

Atherosclerotic Vascular Calcification Detection and Segmentation on Low Dose Computed Tomography Scans Using Convolutional Neural Networks

Karthik Chellamuthu, Jiamin Liu, Jianhua Yao, Mohammadhadi Bagheri, Le Lu, Veit Sandfort, Ronald M. Summers

Imaging Biomarkers and Computer-aided Diagnosis Laboratory
Radiology and Imaging Sciences, National Institutes of Health Clinical Center
Building 10 Room 1C224 MSC 1182, Bethesda, MD 20892-1182

ABSTRACT

We propose an automated platform for extra-coronary calcification detection on low dose CT scans. We utilize faster regional convolutional neural networks (R-CNN) to directly detect calcifications at the lesion-level without performing vessel extraction. To segment detected calcifications at the voxel-level, we employ holistically nested edge detection (HED). CT scans of 112 vasculitis patients and 3219 images with labeled calcifications were used to develop and evaluate our method. By employing a two-class faster R-CNN, the average precision (AP) increased from 49.2% to 84.4% for calcification detection. In addition, sensitivity of 85.0% at 1 false positive per image was observed. The Dice Similarity Coefficient (DSC) for calcification segmentation using HED (0.83 ± 0.08) was significantly better ($p < 0.01$) than the traditional threshold-based method (0.59 ± 0.26).

Index terms – Calcification, plaque, region proposal, CNNs, HED

1. INTRODUCTION

Vascular calcification is an important clinical manifestation of atherosclerosis, which is the primary cause of heart disease and stroke and is implicated in about 50% of all deaths in westernized societies [1].

The disease is characterized by a progressive thickening of the arterial wall due to the buildup of lipids, macrophages, fibrous elements, and debris. As the condition progresses, the arterial wall will present with calcific plaques [2]. These calcified plaques can obstruct blood flow and signify the later stages of atherosclerosis [3]. In clinical practice, there is a need to characterize atherosclerotic plaques to determine the status and disease progression.

Many works have been proposed for coronary artery calcification detection in cardiac CT scans. However, there are few works about extra-coronary calcification detection [4, 5]. In addition to serving as a way to identify and classify the progression of coronary disease, extra-coronary plaque detection and quantification may also be

a relevant biomarker in association studies with other disorders such as kidney disease and cancers [6, 7]. In the current workflow, each CT image is manually examined to detect plaque instances, making the process tedious and time consuming. Therefore, computer-aided calcification detection could greatly reduce the radiologists' workload and could be employed as a first or second reader for improved disease assessment.

In the past, vessel extraction was commonly required for computer-aided calcification detection. However, automatic vessel extraction is challenging on non-contrast or low dose CT scans. Figure 1 shows two calcified plaques on low dose CT scans.

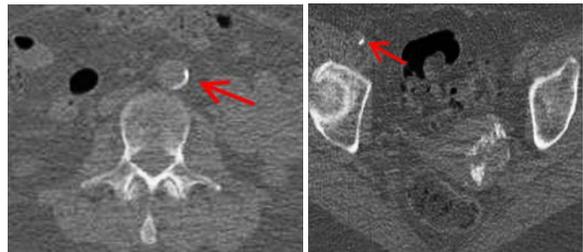


Figure 1. Examples of calcified plaques (red arrows) on abdominal (left) and pelvic (right) CT scans.

Recent breakthroughs in object detection in computer vision have been driven by the success of deep convolutional neural networks (CNNs) [8]. These hierarchical features have been found to be more efficient for object detection and recognition than hand-crafted features such as SIFT [9] and HOG [10]. CNNs do not require hand-crafted features, making them good candidates for sophisticated clinical applications such as lymph node detection in CT images [11], pancreas segmentation in CT images [12], and brain segmentation in MR images [13].

This paper presents an automated method for extra-coronary calcification detection on low dose CT scans. We leverage the development in faster regional convolutional neural networks (R-CNN) [8] to directly detect calcifications on low dose CT scans. Furthermore, we accomplish this without needing to perform vessel extraction. For detected calcifications, we utilize the holistically nested edge detection (HED) [14] for

calcification segmentation, a crucial component of the subsequent quantitative analysis.

