

This submission is inspired by Fabian Isensee's nnU-Net. We utilized two cascaded 3D U-Net to segment the prostate, with the 1st 3D U-Net predicting the approximate bounding box of the prostate regions, which then served as input of the 2nd 3D U-Net.

Data preprocessing included: a) Undersampling for 1st U-Net, and zero padding for 2nd U-Net. Since the output of the 2nd U-Net was adopted as final result, the principle is to avoid any resampling of the bounded regions containing the segmentation target as much as possible. b) Normalization to zero mean and unit variance.

For data augmentation, we adopted data flipping, rotation and deformation. All augmentations were performed on the fly during training.

After training, a simple strategy of largest connected component was used to remove potential tiny objects surrounding the predicted prostate regions.

Overall, the current model achieved a dice score of 88.6 during training phase. We'd anticipate a slight drop of the dice score for the unseen testing data.