image is run through a series of hidden layers, composed of repeated structures of convolution and pooling layers, and scaled to hyperparameters that match the intent of the study. The purpose of the convolution layer is to identify certain characteristics of the image, called features, which are crucial to pattern recognition. This is done by the use of many kernels, each of which is a matrix weighted by the specific pattern that it attempts to discern, of dimensions orders of magnitude lower than that of the original images. A kernel, with three channels for red, green, and blue, identifies patterns through the dot product, extracting higher values in sections that reflect its designated pattern and lower values for those that don't.

After these kernels are employed to identify horizontal lines, vertical lines, diagonals, and more in the image, the products are then passed through an activation function such as a rectified linear unit (or ReLU for short). This function is necessary because it introduces non-linearity into the network, handling the complex relationship between the input and output nodes that cannot be accounted for by a linear combination of kernel products. That is, if an activation function were not introduced, then the neural network could be condensed into a single layer consisting of repeated matrix multiplication, resulting in a nullified increase in accuracy for each additional layer.

The final layer in this repeating structure is known as the pooling layer, which reduces the dimensionality of the input matrix while preserving the number of channels using a two-dimensional filter for condensation. This effectively reduces the complexity of the entries in each node of the layer, producing a model robust to overfitting or minor image distortion.

By replicating this repeated structure of convolution, activation, and downscaling, the neural network picks up on nuances in the image structure, resulting in refined classification. Finally, the output of the hidden network is flattened into a single dimension, interconnected with all potential relative weights of the final outputs, and standardized into a probability distribution using the softmax function. Rather than aggregation, which may occur in binary regression algorithms, these probabilities simply produce the model's most confident prediction for the class of tumor in the image.

2.2 Analysis and Results

We performed four-class classification on a public brain tumor MRI dataset from Kaggle with 7023 images. The dataset was pre-split into training and testing images, and brain scans were labeled. Between 80-85% of the total data was used for testing, depending on the class of tumor, while the remaining data was allocated for testing.

We used the Keras API in Python with the standard categorical cross-entropy loss function and an Adaptive Moment Estimation (Adam) optimizer to create the model, executing with 15 epochs on the training data and computing accuracy on pre-split testing images. The hyperparameters of our model are depicted pictorially in Figure 1 below.

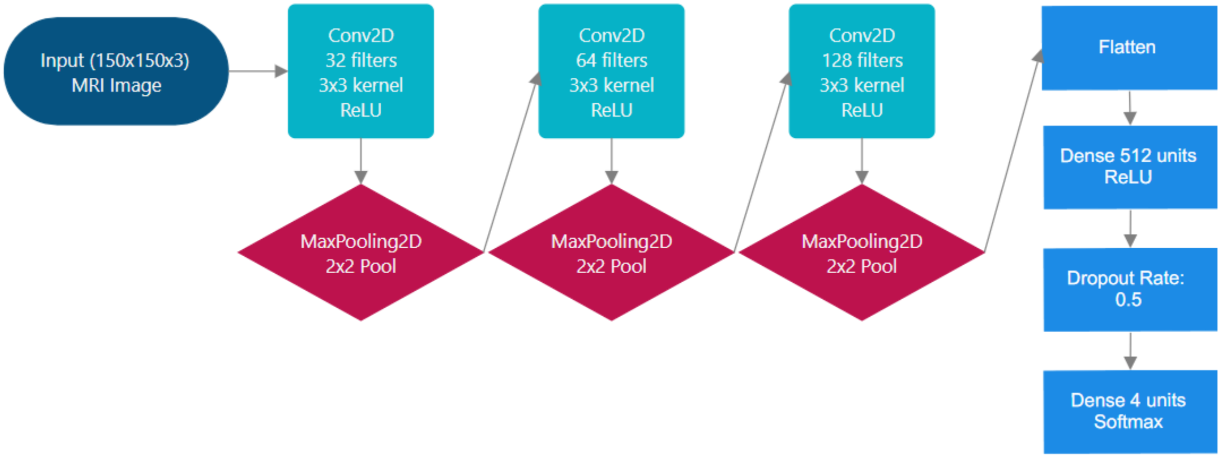

Figure 1: CNN model layers and their parameters.

