Long-Range 3D Self-Attention for MRI Prostate Segmentation

Federico Pollastri, Marco Cipriano, Federico Bolelli, and Costantino Grana

Università degli Studi di Modena e Reggio Emilia, Italy

Abstract

The problem of prostate segmentation from Magnetic Resonance Imaging (MRI) is an intense research area, due to the increased use of MRI in the diagnosis and treatment planning of prostate cancer.

The complexity of the task and the 3D nature of the data make 2D segmentation algorithms suboptimal for the task.

- We design a novel 3D segmentation Convolutional Neural Network (CNN).
- We present a data pre-processing algorithm, to extrapolate a "deep learning-functional" dataset from a publicly available collection of MRI scans.
- In order to grant reproducibility and encourage further advance on this subject, the code for both the data refinement process and the proposed model is publicly available[1].

[1] github.com/PollastriFederico/3D-self-attention

Datasets

- PROMISE12 Dataset

  - T2-weighted MR images of the prostate
  - 50 total volumes
  - Data in MHD/RAW format
  - Non-consistent slices number
  - Squared slices

- MRI-US 2021 Dataset

  - T2-weighted MR images of the prostate
  - 911 total volumes
  - 60 slices for volume
  - Squared slices
  - MRI scans in DICOM file format
  - 711 volumes for training
  - 160 volumes for testing
  - 40 volumes for validation
  - Prostate annotation in STL file format

