Data-driven Methods for Automated Assessment of Coronary Artery Disease

by

Kritika Iyer

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Biomedical Engineering)
in the University of Michigan
2022

Doctoral Committee:

Professor C. Alberto Figueroa, Co-Chair
Professor Brahmajee K. Nallamothu, Co-Chair
Professor Raj R. Nadakuditi
Professor David Nordsletten

Kritika Iyer
kritiyer@umich.edu
ORCID iD: 0000-0002-3993-3348

© Kritika Iyer 2022

ACKNOWLEDGMENTS

I would like to first thank my family for their unconditional support during the toughest times. My mother especially flew across the country on short notice multiple times during my PhD simply because I said I needed her, and I am so grateful and thankful to know that my family will always be in my corner. When I was in high school, I swore that I would never become a computer scientist and sit in front of a computer all day like my dad. Joke's on me, and I would like to thank my dad for all his help from teaching me about bash scripting to offering to help me with coding logic even if he had no idea what the purpose of my code was. And of course I would like to thank my brother, my biggest cheerleader. Arjun, I couldn't have done this without you!

Next I would like to thank my advisors, Dr. C. Alberto Figueroa and Dr. Brahmajee K. Nallamothu. Thank you both for taking a chance on someone with no experience in machine learning or computational modeling and trusting me with this project. I have learned so much in the past few years, both in terms of technical skills and life skills. I would also like to thank my committee members, Dr. David Nordsletten and Dr. Raj Rao Nadakuditi. They both went above and beyond their role as committee members to support me and I've learned a lot from our scientific discussions.

I am also extremely thankful for all the members of the CVBL lab, past and present: Fede, Sabrina, Jonas, Chris, Theo, Taeouk, Liz, Jonathan, Hamid, Vasilina, Nitesh, Abhilash, Diederik, Ignas, Yunus, Pieter, and Nic. My lab members make up one of the most supportive, thoughtful, and fun groups I've been a part of. I have used almost all of you as my sounding board when I was stuck in my research, so thank you for always agreeing to help me even when you didn't always know what I was talking about! I'm really going to miss all our shenanigans, from having a cake in the lab almost every week to dinners together and game nights. I would also like to thank all the undergrads and post-grad students who helped me with different parts of the project: Cyrus Najarian, Aya Fattah, Mohammed Fattah, Ricky Hong, Jiayun Zheng, Dalia Kaufman, Noah Stevens, and Nathan Li. I couldn't have completed this work without their help.

To all the friends I have made in Ann arbor, thank you all for being my family away from home. To all the friends who have stayed in touch from undergrad and beyond, thank you for still being there for me. There are way too many of you to list, but thank you for everything: long walks and hikes, ranting sessions, virtual movie nights, surprising me with strawberry shortcake and care packages, exploring Michigan together rain or shine, and making sure I never felt alone.

TABLE OF CONTENTS

ACKNOWLEDGMENTS ii

LIST OF FIGURES vi

LIST OF TABLES xiii

LIST OF ACRONYMS xiv

ABSTRACT xv

CHAPTER

1 Introduction 1

1.1 Motivation 1
1.2 Anatomical and Functional Assessment of CAD 3
1.2.1 Anatomical Assessment 3
1.2.2 Functional Assessment 9
1.3 Contributions and Outline of Thesis 12

2 Background 14

2.1 Coronary Anatomy and Physiology 14
2.2 Background on X-ray Angiography 15
2.3 Background on Machine Learning and Neural Networks 18
2.4 Projective Geometry and Stereovision 21
2.5 Hemodynamic Modeling in the Coronary Arteries 23

3 AngioNet: A Convolutional Neural Network for Vessel Segmentation in X-ray Angiography 25

3.1 Introduction 25
3.2 Methods 26
3.2.1 Datasets 26
3.2.2 CNN Design and Training 29
3.3 Results 32
3.3.1 Learned Filters using the Angiographic Processing Network 32
3.3.2 Segmentation Accuracy Metrics 33
3.3.3 Comparison of AngioNet versus Current State-of-the-art Semantic Segmentation Neural Networks 34

3.3.4 Evaluation of Vessel Diameter Accuracy versus QCA . . . 37
3.4 Discussion . . . 38
3.4.1 Learned Filters using the Angiographic Processing Network . . . 38
3.4.2 Comparison of AngioNet to Current State-of-the-Art Semantic Segmentation Neural Networks . . . 39
3.4.3 Evaluation of Vessel Diameter Accuracy . . . 41
3.5 Conclusion . . . 42

4 Stenosis Characterization . . . 43

4.1 Introduction . . . 43
4.2 Methods . . . 44
4.2.1 Stenosis Localization . . . 45
4.2.2 Stenosis Segmentation . . . 46
4.2.3 Stenosis Severity Quantification . . . 47
4.2.4 Statistics . . . 48
4.3 Results . . . 48
4.3.1 Study Population . . . 48
4.3.2 Stenosis Localization . . . 49
4.3.3 Stenosis Severity Quantification . . . 51
4.4 Discussion . . . 55
4.4.1 Stenosis Localization . . . 55
4.4.2 Stenosis Severity Quantification . . . 56
4.5 Conclusions and Future Directions . . . 56

5 3D Reconstruction of the Coronaries from 2D X-ray Angiography Images . . . 58

5.1 3D Reconstruction using Projective Geometry . . . 58
5.1.1 Forward Projection . . . 59
5.1.2 Epipolar Matching . . . 62
5.1.3 Point Cloud Method with Ordering Constraint . . . 65
5.2 A Machine Learning Approach for Coronary 3D Reconstruction from Uncalibrated X-ray Angiography Images . . . 67
5.2.1 Introduction . . . 67
5.2.2 Methods . . . 68
5.2.3 Results . . . 72
5.2.4 Discussion . . . 76

6 Relative Flow Extraction from Angiographic Time Series . . . 81

6.1 Introduction . . . 81
6.2 Methods . . . 83
6.2.1 Synthetic Angiography Series Generation . . . 83
6.2.2 Relative Flow Estimation Algorithm . . . 85
6.3 Results and Discussion . . . 86
6.4 Future Work: Application to Clinical Angiograms . . . 89

7 Hemodynamic Studies . . . 90

7.1 Effect of Reconstruction Error on FFR Computation for Varying Levels of Stenosis 90

7.1.1 Methods . . . 91
7.1.2 Results . . . 91
7.1.3 Discussion . . . 95
7.2 FFR Verification in Clinical Case . . . 96
7.2.1 Vessel Segmentation . . . 96
7.2.2 Stenosis Detection and Segmentation Refinement . . . 97
7.2.3 3D Reconstruction . . . 98
7.2.4 FFR Computation . . . 99
7.2.5 Discussion . . . 101
8 Discussion and Future Directions . . . 103
8.1 Discussion . . . 103
8.2 Future Directions . . . 106
8.3 Summary of Ongoing Challenges and Potential Solutions . . . 108
8.3.1 Vessel Segmentation . . . 108
8.3.2 Stenosis Localization . . . 108
8.3.3 3D Reconstruction . . . 109
8.3.4 Relative Flow Extraction . . . 109
APPENDICES . . . 110
BIBLIOGRAPHY . . . 114

LIST OF FIGURES

FIGURE

1.1 X-ray angiography is the most commonly used diagnostic method, followed by QCA and FFR; however, the less widely-adopted diagnostic modalities lead to better patient outcomes. . . . 3
1.2 Examples of filter-based methods of vessel segmentation. These methods are susceptible to including the catheter and background artifacts in the segmented images. . . . 5
1.3 Schematic of the reconstruction method by Galassi et al. [38] Centerlines in 2D angiography planes are projected as surfaces and the 3D centerline is identified at their intersection. . . . 8
1.4 Example of $\mathrm{FFR_{CT}}$ computation from Taylor et al. [137]. This method uses high fidelity 3D models of patient anatomy and flow. . . . 10
1.5 Schematic of QFR workflow. A and B show examples of angiograms and measured FFR. C is the reconstructed model of the vessel using 3D QCA. D shows the computed FFR in the reconstructed vessel. E is the series of angiographic frames used to extract mean flow. [140] . . . 11

2.1 Schematic of the major coronary branches [83] . . . 15
2.2 Schematic of X-ray Angiography machine. A) The X-ray angiography system is composed of a patient table, X-ray detector, and X-ray source. The X-ray detector and source rotate about a fixed point known as the isocenter. B) The position of the X-ray detector is given by two angles, the angulation and elevation. The angulation is expressed in terms of RAO/LAO and the elevation is expressed as CRA/CAU. Figures adapted from [92]. . . . 16
2.3 Selected frames from an angiographic time series. Radio-opaque dye is first injected into the arteries via a catheter. As the heart pumps, the coronaries fill with dye until the entire coronary tree can be visualized. Eventually, the dye is washed out of the vessels of interest. . . . 17
2.4 Schematic of Multi-layer perceptron . . . 19
2.5 Schematic of convolutional neural network . . . 19
2.6 Visualization of attention weights for image captioning. The words in the caption are color-coded to correspond to the region of the image that the network associates with that word. Image reproduced from [75]. . . . 20
2.7 Diagram of 3D point triangulation using stereovision. The corresponding point x_R for a given reference point x_L must lie on the epipolar line e_R. Adapted from [95]. . . . 22
2.8 The intersections between two vessel surfaces from 2D images can produce several possible 3D centerlines. Reproduced from [10] . . . 23