2. METHODS

2.1 Plaque detection by faster R-CNN

In this work, we use the recently advanced faster R-CNN for plaque detection. The faster R-CNN employs an efficient single-stage training process and jointly learns to classify plaque and refine their spatial locations.

Calcification detection by faster R-CNN has two steps (Figure 2). For each 2D image, region proposals are first generated by region proposal networks (RPN). A forward pass is performed on the entire input image to create one feature map. The features of region proposals are computed by the region-of-interest (ROI) projection in the feature map. Then, each region proposal is jointly classified and refined by softmax classifier and bounding box regressor.

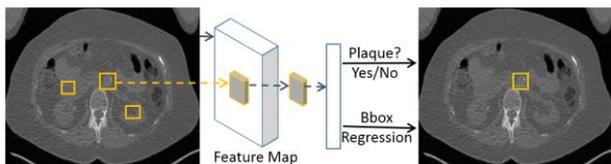


Figure 2: Calcification detection by faster R-CNN.

2.1.1 Region proposal network

The RPN was trained end-to-end using backpropagation and optimized with stochastic gradient descent (SGD). In each SGD iteration, a mini-batch was constructed for loss optimization. Each mini-batch has 128 positive samples and 128 negative samples from the same image. Positive samples are defined as all ground-truth bounding boxes of the plaque and the region proposals with IoU (intersection over union) ≥ 0.7 . Negative samples are defined as the region proposals with $\text{IoU} \leq 0.1$. Caffe [15] implementation of a 16 layer VGG net [16] was used in this work. The RPN is initialized by an ImageNet pre-trained model [17]. VGG net is fine-tuned with a learning rate of 0.001, a weight decay of 0.0005, and a momentum of 0.9 for 10k mini-batches on our plaque dataset.

2.1.2 Detection networks

We use Fast R-CNN [18] as the detection network in our algorithm. A 4-step training method [8] is utilized to learn the shared convolutional layers between the RPN and detection network. (1) The RPN is initialized with a pre-trained model (using ImageNet). All layers are fine-tuned for region proposal generation. (2) A separate detection network is also initialized with a pre-trained model (using ImageNet) and is trained using the proposals generated in the first step. (3) The RPN was initialized by the detection network. The shared convolutional layers are fixed and the layers of RPN are fine-tuned. (4) Retaining the shared convolutional layers, the fully connected layers of the

detection network are fine-tuned. In this way, RPN and detection networks form a unified network by sharing the same convolutional layers.

2.1.3 Training strategy

Without prior anatomical constraints, i.e. vessel extraction, other negatives such as bones (e.g. ribs) and kidney stones are challenging to distinguish from calcified plaques. To reduce these false positives, we use two sequential faster R-CNNs for calcification detection. The first faster R-CNN (Faster R-CNN₁) focuses on detecting and separating plaque from all negatives. The second faster R-CNN (Faster R-CNN₂) distinguishes among plaque, plaque-like, and all other negatives. Figure 3 shows the overview of sequential faster R-CNNs trained for plaque detection.

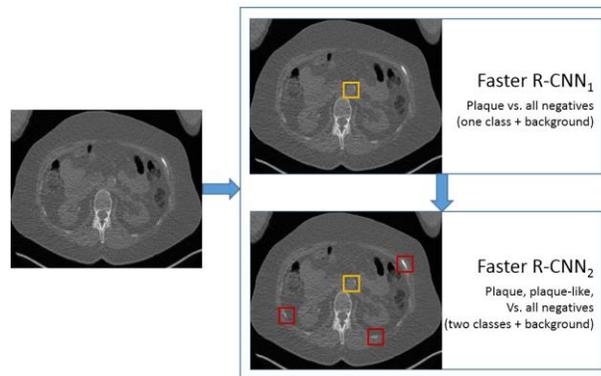


Figure 3: Faster R-CNN₁ is trained first, using the manually labeled plaques only (orange boxes). This CNN discards the negative objects such as muscle and fat. Faster R-CNN₂ is trained using two classes, the manually labeled plaques and the plaque-like objects (red boxes) those have high prediction value from Faster R-CNN₁.

