2D-3D Fully Convolutional Neural Networks for Cardiac MR Segmentation

Jay Patravali, Shubham Jain, Sasank Chilamkurthy | Qure.ai

MOTIVATION AND OBJECTIVE

We developed a fully-automated 2D and 3D CNN models designed to segment the Left Ventricle, Right Ventricle and Myocardium.
We design and compare the performance of 2D and 3D segmentation models such that they are:

- Light-weight and modular in architecture
- Easy to implement
- Ease of deployment in clinical settings.

We show that both our segmentation models achieve near state-of-the-art performance scores in terms of distance metrics and have convincing accuracy in terms of clinical parameters.

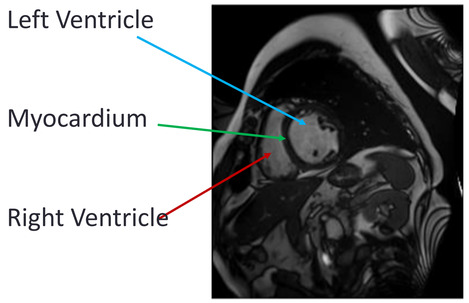

DICE LOSS OPTIMIZATION

Compared the performance of three loss functions:

- Cross Entropy Loss (weighted)
- Log Dice Loss (weighted)
- Hybrid Combination of both

Log Dice Loss:

Dice loss for class i is given by,

$$L_{\text{dice}} = \sum_i w_i l_i \quad l_i = \log\left(2 - \frac{\sum_x t_i(x) p(x,i) + \epsilon}{\sum_x t_i(x) + p(x,i) + \epsilon}\right)$$

SEGMENTATION RESULTS

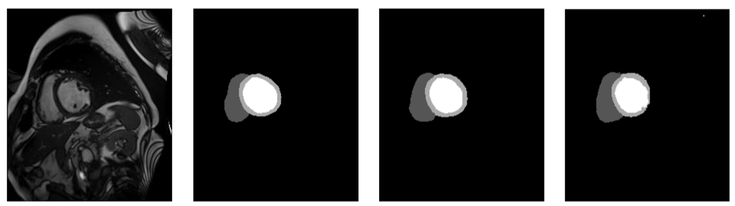

	2D Model	3D Model
Inference Stage (complete phase)		
Time	0.3s	0.3s
GPU Memory	1.2 GB	2 GB
Training Stage		
Time	2.9 hours	2.6 hours
GPU Memory	4 GB	4 GB

Complete Phase

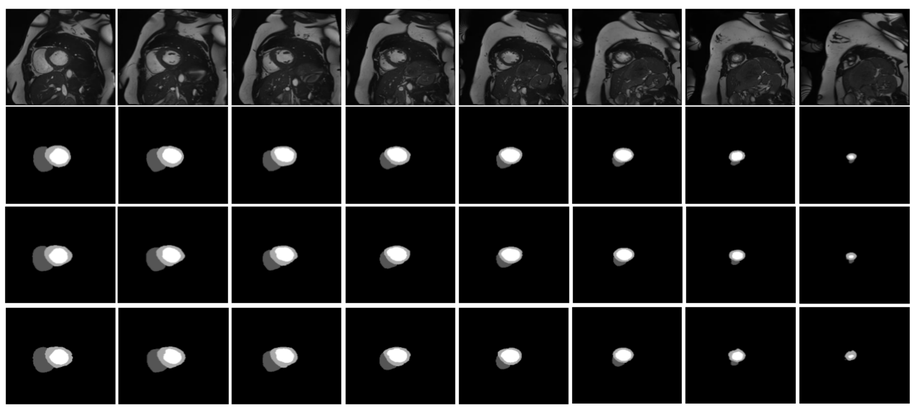

- Raw MR Scan, slice 0-8
- Ground Truth
- 2D model output
- 3D model output

2D SEGMENTATION

- Similar to U-Net architecture
- Option to feed *varying number (N)* of image slices that can be passed as input channels to the model
- On left side is the “contracting” stage and on the right side is “expanding” stage. At the bottom is a base layer.
- Each contraction block doubles the feature maps and downsamples by 2
- Each expanding block concats feature maps, reduces number of feature maps and upsamples by 2.
- Base layer has 512 feature maps

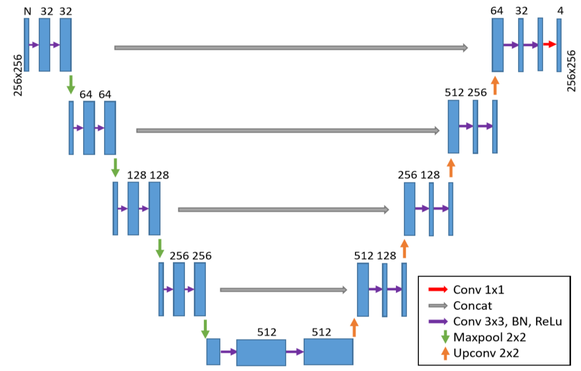

Fig. 1. 2D Model Architecture.

	Dice Score						Hausdorff Distance					
	LV		RV		MYO		LV		RV		MYO	
	ED	ES	ED	ES	ED	ES	ED	ES	ED	ES	ED	ES
CE Loss	0.95	0.90	0.87	0.76	0.79	0.82	13.92	17.67	27.40	27.73	24.81	22.11
Dice Loss	*0.95*	*0.90*	*0.90*	0.79	*0.86*	*0.88*	9.51	12.29	16.1	20.38	*13.45*	*14.88*
Dice-CE Loss	0.95	0.90	0.89	*0.81*	0.83	0.84	*9.15*	*11.7*	*16.0*	*18.22*	13.87	15.35

3D SEGMENTATION

- Similar to 3D U-Net architecture
- An extension to our 2D model with few modifications
- Replace all 2D operations with its 3D counterparts
- We apply 3D 1*2*2 maxpooling operation only in X and Y leaving out the Z dimension. It allows our 3D model to accept input volumes of varying slices at inference stage.
- Base layer has 256 feature maps

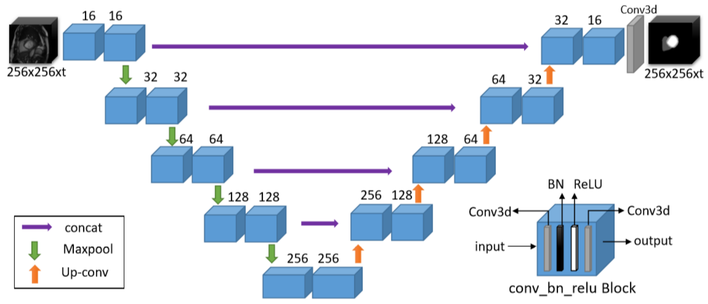

Fig. 2. 3D Model Architecture.

	Dice Score						Hausdorff Distance					
	LV		RV		MYO		LV		RV		MYO	
	ED	ES	ED	ES	ED	ES	ED	ES	ED	ES	ED	ES
CE Loss	0.94	0.89	0.86	0.73	0.76	0.81	12.36	14.11	25.85	29.57	43.47	43.82
Dice Loss	*0.95*	*0.90*	*0.91*	*0.83*	*0.85*	*0.86*	14.95	14.35	*23.15*	*22.14*	*37.75*	*38.50*
Dice-CE Loss	0.94	0.89	0.91	0.81	0.83	0.85	*10.71*	*11.52*	38.01	32.26	43.28	44.98

Contact: shubham.jain@qure.ai