3.1 Diagram of datasets for CNN training and evaluation. AngioNet's performance was compared against state-of-the-art neural networks, all trained on the UM Dataset. The MMM QCA dataset was used to quantify segmentation diameter accuracy by comparing AngioNet's results against the diameters reported in QCA. Acronyms: left coronary artery (LCA); right coronary artery (RCA); Madras Medical Mission (MMM). 27
3.2 a) Annotated QCA report, along with the corresponding diameters in the report table. Highlighted values correspond to maximum (proximal) and minimum (distal) diameters (segments 1 and 9, respectively). b) Schematic of how a MATLAB script was used to delineate regions in the neural network segmentation corresponding to the regions measured in the QCA report, along with the computed proximal and distal diameters. 29
3.3 AngioNet Architecture Diagram. AngioNet is composed of an Angiographic Processing Network (APN) in tandem with Deeplabv3+. The APN is designed to improve local contrast and vessel boundary sharpness. The output of the APN, a single channel filtered image, is copied and concatenated to form a 3-channel image which is input into the backbone network. 30
3.4 Examples of learned filters when the APN is fine-tuned together with Deeplabv3+ and U-Net. The APN was initialized with the combination of unsharp mask filters shown above, and learned new filters to aid segmentation. Each example image is the output of the APN after training with different data partitions. 32
3.5 Examples of AngioNet segmentation on left coronary tree, taken at two different angles (1,2), and right coronary tree (3). AngioNet does not segment the catheter (red arrows), despite its similar diameter and pixel intensity as the vessels (2,3). It also ignores bony structures such as the spine in (3) and ribs in (1). 34
3.6 Summary of APN performance. All results are derived from the networks trained on all five partitions of the UM training set, unless otherwise noted as a k-fold result. a) Comparison of Dice score distribution on test set. AngioNet has the highest average Dice score, with scores ranging from 0.737 to 0.946. Adding the APN improves the Dice scores of all backbone networks (Deeplabv3+, U-Net, and Attention U-Net). Dashed lines correspond to the Test Dice in the table below. clDice is also summarized in the table, with the Attention U-Net backbone demonstrating highest topology preservation. b) Segmentation comparison on challenging images with low contrast, faint vessels, and a curved catheter. AngioNet can segment more vessels in these images without segmenting the catheter (red arrows). The Dice scores of all predicted images are indicated in the corner of each image. 36
3.7 a) Correlation plot of QCA and AngioNet derived vessel diameters. b) The Bland-Altman plot demonstrates that AngioNet's segmentation and QCA are interchangeable methods to determine vessel diameter since more than 95% of points lie within the limits of agreement. The red error bars represent the 95% confidence interval containing the limits of agreement. The mean difference in diameter between methods is 0.24mm or 1.1 pixels. 37

4.1 Diagram of training and validation datasets. Angiograms from UM were used to train neural networks for the stenosis localization and vessel segmentation tasks. Validation angiograms and their corresponding QCA reports were obtained from ICRF. The ground truth for stenosis localization was lesions of ≥50% diameter reduction according to two board-certified cardiologists, and the ground truth for vessel diameter and percent stenosis calculation was from QCA analysis of vessels containing >20% stenosis. 45
4.2 Diagram of the proposed neural network stenosis characterization pipeline (top) and the QCA workflow (bottom) to determine lesion diameter and stenosis severity. 46
4.3 Architecture of our stenosis patch segmentation network, based on AngioNet [55]. The angiographic processing network (APN) learns filters that improve segmentation performance, and the VGG U-Net backbone performs the segmentation task. 47
4.4 Examples of stenosis localization by the YOLOv5s network (top) and their corresponding ground truth as determined by 2 cardiologists (bottom). In these examples, correctly identified regions containing stenoses are given by the solid red bounding boxes. False positive neural network detections are shown using black dashed boxes, and false negative detections are shown using blue dashed boxes. Some of the false positive detections do contain stenoses; however, these were determined to be <50% stenoses by the cardiologists. 49
4.5 Precision-Recall curve for YOLOv5 stenosis localization. There is a trade-off that occurs between precision and recall, since minimizing false positives by increasing the confidence threshold for stenoses can lead to an increase in false negatives and vice versa. The mean average precision was 0.65 across all levels of recall, and the best performance occurs at a precision of 0.63 and recall of 0.68 (red circle). 51
4.6 Variability in ground truth stenosis localization. A total of 141 stenoses ≥50% were localized by both cardiologists. Both cardiologists independently identified the same lesion 65 times, and disagreed on 76 lesions. 52
4.7 Examples of stenosis segmentation and their corresponding measurements from our pipeline and QCA. 52
4.8 Agreement between lesion diameters measured by QCA and our pipeline. A) Box plot of relative error in millimeters (QCA – Pipeline diameters), with a median value of 0.006mm. B) Correlation plot of diameters measured using both methods and the line of best fit, with a Pearson's r value of 0.9098. 53
4.9 A Bland-Altman plot comparing vessel diameter measured by QCA and our stenosis characterization pipeline in vessels containing stenoses greater than 20% diameter reduction. The gray dashed line is the mean relative error, 0.045mm. The limits of agreement are at 0.52 and -0.49mm 54
4.10 A Bland-Altman plot comparing percent diameter reduction measured by QCA and our stenosis characterization pipeline. The mean relative error of 7% is given by a gray dashed line. The limits of agreement are given by solid lines at 34.3% and -23%. 54

5.1 3D right coronary tree from CCTA and its corresponding projection images. These images were used as the input for all projective geometry reconstruction methods. . . . 59

5.2 This figure demonstrates how 2D projection images were mapped into 3D space to mimic the geometry of image acquisition. The X-ray source and detector/image plane positions are given by their angle of rotation about the isocenter (ISO). The vessel pixels on each image can then be mapped into coordinates on the 3D image plane. The 3D object is assumed to be centered on the isocenter. 60
5.3 Forward projection results. When using 2 images for reconstruction, there is uncertainty about which of the many possible branches are correct. When 3 images are used, the reconstruction time increases by a factor of 2.2. 61
5.4 An example of forward projection of centerline points. There are several possible centerlines since we only have a limited number of projection images. The color of the points reflects their depth into the page, where the darkest points are the furthest away. 62
5.5 Epipolar matching method. The algorithm first automatically identifies vessel centerlines and breaks them up into segments. The user chooses which segments belong to the same vessel branch. Reference points on the branch from one image are considered sequentially, and matching points are identified using a combination of epipolar line intersections and an ordering constraint. Finally, an example set of corresponding point is shown. 63
5.6 The two main challenges of this method are foreshortening in the input images and choice of algorithm parameters. Foreshortening in one or more images can lead to false correspondences using epipolar matching. Other parameters such as the number of centerline points to sample can affect accuracy in highly curved vessels. In this example, minimizing the distance between the previously identified point p_{i-1} and the two epipolar line intersections does not identify the correct intersection. 64
5.7 An example of a right coronary tree reconstructed using the point cloud approach. Depending on the relative orientation of different vessels in the tree, there can be gaps in the centerline due to foreshortening or regions of uncertainty with many possible 3D points corresponding to multiple epipolar line intersections. 66
5.8 Diagrams of coronary tree generation. A: Centerlines derived from patient image data are aligned to determine the distribution of their control points. New control points are randomly sampled from this distribution to produce new centerlines. B: Schematic of how the gaussian stenosis profile is introduced in a tapering vessel C: Examples of a real patient coronary tree and synthetically generated coronary trees 68
5.9 Schematic of Synthetic X-ray Projection Generation. Projections were modeled as a cone beam of X-rays from a point source hitting a detector at a series of sampled random positions. 69
5.10 Neural network architecture. The input is binarized X-ray angiography images of the ground truth vessel geometry with a Euclidean distance transform applied for mathematical smoothness. The network is composed of a backbone network and multi-layer perceptrons which predict vessel centerlines and radii. 70

5.11 Single- and multi-stage neural network reconstruction results. 5.11A shows an example vessel with a 70% stenosis and the corresponding reconstructions. The single-stage network fails to predict the stenosis while the multi-stage network prediction does capture the stenosis and more closely follows the vessel centerline path. 5.11B shows the reconstruction error for the vessel centerlines (root mean squared error) and radii (mean absolute error at the stenosis) in 10 synthetically generated vessels, each with varying levels of stenosis from 20-90%. The multi-stage approach had a 52% lower centerline error compared to the single-stage approach across all levels of stenosis. The single-stage approach cannot predict the presence of stenosis, demonstrated by the increasing radius error for increasing stenosis severity. 73
5.12 Comparison of neural network and projection-based reconstruction of a single vessel from different input images. A) Sample 3D RCA geometry and corresponding projection images. B) The neural network has on average 38% lower reconstruction error for all values of $\Delta\theta$, the angle between input projection images. C) Examples of the same 3D geometry reconstructed from different input images. Red arrows indicate regions where the neural network centerline does not follow the same bends as the ground truth while brown arrows indicate regions of uncertainty and gaps in the projection-based centerline. 74
5.13 Left: Examples ground truth centerlines (blue), reconstructed centerlines (red), and their corresponding input images. As in the single vessel case, the neural network predictions follow the average path of all vessels but do not have the same tortuosity as the ground truth centerlines. Right: Examples of radius along the centerline in vessels with varying numbers of stenoses. When no stenosis is present, the neural network captures the linear tapering of the vessel well. The neural network is able to accurately predict the severity and location of a single stenosis, but in the multi-stenosis case, the neural network approximates serial stenoses as a single stenosis with an error of 0.3mm in stenosis severity. 75
5.14 Comparison of pressures in ground truth and reconstructed coronary trees. When no stenosis is present, the error in pressure drop ΔP along the main RCA vessel is small (1.4mmHg) despite the error in geometric reconstruction. In the stenosis case, subpixel errors in stenosis severity (0.07mm) can lead to a larger error in pressure drop (4.5mmHg) along the vessel. 77

