Automated Volumetric Intravascular Plaque Classification Using Optical Coherence Tomography (OCT)

Ronny Shalev,[a] Daisuke Nakamura,[b] Setsu Nishino,[b] Andrew M. Rollins,[c] Hiram G. Bezerra,[b] David L. Wilson,[c,d] Soumya Ray[a]

aDepartment of Electrical Engineering & Computer Science, Case Western Reserve University, Cleveland, OH, 44106, USA

bCardiovascular Imaging Core Laboratory, Harrington Heart & Vascular Institute, University Hospitals Case Medical Center, Cleveland, OH, 44106, USA

cDepartment of Biomedical Engineering, Case Western Reserve University, Cleveland, OH, 44106, USA

dDepartment of Radiology, Case Western Reserve University, Cleveland, OH, 44106, USA

Abstract

An estimated 17.5 million people died from a cardiovascular disease in 2012, representing 31% of all global deaths. Most acute coronary events result from rupture of the protective fibrous cap overlying an atherosclerotic plaque. The task of early identification of plaque types that can potentially rupture is, therefore, of great importance. The state-of-the-art approach to imaging blood vessels is intravascular optical coherence tomography (IVOCT). However, currently, this is an offline approach where the images are first collected and then manually analyzed a frame at a time to identify regions at risk of thrombosis. This process is extremely laborious, time consuming and prone to human error. We are building a system that, when complete, will provide interactive 3D visualization of a blood vessel as an IVOCT is in progress. The visualization will highlight different plaque types and enable quick identification of regions at risk for thrombosis. In this paper, we describe our approach, focusing on machine learning methods that are a key enabling technology. Our empirical results using real OCT data show that our approach can identify different plaque types efficiently with high accuracy across multiple patients.

Introduction

Cardiovascular diseases are the leading cause of death worldwide. An estimated 17.5 million people died from a cardiovascular disease in 2012, representing 31% of all global deaths. Of these deaths, an estimated 7.4 million were due to coronary heart disease and 6.7 million were due to stroke (Mendis 2011). The underlying disease process in the blood vessels that results in coronary heart disease (heart attack) and cerebrovascular disease (stroke) is known as *atherosclerosis*. It is a complex pathological process where fatty material and cholesterol are deposited inside the lumen of medium and large-sized blood vessels (arteries). These deposits (plaques) cause the inner surface of the arteries to become irregular and the lumen to become narrow, making it harder for blood to flow through. Further, the plaque can rupture, triggering the formation of a blood clot, which may eventually lead to disease.

To treat atherosclerosis, we must first have access to an imaging technique with suitable resolution. As we describe in the next section, the recently approved intravascular optical coherence tomography (IVOCT) approach fulfills this need. In IVOCT, a probe is inserted into a blood vessel. As the probe moves through the vessel, it collects images of the vessel wall. These images are subsequently analyzed by experts to identify at-risk regions.

A major issue when working with IVOCT, however, is that it can produce more than 500 image frames in a single scan, resulting in an explosion of image data. This can be difficult and labor-intensive to analyze manually, taking up to one hour of examination for each frame by a trained analyst. This often precludes measurements from every frame, and plaque classification is not even done because it is infeasible in terms of time. Further, this manual process is also prone to error. In prior work (Lu et al. 2012), our group has found evidence of up to 5% intra and 6% inter-rater variability among analysts looking at these images.

The goal of our work is to enable an effective detection and diagnosis of atherosclerosis, which is a necessary precursor for effective treatment. We do it in three ways: (i) reduce the effort involved, (ii) improve the accuracy of high-risk plaque identification and (iii) make the diagnosis available as early in the process as possible. The prevalence of atherosclerosis means achieving these goals can have a major impact on health worldwide.

