

HALE: Healthy Area of Lumen Estimation for Vessel Stenosis Quantification

Sethuraman Sankaran¹(✉), Michiel Schaap¹, Stanley C. Hunley¹, James K. Min², Charles A. Taylor¹, and Leo Grady¹

¹ HeartFlow Inc., Redwood City, CA, USA
ssankaran@heartflow.com

² Department of Radiology, Weill-Cornell Medical College, New York, NY, USA

Abstract. One of the most widely used non-invasive clinical metric for diagnosing patients with symptoms of coronary artery disease is %stenosis derived from cCTA. Estimation of %stenosis involves two steps - the measurement of local diameter and the measurement of a reference healthy diameter. The estimation of a reference healthy diameter is challenging, especially in diffuse, ostial and bifurcation lesions. We develop a machine learning algorithm using random forest regressors for the estimation of healthy diameter using downstream and upstream properties of coronary tree vasculature as features. We use a population-based estimation, in contrast to single patient estimation that is used in the majority of the literature. We demonstrate that this method is able to predict the diameter of healthy sections with a correlation coefficient of 0.95. We then estimate %stenosis based on the ratio of the local vessel diameter to the estimated healthy diameter. Compared to a reference anisotropic kernel regression method, the proposed method, HALE (Healthy Area of Lumen Estimation), has a superior area under curve (0.90 vs 0.83) and operating point sensitivity/specificity (90%/85% vs 82%/76%) for the detection of stenoses. We also demonstrate superior performance of HALE against invasive quantitative coronary angiography (QCA), compared to the reference method (mean absolute error: 14% vs 31%, $p < 0.001$).

Keywords: Healthy lumen diameter · Stenosis detection · Coronary artery disease

1 Introduction

Coronary artery disease (CAD) is one of the leading causes of death and may result in acute events such as plaque rupture which demands immediate care or gradual events such as accumulation of plaque which leads to progressive anatomic narrowing resulting in ischemia. Coronary computed tomography angiography (cCTA) provides information on the degree of anatomical narrowing (stenosis) in different regions of the coronary artery tree. The degree of stenosis, called %stenosis, is a widely used clinical measure to decide between performing invasive angiography and pressure measurements or deferment of invasive measurements. The estimation of %stenosis is usually performed categorically

(e.g. 0%, 1–30%, 31–49%, 50–69%, 70–100%) in the clinic, or, less frequently, sent to a core lab with expert readers for analysis. Quantitative computed tomography (QCT) and QCA are methods where %stenosis is estimated as a number between 0% (healthy vessel) and 100% (complete obstruction). QCA is invasive, while QCT, evaluated on cCTA, is time consuming and generally performed in a core lab. The method we present enables automatic measurement of %stenosis given a lumen segmentation. The main novelty is that we estimate the healthy diameter with respect to a database of healthy sections from other patients, in contrast to regressing from individual patient data.

Different methods for stenosis detection and quantification were compared on a set of cCTAs in the 2012 MICCAI stenosis challenge [1]. The goal was to quantify results in a set of patients who were evaluated by expert readers and QCA, and to report diagnostic performance and differences in stenosis grade. For most of the algorithms, this process had, built into it, (i) a centerline detection algorithm, (ii) a lumen segmentation algorithm and (iii) a stenosis detection step. We believe that (iii) by itself has scope for improvement based on the best performing methods in literature, as shown later. Therefore, we use the same lumen segmentation to analyze the performance of all stenosis detection algorithms. Towards this, a lumen segmentation is read by expert readers and stenosed sections are annotated, which are then compared to the proposed method to quantify performance. We use the acronym HALE for the proposed method (Healthy Area of Lumen Estimation). HALE is designed to predict healthy lumen diameter, and thereby also enables the prediction of healthy lumen area and evaluation of %stenosis. We believe that HALE, when used with a state-of-the-art lumen segmentation algorithm, can be used as an accurate and efficient QCT tool. Kirisli et al. [1] conclude the MICCAI stenosis challenge by stating that “Given a similar accurate lumen segmentation, robust kernel regression outperforms the other approaches and is a good approach to quantify lesions from accurate lumen segmentation”. The quoted method [2] uses a robust kernel regression approach with a radial basis function applied on the segmented lumen radius profile. Later, it was suggested that natural discontinuities in lumen radii at bifurcations can be accounted for by using an anisotropic kernel [3].