The two faster R-CNN are trained in sequence. Faster R-CNN₁ is trained first, using one class, the manually labeled plaques (orange boxes in Figure 3). This CNN discards the negative objects such as muscle and fat and detects plaque and plaque-like objects. Subsequently, Faster R-CNN₂ is trained using two classes: the manually labeled plaques and the plaque-like objects (red boxes in Figure 3) which have high prediction value from Faster R-CNN₁.

During testing, the trained Faster R-CNN₂ is used for plaque detection. We run RPN on the test image to extract around 200 region proposals. For each region proposal, the forward pass outputs a class posterior probability and a set of predicted bounding-boxes. We then perform non-maximum suppression independently for each class for final detections.

2.2 Plaque segmentation

Once the location of the plaque has been detected from the prior step, segmentation is performed to isolate the plaque voxels.

2.2.1 Threshold-based approach

In clinical practice, voxels exceeding 130 Hounsfield units (HU) are classified as calcified tissue [19]. Regions with higher HU values will be correspondingly weighted to represent the grade of the plaque.

With this in mind, we considered employing a threshold-based approach for calcification segmentation in voxel-level. In this method, a global threshold (130 HU) was first applied to the detected region. The resulting image was subsequently processed with Sobel filtering and border removal. Finally, the image underwent morphological filling and erosion to yield a segmentation of calcification.

2.2.2 HED

For comparison, we also utilize the deep learning-based method, HED, for calcification segmentation. HED utilizes CNNs to hierarchically learn the relevant features for segmentation. Furthermore, these features are used to produce edge maps which are progressively refined in a structured manner. As a result, the deeply supervised HED framework allows for better discrimination in edge and object boundary detection problems. This technique has proven to be successful in advancing the state-of-the-art in the BSD500 and NYU Depth segmentation benchmark datasets [17].

We implement our framework using the Caffe library and build from the default HED implementation. Our model parameters follow the configuration used in [18]. HED is fine-tuned with a learning rate of $1e-6$, a weight decay of 0.0001, and a momentum of 0.9 for 10k mini-batches on our plaque dataset. The training phase takes only 30 minutes on a K40 GPU since we only trained HED using the cropped images (plaques within the bounding box).

3. RESULTS

3.1 Dataset

CT scans of 112 patients with vasculitis are used to develop and evaluate our method. For each patient, vascular calcifications were considered along the aorta, aortic arch, coronary, carotid, subclavian, innominate, common iliac, and common femoral arteries. These lesions were manually labeled on CT scans by a radiologist. Thus our dataset has 3219 images with labeled calcifications in total. 75% of this dataset (84 patients with 2119 plaques) is randomly selected. These patients constitute a training set for faster R-CNN for plaque detection and for HED for plaque segmentation. The remaining 28 patients with 1100 plaques form the test set for evaluation.

3.2 Plaque detection

Figure 4 shows the precision-recall curve and FROC curve of our calcification detection on 28 patients. The detection average precision (AP) using Faster R-CNN₂ is increased to 84.4% from 49.2% using Faster R-CNN₁. The system using Faster R-CNN₂ can achieve an 85.0% sensitivity at 1 false positive per image. Some true cases, false positives, and missed detections using Faster R-CNN₂ are shown in Figure 5.