We anticipate fulfilling our goals in two steps. In the first step, reported in this paper, we develop an automated method to process single frames generated by IVOCT scans. We demonstrate that it is accurate and efficient on real IVOCT data, it outperforms a previously published baseline (Ughi et al. 2013), and the output can be used by analysts to greatly reduce their annotation effort. In the second step, our goal is to integrate this approach into a real time visualization that accompanies an IVOCT scan. We would like to produce 3D images as in Figure 1, by stacking the output of multiple 2D frames. These images will be annotated with different detected plaque types, and will be used for rapid identification of high-risk regions for intervention and management and guidance.

In the rest of the paper, we first describe IVOCT. Next, we describe how we extract meaningful features from IVOCT images for our automated analysis, followed by a description of our classifier. We then describe empirical results that illustrate the performance characteristics of our approach, and discuss current limitations and future work.
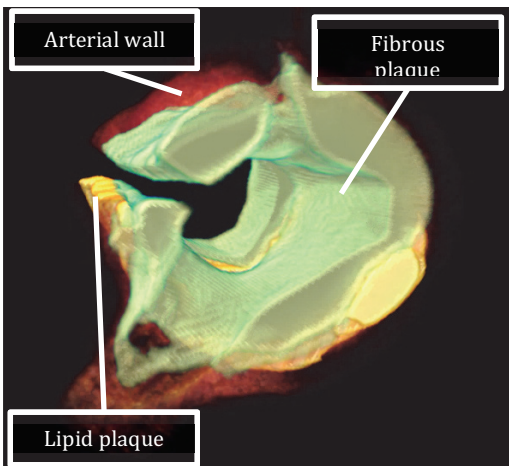

Figure 1: Example of desired 3D output (using current automated approach's output). This will greatly enhance the physician's ability to make (pre)treatment decisions.

Optical Coherence Tomography (OCT)

Before OCT was approved, Intravascular Ultrasound (IVUS) was used to identify plaques. IVUS has a resolution of ~200 μm and is able to detect for example, calcium. However, it cannot measure the distance between the superficial calcium and the lumen, nor can it assess the thickness of calcium due to acoustic distortion (Mintz et al. 1995). The underlying concept of OCT is similar to that of ultrasound; by measuring the delay time of optical echoes reflected or backscattered from subsurface structures in tissues, we can obtain structural information as a function of depth within the tissue (Tearney et al. 2012). However, OCT does not possess the limitations of IVUS.

In IVOCT we obtain cross-sectional images by inserting a flexible imaging probe (catheter) into the blood vessel to be imaged. The catheter has an optical fiber coupled to a lens and micro-prism. The micro-prism reflects the OCT beam perpendicular to the catheter longitudinal direction and captures the light that is back-scattered from that tissue (the reflected beam is referred to as an *A-Line*). The probe is then rotated and pulled back. This *pullback* creates a two-dimensional image (referred to as *polar* or *r-θ* image) by assembling successive A-lines next to each other resulting in an image shown in Figure 2b. This image is then transformed to Cartesian coordinates to produce the image shown in Figure 2c. A typical pullback contains 271 images covering 54*mm* and an image contains 504 A-lines.

Different tissues have different qualities that influence the back-reflectance. The longer the distance traveled, the longer the delay in returning to a detector. The delay in the returning light from deeper structures compared with shallow structures is used to reconstruct images.
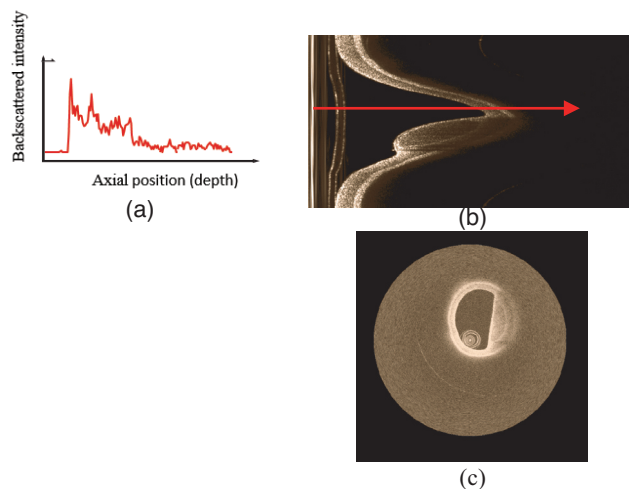

Figure 2: (a) Backscattered intensity of a single A-line (b) polar (r-θ) image (the red line is the A-line in (a)). (c) the polar image converted to the more human readable x-y.