While focal coronary disease is captured well by many methods in literature, diffuse and ostial coronary artery disease are difficult to diagnose due to the absence of a clear reference lumen diameter. Huo et al. [4] suggested that epicardial volume, length and lumen area are related by a power-law and that the coefficient of power-law separates subjects with and without diffuse disease, indicating only the presence or absence of diffuse lumen narrowing without specifying where the disease is present. In this work, we present a general framework that can identify regions of lumen narrowing in (coronary) arteries, including focal, diffuse, ostial and bifurcation disease. The (coronary) arteries are split into sections or stems, and each stem is associated with features corresponding to its crown (downstream vasculature), root (upstream vasculature) and sibling (the other child vessel of its parent, if available). We predict the healthy diameter of the stem using a machine learning method trained on these features on a

database of 6000 stems (from 200 patients) and tested on 4697 stems (150 patients). We demonstrate that HALE performs better than state-of-the-art techniques over different lesion characteristics including the challenging cases of diffuse and ostial disease.

2 Methods

The first step in our process is the extraction of a coronary centerline tree and lumen segmentation. Following this, trained CT readers evaluate the lumen segmentation and make corrections, as necessary. Since manual annotation of diseased sections is performed on the lumen segmentation rather than the cCTA, performance does not depend on the algorithm used for centerline detection and lumen segmentation and many available methods may be used [5,6]. In the section below, we first describe the process of manual annotation of sections of disease. Then, we describe the proposed method, HALE, including definition of features and estimation of healthy lumen diameter. We then discuss how we evaluate %stenosis and the metrics used for validation.

2.1 Manual Annotation

Trained readers of cCTA assess lumen segmentation on a cohort of patients and identify locations of lumen narrowing (i.e. %stenosis $\geq 50\%$). This process mimics the process of reading %stenosis from CT scans in the clinic, i.e. estimated visually rather than assessing a reference diameter and evaluating the ratio of minimum lumen diameter to the reference diameter. To provide confidence in the readings, each patient is assessed by three readers and only sections that have a consensus read are used for training and testing. For convenience, the coronary trees are split into sections, where each section is marked either as diseased or healthy. Sections are split using locations of bifurcations as separators. The rationale for using bifurcation as separators is that flow rate in a given section being constant, a healthy vessel maintains its radius within a section to preserve a homeostatic state of wall shear stress.

2.2 Data-Driven Estimation of Healthy Radius

The approach we outline here aims to estimate healthy vessel diameter. Regions of disease are evaluated by dividing the difference between the estimated healthy diameter and local diameter with the estimated healthy diameter, and comparing it to a diagnostic threshold of 50%. Our approach maps metrics derived from the epicardial vasculature to a healthy diameter using a machine learning approach, which is trained on a database of sections annotated as healthy.

In contrast to previously published methods, we do not solely rely on the patient's vasculature to determine the healthy diameter at a given location. We use a machine-learning approach relying on a population of 6000 healthy stems from 200 patients. This enables a better identification of non-focal stenosis

morphologies such as long diffuse lesions, ostial lesions, or lesions which are present along an entire section. To determine features of the machine learning algorithm, we first evaluate local diameter using maximum inscribed spheres. An alternative approach is to derive average diameter from the area of lumen along the normal to centerlines. This approach was not used as it provides inaccurate values near bifurcations. To assess the healthy lumen diameter, we first split the lumen segmentation into stem-crown-root units. A stem is the section for which we are evaluating healthy diameter. A crown refers to its downstream vasculature and a root refers to its upstream vasculature. We also identify a sibling vessel which is the other child of the parent vessel. By definition, all the features are not available for all stems (e.g. ostial sections do not have a root unit and terminal sections do not have a crown unit) and are given a default special value of -1 . The crown, root and sibling units for a stem are shown in Fig. 1.