The Proposed Architecture

- The proposed 3D architecture can be divided in three main components: the feature extraction CNN, the Long-Range 3D Self-Attention Block, and the decoder.
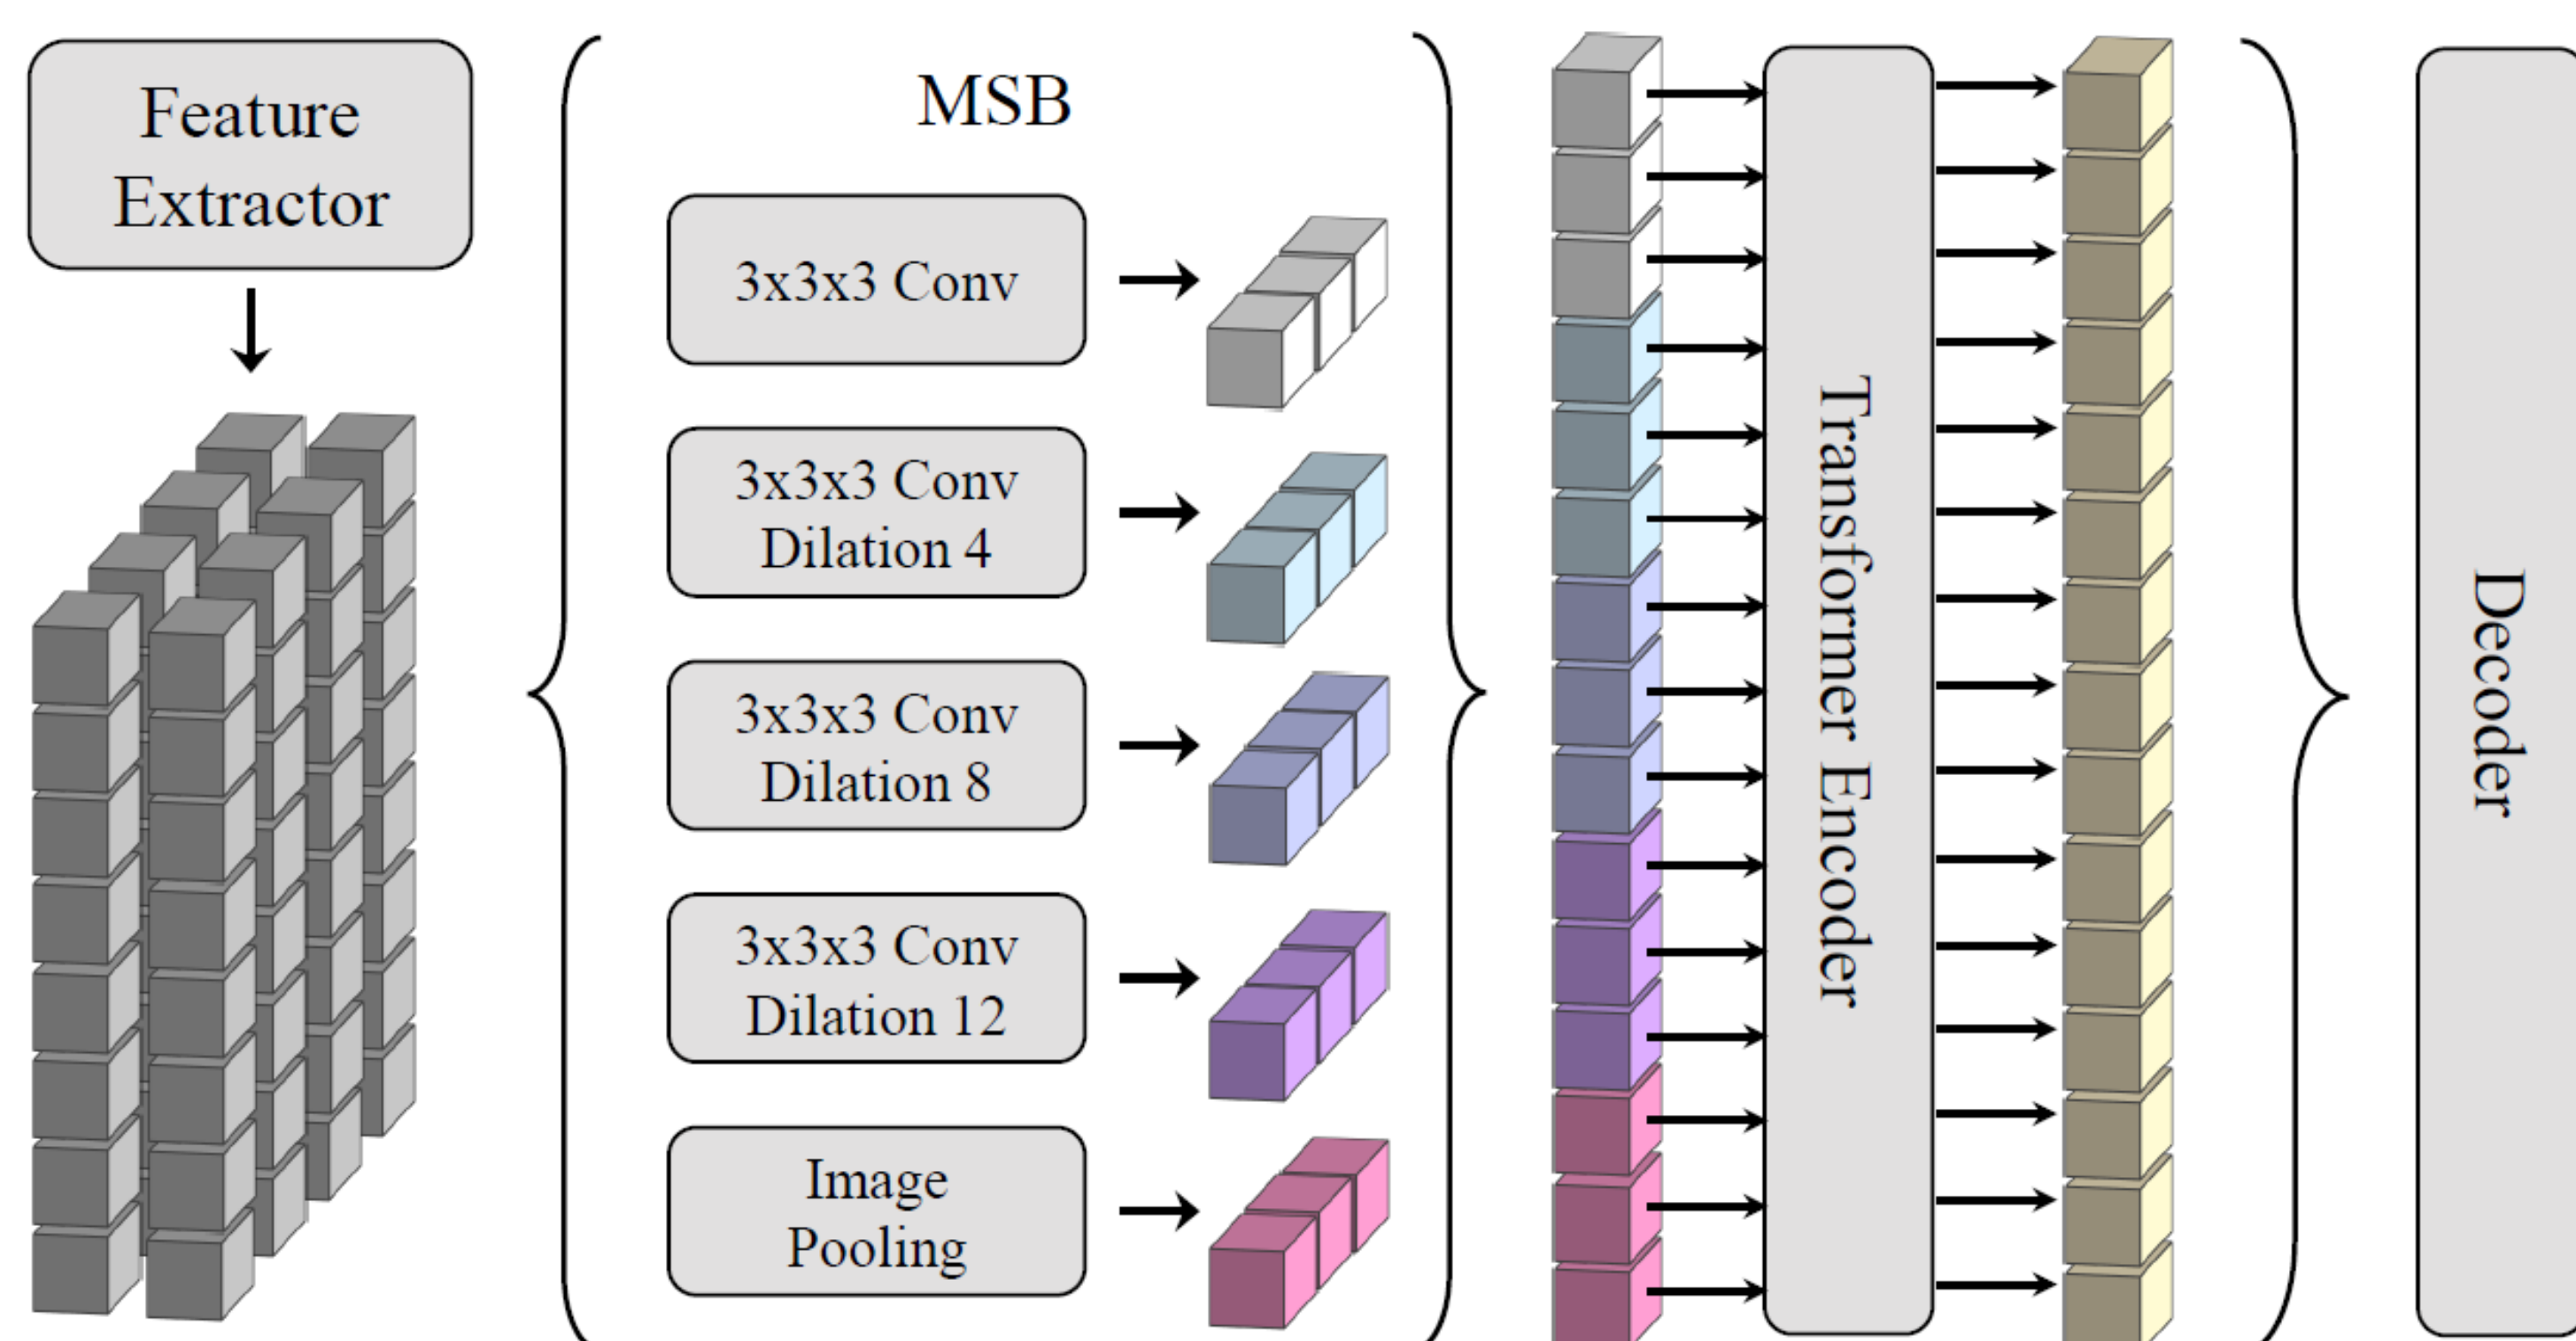
Qualitative Results
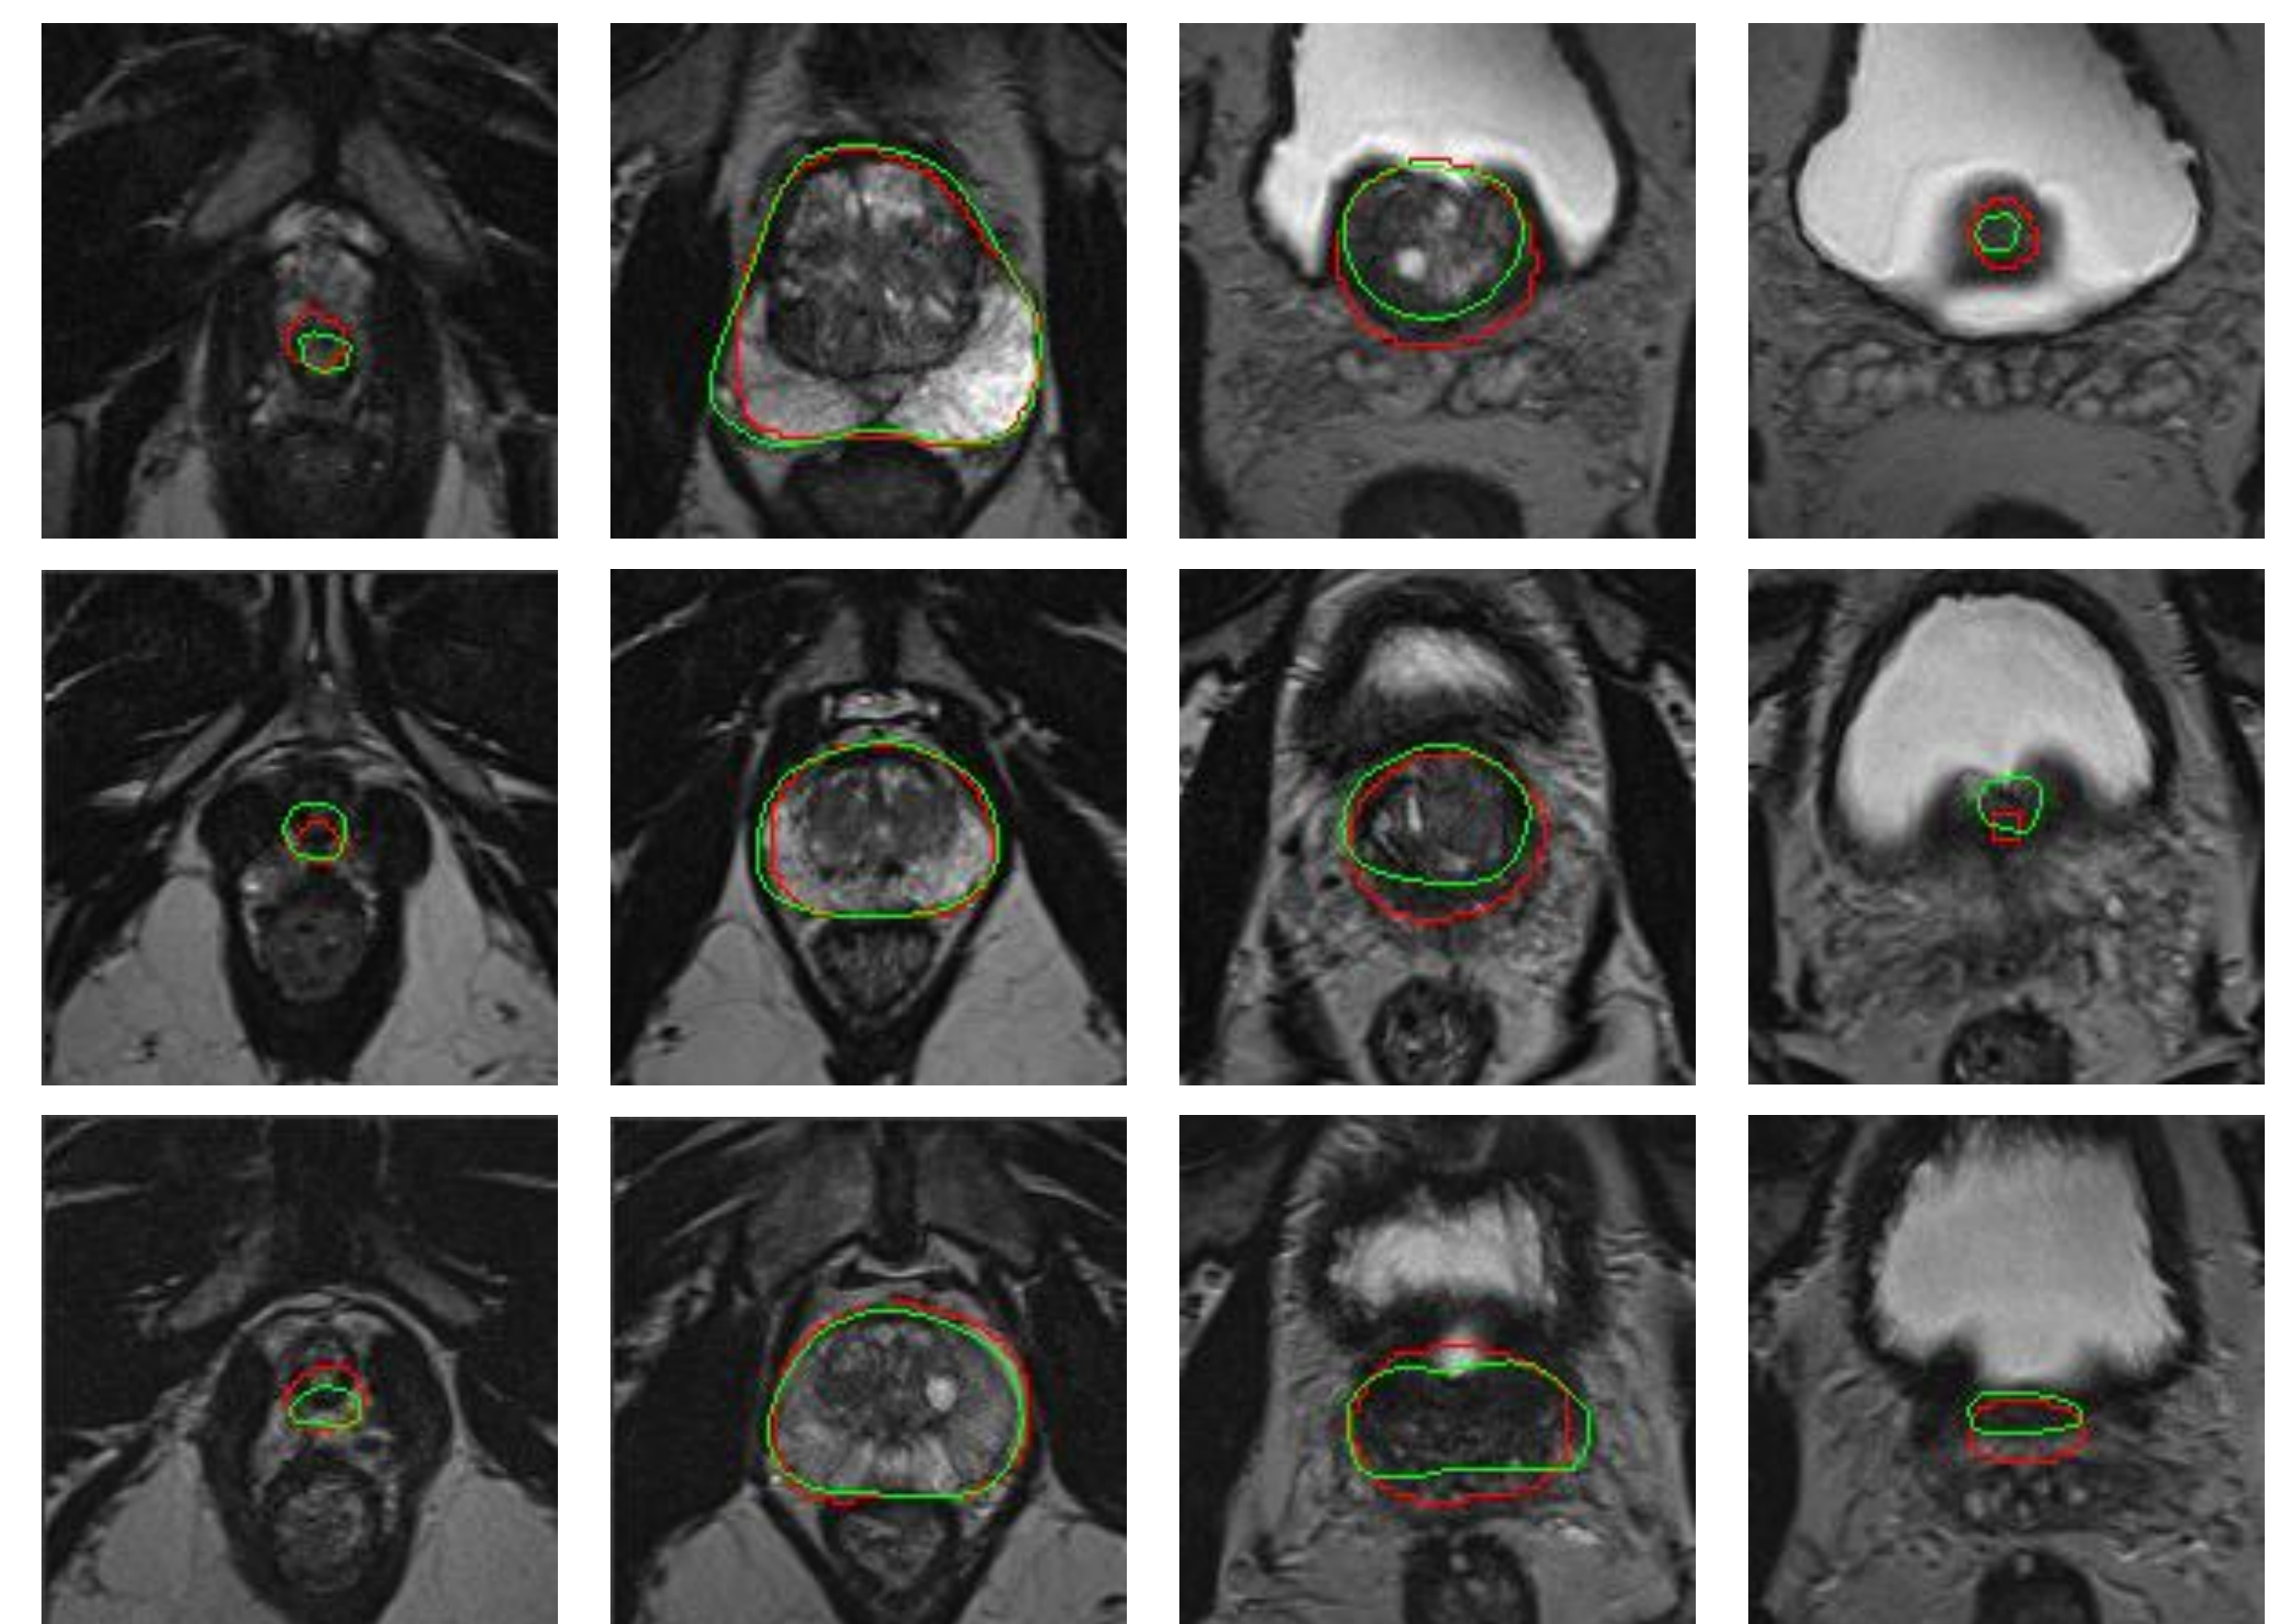
- Results visualization for three volumes, at four different depths. The ground truth is depicted in green, model prediction in red. The first and last column show, respectively, the most inferior and superior appearance of the prostate in the MRI scan.