Since its approval for clinical use, IVOCT has become an invaluable tool for vascular assessments due to its high contrast and microscopic resolution (5-15 μm), which is superior to other *in-vivo* imaging modalities such as IVUS. It has been shown that IVOCT is able to distinguish between key types of plaque (Yabushita et al. 2002), and aid in assessment of new coronary artery stent designs (Lu, Gargesha et al. 2012, Wang 2012). These characteristics make it ideal for our purposes.

Our group has access to a large database of manually analyzed OCT images obtained in a clinical setting. Images were collected on the C7-XR system from St. Jude Medical Inc., Westford, MA. It has an OCT Swept Source having a 1310 *nm* center wavelength, 110 *nm* wavelength range, 50 kHz sweep rate, and ~12 *mm* coherence length. The

pullback speed was 20 mm/s and the pullback length was 54 mm. Images from this source were used in training our machine learning approaches, described below.

Representing an OCT Image

In order to build our system, we need to automatically and accurately identify different plaque types in OCT images. In this section, we describe image characteristics that are key to identifying different plaque types. In constructing our features we use the qualitative description of the different plaques' characteristics in prior work (Yabushita, Bouma et al. 2002) described below. This also provides the ability to interpret results in a meaningful way.

A **fibrous plaque** (Figure 3A) has high backscattering and the region has relatively homogeneous intensity values. We see that the average intensity is high (bright). Likewise, the intensity is not attenuated much along the A-line.

A **lipid plaque** (Figure 3B) is a low intensity region with poorly delineated borders, a fast IVOCT signal drop-off, and little or no OCT signal backscattering, within a lesion that is covered by a fibrous cap. We see that the intensity starts very bright and decreases quickly along the A-line.

A **calcified plaque** (Figure 3C) appears as a low intensity or heterogeneous region with a sharply delineated border (leading, trailing, and/or lateral edges). Calcium is darker than fibrous plaque with greater variation in intensity level.
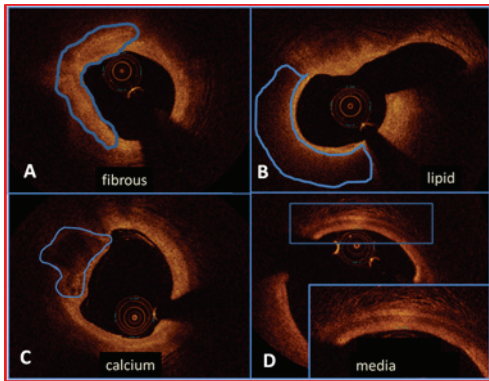

Figure 3: Appearance of plaque types in clinical images. A is fibrous, B is lipid and C is calcium. D shows the appearance of a normal blood vessel wall which has layered structure.

Based on this description, we construct a set of eight (real-valued) features for each pixel in the image. We compute these features using a three-dimensional (3D) neighborhood centered on the pixel of interest. The third dimension comes from neighboring frames (human analysts will often use adjacent frames when annotating a frame). In these features, σ represents the standard deviation of the intensity values within a 3D neighborhood.

Distance to lumen (D_l): This is a measure of the distance of the center pixel from the lumen border (i.e. the wall of the blood vessel). This feature helps identify lipid plaques since they are typically within a fibrous plaque as above.

Beam penetration (D_d) – This is a measure of the length of the beam from the lumen border to the back-border (the border beyond which the near infrared beam does not reach and the signal is at baseline). It depends on tissue type, thus can distinguish between plaques. This feature is invariant for pixels across an A-line but varies across A-lines.

Mean Intensity ($\bar{I}$): This represents the average signal intensity of the different plaque types within the 3D neighborhood. As can be seen from Figure 3, this is a very distinctive feature.

