

CHAOS challenge

I am sending our results in the CHAOS challenge for evaluation in tasks 3 and 5. We are part of the Biomedical Computer Vision group at Universidad de los Andes, Colombia. I have a M.Sc. in Biomedical Engineering oriented to image analysis and machine learning, and my team mate is doing a Ph.D. in Engineering. Our method is fully automatic, and consists of a single model with less than 4.5 million parameters. Our network is based on DeepLab V3+. It has an Xception encoder with five stages separated by strided convolutions. As the network goes deeper, the amount of layers in each stage increases by 1 and the number of feature maps in the layers doubles. The decoder uses trilinear upsample to match the size of the previous stage, concatenates the results from both stages and then combine the outputs with a 3x3x3 convolutional layer. This process is repeated until the original resolution is recovered. As a preprocessing step, we adjusted the (x, y, z) spacing of the volumes to (1.66, 1.66, 3), and adjusted the size of the volumes to the size of T2 sequences to train the model with multimodal inputs. We merged the labels of T1 and T2 and defined their union as the ground truth. As the loss function, we combine the Dice Loss and Cross Entropy. We split the train patients into training and validation sets, to adjust the learning rate schedule during the training phase.