Our approach omits one stem at a time and uses features from the rest of the vascular tree to infer the healthy lumen diameter of the stem under consideration. For each stem in a given coronary vasculature, the following features, hypothesized to be relevant, are extracted for the corresponding crown, root and sibling vessels (when available) - average, maximum and minimum lumen area (A), volume (V), length (L), V/A , and V/L . We also use a feature (d_m) derived from Murray’s law [7], which is a physiologic model of how the diameters of parent (d_p) and daughter vessels are related. This feature is calculated as $d_m = (d_p^3 - d_s^3)^{1/3}$, where d_s is the diameter of its sibling vessel.

We use random forest regression [8] to predict the healthy lumen diameter. Random forests are known to be effective and powerful for high dimensional and heterogeneous features. Random forests employ an ensemble of decision trees, each of which is composed of a random subset of features and training data. The values from the decision trees are pooled together and averaged to compute the final predictor of healthy lumen diameter (d_p). Percent stenosis is evaluated from the ratio of the local lumen diameter (d_l) to the predicted healthy lumen diameter as $\alpha = \max\left(0\%, \left(1 - \frac{d_l}{d_p}\right) \times 100\%\right)$. We use 50 trees with an average of 5 features per tree (chosen based on a 5-fold cross validation).

Results of the random forest regressor on a test set of 4697 stems (from 150 patients) are evaluated by assessing sensitivity, specificity and area under the receiver-operator characteristic (ROC) curve. Sections annotated by readers as




Fig. 1. A schematic of the method used is shown. The coronary tree is split into many stem-crown-root units. Stems are defined based on branch points as separators with the corresponding crown and root being the downstream and upstream vasculatures respectively. Features are derived from epicardial volume, length and lumen diameter.

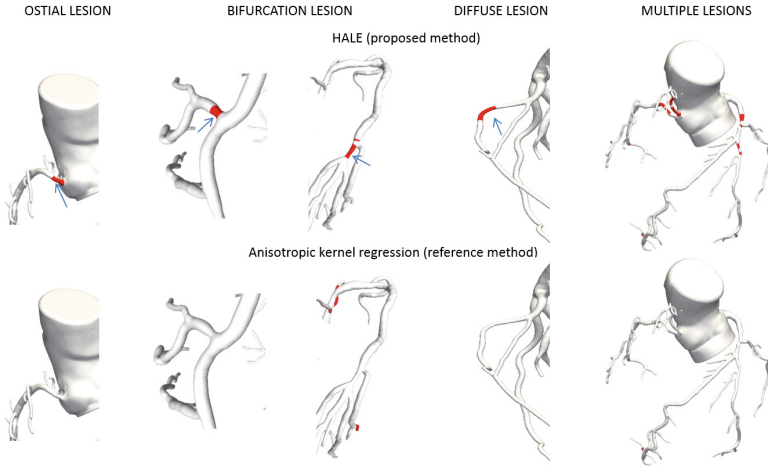


Fig. 2. Comparison of the machine learning method (top) with an anisotropic kernel method [3] (bottom) on patients with (from left) ostial, bifurcation, diffuse and multiple lesions. Anisotropic kernels can minimize the effect of steep changes in healthy lumen area across bifurcations [3] but is unable to detect lesions in the examples shown above in contrast to HALE. Regions with $\alpha \geq 50\%$ are shown in red. (Color figure online)

diseased are considered positive, which are further classified as “true positive” if the random forest predicts $\%stenosis \geq o\%$ or “false negative” otherwise, where the operating point o can be different for each method. Similarly, sections which are annotated as healthy are classified as “true negative” if the random forest predicts $\%stenosis < 50\%$ and “false positive” otherwise. Sensitivity (Se) and specificity (Sp) are defined as

$$Se = \frac{TP}{TP + FN}, Sp = \frac{TN}{TN + FP}$$

An ROC curve is plotted by evaluating the sensitivity and specificity for different value of cutoffs used to define sections of disease, (i.e.) $\alpha \leq x \quad \forall x \in [0\%, 100\%]$.