Homogeneity (H): This is a local coefficient of variation, ($\sigma / \bar{I}$). It helps in distinguishing between heterogeneous intensity regions and homogeneous intensity regions.

Relative Smoothness of Intensity (S): This is defined as: $S = 1 - 1 / (1 + \sigma^2)$. S is 0 for constant intensity regions and it approaches 1 for large deviations in intensity values.

Entropy (E): Entropy is another measure of the variability of the signal intensity within the respective plaque type regions. To compute it, we construct a histogram of the intensity distribution within a 3D-neighborhood. Then entropy is defined as:

$$E = -\sum_{i=0}^{L-1} p(z_i) \log_2 p(z_i),$$

where $p(z_i)$ is the probability of the intensity level z_i, $i=1,\ldots,L$ for L bins in the histogram of intensity levels. It is expected that within homogeneous regions the entropy will be low and within heterogeneous regions it will be high.

Similar features as these are often used in image processing applications (Gonzalez et al. 2009).

The final two features we use are *optical parameters*. These features are based on models of light transmission and reflectance described below.

Attenuation coefficient, μ_t – This feature measures the rate at which the signal intensity drops off within the tissue. Calcified plaque has lower attenuation, and as a result, IVOCT can see deeper into these tissues, compared to lipid where IVOCT does not see as deeply. For this reason, the attenuation coefficient (or penetration depth) gives useful information about plaque types.

Incident intensity, I_0 – This represents the backscattering characteristics of the plaque at the point where the light touches it. This feature is excellent at distinguishing fibrous plaques, which are very reflective.

In order to estimate I_o and μ_t, we modeled the OCT signal as Lambert-Beer exponential decay function (Wang et al. 2007), with the addition of baseline $I_{baseline}$ to account for noise and other sources that elevate the expected signal.

$$I_d(r) = I_0 T(r) S(r) \exp(-\mu_t r) + I_{baseline} \qquad (1)$$

Here $T(r)$ is the longitudinal point spread function which describes the shape and focal point of the beam and thus affects the contrast of the image, and $S(r)$ is the Gaussian coherence function which describes the signal roll-off with depth (Van Soest et al. 2010). In order to identify the unknown parameters in these two functions, we use nonlinear optimization over their unknown variables. After estimating $T(r)$ and $S(r)$, and subtracting the baseline intensity, we use least squares estimation to find the unknowns I_o and μ_t for each separate A-line (Gargesha et al. 2015). We verified these estimates by fabricating *phantom* (realistic imitations) blood vessels with known plaque types and checking the estimates against measured values in these cases.

The Plaque-Type Classifier

After extracting features from pixels in our OCT images, we then train a support vector machine (SVM) (Cristianini et al. 2000) with a radial basis function (RBF) kernel for classification of the individual pixels. The SVM is a state-of-the-art classification method. It is widely used due to its high accuracy, ability to deal with high-dimensional data, and flexibility in modeling diverse data sources. We use a standard SVM formulation (omitting details due to space). Given that we are interested in classifying three different plaque types, we use a one-versus-rest (OVR) approach for multi-class classification. There are two parameters which must be input to the SVM: C, the regularization parameter that trades off margin size and training error, and γ, the RBF kernel's bandwidth. In our experiments, we select these parameters using an internal 5-fold stratified cross validation loop and a two-dimensional grid search.

Empirical Evaluation

We now describe experiments to test our hypothesis that the system we described will be able to accurately and efficiently classify different plaque types from OCT images.

The clinical images (*in-vivo*) that we use were selected from the database available at our institutions. The images consist of 35 IVOCT pullbacks of the Left Anterior Descending (LAD) and the Left Circumflex (LCX) coronary arteries of patients acquired prior to stent implantation, with a total of 287 images across 35 patients. An expert cardiologist on our team then labeled volumes of interest (VOIs) as belonging to one of the three plaque types in the images. The expert marked the VOIs of a particular plaque type using freehand brush strokes. On the clinical images the expert annotated 311 VOIs (roughly equal number from each plaque type). VOIs were of various sizes and shapes. Most consisted of 2-5 image frames, 50-200 A-lines, and 20-50 sample points in each A-line.

