MACHINE-LEARNING BASED CLINICAL PLAQUE DETECTION USING A SYNTHETIC PLAQUE LESION MODEL FOR CORONARY CTA

Nikolas D. Schnellbächer[a+], Haissam Ragab [b+], Hannes Nickisch [a+], Tobias Wissel[a], Clemens Spink[b], Gunnar Lund [b], Michael Grass[a]

[a]Philips Research – Hamburg, Germany; [b]Department of Diagnostic and Interventional Radiology and Nuclear Medicine, University Hospital Hamburg Eppendorf (UKE), Hamburg, Germany;
[+]Authors contributed equally to this work

ABSTRACT

Coronary computed tomography angiography (coronary CTA) is a robust and well-established non-invasive diagnostic tool to detect and assess coronary artery disease (CAD). The accurate detection, quantification and characterization of the coronary plaque burden has become an important part of this imaging modality.
The quality and performance of modern machine-learning-based data-driven learning approaches is often impacted by either insufficient or inconsistently-labeled training data and is further subject to additional bias from human annotators. To address these shortcomings for coronary plaque characterization, we have developed a synthetic lesion generating framework for CTA applications, which can produce accurate and high-quality labeled training data for data-driven learning approaches.

This approach can help to ease the manual annotation burden, which is often the limiting factor in data-driven learning algorithms and instead provides reliable ground truth data for modern deep learning approaches. Furthermore, this framework can easily be used to create custom tailored training data that can be used for pre- or post-training steps of already existing machine learning approaches for CTA applications.

We tested this data generation framework by inserting synthetic lesions in 11 clinical CTA scans of healthy patients resulting in a data set of ~7000 annotated 2D slices. With this data we performed several plaque detection experiments using a data-driven machine learning approach with a neural encoder architecture. In this plaque classification task we first demonstrate that the synthetic lesion generation module can consistently perform well in recognizing unseen synthetic test data with an overall classification accuracy of 93%. Next we apply the synthetic lesion framework in a transfer learning experiment, where we demonstrate the feasibility to learn to classify real clinical plaque lesions with a purely synthetic model (overall classification accuracy 84%) that never saw real clinical lesions during model training. Second, we show that using synthetically data for pre-training with a subsequent training on clinical data can enhance the overall classification accuracy (from 91% to 92%) while strongly increasing the true positive count.

We conclude that the synthetic plaque lesions model faithfully covers many important image characteristics of real plaque lesions in coronary CTA imaging and can thus help reduce the annotation burden for data-driven predictive vascular systems in this domain. This allows the creation of exhaustively annotated and site-specific customizable training data with a computationally fast forward model.

Keywords: coronary CTA (CCTA), CT angiography, coronary plaque, machine learning customization, data synthesis, forward model, convolutional neural networks, classification

1. INTRODUCTION

Coronary computed tomography angiography (coronary CTA) is a robust and well-established non-invasive diagnostic tool to detect and assess coronary artery disease (CAD). The accurate detection, quantification and characterization of the coronary plaque burden has become an important part of this imaging modality [1], [2].

