Automatic Detection of Brain Aneurysms: segmenting CT scans using CNNs

Jason Qin (jzqin), Harry Emeric (harryem)
CS230: Deep Learning

MOTIVATION

Brain aneurysms affect 1-3% of adults, and rupture is often fatal. CT scans can contain hundreds of images so detection is a lengthy process; in an emergency setting this delay may prove deadly. Automatic detection with a prediction model would be a valuable tool to address this.

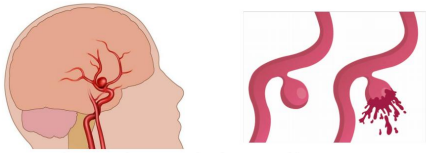

Figure 1: Schematic of aneurysm localization and bursting. Sources: Shutterstock, Mayo Clinic

DATASET

- DICOM images from 57 patients' CT scan, ~500 slices each
- Images were converted to 512 x 512 png
- Augmentation: flips, rotations, crops
- Class imbalance in original data: only ~1% of slices images contain aneurysm

Label	Pixels	Percentage
Background	10^9	99.9976%
Aneurysm	~160K	0.00235%

- Images with aneurysm resampled 40x

METHODS

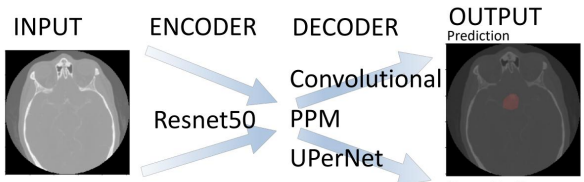

- Objective: pixel-wise labels
- Vary decoder, fine-tune encoder
- Training/test split of 51/6 patients
- Seek to minimize NLLL over 2 classes (aneurysm, background)

$$\ell(x, y) = L = \{l_1, \ldots, l_N\}^\top, \quad l_n = -w_{y_n} x_{n,y_n} \qquad w_{y_c} = \frac{1}{pixels\ per\ class\ c}$$

Evaluation metric:

$$IOU = \frac{|X \cap Y|}{|X \cup Y|}$$

Parameters:

- Adam, $\beta_1 = 0.9$, $\alpha = 0.02$
- 20 epochs, 500 iterations each

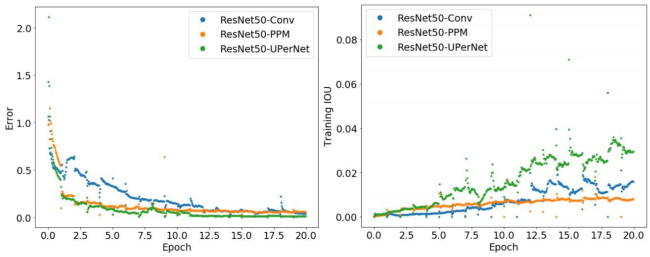

Figure 2: Training error (negative loss likelihood) and training aneurysm IOU over 20 epochs of training across three architectures of interest

DISCUSSION OF RESULTS

Encoder	Decoder	Eval IOU	Time per Epoch (s)
Resnet50	Conv	0.0018	270
	PPM	0.0013	360
	UPerNet	0.0141	540

- Random chance IOU: $\frac{\#\ Aneurysm\ Pixels * P(Aneurysm)}{2(\#\ Aneurysm\ Pixels)} = \sim 10^{-5}$
- UPerNet ~10x better than conv, PPM
- We are overlabeling aneurysm pixels
- In practice, would rather overdetect than miss aneurysm

CONCLUSIONS AND FUTURE WORK

- Semantic segmentation ~1000x better than random at detecting aneurysms
- Significant resampling of aneurysms required to correct class imbalance
- In future, can further tune resampling level, loss function weights
- More data to prevent overfitting

REFERENCES & ACKNOWLDGEMENTS

[1] He, Kaiming; Zhang, Xiangyu; Ren, Shaoqing; Sun, Jian (2015-12-10). "Deep Residual Learning for Image Recognition"
[2] Zhao, Hengshuang, et al. "Pyramid scene parsing network." *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. 2017.
[3] Xiao, Tete, et al. "Unified perceptual parsing for scene understanding." *arXiv preprint arXiv:1807.10221* (2018).

Thanks to Yeom Lab for providing the labelled data, and Aarti Bagul for project guidance.

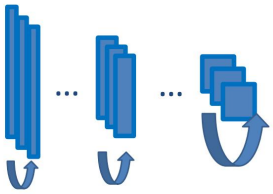

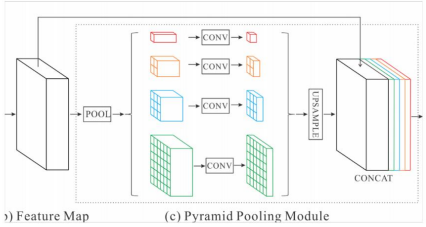

) Feature Map (c) Pyramid Pooling Module