Automated Detection of Left Ventricle in Arterial Input Function (AIF) Image Series for Cardiac MR Perfusion Imaging: A Large Study on 13K Patients

CS 230
Spring 2019

Hui Xue
hxue@Stanford.edu
National Heart Lung and Blood Institute, National Institutes of Health

NIH National Heart, Lung, and Blood Institute

Cardiac MR perfusion imaging

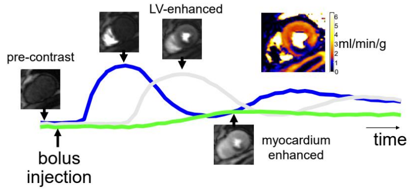

Quantify the blood flow in myocardium requires arterial input function

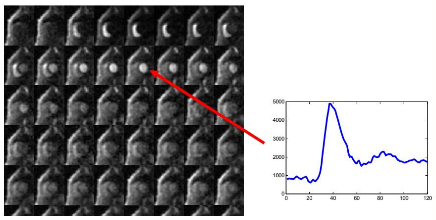

High accuracy to detect of AIF LV

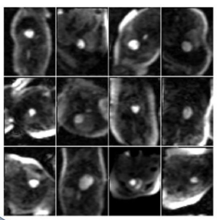

- Anatomical and imaging variation
- High accuracy required by stress test
- Run in hospitals →Deploy the NN

Large data cohort

- Three hospitals: Barts Heart Center, London, UK, Barts; Royal Free Hospital, London, UK, RFH; University of Leeds Hospital, Leeds, UK, ULH

Centers	#Pts	#scans	#MR	Duration
Barts	10,128	20,346	5	20160429 - 20190215
RFH	2,273	3,267	2	20160609 – 20190214
ULH	583	1,414	1	20160524-20190214
Total	**12,984**	**25,027**	8	(22,941 for Tra; 2086 for Dev)

- Test set: 429 patients , 880 perfusion scans 20190216-20190314

Speed up data labelling

- Heuristic method was implemented to help mark LV
- **Manual check for failed detection and correct if needed**
- ~200hrs to label data cohort

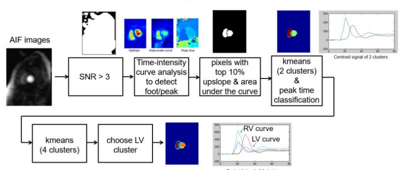

Model and loss function

- A variation of U-Net [1]
- ResNet module as building block [2]
- Both binary and 3 class segmentation (Background, LV, RV)
- Loss function: combined cross-entropy and log IoU [3]

$$J_{IoU} = \frac{1}{N}\sum_{i=0}^{N-1}\frac{y^i\hat{y}^i}{y^i+\hat{y}^i-y^i\hat{y}^i} \qquad J = J_{cross-entropy} - \lambda \cdot \log(J_{IoU})$$

N is the number of pixels

ADAM optimization(0.9, 0.999, 1e-8), λ=0.5, learning rate 1e-3, reduce by x2 every 10 epochs

- Hyper-parameter search for:
  Number of CONVs in ResNet Module (5 – 9)
  Number of filters for each CONV (64 – 256)

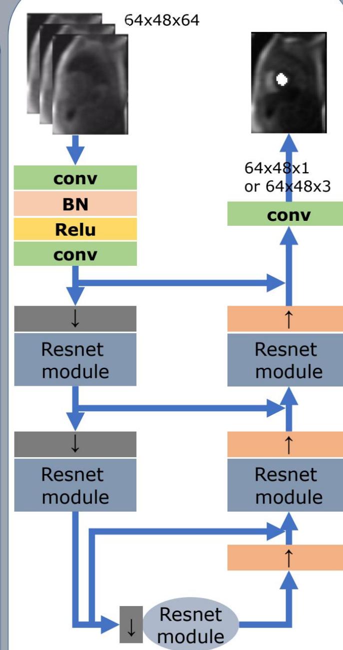

Residual U-Net architecture

Architecture of a single Resnet module

Results

- Dice > 0.5 → success
- 3-class seg
  Tra: 99.956% (10 out of 22,941);
  Dev: 99.856% (3 out of 2,086)
- Binary seg
  Tra: 99.865% (30 out of 22,941)
  Dev: 99.808% (4 out of 2,086)
- 3-class seg, Test performance
  99.77% success (878 out of 880 cases)

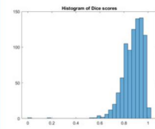

93.1% cases with dice higher than 0.75

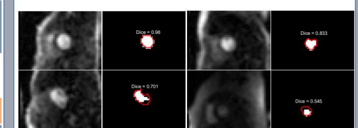

Deployed to hospitals

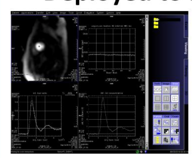

- PyTorch 1.0
- Model loading ~250ms
- Apply model ~150ms
- CPU only inference

Next steps: a) Continue to collect test data; b) retrain/redeploy with failed cases added; c) collect end-user feedback

References

[1] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in MICCAI, pp. 234–241, Springer, 2015.
[2] Z. Zhang, Q. Liu, and Y. Wang. Road extraction by deep residual u-net. CoRR, abs/1711.10684, 2017.
[3] Shvets A, Rakhlin A, Kalinin AA, Iglovikov V. Automatic Instrument Segmentation in Robot-Assisted Surgery Using Deep Learning. 2018