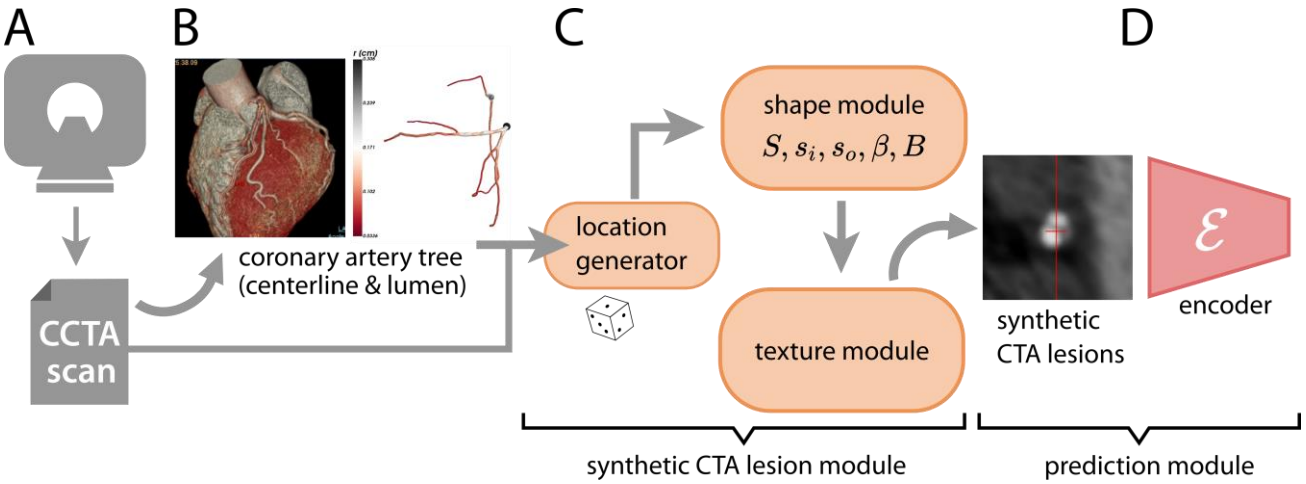

Figure 1 High-level schematic depiction of the presented synthetic plaque lesion module. Starting from a clinical CCTA scan (**A**) together with a coronary artery tree segmentation (centerline and lumen segmentation) (**B**) the synthetic lesion generator can inpaint synthetic lesions in healthy coronary vessels (**C**) and therewith create high-quality annotated training data for predictive machine learning applications (**D**). The synthetic lesion generation module itself consists of three modular parts. A stochastic location generator samples lesion position and lesion length in the coronary artery tree. This step is then followed by a shape module which creates a voxelized lesion volume according to a handful of geometric parameters and a texture module which adds realistic intensity grey values to arrive at a final synthetic lesion, as exemplary shown in part (**D**).

The detection and characterization of such lesions are often used as risk markers to predict major adverse cardiac events (MACE). The quality and performance of modern machine-learning-based data-driven learning approaches is often impacted by either insufficient or inconsistently-labeled training data and is further subject to additional bias from human annotators. To address these shortcomings for coronary plaque characterization, we have developed a synthetic lesion generating framework for CTA applications, which can produce accurate and high-quality labeled training data for data-driven learning approaches.

The early success of many modern machine learning methods can largely be attributed to the successful application of deep convolutional neural networks in the field of computer vision, showing excellent performance in particular in the field of image recognition. The early paradigm in the field was, that an enormous amount of labeled training data is a necessary requirement for the successful application of such models. However, more elaborate training routines combined with efficient data augmentation strategies allowed their application in the biomedical domain [3], where good results on common computer vision tasks were within reach using only on the order of thousand annotated images.
However, for many important clinical applications, such as coronary CTA, the manual labor to obtain such curated and high-quality annotated data sets is still very cumbersome and manual labor intensive.
Our herein presented method can help to ease the manual annotation burden, which is often the limiting factor in data-driven learning algorithms and instead provides reliable ground truth data for modern deep learning approaches. Furthermore, this framework can easily be used to create custom tailored training data that can be used for pre- or post-training steps of already existing machine learning approaches for various CTA applications. Fig. 1 shows a high-level description of our method. Starting from a given CCTA scan together with a coronary artery segmentation (centerline and lumen segmentation), the synthetic CTA lesion generator module can *synthesize* and exhaustively annotate training data for CTA prediction tasks.

The next section first describes our used data sources and the synthetic lesion generation step. Then we provide more details on the used learning approach. Finally results from various plaque classification experiments are presented and discussed.

2. MATERIAL AND METHODS

2.1 Dataset

We collected and curated 92 clinical CTA scans acquired at the Department of Diagnostic and Interventional Radiology and Nuclear Medicine, University Hospital Hamburg Eppendorf, Germany. These cases remained after filtering out cases with poor image quality due to strong motion or step artifacts or with too large image slice thickness. For all these we performed a coronary artery tree segmentation to obtain the centerline position and lumen segmentation. Within this set we also identified 11 healthy patients without any clinical findings that were used for our synthetic forward model in this work.