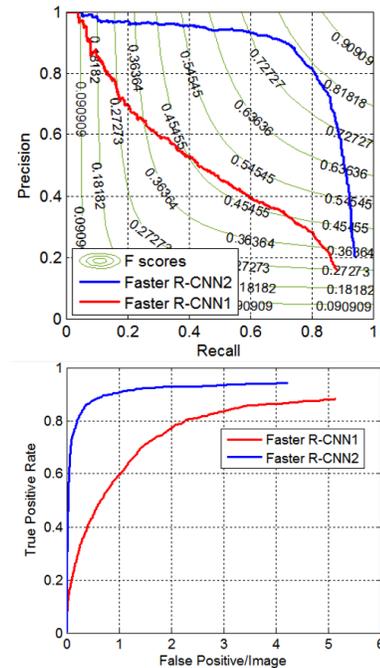


Figure 4. Precision-recall curves (top) and FROC curves of plaque detection on test set (bottom). Isobars in (top) are F scores.

3.3 Plaque segmentation

Dice Similarity Coefficient (DSC) [20] is used as a measure of segmentation accuracy for evaluation. The DSC for plaque segmentation using HED (0.83 ± 0.08) is significantly better ($p < 0.01$, paired t-test) than the segmentation using the threshold-based method (0.59 ± 0.26). Figure 6 shows one example of plaque segmentation by using two approaches. The DSC using HED (0.79) is much higher than the threshold-based approach (0.57). In this case, thresholding could not effectively separate the plaque and the adjacent vertebra. In such situations, plaque segmentation via deep learning based methods such as HED can significantly outperform conventional techniques.

4. CONCLUSION

Convolutional neural networks have been applied to detect and segment the vascular calcifications on low dose CT scans without vessel extraction. The experimental

results reveal high detection sensitivity at a low false positive rate and accurate segmentation.

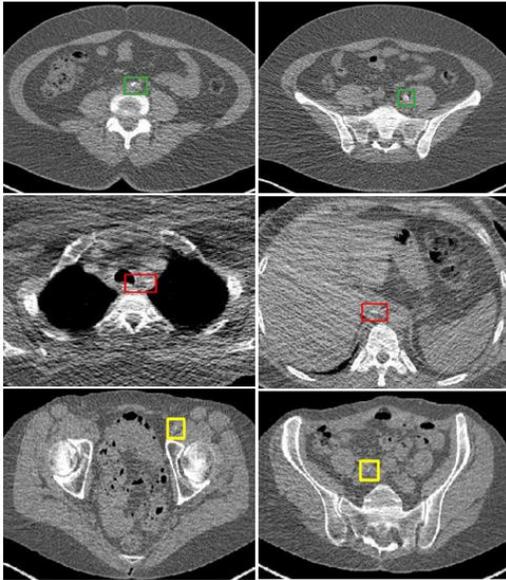


Figure 5. Calcification detection examples: true positives (green), false positives (red) and missed detections (yellow).

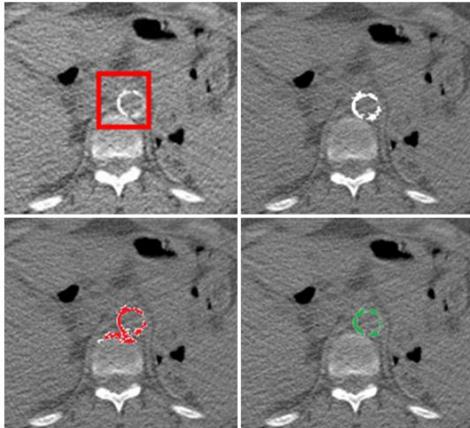


Figure 6. Example of plaque segmentation. Detected plaques (red boxes), manually-labeled plaque segmentations (white), plaque segmentations by threshold-based approach (red) and segmentation by HED (green).

ACKNOWLEDGMENT

This research was supported by the National Institutes of Health, Clinical Center. The authors thank Nvidia for the K40 GPU donation.