We also acquired a second set of 106 images from blood vessels used in cadaver studies. Since in this case the blood vessels can be extracted and cryogenically frozen and imaged, they are much easier to label very accurately for the expert. However, since this is *ex-vivo*, we do not use these images for training our classifiers, but use them to validate the [I]results. We call these images "Cryo-images" below to distinguish them from the previous set.

Next, we preprocess all images for speckle noise reduction, baseline subtraction, catheter optical system correction, and catheter eccentricity correction. We segment the lumen and the back-border using dynamic programming. To do this, we use a cost function from prior work (Wang 2012). An example of the results of the back-border segmentation is shown in Figure 4 in both (r-θ) view and (x-y) view. Segmenting the image in this way is important because (i) the regions of interest are contained between these borders and the rest of the pixels do not contain any relevant information, and (ii) it enables us to properly compute the distance to the lumen and the beam penetration depth discussed above, which are important signals for different plaque types.
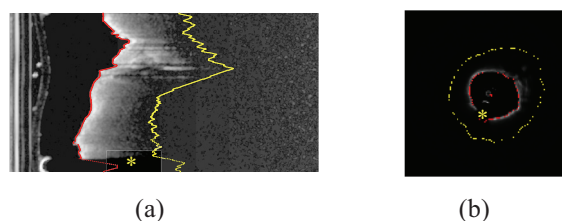

| (a) | (b) |

Figure 4: An illustration of back border segmentation (yellow line) along with lumen segmentation (red line) in a typical clinical image in both views. (a) is the polar image and (b) is the x-y image. The yellow line is broken due to view conversion). Asterisk marks the guide-wire shadowing artifact.

Next, we generate features by scanning the annotated VOIs in the image pixel by pixel. For each pixel, we construct a 7 x 11 x 3 neighborhood (0.035mm x 0.055mm x 0.6mm) around it. As long as the neighborhood is within the VOI, the features of the box are computed as explained above and the values are assigned to the pixel. In the cryo-images (images that did not take part in the training at all), we generated features for *all pixels* between border regions in a similar way.

For cross validation we use the processed images with a *leave-one-pullback-out* strategy. Here, in each iteration, we hold out all the data from one pullback as the test set and use the remaining 34 pullbacks as the training set. This mimics practical usage where the system will operate on novel pullbacks and is more stringent than using random folds. In a second experiment, we ran the trained classifiers on the cryo-images (these were not used at all during training/cross validation). We ran our experiments on a 64-bit Windows 7 machine with 3rd generation Intel Core i7 and 16 GB RAM

Results and Discussion

The Receiver Operating Curves (ROC) for each OVR classifier from the cross validation experiment is shown in Figure 5. The summary statistics are shown in Table 1, where the accuracy, sensitivity and specificity are noted at the optimal operating point along the curves. The ROC describes the system's behavior for a range of confidence threshold settings and enables the cardiologist (the end user) to decide on weighting the false positives (FP) and false negatives (FN) unequally (a very desirable property according to our expert).

	CALCIUM	LIPID	FIBROUS
ACCURACY	92.2±6.28%	96.95±2.79%	96.17±4.0%
SENSITIVITY	93.0±2.58%	98.95±2.35%	94.28±5.23%
SPECIFICITY	96.5±3.39%	93.65±2.77%	95.89±2.18%
AUC	0.9837	0.9947	0.9959

Table 1: Performance measures: Area under ROC and the accuracy, sensitivity and specificity at the optimal operating point on the ROC curves.