All cases were prepared for manual annotation by a human expert reader based on a stretched multi-planar reformatted (MPR) image reconstruction based on the underlying vessel segmentation. This allows to view all vessels both in a longitudinal view and on a cross-sectional level. In this detailed annotation process we labeled various plaque types as well as other high-risk markers such as positive remodeling and napkin-ring-sign alongside to other characteristic findings, such as stents, side branches, neighboring vessels, bifurcations or the lumen, to arrive at a complete picture for a given CTA case. In this process labels can be created either on a per cross-section level, by annotating a certain interval in the stretched MPR view, or on a per-pixel level on a specific cross section by creating a corresponding label bitmask. For this work we focused on calcified plaque labels and pooled all labels from both label types accordingly to obtain consistent per-frame labels for calcified plaque.

2.2 Synthetic Lesion Module

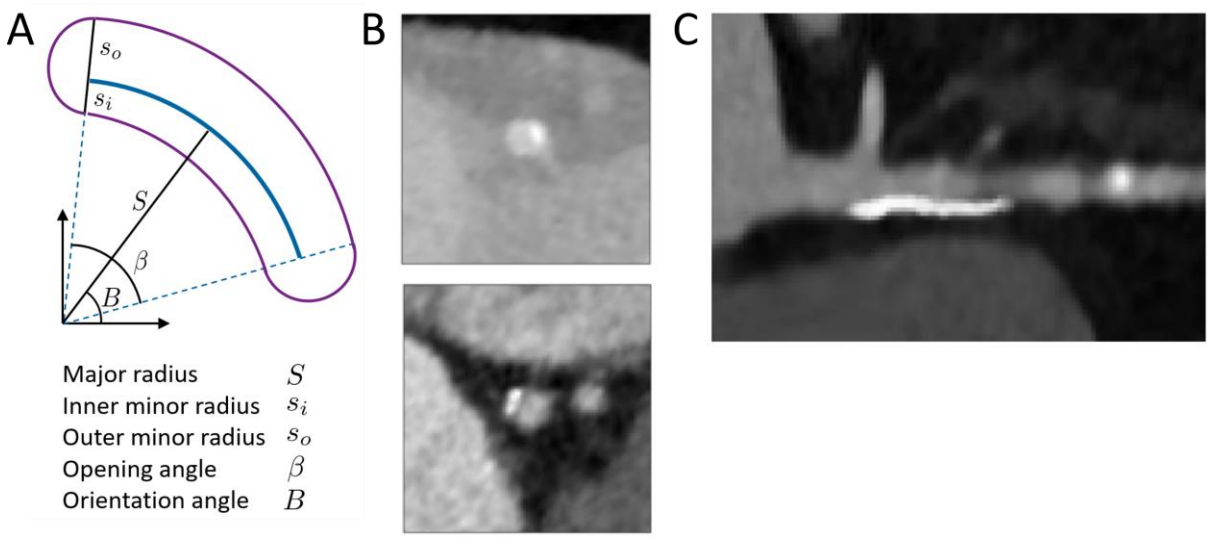

Figure 2 Synthetic plaque lesion model. (**A**) depicts the toroidal plaque shape model which creates the lesion shape based on the depicted geometrical shape parameters. Shown is a cross-sectional view with the center of origin in the vessel center. These five shape parameters can vary along the vessel's centerline and are sampled per inserted lesion. This lesion shape model produces a smooth motion vector field (with zero values outside of the purple delineated lesion shape) which describes the lesion growth. (**B**) shows two synthetically created lesions in cross-sectional view, while (**C**) illustrates the extent and lesion shape in a stretched MPR view.

We are interested in modeling the image characteristics of coronary lesions as detected in contrast-enhanced CT angiography scans (CTA). In this context, the lesion shape, their location, plaque volume, location of calcifications have to be faithfully captured, in order to model them synthetically. We propose a model that is implemented as an explicit forward algorithm. This synthetic lesion growth model allows additions of artificial lesions to healthy coronary vessels in CTA scans and thus creating precisely labeled training data for data-driven learning algorithms.

The lesion model consists of a shape and a texture module (see Fig. 1 and Fig. 2 (A)) which jointly create an artificial lesion using only a few geometric parameters. Based on a coronary artery tree segmentation (centerline and lumen), the lesion shape is modulated by varying radii, angles and length parameters of a given toroidal geometry along the coronary's centerline (compare Fig. 2). Based on the lesion's shape, a motion vector field (MVF) is generated to capture

the lesion growth process. After creating this shape and applying the growth model, a given lesion is inpainted into a given healthy vessel using a texture module. Exemplary results of this approach are depicted in Fig. 2 (B and C).