6.1 A snapshot of dye transport in a healthy and stenosed coronary tree, where dye concentration is highest in the red regions. At the same point in time, the dye has traveled farther down the RCA in the healthy coronary tree compared to the diseased case due to the flow limiting severe stenosis in the mid RCA. 84
6.2 Examples of synthetic X-ray angiography projections acquired at different angles by projecting regions containing dye concentration above 190 mol/mm^3. Selected frames from the time series of projection angle II are shown below. 85

6.3 Overview of flow extraction algorithm. The centerline of the vessel of interest is first specified. Radii along the centerline are used to compute the area of the target vessel. This process is repeated for all frames, and a plot of area per frame is generated. The first differences method is used to find the derivative of the graph, which can be normalized by mean vessel radius to get axial velocity. The velocity is then multiplied by cross sectional area to compute the flow. 86
6.4 The computed flow between frames is shown for each vessel branch in the healthy coronary tree (solid lines) and stenosed coronary tree (dashed lines). For each coronary tree, the flow computed using each of the 2 projections is shown. Although the values of flow in both projections differ, the average values of flow in each branch are similar. 87
6.5 Comparison of ground truth flow splits and computed flow splits in the healthy (blue) and stenotic (red) coronary trees. The average difference in % flow across all branches was 5.15% in the healthy tree and 3.68% in the diseased tree. 88

7.1 Ground truth RCA vessel geometries with 20-90% stenoses (gray arrows). Geometry A had the lowest reconstruction error out of the 10 geometries that were reconstructed in subsubsection 5.2.3.1 while Geometry B had the maximum reconstruction error. FFR computed in each of these geometries was compared to FFR in their reconstructions. 90
7.2 Comparison of ground truth and reconstructed geometries with 70% stenosis. The reconstruction in A underestimated stenosis diameter by 0.007mm, leading to slightly higher pressure at the inlet. Meanwhile, reconstruction B overestimated stenosis diameter by 0.203mm and underestimates pressure at the inlet. 92
7.3 Comparison of ground truth vs reconstruction pressure along the centerline in cases with 50,70, and 90% stenosis. 93
7.4 Comparison of stenosis radius error and error in FFR for varying levels of stenosis. In the radius graph (left), the orange series summarizes the mean and standard deviation of the set of 10 geometries from which geometry A and B were selected. Geometry A has the minimum error, in green, and Geometry B has the maximum error, in purple. The FFR graph (right) demonstrates the error in FFR in the ground truth (solid lines) and reconstructed (dashed lines) geometries A and B. Error in FFR increases sharply as stenosis severity increases past 50%. 94
7.5 Input angiography frames (top) and their corresponding segmentations (bottom). The main branches of the coronaries are included in the segmentation images, besides the SA branch, and the catheter is excluded. 97
7.6 Stenosis segmentation and refinement. The top row shows the regions of the image that were localized as stenoses and their segmentation overlaid on the input image. The bottom row shows the refined segmentation image where the results of the patch segmentation replace the region in the original image segmentation. Blue arrows point to areas of diameter mismatch, gaps, or missing vessels. 98
7.7 Reconstructed coronary tree and the predicted radii along each branch. The network identified a stenosis in the RCA and did not predict a stenosis in the PLV (distal half of the RCA vessel in this graph). 99

7.8 FFR computation in the ground truth geometry from CCTA compared with the reconstructed geometry. The reconstruction underestimates stenosis severity and has the RCA stenosis in the wrong position relative to the AM branch, leading to a lower pressure drop across the stenosis and higher FFR. 100

LIST OF TABLES

TABLE

3.1 Accuracy of AngioNet using K-Fold Cross Validation 35

7.1 FFR values in ground truth and reconstructed geometries for varying levels of stenosis 94

A.1 Dimensions of RCA tree branches for synthetic vessel generation 111

LIST OF ACRONYMS

CAD Coronary Artery Disease

CCTA Coronary Computed Tomography Angiography

CFD Computational Fluid Dynamics

CNN Convolutional Neural Network

FEM Finite Element Method

FFR Fractional Flow Reserve

iFR instantaneous wave-Free Ratio

LAD Left Anterior Descending

LCx Left Circumflex

MLP Multi-layer Perceptron

PCA Principal Component Analysis

QCA Quantitative Coronary Angiography

QFR Quantitative Flow Ratio

RCA Right Coronary Artery

ABSTRACT

There are several diagnostic tools to assess stenosis severity in coronary artery disease. The gold standard is anatomical assessment via cardiologists' visual inspection of X-ray angiography (XRA), which can be subjective. Quantitative Coronary Angiography (QCA) is a less subjective, semi-automated method to quantify stenosis severity. Alternatively, functional metrics that estimate how a stenosis affects blood flow, such as Fractional Flow Reserve (FFR), have demonstrated better diagnostic outcomes than anatomical assessment; however, they are not widely used due to their cost and risk. Ideally, quantitative and functional information could be derived directly from XRA images without the additional operator dependence, time, and cost associated with performing FFR or QCA.

This project aims to develop automated approaches for anatomical and functional quantification of disease severity using XRA images. To this end, we have developed algorithms for 1) coronary vessel segmentation, 2) stenosis detection and characterization, 3) 3D reconstruction of coronary trees, and 4) image-based flow extraction. These algorithms can be used in conjunction with computational fluid dynamics modeling to assess the functional significance of disease.

We first present AngioNet, a neural network for coronary segmentation from XRA images. AngioNet's key innovation is an Angiographic Processing Network which learns the best preprocessing filters to improve segmentation performance. AngioNet demonstrates state-of-the-art segmentation accuracy (Dice score=0.864) and does not segment the catheter in challenging cases where other neural networks fail.

Building upon AngioNet, we developed a neural network and image processing pipeline to automate the tasks of QCA: localizing, segmenting, and measuring stenoses. This pipeline was able to measure stenosis diameter within 0.206 ± 0.155mm or approximately 1 pixel of ground truth measurements from QCA. It is also the first automated pipeline to quantify rather than categorize disease severity. To assess functional disease severity, we require the patient's 3D coronary geometry to perform hemodynamic simulations and compute FFR. To this end, we developed a machine learning approach for 3D vessel reconstruction from a series of uncalibrated 2D XRA images. Our neural network has demonstrated sub-pixel error in radius reconstruction (0.16 ± 0.07mm) and 1% error in FFR computed in a reconstructed coronary tree.

In addition to the 3D coronary geometry, information about patient-specific flow or pressure is

required to perform a hemodynamics simulation for FFR computation. We developed an algorithm that tracks vessel area in sequential XRA frames to estimate relative flow in each branch. We validated the algorithm using a simulation of dye transport under steady flow conditions as the ground truth. On average, the error in relative flow per branch was 5.15% for a healthy coronary tree and 3.68% in a coronary tree with stenosis.

We finally combined the methods developed in this thesis to compare stenosis severity and FFR using the above algorithms against patient-specific clinical measurements in a proof-of-concept analysis. In this patient, AngioNet and the stenosis characterization pipeline accurately captured the vessel anatomy and stenosis diameters. There was a substantial 11% error in computed FFR (0.98) compared to clinically measured FFR (0.88), mainly due to errors in 3D reconstruction. Reconstruction accuracy could be improved by retraining the network on a wider range of anatomical variations and incorporating patient motion between input images to match clinical practice. Another improvement would be to automate and incorporate motion compensation in the flow extraction algorithm so that it can be applied to clinical angiograms.

CHAPTER 1

Introduction

1.1 Motivation

Coronary Artery Disease (CAD) is a condition which affects the vessels that supply blood to the heart muscle. The disease presents as a buildup of plaque in the coronaries leading to a narrowing of the vessel lumen known as stenosis. As the disease progresses, stenoses can limit blood flow and cause an oxygen deficit in the heart tissues known as ischemia. According to the American Heart Association, CAD affects over 20 million adults in the United States alone [5]. Globally, CAD disproportionately affects middle and low-income countries and accounts for over 7 million deaths daily [107]. Although improved diagnostic procedures and treatments have decreased mortality over the last few decades, CAD continues to account for over 25% of adult mortality worldwide [107].

There are several diagnostic modalities to assess whether a patient has CAD and requires intervention. When a patient is suspected to have coronary artery disease, Coronary Computed Tomography Angiography (CCTA) can be performed to non-invasively identify regions of stenosis [94]. However, this imaging modality is sensitive to artifacts created by calcified plaque which can make it difficult to assess disease severity. CCTA also only provides anatomical information and does not provide any information on possible ischemia. Since functional information is a better indicator of disease severity, stress tests are more commonly performed to determine whether catheterization is required [80].