Quantitative Results

- The PROMISE12 section display the results obtained by the networks when trained and tested using only the PROMISE12 dataset

- The last 4 columns evaluate the models when pre-trained with the Prostate-MRI-US-Biopsy and fine-tuned for the PROMISE12 dataset.

- When processing volumes with variable resolutions, slice metrics favor 2D architectures over 3D ones, since the resampling operations along the z axis can generate minor errors in slices with no foreground

The proposed method outperforms every competitor

- 3D Competitors & 2D Competitors

  - V-Net
  - Med3D
  - U-Net3D
  - DeepLabv3+
  - U-Net

- Two Metrics

  - DICE score
  - Intersection over Union (IoU)

- Two Metric-Aggregation Strategies

  - Per Volume
  - Per Slice

Method	Prostate-MRI-US-Biopsy				PROMISE12				Fine-Tuned for PROMISE12			
	Volume IoU	Slice IoU	Volume DICE	Slice DICE	Volume IoU	Slice IoU	Volume DICE	Slice DICE	Volume IoU	Slice IoU	Volume DICE	Slice DICE
Ours	0.846	0.859	0.916	0.895	0.716	0.726	0.834	0.775	0.785	0.807	0.880	0.847
V-Net	0.822	0.840	0.901	0.880	0.390	0.463	0.551	0.534	0.692	0.761	0.815	0.811
Med3D	0.822	0.840	0.901	0.880	0.653	0.714	0.787	0.762	0.736	0.776	0.847	0.821
U-Net3D	0.822	0.840	0.901	0.880	0.482	0.519	0.635	0.584	0.704	0.740	0.824	0.790
DeepLabv3+	0.826	0.841	0.904	0.880	0.701	0.735	0.821	0.782	0.759	0.803	0.862	0.848
U-Net	0.776	0.810	0.871	0.855	0.699	0.779	0.820	0.822	0.763	0.814	0.865	0.857