Our model performed relatively well with respect to the time to run the code, which was approximately 72 minutes total with 277 seconds per epoch. The average accuracy of the model was 93%. Unlike for binary classification problems, further quantifying our results by means of a receiver-operating characteristic curve is not effective or purposeful, and such calculations have been deliberately omitted. Instead, we analyze and account for the regular inaccuracies that occur during repeated execution of the model.

The most frequent error made was incorrectly classifying an MRI image of a meningioma patient as not having meningioma. One such instance where this was detected is shown in

Figure 2 below.

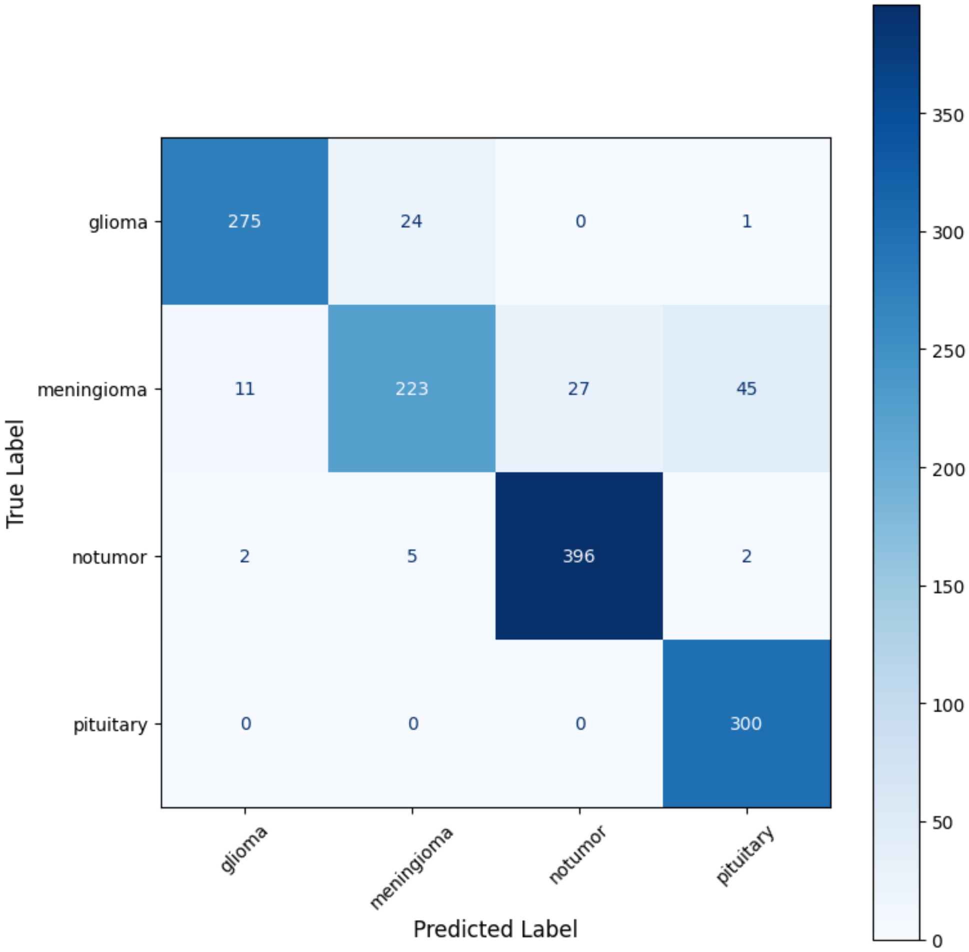

Figure 2: Confusion matrix summarizing CNN model's performance on test set.

To visualize possible sources of this error, we showcase three randomly selected MRI images from the test set which were incorrectly classified as having non-meningioma classes. The images and their true classes are displayed in Figure 3 below.