Although stress tests and CCTA are useful for early-stage diagnosis, cardiac catheterization remains the gold standard for CAD diagnosis and treatment. Cardiac catheterization has been used since the 1950s to assess disease severity [47]. In this procedure, a series of X-ray angiograms are acquired as radio-opaque dye is injected into the coronaries via catheter. Cardiologists visually inspect the acquired images to identify stenoses and estimate diameter reduction. The advantages of X-ray angiography are that it has higher resolution than CCTA and is not sensitive to calcification artifacts. Another advantage of X-ray angiography is that an intervention can be immediately per-

formed while in the cardiac catheterization laboratory if necessary. Furthermore, an angiographic time series of dye being flushed through the coronaries contains implicit information about flow. Although this is not a direct measurement of coronary function like a stress test, it can provide more insight into flow-limiting lesions or microvascular dysfunction than CCTA.

One disadvantage of X-ray angiography is that visual estimation of disease severity is semi-quantitative and can be subjective. Several studies have demonstrated that visual inspection leads to over-estimation of stenosis severity, particularly for intermediate lesions. A more objective and accurate alternative is Quantitative Coronary Angiography (QCA), which is a semi-automated software tool used to measure vessel diameter and diameter reduction. QCA has has been validated in numerous clinical and experimental studies [45, 90, 123] and has lower inter-operator variability than visual estimation [159, 39, 132]. Despite these advantages, QCA only provides anatomical and not functional information about disease severity.

The most accurate measure of CAD severity is also the most invasive. Fractional Flow Reserve (FFR) is a quantitative functional metric defined as the ratio of pressures on either side of a stenosis [101]. In order to measure FFR, adenosine is administered to the patient to induce maximum blood flow to the heart, which is known as hyperemia. A pressure guidewire is then pulled through the stenosis to obtain a time-averaged measure of pressure on either side of the stenosis. Studies have shown that FFR can lead to better long-term patient outcomes [27, 103] and up to 30% fewer adverse events compared to diagnostic decisions guided by angiographic estimates of disease severity [34].

Despite its clear benefits, the cost of performing the measurement and increased risk to the patient have limited the adoption of FFR in practice. There is an inherent risk to the patient when passing a guidewire through a stenosis; another potential risk is an adverse reaction to induced hyperemia in those with pulmonary issues such as asthma [102]. Another barrier to widespread use of FFR measurement is the challenges of inducing maximum hyperemia using pharmacological agents. Pijls and Tonino suggest that maximum hyperemia may not be achieved depending on the method and dosage of adenosine administration [102]. This can lead to overestimation of FFR, which is only accurate when maximum hyperemia is achieved.

From this overview of diagnostic techniques, we observe two trends (Figure 1.1). The first is that quantitative and functional metrics lead to better patient outcomes compared to estimates based on anatomy. The second is that anatomical assessment remains the most prevalent diagnostic modality. One barrier for wide-scale adoption of quantitative measures such as QCA and FFR is their high cost. Performing these measurements also adds time to decision making, which is not always feasible in patients with acute disease. In an ideal scenario, quantitative and functional information could be derived directly from X-ray angiography images. Computational modeling has the potential to do exactly this.

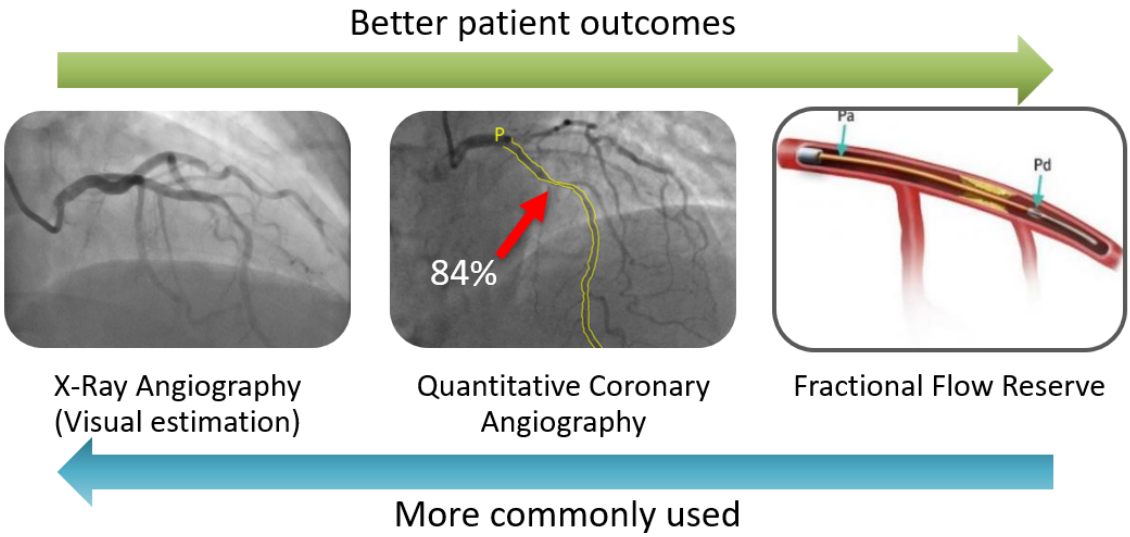


Figure 1.1: X-ray angiography is the most commonly used diagnostic method, followed by QCA and FFR; however, the less widely-adopted diagnostic modalities lead to better patient outcomes.

The goal of this work is to develop automated algorithms to extract quantitative anatomical and functional information from X-ray angiography images. Not only do computational methods have reduced subjectivity and inter-operator variability compared to operator-derived measurements, they also can reduce time to a diagnostic decision. By supplementing X-ray angiography with computationally-derived metrics, diagnostic outcomes would improve without the additional risks, time, and cost associated with FFR or QCA.

1.2 Anatomical and Functional Assessment of CAD

In this section, we describe previous work in the field of computational assessment of CAD severity. Broadly, these methods can be classified as either anatomical or functional assessment of disease.

1.2.1 Anatomical Assessment

There have been two major efforts to quantify disease severity based on coronary anatomy. The first is to develop algorithms that can fully automate the same task that QCA performs: stenosis characterization. Stenosis characterization in X-ray angiography images is comprised of three steps: 1) vessel segmentation, 2) stenosis localization, and 3) stenosis severity quantification.

The second set of methods focus on developing a more robust measurement of diameter reduc-

tion – one that is not dependent on a single projection of a 3D structure – by reconstructing the 3D coronary geometry from several 2D X-ray angiograms. Not only do 3D anatomical models provide a more accurate picture of diameter reduction, they can also be used to derive functional metrics.

1.2.1.1 Vessel Segmentation

Automatic segmentation of the coronaries in X-ray angiography images may seem like a trivial task. Since the coronaries are visualized by imaging the contrast agent as it flushes through the vessels, it may seem as if a simple thresholding filter could separate the vessels from the background. However, here are several aspects of this imaging modality that make segmentation challenging.

The catheter that is used to inject contrast into the coronaries during image acquisition has a similar diameter and curvature as the vessels themselves. The catheter also has the same pixel intensity as the coronaries due to the contrast agent. These characteristics make it difficult to distinguish the catheter from the vessels from a computer vision perspective. There are other imaging artifacts that further complicate the task of segmentation. For example, structures such as the patient's ribs or implanted devices such as stents attenuate X-rays and appear in the angiography image. There is also large variability in local contrast throughout the X-ray image due to tissue attenuation of the X-rays and a poor signal-to-noise ratio [73]. Due to these challenges, vessel segmentation continues to be an area of great interest and algorithm development.

Several filter-based methods for vessel segmentation have been adopted over the last 10 years. The Frangi vesselness filter is a Hessian-based filter designed to identify tubular structures in images [36]. Many groups have used the Frangi filter in combination with other image-processing techniques such as fast-marching algorithms, edge-detection filters, or region-growing to segment the coronaries in angiography images and remove spurious objects [60, 127, 77, 135]. Another set of segmentation methods involve using Principal Component Analysis (PCA) to separate foreground and background components of the image [78] [105]. These methods applied specialized filtering algorithms to the PCA-separated layers to segment the coronaries.

While these filter-based methods are effective at segmenting the coronaries, they have a major limitation. Due to the similarity of the catheter and coronary vessels in an angiography image, these methods often include the catheter in the segmentation (Figure 1.2). Many of these methods require user input to eliminate the catheter or correct the vessel boundaries, which can be both subjective and time consuming. More recently, machine learning approaches have demonstrated greater robustness to image artifacts in coronary segmentation.