In a separate annotation task we verified the synthetically generated lesions both qualitatively and quantitatively by asking a human expert reader to assess lesion appearance under the same conditions as the clinical annotations were performed. In a qualitative per case assessment the majority of all synthetically created lesions were rated as highly realistic with regard to their image appearance. In some cases a too eccentric lesions shape was criticized as artificial. Complementary, we verified the exhaustive ground truth annotation of the synthetic lesions per cross-section and compared it against a human expert annotation on a per-pixel level. For a given test case with manually annotated cross sections, we found them to be in good agreement resulting in an average dice overlap score of 74%, where outliers could be identified to be largely occurring at the beginning or the end of a lesions.

2.3 Learning Approach

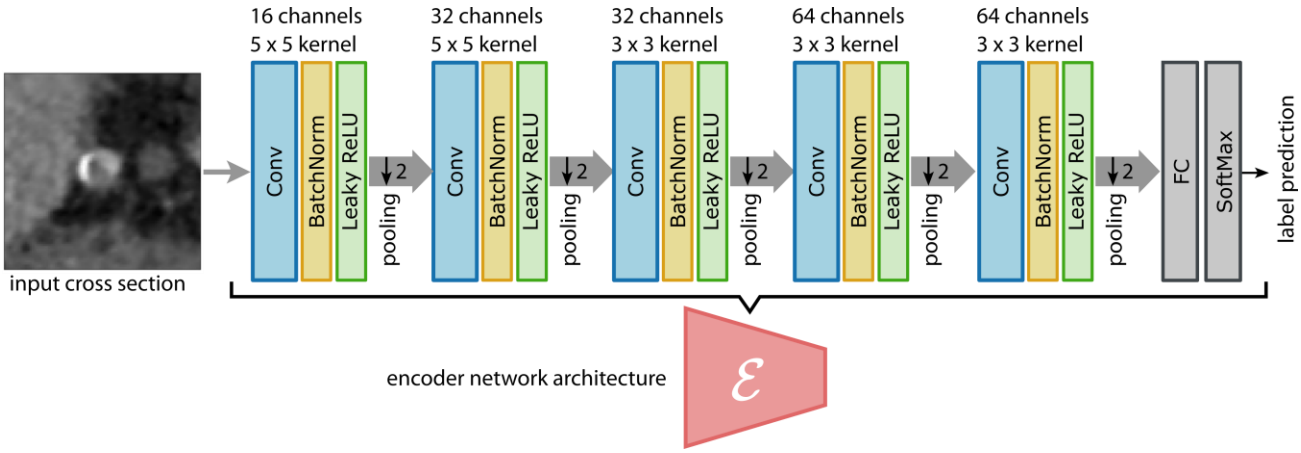

Figure 3 Encoder neural network architecture. Starting from an input cross-section at the very left, a forward pass consists of five VGG-like layers (2D convolution, batch normalization and nonlinear activation function) with a subsequent average pooling step after each layer. Above each VGG-like block we note the number of feature channels and the 2D convolution kernel size. At the very end, the output from the last pooling layer is passed into a fully connected (FC) layer followed by a softmax layer outputting the corresponding predicted class probabilities.

In this work we use a neural encoder architecture similar to the encoder branch of Noh et al. [4]. This is a deep CNN architecture where VGG-like neural network layers are combined [5]. A graphical depiction of the encoder architecture is shown in Fig. 3. The encoder consists of five sub-blocks, each containing a convolutional layer (twice kernel size 5x5 and 3x3 in the last three blocks, all with stride 1), batch normalization [6] and leaky rectified linear unit (ReLU) as nonlinear activation functions. The single channel raw input is increased to 16 feature channels in the first sub-block and then further increased to 32 and 64 as depicted in Fig. 3. After each VGG-like sub-block an average pooling operation is performed with kernel size 2x2 and stride 2. At the bottleneck, the output of the last sub-block is fed into a fully connected layer (FC) followed by a softmax operation to output the class probabilities for the classification task.

In this work we looked at the binary classification task of predicting calcified plaque on a per-frame level and trained this encoder structure in a supervised manner. For this process we used the ADAM optimizer [5] for training and ran all models for 200 epochs using a cross-entropy loss function.