References

- Lusis, A.J., *Atherosclerosis*. Nature, 2000. 407(6801): p. 233-41.
- Kriszbacher, I., M. Koppan, and J. Bodis, *Inflammation, atherosclerosis, and coronary artery disease*. N Engl J Med, 2005. 353(4): p. 429-30; author reply 429-30.
- Wong, K.K., P. Thavornpattanapong, S.C. Cheung, Z. Sun, and J. Tu, *Effect of calcification on the mechanical stability of plaque based on a three-dimensional carotid bifurcation model*. BMC Cardiovasc Disord, 2012. 12: p. 7.
- Isgum, I., A. Rutten, M. Prokop, M. Staring, S. Klein, J.P. Pluim, M.A. Viergever, and B. van Ginneken, *Automated aortic calcium scoring on low-dose chest computed tomography*. Med Phys, 2010. 37(2): p. 714-23.
- Isgum, I., B. van Ginneken, and M. Olree, *Automatic detection of calcifications in the aorta from CT scans of the abdomen. 3D computer-aided diagnosis*. Academic Radiology, 2004. 11(3): p. 247-57.
- Hager, M., G. Mikuz, G. Bartsch, C. Kolbitsch, and P.L. Moser, *The association between local atherosclerosis and prostate cancer*. Bju International, 2007. 99(1): p. 46-48.
- Budczies, J., M. von Winterfeld, F. Klauschen, A.C. Kimmritz, J.M. Daniel, A. Warth, V. Endris, C. Denkert, H. Pfeiffer, W. Weichert, M. Dietel, D. Wittschieber, and A. Stenzinger, *Comprehensive analysis of clinico-pathological data reveals heterogeneous relations between atherosclerosis and cancer*. Journal of Clinical Pathology, 2014. 67(6): p. 482-490.
- Ren, S., K. He, R. Girshick, and J. Sun, *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks*. IEEE Trans Pattern Anal Mach Intell, 2016.
- Dalal, N. and B. Triggs, *Histograms of oriented gradients for human detection*. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol 1, Proceedings, 2005: p. 886-893.
- Lowe, D.G., *Distinctive image features from scale-invariant keypoints*. International Journal of Computer Vision, 2004. 60(2): p. 91-110.
- Roth, H.R., L. Lu, A. Seff, K.M. Cherry, J. Hoffman, S. Wang, J. Liu, E. Turkbey, and R.M. Summers, *A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations*. Med Image Comput Assist Interv, 2014. 17(Pt 1): p. 520-7.
- Roth, H.R., L. Lu, A. Farag, H.C. Shin, J.M. Liu, E.B. Turkbey, and R.M. Summers, *DeepOrgan: Multi-level Deep Convolutional Networks for Automated Pancreas Segmentation*. Medical Image Computing and Computer-Assisted Intervention - Miccai 2015, Pt I, 2015. 9349: p. 556-564.
- Havaei, M., A. Davy, D. Warde-Farley, A. Biard, A. Courville, Y. Bengio, C. Pal, P.M. Jodoin, and H. Larochelle, *Brain tumor segmentation with Deep Neural Networks*. Med Image Anal, 2016. 35: p. 18-31.
- Xie, S.N. and Z.W. Tu, *Holistically-Nested Edge Detection*. 2015 Ieee International Conference on Computer Vision (Iccv), 2015: p. 1395-1403.
- Jia, Y., E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Suadarrama, and T. Darrell, *Caffe: Convolutional Architecture for Fast Feature Embedding*. arXiv preprint arXiv:1408.5093, 2014.
- Simonyan, K. and A. Zisserman, *Very Deep convolutional networks for large-scale image recognition*, in *ICLR2015*.
- Russakovsky, O., J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z.H. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, and L. Fei-Fei, *ImageNet Large Scale Visual Recognition Challenge*. International Journal of Computer Vision, 2015. 115(3): p. 211-252.
- Girshick, R., *Fast R-CNN*. 2015 Ieee International Conference on Computer Vision (Iccv), 2015: p. 1440-1448.
- Yamak, D., B. Pavlicek, T. Boltz, P. Panse, D. Frakes, and M. Akay, *Coronary calcium quantification using contrast-enhanced dual-energy computed tomography scans (vol 14, pg 4014, 2013)*. Journal of Applied Clinical Medical Physics, 2015. 16(2): p. 465-465.
- Dice, L.R., *Measures of the Amount of Ecologic Association between Species*. Ecology, 1945. 26(3): p. 297-302.