

Convolutional Neural Networks for coronary plaque classification in intravascular optical coherence tomography images

Chaitanya Kolluru¹, David Prabhu¹, Yazan Gharaibeh¹, David L. Wilson^{1,2}, Sanjaya Gajurel³

1. Department of Biomedical Engineering, Case Western Reserve University, Cleveland, OH, 44106, USA



2. Department of Radiology, Case Western Reserve University, Cleveland, OH, 44106, USA

3. RCCI [U]Tech, Case Western Reserve University, Cleveland, OH, 44106, USA



Introduction

- Coronary artery disease (CAD) is the leading cause of death worldwide, resulting in more than 17M deaths in 2015.
- Percutaneous coronary interventions (PCIs) require identification of both plaque type and location within the coronary vessel wall.
- Intravascular Optical Coherence Tomography (IVOCT) allows for the identification and quantification of calcium and lipid plaques in coronary arteries.
- Automatic segmentation of calcific and lipid plaques is necessary due to the image overload in IVOCT (> 500 images) generated in a single pullback.
- Additionally, IVOCT images contain artifacts such as speckle, tangential signal dropout and proximity artifact making image interpretation challenging to a novice reader.



(c)

Fig 1. (a) IVOCT data in cartesian coordinates (XY domain). (b) Images can be converted to polar domain. (r-θ) (c) Depth profile at a given angle θ is referred to as an A-line.

- To reduce the need of manual analysis of IVOCT images, we are developing deep learning approaches to perform plaque characterization of A-lines in IVOCT images.

Datasets

IVOCT imaging

- 48 in-vivo clinical IVOCT pullbacks were obtained prior to stent treatment.
- Images were collected using a frequency-domain OCT system, Illumien Optis, St. Jude Medical, St. Paul, Minnesota.
- Select regions of pullbacks were annotated by two expert IVOCT readers, leading to an annotated dataset of 4,469 image frames.
- A-lines were annotated to three classes, namely fibrocalcific, fibrolipidic and other.
- A-line segmentations would blend seamlessly into previously described segmentation techniques [1] for calcium segmentation and fibrous cap thickness determination for lipid.




Fig 2. Manual A-line segmentation of plaques in XY images (labels indicated in a ring around image). Fibrocalcific A-lines are indicated in red color and fibrolipidic plaques are shown in green, other A-lines shown in blue and A-lines within guidewire shadow are colored black.

Acknowledgements

This work was supported by the National Institutes of Health through grant number R01HL114406. This work made use of the High Performance Computing Resource in the Core Facility for Advanced Research Computing at Case Western Reserve University.

Methods

Preprocessing

Convolutional Neural Network (CNN)

Dense Conditional Random Field

Preprocessing

- Lumen and guidewire regions are segmented using a dynamic programming approach. A-lines containing guidewire shadow are set to zero.
- A fixed back border distance of 1 mm (200 pixels) is considered.
- Image is pixel shifted to reduce the effects of large lumen eccentricity.
- Speckle noise is reduced by performing a log transform followed by Gaussian filtering with kernel size (7, 7) and sigma 1.




Fig 3. Preprocessing steps. (a) Example lumen boundary segmentation overlaid (in red) on OCT image. (b) Pixel shifted version of the image. (c) Filtered image after performing Gaussian filtering.

CNN architecture

- We develop a convolutional neural network (CNN) that takes 1D A-lines as input.
- The network consists of two convolutional layers, two max-pooling layers and two fully connected layers.
- Batch normalization was used to allow for a high learning rate.
- Categorical cross-entropy was used as the loss function for network optimization.
- Training was stopped when the validation loss did not decrease over five consecutive epochs.




Fig 4. CNN architecture.

Layer	Output Shape	Comments
Input	(200, 1)	Individual A-lines of length 200 pixels
Padding layer	(210, 1)	Pixels replicated at corners
Convolutional layer	(200, 1, 32)	Kernel size of (11, 1) and 32 filters
Max-pooling layer	(100, 1, 32)	Pool size of (2,2) and stride 2 pixels
Convolutional layer	(100, 1, 64)	Kernel size of (9, 32) and 64 filters
Max-pooling layer	(50, 1, 64)	Pool size of (2,2) and stride 2 pixels
Fully connected layer	100	
Fully connected layer	3	Output layer, softmax activation

Dense Conditional Random Field

- Since individual A-line classification can lead to noisy segmentations, we use a dense conditional random field to clean classification results obtained from the deep neural networks.
- A dense conditional random field [2] is used over the en-face view (θ -z) to collect neighborhood information from consecutive frames and neighboring A-lines.
- A closing operation is performed to remove any islands less than 10 pixels in area.

References

- Z. Wang *et al.* 2010. Semiautomatic segmentation and quantification of calcified plaques in intracoronary optical coherence tomography images, *Journal of biomedical optics*, vol. 15, no. 6, p. 061711, 2010.
- Philipp Krähenbühl, Vladlen Koltun. 2011. Efficient Interference in Fully Connected CRFs with Gaussian Edge Potentials Advances in Neural Information Processing Systems 24 (2011) 109-117.
- Sanjaya Gajurel. 2017. Singularity. GitHub. <https://github.com/sxg125/singularity>.

Results

Experiments

- For the experiment, Tensorflow image is imported from Docker using Singularity [3] in CWRU HPC.
- Dataset containing 48 pullbacks was split into 10 folds.
- Fold consisted of roughly 38 pullbacks for training, 5 pullbacks for validation and 5 pullbacks for testing.
- Confusion matrix is reported across all A-lines in the dataset after noise cleaning is performed by the dense conditional random field.




Fig 5. Ground truth and prediction result for two calcium lesions ((a) & (b)) in en-face view (θ -z) after noise cleaning by the dense conditional random field. Red indicates calcium, green is lipid, blue is other and black is guidewire shadow.

	Predicted Fibrocalcific	Predicted Fibrolipidic	Predicted Other
True Fibrocalcific	243821 (76%)	33753 (10%)	43753 (13%)
True Fibrolipidic	19535 (4%)	362332 (84%)	48144 (11%)
True Other	67455 (5%)	126298 (10%)	1101586 (85%)

Table 1. Confusion matrix for ten-fold cross validation on the dataset.

- Network training time was found to be less than 20 minutes per fold.
- Trained neural network classifier takes less than 10 seconds to predict a complete test pullback (500 images * 496 A-lines per image) when run on a notebook GPU (NVIDIA GTX 950M) indicating possible real-time and offline use.




Fig 6. Example predictions. OCT image overlaid with ground truth (outer ring) and predictions from the classifier (inner ring). Red indicates calcium, green indicates lipid, blue indicates other and black indicates guidewire shadow A-lines.

Discussion

- A method to classify individual A-lines in IVOCT images is presented.
- We employ a ten fold cross validation procedure and achieve good classification accuracy over 75% across three classes.
- Multiple GPU nodes in the HPC will be considered to scale up the performance using larger training datasets in the future.