3. RESULTS AND DISCUSSION

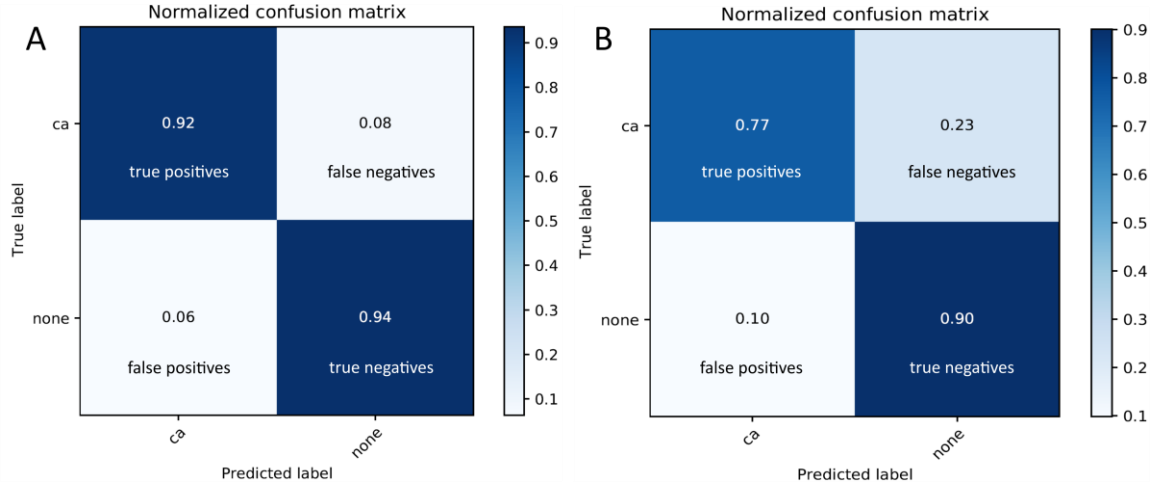

Figure 4 Normalized confusion matrices summarizing the classification performance. Rows are normalized to unity. **(A**) Results for a purely synthetic training tested on purely synthetic test data. (**B**) Results for the transfer learning experiment by applying the purely synthetically trained model on a test set with clinical plaque findings. Ca denotes the calcified plaque label in the binary classification task.

The first classification experiment was run purely on synthetic training data using a 80% to 20% train/test data split and served the purpose to assess the self-consistency, meaning that the synthetically generated lesions can be used to train a classifier to recognize data which contain synthetic lesions, not seen during training, but also exclusively generated using our synthetic plaque lesions module as described above. For this experiment we used 11 CTA scans with healthy coronary vessels, resulting in 7828 2D cross sections in total. Fig. 4 (A) summarizes the classification performance by showing the normalized confusion matrix. The overall classification accuracy here is 93%. This illustrates that lesions that are generated using our synthetic forward model can be easily learned with high accuracy, demonstrating self-consistency.

In a second experiment we used a purely synthetically trained plaque model to predict plaque in unseen clinical test data. The classification performance is shown in the normalized confusion matrix as depicted in Fig. 4 (B) with an overall classification accuracy of 84%. While this is noticeably worse than the dedicated training shown before, this still demonstrates that the synthetically generated lesion can also be used to solve the much harder learning task to recognize real clinical calcified plaque lesions, although the used predictor has never seen any real lesion in the training process (transfer learning).

Using the clinical ground truth annotations obtained by a human expert reader, we are of course also able to run a benchmark experiment with the same settings on purely clinical data, solving the identical plaque detection task as before. The baseline results from this experiment are shown in Fig. 5 (A), giving on overall a classification accuracy of 91%. Using a synthetically trained model to initialize the neural network weights for a clinical training (synthetic model as pre-training), we can show, that we can improve the overall classification accuracy for this task to 92%, mostly accredited to a much higher true positive count (increased to 92% from an 85% baseline value, compare Fig. 5).
Fig. 6. shows sample predictions from this final clinically trained model with a synthetically trained model as neural network weight initialization (pre-training).