3 Results

We implemented the robust kernel regression [2] and anisotropic kernel regression methods [3] to serve as a baseline. We refer to anisotropic kernel regression as the “reference” method since it outperformed the robust kernel regression method in our dataset (shown later). First, we use HALE and evaluate healthy lumen diameter and subsequently $\%stenosis$ on five patients with either ostial lesions, bifurcation lesions, diffuse lesions or a combination thereof. Figure 2 shows a map of regions identified as diseased using HALE and the reference anisotropic kernel regression method, suggesting that the latter fails to capture non-focal lesions.

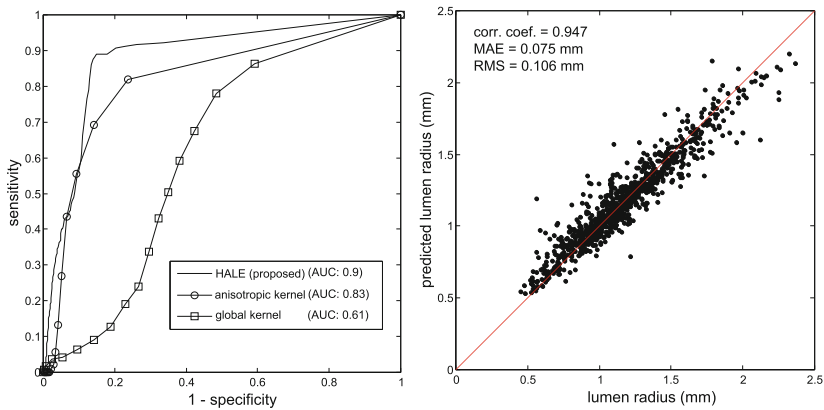


Fig. 3. The figure on the left demonstrates superior performance of HALE (area under the curve (AUC) : 0.90) compared to anisotropic kernel (AUC: 0.83) and global kernel methods (AUC:0.61). The figure on the right shows a comparison of the predicted and measured lumen radius in vessel sections marked healthy by human evaluation.

To quantify the overall performance of the algorithm, we predict healthy diameter and assess presence of stenosis on 150 patients (4697 sections) distinct from those used for training. A scatterplot of average radius of sections annotated as healthy compared to their estimated healthy radius is shown in Fig. 3. The correlation coefficient between the predicted and the measured healthy lumen diameter is 0.947, with a mean absolute error of 0.150 mm and a root mean squared error of 0.212 mm. The operating point sensitivity and specificity for detecting %stenosis using HALE is 90%/85% (operating point, σ is 48%), compared to 77%/52% using a global kernel regression method and 82%/76% using the reference anisotropic kernel regression method (operating point, σ is 32%). The receiver operator characteristic (ROC) curves for the three methods are compared in Fig. 3 with the area under curve being 0.90 (HALE), 0.83 (anisotropic kernel regression) and 0.61 (global kernel regression).

Next, we apply the method on patients who underwent a coronary angiography and the corresponding diseased locations were identified and quantified using QCA by an independent expert at a core laboratory. Coronary QCA data from a subset of the DeFACTO clinical trial [9] (Clinicaltrials.gov # NCT01233518) is used as the reference ground truth data. Performance of HALE and the reference method against QCA data on 69 measured vessel segments is tabulated in Table 1. There is a significant improvement in mean absolute error (MAE) and an insignificant difference in correlation coefficient, evaluated on the same lumen segmentation. Figure 4 compares locations of disease identified by HALE, the reference method, CTA and QCA evaluated by a core-lab for a severely diseased patient with multiple locations of lumen narrowing. The results show that HALE is able to identify four lesions out of five identified by CTA (only four of these are identified by QCA due to distal loss of contrast).

Table 1. The mean absolute error (MAE) of HALE is significantly better ($p < 0.001$, for both healthy diameter and %stenosis) and correlation coefficient (R) is slightly better but not significant ($p = 0.47$ for healthy diameter and $p = 0.36$ for %stenosis) compared to the reference method, both using QCA data as the ground truth.