There are several hybrid methods which employ image-processing filters for an initial segmentation or feature extraction followed by machine learning to refine the segmentation. Nasr-Esfahani et al. [91] used a combination of multiscale top hat filters, edge detection filters, and features ex-

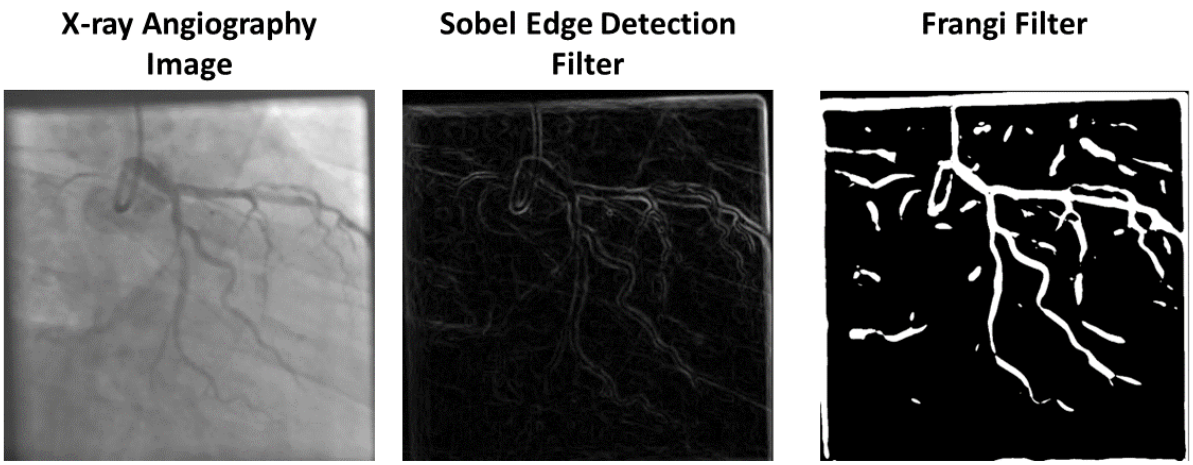


Figure 1.2: Examples of filter-based methods of vessel segmentation. These methods are susceptible to including the catheter and background artifacts in the segmented images.

tracted from a neural network to perform vessel segmentation. Another hybrid method was developed by Cervantes-Sanchez et al. [20]. In this method, Gaussian Matched Filters and Gabor filters were applied to coronary angiograms to enhance the vessel-like structures present in the images. The filtered images were then input into a neural network to obtain a segmented image.

Others have employed pure deep-learning approaches for vessel segmentation. Many convolutional neural networks have been designed specifically for angiographic segmentation, while others add additional stages to state-of-the-art networks such as U-Net [114]. Yang et al. [153] developed a Convolutional Neural Network (CNN) based on U-Net to segment the major branches of the coronary arteries. Like Yang et al., the authors of [57] focused on single-vessel segmentation using selective feature mapping applied to a backbone CNN. Multichannel Segmentation Network with Aligned input (MSN-A)[33], is a CNN based on U-Net designed to segment the entire coronary tree. The inputs to MSN-A were the angiographic image and a co-registered "mask" image taken before the dye was injected into the vessel. Samuel and Veeramalai [121] developed VSSC Net (Vessel Specific Skip Chain Network), which added two stages on top of VGG16, a seminal CNN for image classification [131]. The additional stages refined the feature maps learned by VGG16 to improve vessel segmentation. Shin et al. [128] combined a feature-extraction CNN with a graph convolutional network in order to improve segmentation by learning the tree structure of the vessels.

These deep learning approaches have demonstrated higher accuracy and faster segmentation than the filter-based and hybrid methods. There is still however room for improvement. Some of the deep learning segmentation methods are susceptible to segmenting the catheter or other artifacts in addition to the coronaries. Others were trained on small datasets, therefore their accuracy

on noisy or challenging images and generalizability to external datasets is unknown. The best-performing methods have limited use-cases: the network of Yang et al. [153] was only developed for single vessel segmentation (the main coronary branches) while the multi-input strategy of Fan et al. [33] requires the angiographic sequence to be acquired with minimal patient table motion, which goes against standard clinical practice.

1.2.1.2 Stenosis Localization and Severity Quantification

While algorithms for vessel segmentation have been explored for over a decade, automated detection of stenoses has been a much more recent area of focus. Stenosis localization is essentially an object detection task, and neural networks have demonstrated good performance on object detection in natural camera images since 2014 [43]. In this case, instead of identifying discrete objects such as cars or animals in natural images, the network must identify the regions of interest in a larger structure.

Since 2019, there have been several papers that propose algorithms based on neural networks to automatically localize stenoses. Some use single angiographic frames, while others use the entire angiographic time series to track the location of stenoses for the duration of image acquisition. Stenosis-DetNet was developed by Pang et al [97] to leverage sequential information in an angiographic series to improve stenosis localization. The network uses a classification-regression model to track bounding box movement for each stenosis in the sequence and reduce the incidence of false positive detections. Moon et al. developed a weakly supervised attention-based network [89] which was able to automatically identify key frames in an angiographic series and localize stenoses.

Since the goal of stenosis localization is ultimately to characterize stenosis severity, pipelines that incorporate both stenosis detection and classification have been developed. Au and colleagues [6] proposed a fully automated pipeline for stenosis characterization analogous to QCA, composed of separate CNNs for stenosis localization, segmentation, and classification. Their pipeline used a simplified version of the YOLO object detection network [110] to identify 64x64 pixel regions of stenosis, a variation of U-Net [114] to segment the localized image patch, and a custom fully connected CNN to classify the lesion as greater or less than 70% diameter reduction. Another end-to-end method was developed by Cong et al [7]. This approach bypassed vessel segmentation and used a combination of a CNN and recurrent neural network to a stenosis activation heatmap and corresponding classification as greater or less than 25% stenosis. The authors later expanded this work to include automatic key frame selection using neural networks and improved stenosis localization accuracy using a more sophisticated network [25]. CathAI is another end-to-end pipeline analogous to QCA. It uses neural networks for key frame selection, stenosis localization, and prediction of stenosis severity.

While end-to-end pipelines for stenosis characterization exist, their main limitation is the lack of quantitative metrics for disease severity. Most approaches use a binary classification system for stenoses greater than or less than a threshold value. The approaches that predict continuous values for stenosis severity do not use quantitative measurements such as those from QCA to train the neural network. Instead, they use the visual estimates by cardiologists in the medical chart as the ground truth. Another drawback of directly predicting stenosis severity rather than measuring it is the lack of interpretability; there is no intuitive way to verify the accuracy of the network on unlabeled data or identify cases where the network would fail. There is a need for a stenosis characterization pipeline that measures diameter reduction rather than classifying or estimating stenosis severity.

1.2.1.3 3D Coronary Reconstruction

While characterizing stenosis severity from angiography images is useful for making diagnostic decisions, 2D images have their limitations. Since they are projections of the true 3D geometry, the images as susceptible to vessel foreshortening. A 3D anatomical model would provide a more accurate picture of stenosis severity without any artifacts or distortion.

The two main classes of 3D reconstruction methods are based on epipolar constraints and those based on deformable models. Epipolar constraint methods determine corresponding points or segments in multiple images and triangulate the position of the 3D object using these corresponding structures [10, 13, 17, 22, 38, 72, 145]. In the deformable model approach, mechanical properties are prescribed to a 3D cylinder which is iteratively deformed until its projections match a series of X-ray angiography images. An energy function is minimized to identify the deformation to apply to the 3D snake during each iteration [16, 26, 152, 158].

We first describe epipolar constraint algorithms that reconstruct Non-Uniform Rational B-Spline (NURBS) 3D models of the coronary tree from 2-3 angiography images. Galassi et al [38] applied multiscale Hessian-based filters and a fast-marching level set algorithm to segment the vessels from a pair of angiography images. The centerlines of the vessels were extracted using a subvoxel skeletonization algorithm and corresponding segments in both images were matched. Each segment was projected and the centerline of the 3D coronary vessel was identified by the intersections of the segment surfaces (Figure 1.3). NURBS contours were used to define the luminal contours of the 3D vessel surface. Vukicevic et al. [145] also used a NURBS-based semi-automatic approach to 3D reconstruction of the coronary tree. Their segmentation algorithm involved user specification of the start and end points of each vessel of interest, followed by a Frangi filter to extract the vessel-like structures. The vessel centerlines were extracted and split into branches by a fast-marching algorithm. The 2D centerlines were projected using partial matching to create a 3D centerline, followed by B-spline interpolation. The coronary branch surface was then recon-

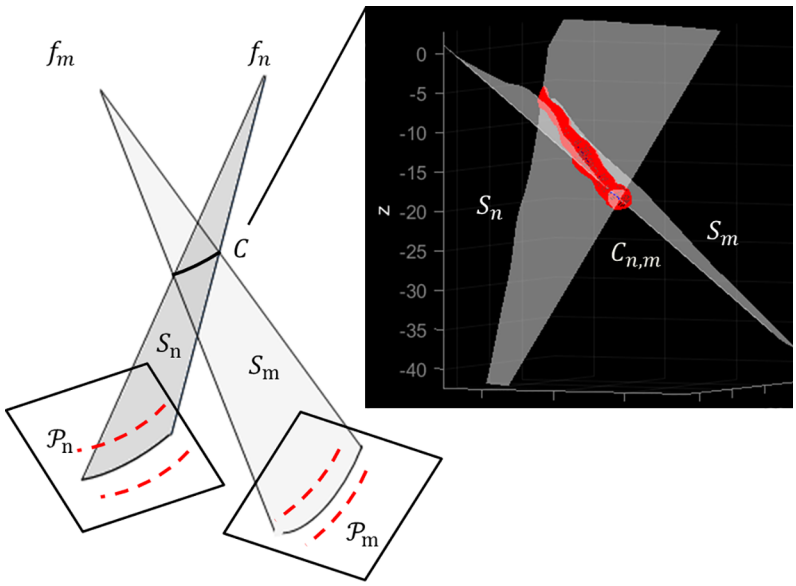


Figure 1.3: Schematic of the reconstruction method by Galassi et al. [38] Centerlines in 2D angiography planes are projected as surfaces and the 3D centerline is identified at their intersection.

structed by projecting vessel border points from the 2D images to create 3D vessel cross-sections. Banerjee et al. [10] used a similar approach as Galassi et al. for the vessel segmentation portion of their method but implemented a fully automated 3D vessel reconstruction algorithm. Once the images were segmented and centerlines were extracted, centerline points from corresponding branches were projected to form an outer hull point cloud containing many possible point correspondences. Several algorithms such as B-spline fitting, distance thresholding, and re-projection error were used to refine the point cloud into the true 3D centerline.