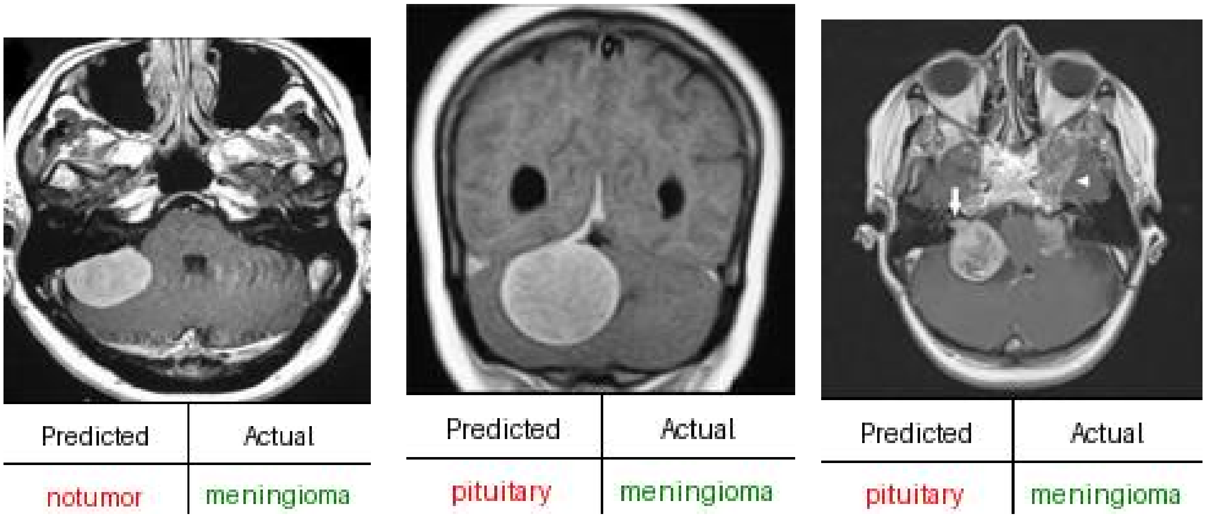

Figure 3: Sample erroneous predictions on test images and their true classes.

The mischaracterization as glioma can be attributed primarily to location, as glioma tumors may form near the meninges region of the brain. This would make it difficult for the model to differentiate patterns on the basis of location alone. Although this can occur for pituitary tumors as well when a meningioma congregates near the sella turcica compartment of the brain, the natural error from signal effects of scans coupled with the homogeneous morphology of the tumors most plausibly gives rise to inaccurate classification of that kind.

Lastly, the misclassification as "no tumor" can be explained by the low color contrast of some meningioma tumors with respect to the surrounding gray matter in the MRI, coupled with the smaller size of tumors considered as benign or low-grade. Conversely, it is worth noting that the label that was classified most accurately and precisely by the model was"pituitary". This tumor class was the most strongly defined by placement of the tumor in the brain, indicating a strong feature importance of location for CNN-based classification models.

3 Summary and Conclusion

In summary, multinomial classification was performed on an MRI dataset containing over 7,000 images using a convolutional neural network (CNN) implemented in Python. The architecture included three convolution-pooling layer sets and a ReLU activation function. The model achieved a test accuracy of approximately 93%, demonstrating strong and consistent performance.

We also found that the most frequent type of misclassification was a type I error (false positive) under the null hypothesis that the tumor corresponds to a meningioma. This is likely due to similarities in location or morphology shared with other primary tumors. Such errors may be mitigated in post-analysis through careful image inspection and by incorporating additional diagnostic metrics (e.g., tumor severity or size).

Our model's accuracy could potentially be increased by incrementing model hyperparameters, which would consequently increase the run time. Additionally, implementing a larger training dataset would help offset the effects of bias or overfitting, having a similar effect on accuracy. These results have instrumental applications in tumor classification and the fields of radiology and oncology, offering a supportive tool for verifying diagnostic accuracy in clinical settings.