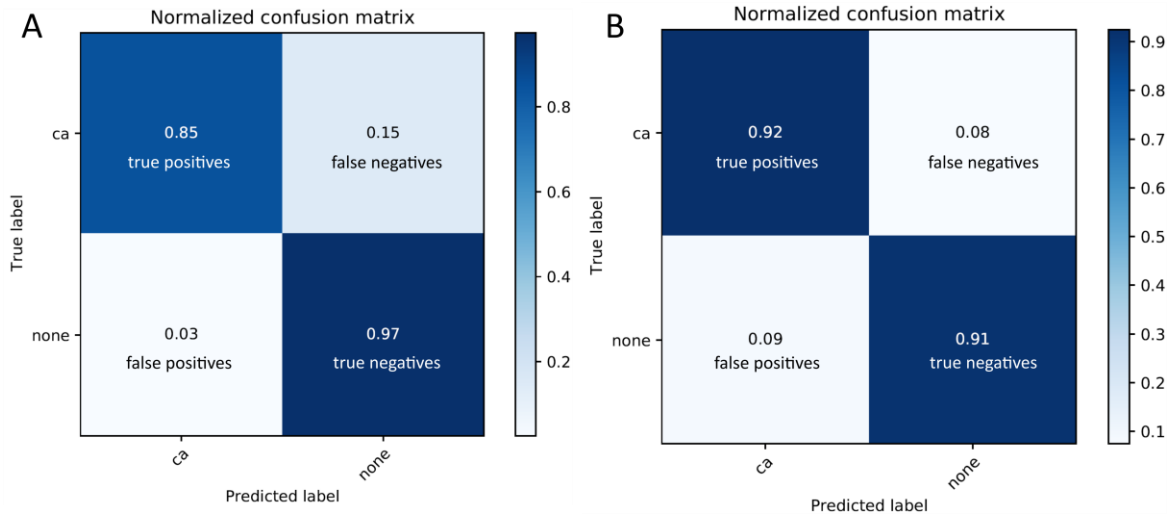

Figure 5 Normalized confusion matrices as in Fig. 4. **(A)** Results for a purely clinical training without any synthetic lesions. (**B**) Results for the same setup as in (A) but this time using a pre-trained model that was trained on purely synthetic lesions.

