

Fig. 1. Example segmentation results comparing modality dropout (MD) with our proposed method on the MSSEG-2016 dataset. True positives and false positive/negatives are overlaid on the FLAIR image. Four sets of segmentation results are shown for different input modality configurations.

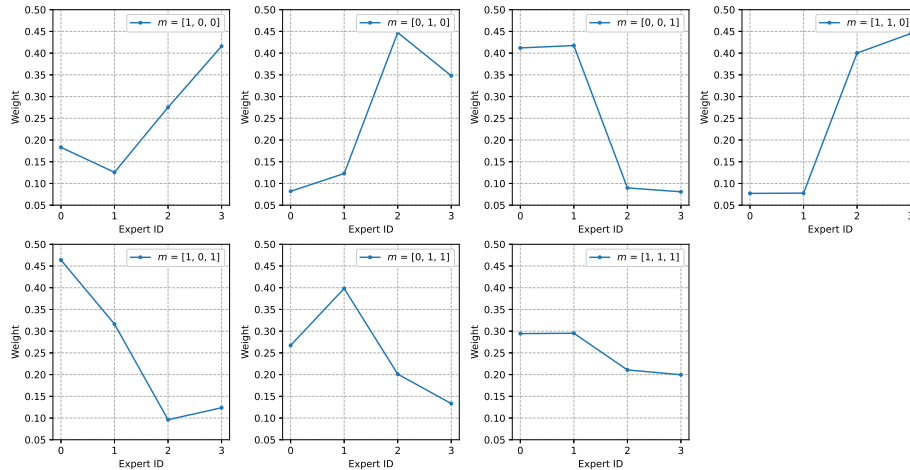


Fig. 2. Example expert activations for an MoE layer replacing the UNet stem convolutional layer. Each input modality configuration m produces a different vector of expert weights. While the experts tend to naturally specialize on different tasks, there is still considerable overlap between the distributions. Note that the template for m is [T1w, T2w, FLR].