4 Future Work

Following a classification task, one natural next step would be to perform localization and instance segmentation of tumors within the MRI image dataset. This can be approached at two levels: drawing bounding boxes and delineating contour lines. While the former is simpler, it offers less precision; both methods, however, can be implemented using convolutional neural networks (CNNs). Furthermore, alternative models such as the Segment Anything Model (SAM) may be explored, with their prediction accuracies compared on a testing set.

5 Supplemental Materials

The code developed for this paper is available in this GitHub repository. To access the data file directly on Kaggle, use this link. To run the code directly, use this Google Colab link.

Acknowledgements

Sincere gratitude is extended to Dr. Olga Korosteleva of California State University, Long Beach, for her invaluable mentorship throughout this project. Her guidance, support, and the resources she provided were instrumental in developing a deeper understanding of statistics and machine learning.

Appreciation is also given to Eric Shulman, mathematics teacher at University High School, for his instruction in calculus and matrix theory, which provided foundational knowledge essential to the research.

Thanks are due to the many other teachers whose instruction was influential, to peers for their support and collaboration, and to family members for their continuous encouragement throughout the course of this project.

References

[1] National Cancer Institute. Cancer of the brain and other nervous system - cancer stat facts, 2018. SEER.

[2] H. S Greenberg, W. F Chandler, and H. M Sandler. *Brain Tumors*. Oxford Academic, 1999.

[3] H. H. Publishing. Brain tumor overview, 2023. Harvard Health.

[4] John Hopkins Medicine. Meningioma, n.d. Johns Hopkins Medicine.

[5] John Hopkins Medicine. Gliomas, 2019. Johns Hopkins Medicine.

[6] Pituitary tumors — national institute of neurological disorders and stroke, n.d. National Institute of Neurological Disorders and Stroke.

[7] Daines, Markey, Rubio, Warren, Sinema, and Warnock. S.res. 229 (118th): A resolution designating may 2023 as “national brain tumor awareness month”, 2023. 118th Congress 1st Session.

[8] P. Tiwari, B. Pant, M. M. Elarabawy, M. Abd-Elnaby, N. Mohd, G. Dhiman, and S. Sharma. Cnn based multiclass brain tumor detection using medical imaging. *Computational Intelligence and Neuroscience*, 2022:1830010, 2022.

[9] Hareem Kibriya, Momina Masood, Marriam Nawaz, and Tahira Nazir. Multiclass classification of brain tumors using a novel cnn architecture. *Multimedia Tools and Applications*, 81(21):29847–29863, 2022.

[10] K. Jaspin and Shirley Selvan. Multiclass convolutional neural network based classification for the diagnosis of brain mri images. *Biomedical Signal Processing and Control*, 82:104542, 2023.

[11] N. Mukkapati and M. S. Anbarasi. Multi-class classification framework for brain tumor mr image classification by using deep cnn with grid-search hyper parameter optimization algorithm. *International Journal of Computer Science and Network Security*, 22(4):101–110, 2022.

[12] Ngangbam Singh, N. Merlin, Thandaiah Prabu, Deepak Gupta, and Meshal Alharbi. Multi-classification of brain tumor by using deep convolutional neural network model in magnetic resonance imaging images. *International Journal of Imaging Systems and Technology*, 34:n/a–n/a, 08 2023.

[13] Suruchi Gautam, Sweety Ahlawat, and Prabhat Mittal. Binary and multi-class classification of brain tumors using mri images. *International Journal of Experimental Research and Review*, 29:1–9, 12 2022.

[14] Sunil Kumar and Dilip Kumar. Human brain tumor classification and segmentation using cnn. *Multimedia Tools and Applications*, 82(5):7599–7620, 2023.

[15] O. Özkaraca, O.İ. Bağrı açık, H. Gürüler, F. Khan, J. Hussain, J. Khan, and U.E. Laila. Multiple brain tumor classification with dense cnn architecture using brain mri images. *Life*, 13(2):349, 2023.

[16] Şaban Öztürk. *Convolutional Neural Networks for Medical Image Processing Applications*. CRC Press, 2022.