The overall accuracy results averaged over 35 folds are shown in Table 2. As can be seen from all of these results, our approach has excellent accuracy for all three plaque types. In fact, across the 35 folds, the median accuracy for all three plaque types is 100%, indicating that our classifiers are (in most cases) able to perfectly separate the plaque types using the features we designed. In a few folds, the accuracy is lower than 100%. We conjecture that this is because some pullbacks have many more images associated with them than others. When such a pullback is held out, the training set size becomes much smaller, and yields a classifier with lower accuracy.
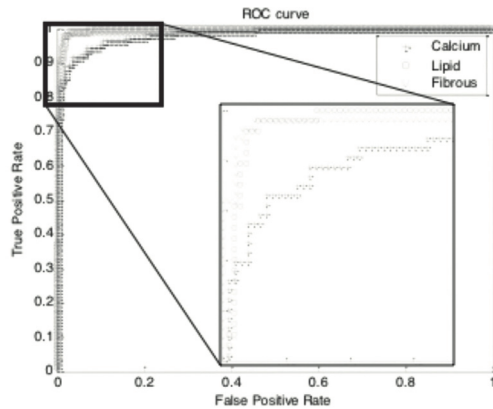

Figure 5: ROC curve for all three plaque types. Area Under the Curve (AUC) values are 0.9837, 0.9947 and 0.9959 for calcium, lipid and fibrous respectively.

In the second experiment, we ran our trained classifier on the cryo-images. We also ran a baseline approach following (Ughi, Adriaenssens et al. 2013). This approach used beam attenuation estimates from a layer model applied to single A-lines and 2D texture and geometric measures as features for classification with the added requirement of manual region of interest selection for analysis. These results are shown in Table 3. Here the "Other" row corresponds to pixels in these images that belong to none of the three plaque types. The accuracy of the approach in this case is lower, possibly because these are ex-vivo images which have somewhat different characteristics from the training set. However, our approach still outperforms the state of the art. Further, these values are still at a very useful level according our expert. In particular, cardiologists now divide an image into quadrants and simply state whether a quadrant contains a certain plaque type. If we use this as a performance measure, our current approach has perfect accuracy on the cryo-images.

	ACCURACY	MEDIAN ACC.
OVERALL	90.70±8.28%	
CALCIUM	92.14±10.74%	100%
LIPID	96.40±8.87%	100%
FIBROUS	100%±0.0%	100%

Table 2: Accuracy results for leave-one-pullback-out experiment

	OUR APPROACH	BASELINE
OVERALL	81.15%	69.4%
CALCIUM	97.62%	66.88%
LIPID	87.65%	67.07%
FIBROUS	97.39%	77.95%
OTHER	77.96%	30.46%

Table 3: Accuracy results for Cryo-images.

The results also indicate that in some cases some plaque types may be confused with others. For example, the average intensity of a lipid region may be very close to that of calcium. However, they may still be separable due to the fact that the lipid's attenuation coefficient is much higher.

To confirm our intuitive understanding of the plaques' characteristics we performed a *leave-one-feature-out* experiment. In this experiment, we ran the classifier using all of the features and noted the accuracy measures (as shown in Table 2). We then removed each feature at a time to see the impact on the accuracy. We found that removing the attenuation parameter had the biggest impact on the lipid accuracy reducing it down to 92.4±8.87% while removing the average intensity feature, had a significant effect on the fibrous' accuracy and uncertainty (down to 95.2%±10.75).

In addition to high accuracy, our approach was also efficient at classification. Each test fold (on average 200,000 datapoints) was classified in 0.0366 seconds by our implementation. This facilitates future real-time usage.

Finally, we consider whether an automatic classification procedure such as this can be useful in reducing the amount of time taken to process images in a clinical setting. In an initial experiment, we found that cardiolo-

