

Biophysics Informed Pathological Regularisation for Brain Tumour Segmentation -Supplementary-

Lipei Zhang, Yanqi Cheng, Lihao Liu,
Carola-Bibiane Schönlieb, and Angelica I Aviles-Rivero

Department of Applied Mathematics and Theoretical Physics,
University of Cambridge, UK
{lz452,yc443,11610,cbs31,ai323}@cam.ac.uk

1 More Details of Experimental Setting

Preprocessing and Augmentation: A variety of augmentation techniques were implemented to enhance model robustness. These included random rotations of 90° in the (x,z) plane with a 70% probability, flips along all axes (70%), intensity shifts (70%), and contrast adjustments using a gamma range of 0.5 to 4.5 (20%). Additionally, Gaussian noise was added based on a normal distribution with a standard deviation of 0.1 (50% probability), and Gaussian smoothing was applied (20% probability). Z-score standardisation was performed on non-zero voxels independently for each channel, clipping outliers exceeding $3\times$ the standard deviation.

Test-Time Augmentation (TTA): To improve segmentation accuracy, each sample underwent TTA using 16 distinct augmentations, including operations in two planes (axial and coronal), two horizontal flips, and four rotations ($0^\circ, 90^\circ, 180^\circ, 270^\circ$).

Postprocessing: potentially presented background voxels in the 3-channel volume (TC, WT, ET) of MRI images were removed. The segmentation map was then reconstructed into a 3D volume using original labels (NCR/NET- label 1, ED - label 2, ET - label 4). Boolean operations were employed to exclude ET from the TC (label 1) and TC from the WT (label 2). Isolated connected components smaller than 20 voxels were eliminated, and any ET regions smaller than 300 voxels were interpolated from the nearest axial plane values using 2D nearest interpolation.

2 More Details of Architecture

Tumour Cell Density Estimator: The biophysics informed regularisation model features three fully connected layers interspersed with two sine activation functions, configured as $nn.linear + sine + nn.linear + sine + nn.linear$. The time matrix T consistently uses a value of 0.1, and the hidden size for all MLPs is set to 128.

Baseline Network Architectures:

- **UNet and R2UNet:** These models start with 32 feature maps and include three down-sampling and up-sampling stages per axis, utilizing ReLU activation functions.
- **nn-UNet:** Reimplemented using MONAI’s DynUNet, this setup follows the original specifications with a kernel size of [3, 3, 3], five down-sampling and up-sampling steps, instance normalisation, and no deep supervision, aligning with the BraTS challenges standards.
- **UNet-TR:** This configuration adheres to the original design with a feature size of 16, a hidden size of 768, an MLP dimension of 3072, 12 heads, perceptron-based positional embedding, and instance normalisation.
- **SegResNet and SegResNet-VAE:** While based on default MONAI settings, instance normalisation was adapted for consistency across all models.