Brain Tumor Detection using CNN

Brain tumor diagnosis from magnetic resonance imaging (MRI) is traditionally a manual, time-consuming process that depends on the availability of trained radiologists. This project presents an automated brain tumor detection system built on convolutional neural networks (CNNs), one of the most widely studied applications of deep learning in medical imaging. The objective is to classify brain MRI scans into tumor and non-tumor categories while making the model's decisions interpretable — a critical requirement for any medical-image analysis system.

The pipeline begins with preprocessing: scans are resized to a uniform input resolution, pixel values are normalized, and data augmentation (random rotations, flips, and zooms) is applied to improve generalization and reduce overfitting. Two model variants are implemented: a custom CNN baseline designed and trained from scratch, and a transfer-learning variant built on the pretrained VGG16 backbone with fine-tuned top layers, enabling a direct comparison of accuracy versus training cost.

A key contribution of this work is explainability. The system integrates Gradient-weighted Class Activation Mapping (Grad-CAM), which back-propagates the predicted class score to the final convolutional layer to produce a heatmap of the image regions that most influenced the classification. This visual explanation lets users verify that the model attends to anatomically plausible tumor regions rather than background artifacts.

The trained models are evaluated on a held-out test set using accuracy, precision, recall, F1-score, and a confusion matrix. Finally, a Flask-based web application — branded NeuroScan — provides an upload-and-diagnose interface: a user submits an MRI scan and receives the predicted class, a confidence score, and a side-by-side Grad-CAM overlay. This project is intended strictly as an academic and research demonstration and is not a validated clinical diagnostic tool.