Method	Metric	Bias	MAE	R	Slope	Intercept
Reference (anisotropic kernel)	healthy dia. (mm)	-0.52	0.58	0.76	0.51	0.90
HALE (machine learning)	healthy dia. (mm)	0.29	0.50	0.77	0.61	1.42
Reference (anisotropic kernel)	%stenosis	-31 %	31 %	0.54	0.98	-0.30
HALE (machine learning)	%stenosis	-13 %	14 %	0.58	0.85	-0.04




Fig. 4. Comparison of (a) HALE and (b) anisotropic kernel method on a severely diseased coronary artery with (c) five identified regions of lumen narrowing in CTA and (d) four stenosed regions identified by QCA, showing that HALE is able to identify four of the lesions while only one was identified by the anisotropic kernel regression method. Regions with $\alpha \geq 50\%$ are shown in red in panels (a) and (b). (Color figure online)

In comparison, the reference method is able to identify only a single lesion. This example demonstrates that with an accurate lumen segmentation algorithm, we are able to evaluate lesions with complex morphologies.

4 Discussion

We proposed a method that quantifies stenoses by first training a machine learning method on healthy vessel sections. We hypothesized that a set of features derived from the geometry of the downstream vasculature, upstream vasculature, and sibling vessel can be used to estimate the healthy vessel dimensions of a given section. To test this, we partitioned the geometry into various “stem-crown-root” units and used metrics such as epicardial vascular volume, lumen area, centerline lengths and other metrics derived from these. We also used a

feature motivated from Murray's law that relates lumen diameter of parent sections to its daughter sections, though the impact of the feature itself was modest. We extracted patient-specific metrics, omitting one section at a time, and using a database of these metrics, mapped them to a healthy lumen diameter. We obtained a correlation coefficient of 0.947 with a mean absolute error of 0.15 mm for predicting lumen diameter of healthy sections. HALE had an operating point sensitivity/specificity of 90%/85% for detecting stenoses. The mean absolute error in quantifying the degree of stenosis reduced from 31% using the reference method to 14% using HALE, with QCA being the ground truth.

The kernel-regression based methods are able to capture regions of focal narrowing but not the other disease morphologies, likely because they rely heavily on local patient-specific data without accounting for population data. The method for detection of diffuse lesions by Huo et al. [4], on the other hand, uses a population-based cutoff tailored to, solely, the presence or absence of diffuse lesions. We posit that the higher mean absolute error and lower bias in quantifying degree of stenosis by the reference method is because regressing on a single patient lumen area results in underestimation of healthy lumen radius.

The proposed method, HALE, can be used with any lumen segmentation algorithm, and is fully reproducible and takes only a few seconds. Depending on the application, HALE can be used with an automated lumen segmentation algorithm for on site evaluation of %stenosis, or be used with a semi-automated method offline or in a core-lab setting. It is our belief that an accurate QCT assessment tool would involve the coupling of an accurate lumen segmentation algorithm with an accurate algorithm for evaluation of %stenosis.

References

1. Kirisli, H.A., et al.: Standardized evaluation framework for evaluating coronary artery stenosis detection, stenosis quantification and lumen segmentation algorithms in computed tomography angiography. *Med. Image Anal.* **17**(8), 859–876 (2012)
2. Shahzad, R., et al.: Automatic segmentation, detection and quantification of coronary artery stenoses on CTA. *Int. J. Cardiovasc. Imaging* **29**(8), 1847–1859 (2013)
3. Sankaran, S., Grady, L., Taylor, C.A.: Fast computation of hemodynamic sensitivity to lumen segmentation uncertainty. *IEEE TMI* **34**(12), 2562–2571 (2015)
4. Huo, Y., et al.: CT-based diagnosis of diffuse coronary artery disease on the basis of scaling power laws. *Radiology* **268**(3), 694–701 (2013)
5. Lesage, D., et al.: A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. *Med. Image Anal.* **13**(6), 819–845 (2009)
6. Schaap, M., et al.: Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Med. Image Anal.* **13**(5), 701–714 (2009)
7. Sherman, T.F.: On connecting large vessels to small. The meaning of Murray's law. *J. Gen. Physiol.* **78**(4), 431–453 (1981)
8. Breiman, L.: Random forests. *Mach. Learn.* **45**(1), 5–32 (2001)
9. Min, J.K., et al.: Diagnostic accuracy of fractional flow reserve from anatomic CT angiography. *JAMA* **308**(12), 1237–1245 (2012)