Tri-VAE: Triplet Variational Autoencoder for Unsupervised Anomaly Detection in Brain Tumor MRI

Supplementary Material

1. Additional Ablation

The risk of adding healthy L1 reconstruction loss. In our experimental investigation, we observed that the inclusion of the reconstruction loss for healthy brain slices at full scale led to an unintended outcome - the network learned an identity mapping. This essentially resulted in the reconstruction of identical images, rendering the model ineffective in detecting anomalous regions. The subsequent Fig. 1 illustrates the differences between the reconstructed images and segmentation results of our current model compared to the one incorporating the healthy brain slices reconstruction loss.

The introduction of the reconstruction loss for healthy brain slices compelled the network to map these already healthy slices back to themselves. This, in turn, exerted a strong constraint on the learning process, with the loss signal favoring identity mapping. Consequently, the network might opt for the simpler route of copying the input to the output (identity mapping) instead of acquiring a more intricate representation capable of distinguishing between healthy and noisy brain slices. If a substantial portion of the model's capacity is allocated to minimizing the reconstruction loss for healthy samples, there might be insufficient capacity remaining to effectively focus on discerning the semantic disparities between healthy and noisy brain slices. This underscores the importance of careful loss function design to strike an appropriate balance in the model's learning objectives.

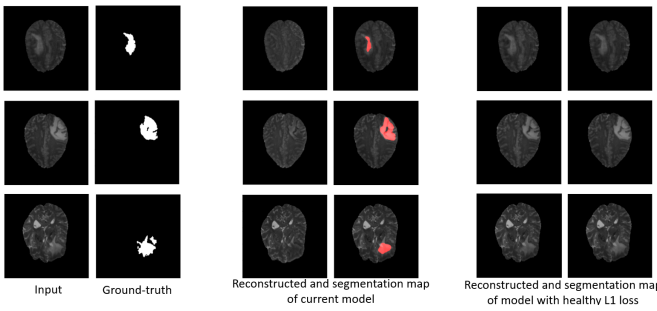

Figure 1. The risk of adding healthy reconstruction loss.