We now present several recent examples of deformable models and energy-based methods for coronary tree reconstruction. The method of Yang et al. [152, 26] did not directly generate a segmentation image; however, active contour models were used to identify the vessel boundaries and find the centerline of the vessels in each 2D image plane. External energy terms derived from generalized gradient vector flow [151] were used to deform a 2D parametric curve until it matched the centerline of the image. The endpoints of the curve were back-projected into 3D and a cylindrical vessel with assigned internal energy was initialized between these endpoints. The 2D external energy forces were then back-projected and applied to this 3D vessel while the internal energy governed its deformation; the 3D vessel was iteratively reprojected and deformed until convergence. Zheng et al. [158] proposed a semi-automatic method to reconstruct the coronaries using deformable snakes. Like the method of Yang et al., this method did not directly segment the coronary vessels; however, pixel intensity and pixel intensity gradients were used to identify vessel boundaries and inform the energy function used to determine the vessel centerline. First, the user

determined corresponding points along the vessel using epipolar constraints to initialize the snake. These points were back-projected to create an initial 3D snake. Energy functions relating the 2D and 3D snakes were used to deform the 3D centerline and maintain continuity, smoothness, and accurate reprojection onto the 2D planes.

Both epipolar constraint-based methods and deformable models have been shown to faithfully reconstruct the 3D geometry of coronary arteries from X-ray angiography images. The main advantage of the epipolar constraint methods are their relatively low computational cost and speed, although these methods sometimes require manual interaction to match vessel branches, endpoints, or bifurcation points in multiple images. Even when fully automated, the accuracy of projection-based methods depends on the angles of images used and the prevalence of vessel foreshortening and overlap. On the other hand, deformable models require less manual interaction, but these methods still require an operator to identify endpoints of the vessel branches and are more time consuming due to their computational complexity. The typical time for reconstruction is on the order of minutes per branch, which can be problematic in coronary trees with many bifurcations.

1.2.2 Functional Assessment

The goal of computational methods for functional assessment of CAD is to provide clinicians with information about potential flow-limiting lesions without having to perform any invasive assessments that would require instrumentation of the stenosed artery such as FFR. FFR is typically only measured where it is difficult to visually estimate ischemic potential, as is the case for intermediate stenoses. An experimental study by De Bruyne et al. [28] validated that FFR can accurately determine the ischemic potential of a series of mild to moderate stenoses as well. Even so, FFR is only measured in less than 30% of eligible patients due to the risks described in section 1.1 and the associated cost [49]. Computational methods would make deriving FFR and related functional metrics more accessible and possibly more widely applied.

There are several approaches to derive computational FFR, using a variety of medical image modalities and hemodynamic modeling frameworks. One of the first methods proposed to compute FFR, known as $\mathrm{FFR_{CT}}$, was developed by Taylor et al. in 2013 [137]. In this method, a 3D model of the patient's coronary anatomy was first extracted from CCTA volumes using a combination of image processing algorithms and manual correction. A finite element mesh was generated based on the patient-specific geometry, and the 3D Navier-Stokes equations were solved on the nodes of the mesh. The boundary conditions were as follows: Coronary inflow was determined by allometric scaling based on the segmented volume of the left ventricle. Outlet resistances were similarly tuned based on morphological scaling until the mean computed pressure matched the measured mean brachial pressure. No slip conditions were applied at the walls. A 3D computational fluid dynamics

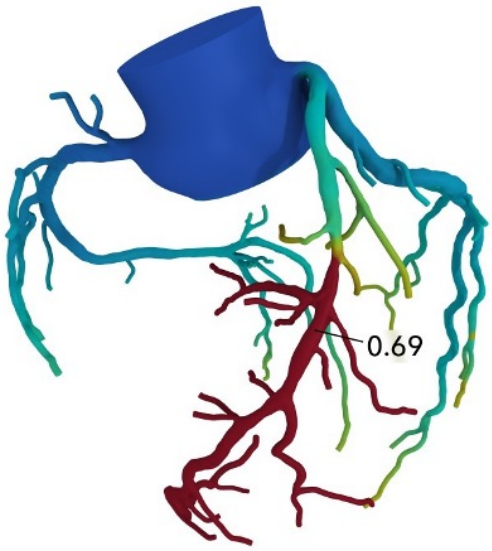


Figure 1.4: Example of $\mathrm{FFR_{CT}}$ computation from Taylor et al. [137]. This method uses high fidelity 3D models of patient anatomy and flow.

simulation was used to compute pressure in the coronary tree and FFR was calculated along the centerline of the target vessel, shown in Figure 1.4. This method has been further developed and validated using several clinical studies and demonstrated greater diagnostic accuracy than visual assessment of CTA [86, 69, 87]. Others have validated a similar commercially available software tool to derive FFR from CTA images known as cFFR [23, 63, 136]. The methods for deriving the 3D coronary geometry and boundary conditions are similar to $\mathrm{FFR_{CT}}$. The main difference is that cFFR uses a simplified 1D model of flow, with an additional correction term for branches containing stenoses. The advantage of this approach is that computation can be performed on a desktop computer rather than a high-performance computing cluster.

An alternative functional metric to FFR is instantaneous wave-Free Ratio (iFR). While both metrics are a ratio of pressures on both sides of a stenosis, iFR measured during the "wave-free" component of diastole, when microvascular resistance remains mostly constant. The main advantage of iFR is that it does not require hyperemia to be induced [124]. iFR has demonstrated similar diagnostic accuracy as FFR, but its integration into clinical practice has been limited. Carson et al developed a method to compute both iFR and FFR using a CCTA-derived geometry and a 1D formulation of the Navier-Stokes equations [19]. Inlet and outlet boundary conditions were determined using population estimates of flow and patient-specific scaling of resistance using a modification of Murray's law for the coronaries[143]. Another study aimed to compute iFR in idealized coronary geometries using the 3D Navier-Stokes equations, a pulsatile inflow waveform, and 5-element Windkessel models at the coronary outlets [41]. Resistances were once again as-

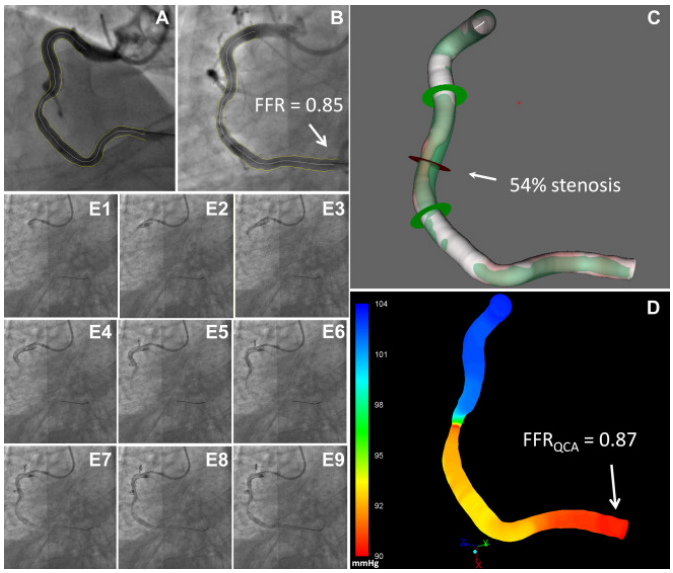


Figure 1.5: Schematic of QFR workflow. A and B show examples of angiograms and measured FFR. C is the reconstructed model of the vessel using 3D QCA. D shows the computed FFR in the reconstructed vessel. E is the series of angiographic frames used to extract mean flow. [140]

signed using the modified form of Murray's law and compliance was determined using population estimates. While more sophisticated than other computational approaches to derive functional metrics, this method requires many iterations of simulations to perform parameter tuning. The time and computational cost of such a method is not feasible for clinical implementation.

While methods to derive functional metrics from CCTA volumes have shown good agreement with clinically measured FFR or iFR, they have two main disadvantages. The first is that only a small proportion of patients with suspected CAD undergo CCTA before catheterization, which limits widespread integration of these methods into the standard diagnostic workflow. Another disadvantage is that CCTA does not provide any information about coronary flow. As discussed in section 1.1, FFR is only accurate when maximum hyperemia is achieved, and it is difficult to assess how closely allometric scaling matches patient-specific flow during pharmacologically-induced hyperemia. Due to these limitations, there have been a large number of studies that use X-ray angiography images to derive functional metrics instead.