gists would spend approximately 5 hours analyzing a section of a blood vessel. We then created a tool (Figure 6) with our classifier built in. The cardiologist would run the classifier for a new image and then, using the tool, analyze the results and correct some of the errors in the predictions. We found that this process took at most an hour, a reduction of 80%. This effort reduction indicates that improving the tool will make it deployable in the near future.
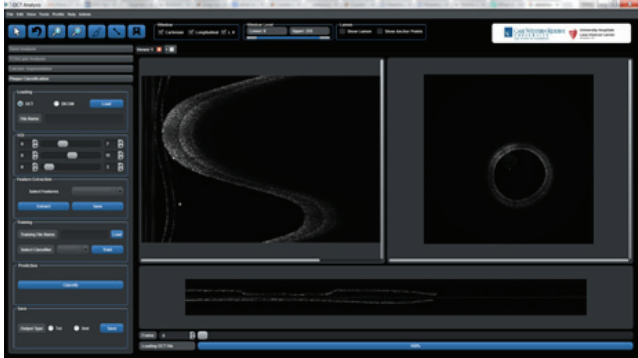

Figure 6: User Interface of editing tool.

Conclusion

In this paper, we have discussed an important emerging application: an automated approach to early plaque detection in blood vessels. Our approach analyzes OCT images to solve this task. Using a carefully designed feature set, we show that an SVM with an RBF kernel is a high-accuracy classifier for this task. Our results are of significant impact on this important problem (Wagstaff 2012) with implications for early diagnosis of cardiovascular disease. In future work, we plan to work on further improving our classification tool and integrating it with a real-time 3D visualization module which will be able to quantify (volume, area covered, etc.) the presence of calcified regions. This can help in decision making regarding stent implantation and pre-implantation treatment (e.g. directional atherectomy). We also plan to add an explanatory module to help explain the automated classification process to the cardiologists, and accept feedback in an active learning environment.

Acknowledgment

References

Cristianini, N. and J. Shawe-Taylor (2000). An introduction to support vector machines : and other kernel-based learning methods. Cambridge ; New York, Cambridge University Press.

Gargesha, M., R. Shalev, D. Prabhu, et al. (2015). "Parameter Estimation of Atherosclerotic Tissue Optical Properties from 3D Intravascular OCT " SPIE Journal of Medical Imaging **2**(1): 14.

Gonzalez, R. C., R. E. Woods and S. L. Eddins (2009). Digital Image processing using MATLAB. S.I., Gatesmark Pub.

Lu, H., M. Gargesha, Z. Wang, et al. (2012). "Automatic stent detection in intravascular OCT images using bagged decision trees." Biomedical Optics Express **3**(11): 2809-2824.

Mendis, S., Pekka Puska, and Bo Norrving., (2011). "Global atlas on cardiovascular disease prevention and control." World Health Organization.

Mintz, G. S., J. J. Popma, A. D. Pichard, et al. (1995). "Patterns of calcification in coronary artery disease. A statistical analysis of intravascular ultrasound and coronary angiography in 1155 lesions." Circulation **91**(7): 1959-1965.

Tearney, G. J., E. Regar, T. Akasaka, et al. (2012). "Consensus Standards for Acquisition, Measurement, and Reporting of Intravascular Optical Coherence Tomography Studies." Journal of the American College of Cardiology **59**(12): 1058-1072.

Ughi, G. J., T. Adriaenssens, P. Sinnaeve, et al. (2013). "Automated tissue characterization of in vivo atherosclerotic plaques by intravascular optical coherence tomography images." Biomedical Optics Express **4**(7): 1014-1030.

Van Soest, G., E. Regar, S. Koljenovi+,,, et al. (2010). "Atherosclerotic tissue characterization in vivo by optical coherence tomography attenuation imaging." Journal of biomedical optics **15**(1): 011105-011105.

Wagstaff, K. (2012). "Machine learning that matters." arXiv preprint arXiv: 1206.4656.

Wang, L. V. and H.-i. Wu (2007). Biomedical optics : principles and imaging. Hoboken, N.J., Wiley-Interscience.

Wang, Z. (2012). "Volumetric quantification of fibrous caps using intravascular optical coherence tomography." Biomedical Optics Express **3**(6): 1413.

Yabushita, H., B. E. Bouma, S. L. Houser, et al. (2002). "Characterization of human atherosclerosis by optical coherence tomography." Circulation **106**(13): 1640-1645.