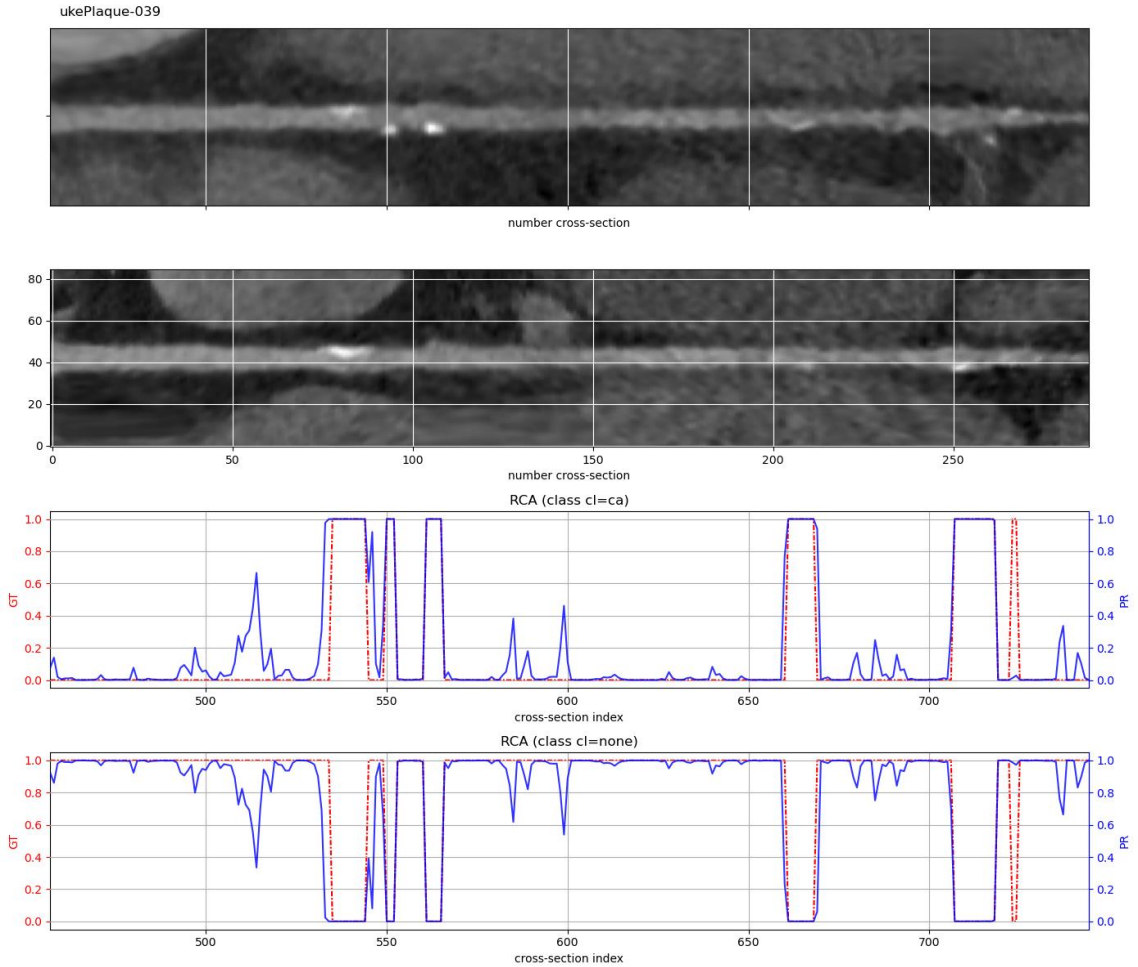

Figure 6 Predicted class probabilities as a function of the cross-sectional position along the RCA vessel of a clinical test case. This shows sample predictions from the clinical training with synthetic pre-training (compare Fig. 5 (B)). Here, the red-dashed lines indicate the ground truth annotation and the blue curves show the network prediction for the calcified plaque detection as probability values as a function of cross-sectional position.

4. CONCLUSIONS

The proposed synthetic lesion framework is a versatile tool that can create high-quality labeled training data for CTA applications. Our experiments demonstrate the transfer capability of the synthetically created plaque lesions for data-driven prediction tasks and thus the quality of the synthetic lesion generation module to capture the essential image characteristics of real lesions. Further extensions to different lesions types such as wall thickenings or spotty calcifications are currently under investigation. The fast algorithmic nature of the implemented forward model and its customization possibilities to site-specific requirements render this a valuable tool for machine-learning-based methods for coronary CTA.

ACKNOWLEDGEMENTS

This work was supported by European Regional Development Fund (ERDF) and the Free and Hanseatic City of Hamburg in the in the Hamburgische Investitions- und Förderbank (IFB)-Program PROFI Transfer Plus under grant MALEKA.

REFERENCES

[1] Maurovich-Horvat P, Ferencik M, Voros S, Merkely B, Hoffmann U. Comprehensive plaque assessment by coronary CT angiography. *Nat Rev Cardiol*. 2014;11(7):390-402. doi:10.1038/nrcardio.2014.60

[2] Kolossváry M, Szilveszter B, Merkely B, Maurovich-Horvat P. Plaque imaging with CT—a comprehensive review on coronary CT angiography based risk assessment. Cardiovasc Diagn Ther 2017;7(5):489-506. doi: 10.21037/cdt.2016.11.06

[3] Ronneberger, O., Fischer, P., Brox, T., „U-Net: Convolutional Networks for Biomedical Segmentation", Medical Image Computing and Computer-Assisted Intervention (MICCAI), Springer, LNCS, Vol. 9351: 234-241, 2015

[4] Noh, H., Hong, S., Han, B, "Learning Deconvolution Network for Semantic Segmentation," Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), IEEE Computer Society, 1520-1528, (2015).

[5] Simonyan, K. and Zisserman, A., "Very Deep Convolutional Networks for Large-Scale Image Recognition", in [3rd International Conference on Learning Representations, ICLR 2015, Conference Track Proceedings], (2015).

[6] Ioffe, S., Szegedy, C., "Batch normalization: Accelerating deep network training by reducing internal covariate shift", arXiv preprint arXiv:1502.03167, 2015

[7] Kingma, D. P. and Ba, J., "Adam: A method for stochastic optimization", in [3rd International Conference on Learning Representations, ICLR 2015, Conference Track Proceedings], (2015).