Quantitative Flow Ratio (QFR) is a method developed by Tu et al. to derive FFR from angiography images without the need to induce hyperemia [140]. In this method, the 3D geometry of the vessel of interest is reconstructed using 3D QCA software and the TIMI-frame counting method is used to compute mean hyperemic flow (Figure 1.5). Allometric scaling is used to determine flow splits at each bifurcation, and a simplified quadratic relation between flow and pressure is used to compute FFR. This method was validated against clinically measured FFR in several studies [150, 142]. More recently, the QFR algorithm has expanded to extract more functional informa-

tion from a single angiography image [141]. Instead of using 3D QCA and TIMI-frame counting, the study used a combination of machine learning (U-Net [115]) and Murray's bifurcation law to determine the dimensions and flow rate through the vessel of interest and its side branches. FFR was computed in the resulting geometry using a 1D model of flow with a correction term for viscous losses across the stenosis. This method has demonstrated high accuracy and can deliver results in under a minute. The main limitations of QFR are the estimation of 3D geometry from a single angiography plane and the simplified relationship between flow and pressure. Both of these methods may not be sufficient to accurately capture the pressure drop across a stenosis in more complex geometries.

Another approach for angiographic FFR estimation was developed by Kornowski et al [65] called $\mathrm{FFR}_{\mathrm{angio}}$. Like QFR, this method uses a reduced order model – in this case, the Hagen-Poiseuille equation – to characterize the relationship between constant flow and pressure. The 3D geometry of the coronary tree is reconstructed from 3 specified angiography planes. The lengths, bifurcations, and radii of each vessel are computed from the reconstructed geometry and used to develop a 0D lumped parameter model of the vasculature. Several multi-center clinical trials have demonstrated the high accuracy of $\mathrm{FFR}_{\mathrm{angio}}$, even within the challenging regime near the clinical cutoff value [100, 35]. Like QFR, however, the simplifying assumptions about the flow and vessel geometry may limit accuracy in more complex anatomies.

1.3 Contributions and Outline of Thesis

In this work, we first describe an automated neural network approach for coronary vessel segmentation in Chapter 3: AngioNet: A Convolutional Neural Network for Vessel Segmentation in X-ray Angiography. We then integrate vessel segmentation into a stenosis localization and severity quantification pipeline, and validate the results against QCA in Chapter 4: Stenosis Characterization. Segmented angiography images are used as the input for several 3D reconstruction algorithms in Chapter 5: 3D Reconstruction of the Coronaries from 2D X-ray Angiography Images. We demonstrate the advantages and disadvantages of projection-based reconstruction algorithms, and propose a new data-driven approach for coronary reconstruction. In addition to extracting anatomical information from X-ray angiography images, we would like to exploit the implicit information about flow in an angiographic time series to derive patient-specific boundary conditions for downstream hemodynamic analysis. In Chapter 6: Relative Flow Extraction from Angiographic Time Series, we describe an algorithm for relative flow extraction and validate its accuracy using simulated data. Finally, we perform hemodynamic analyses in Chapter 7: Hemodynamic Studies. In this chapter, we are interested in understanding how geometric reconstruction error from our segmentation and reconstruction algorithms affects the accuracy of hemodynamic quantities. We also

demonstrate the entire workflow to derive computational FFR from X-ray angiography images in a single patient.

The contributions of this work are as follows:

1. Automatic vessel segmentation: AngioNet, a convolutional neural network that learns the best possible pre-processing filter to improve segmentation.
   (a) This work was published in Scientific Reports (2021). https://doi.org/10.1038/s41598-021-97355-8
   (b) This work was also featured in a publication about the American Heart Association Precision Medicine Platform https://doi.org/10.1161/CIRCOUTCOMES.121.007949
2. Stenosis Characterization: A fully automated end-to-end pipeline for stenosis characterization that quantifies stenosis severity and has been validated against QCA
   (a) A manuscript of this work is in preparation
3. 3D Reconstruction: A machine learning approach for 3D coronary reconstruction from uncalibrated X-ray angiography images
   (a) A manuscript of this work is in preparation.
4. Relative Flow Extraction: Algorithms to compute relative flow in the branches of a coronary tree given an angiographic time series
5. The work presented in this thesis is part of US patent 11,386,563 *Anatomical and Functional Assessment of Coronary Artery Disease Using Machine Learning*.
6. Open Source Software Contributions
   (a) AngioNet github repository: https://github.com/kritiyer/AngioNet
   (b) Synthetic vessel generator github repository: https://github.com/kritiyer/vessel_tree_generator
   (c) I/O pipeline to load 3D reconstructed geometries into CRIMSON [3] for CFD analysis

CHAPTER 2

Background

2.1 Coronary Anatomy and Physiology

The coronary arteries supply blood and oxygenation to the heart muscle, or myocardium. An interesting feature of the coronary circulation is that the contraction of the heart muscle alters the typical blood flow waveform present in the systemic circulation. During systole, when the heart contracts and ejects blood into the pulmonary and systemic circulation, the coronaries are compressed and blood flow is restricted. The majority of coronary blood flow occurs during diastole, when the heart muscle relaxes [93].

The large arteries embedded on the surface of the heart are known as epicardial vessels [93]. The epicardial vessels are divided into the right and left coronary trees, shown in Figure 2.1, each of which supplies blood to a different region of the heart. The three main coronary branches are the Right Coronary Artery (RCA), Left Anterior Descending (LAD), and Left Circumflex (LCx). Each of these main branches bifurcates into several smaller branches. In the right coronary tree, these sub-branches supply blood to the right atrium and ventricle. The distal branches of the right coronary tree perfuse the posterior of the heart and part of the left ventricle, although in some patients these branches can arise from the left coronary tree instead. The LAD supplies blood to the anterior of the heart, while the LCx mainly perfuses the left ventricle [59].

Further down the coronary tree, the arteries eventually bifurcate into capillaries which make up the microvasculature. The microvasculature is embedded within the heart muscle and plays a large role in coronary autoregulation. This is the ability of the coronaries to increase blood flow in response to increased metabolic demand, for example during exercise. The capacity for increasing blood flow to the heart is known as coronary flow reserve, and can be up to 4x greater than blood flow at rest [93]. These conditions of maximum coronary flow are known as hyperemia. In patients with epicardial disease, the microvasculature can compensate for a decrease in flow due to stenosis through autoregulation mechanisms. When patients are affected by microvascular disease, however, coronary flow reserve and the ability to maintain adequate oxygenation of the myocardium is

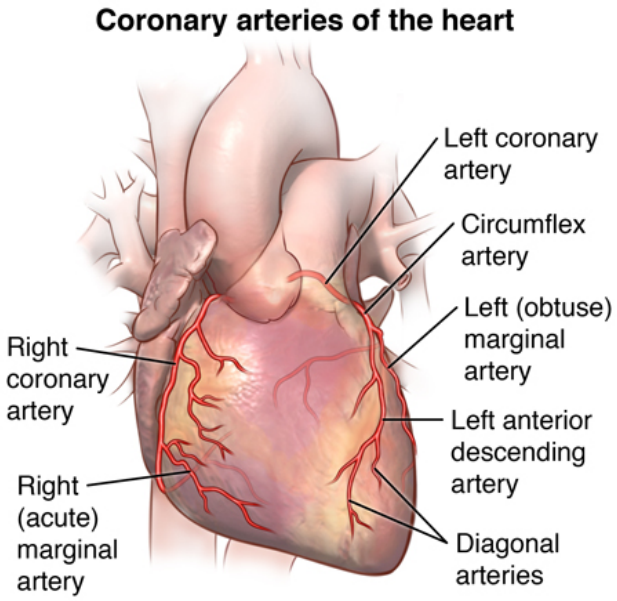


Figure 2.1: Schematic of the major coronary branches [83]

limited. The significance of an epicardial stenosis is therefore dependent on the level of coronary flow, and inducing reactive hyperemia is an effective way to determine the functional significance of a stenosis [46].

2.2 Background on X-ray Angiography

X-ray angiography is a form of fluoroscopy, an imaging modality which produces a sequence of images in real time [14]. It is most commonly used in medical procedures in which immediate feedback is required, such as the placement of a catheter or image-guided coronary angioplasty. Since the 1960s, improvements to angiography systems have improved image quality while decreasing the radiation dose to the patient [14].

A diagram of an X-ray angiography setup is shown in Figure 2.2. The system is composed of the X-ray table, X-ray detector, and X-ray source. The patient lies supine on the X-ray table during image acquisition. The X-ray detector, or imaging plane, can rotate in any direction about a fixed point known as the isocenter. The X-ray source and imaging plane are attached to opposite ends of a C-arm, which ensures that they rotate together and remain perpendicular to each other.

