Brain Tumor Segmentation with Lightweight U-Net

Project Overview

This project applies deep learning to MRI brain tumor segmentation using the LGG MRI Segmentation Dataset from Kaggle. The goal is to detect and delineate tumor regions from FLAIR MRI scans to assist medical diagnosis and treatment planning. The solution uses a Lightweight U-Net model implemented in PyTorch, trained and evaluated on paired MRI and tumor mask images.

Dataset

- Name: LGG MRI Segmentation
- Source: Kaggle
- Format: .tif images, each paired with a corresponding _mask.tif file
- Total pairs: 3,929 (image + mask)
- Image size: Predominantly 256×256 pixels
- Class distribution: 65% of slices contain no tumor (zero-mask slices)

Exploratory Data Analysis (EDA)

Key findings:
- All images are consistently 256×256, simplifying preprocessing.
- Masks are binary (tumor vs. background), with no corrupt files.
- Dataset imbalance exists — majority of slices have empty masks.
- Visual inspection confirmed correct alignment between MRI images and their masks.
- Sample MRI and mask visualizations: images/eda_sample_grid.png, images/val_predictions.png

Model Architecture

Lightweight U-Net featuring:
- Encoder: 4 downsampling blocks (DoubleConv + MaxPooling)
- Bottleneck: DoubleConv with 512 filters
- Decoder: 4 upsampling blocks (ConvTranspose2d + DoubleConv)
- Output: 1-channel convolution for binary mask prediction
- Total parameters: ~7.77 million

Advantages:
- Preserves U-Net's skip connections for fine detail
- Reduced complexity for faster training in Colab

Training Setup

- Loss (Baseline): BCE + Dice
- Loss (Experiment): Weighted BCE + Dice (pos_weight=3.0)
- Optimizer: Adam (lr=1e-3)
- Scheduler: ReduceLROnPlateau (factor=0.5, patience=3)
- Metrics: Dice Coefficient, IoU
- Data Augmentations: Random flip, rotation, normalization
- Splits: Train 2,750 | Val 589 | Test 590

Results

Baseline (BCE + Dice)
- Best Val Dice: ~0.7752
- Best Val IoU: ~0.7358
- Saved model: best_unet.pth
- Curves: images/seg_loss_curves.png, images/seg_val_dice_iou.png

Weighted BCE + Dice
- Best Val Dice: ~0.7296
- Best Val IoU: ~0.6945
- Saved model: best_unet_weighted.pth
- Curves: images/weighted_seg_loss_curves.png, images/weighted_seg_val_dice_iou.png

Qualitative Results

Validation predictions show the model:
- Correctly identifies tumor regions in most positive slices
- Avoids false positives in empty slices
- Weighted loss improves detection of smaller tumors but lowers overall Dice
- Example outputs: images/val_predictions.png

Key Takeaways

- Lightweight U-Net is effective for brain MRI segmentation with limited compute.
- Class imbalance impacts Dice/IoU scores.
- Weighted loss improves tumor pixel focus, but not overall validation Dice.
- Tumor-only metric tracking recommended for future work.

Citations

- Buda, Mateusz, et al. “LGG MRI Segmentation Dataset.” Kaggle, 2019. https://www.kaggle.com/datasets/mateuszbuda/lgg-mri-segmentation
- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. “U-Net: Convolutional Networks for Biomedical Image Segmentation.” MICCAI, 2015. https://arxiv.org/abs/1505.04597
- PyTorch. “BCEWithLogitsLoss.” https://pytorch.org/docs/stable/generated/torch.nn.BCEWithLogitsLoss.html
- Hugging Face. “Medical Imaging Models and Datasets.” https://huggingface.co/models?pipeline_tag=image-segmentation&search;=medical