The position of the detector can be expressed in terms of two angles. The primary angle describes the angulation, or the degree to which the detector has rotated about the axis running longitudinally through the patient (Figure 2.2B). This angle is expressed as a value between 0 and

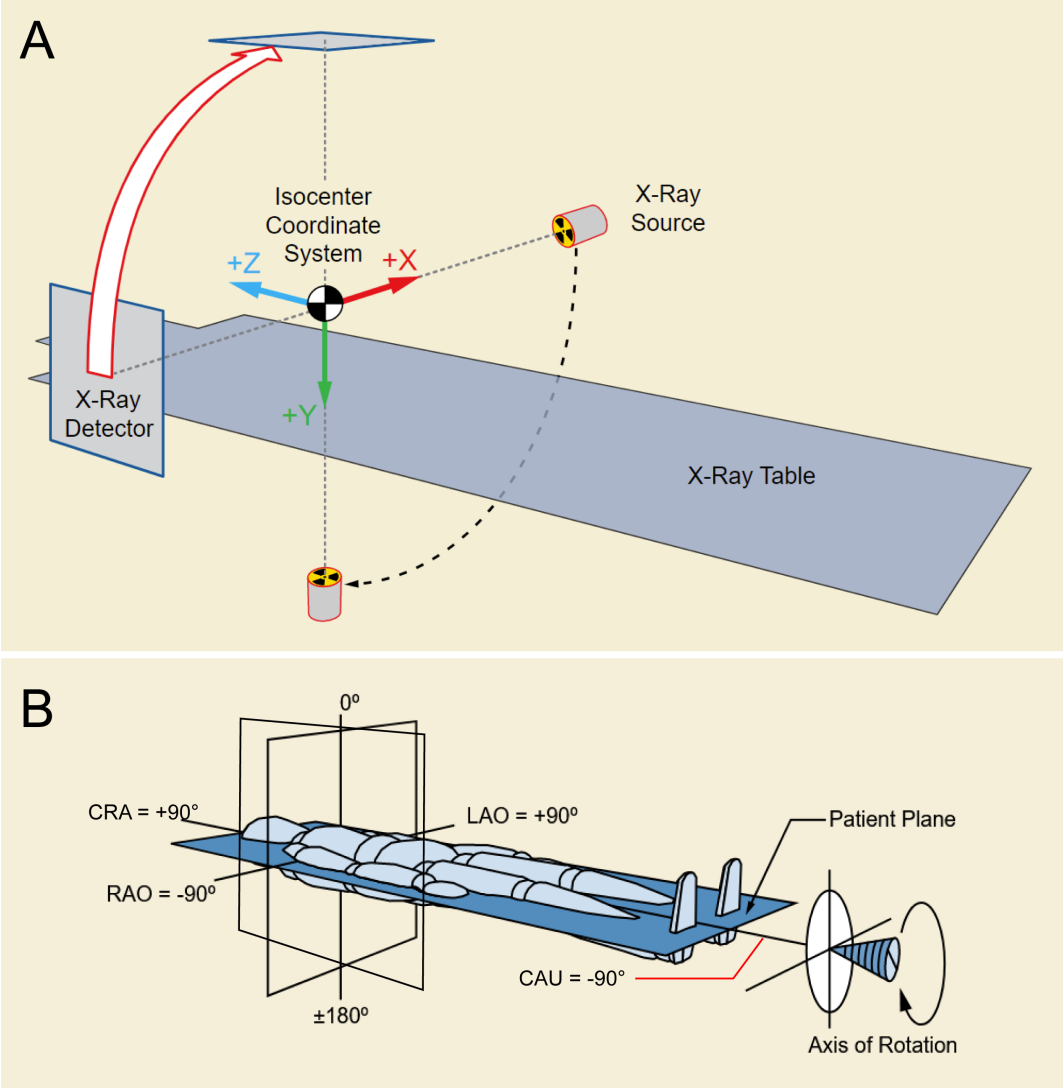


Figure 2.2: Schematic of X-ray Angiography machine. A) The X-ray angiography system is composed of a patient table, X-ray detector, and X-ray source. The X-ray detector and source rotate about a fixed point known as the isocenter. B) The position of the X-ray detector is given by two angles, the angulation and elevation. The angulation is expressed in terms of RAO/LAO and the elevation is expressed as CRA/CAU. Figures adapted from [92].

90, where LAO indicates rotation to the patient's left and RAO indicates rotation to the patient's right. The secondary angle defines the elevation of the detector, where a positive value indicates a rotation towards the patient's head (CRA) and a negative rotation indicates rotation towards the patient's feet (CAU). Together with the distances between the X-ray source, detector, and isocenter, these angles can be used to determine the exact position of the imaging plane during image acquisition.

During an X-ray coronary angiography procedure, a catheter is inserted into the patient and

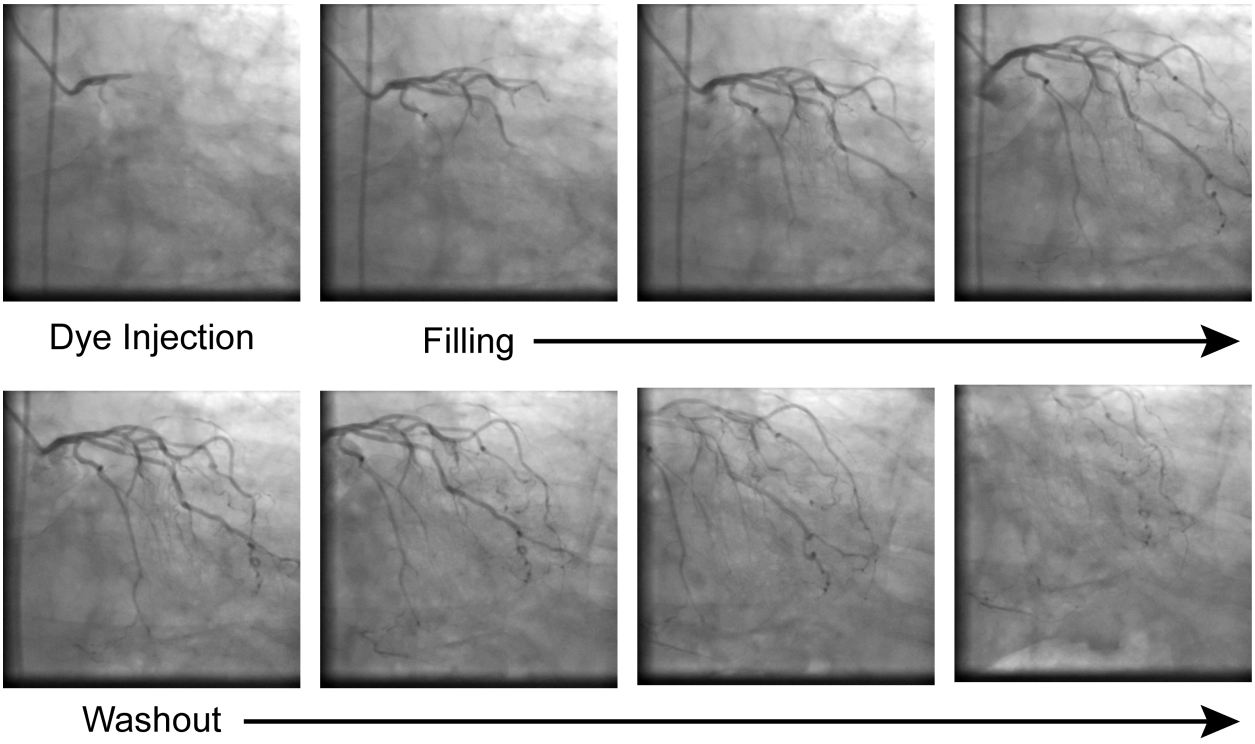


Figure 2.3: Selected frames from an angiographic time series. Radio-opaque dye is first injected into the arteries via a catheter. As the heart pumps, the coronaries fill with dye until the entire coronary tree can be visualized. Eventually, the dye is washed out of the vessels of interest.

positioned at the root of the aorta. Radio-opaque dye is injected into the patient's coronary arteries and a series of X-ray images are acquired as the dye is flushed through the vessels. The position of the X-ray table can be controlled via a joystick mechanism. This allows the cardiologist to move the patient as dye flows through their vessels in order to visualize the entire length of the vessels of interest.

Generally, 2-4 image sequences are acquired from different angles to visualize each of the major branches in the right and left coronary trees. The frame rate of image acquisition is typically 7.5, 15, or 30 frames per second. The typical size of an angiography image is 512x512 pixels, where each pixel corresponds to 0.2-0.3mm. Figure 2.3 contains an example of several frames from an angiographic time series. In this figure, we can see the progression of injected dye as it fills the vessels over several cardiac cycles before eventually being flushed out. Since the coronaries are embedded in the surface of the heart, they undergo deformation and rotation along with the squeezing and twisting motions of the heart as it beats. This motion and deformation is also captured by the series of angiography images.

2.3 Background on Machine Learning and Neural Networks

Machine learning involves using large amounts of data to optimize an algorithm which can predict an outcome or approximate a function. Most machine learning algorithms fall into one of three broad categories: supervised learning, unsupervised learning, and reinforcement learning. For example, logistic regression, support vector machines (SVM), and Naive Bayes are all simple supervised machine learning algorithms for classification. Linear regression and its variations are examples of supervised regression algorithms. K-means clustering is an unsupervised algorithm which can detect patterns in data and sort them into groups. Principle component analysis (PCA) is a linear algebra technique for dimensionality reduction; its extension is an unsupervised machine learning algorithm known as kernel PCA.

Neural networks are a subset of machine learning algorithms that can accomplish more complex tasks than simple machine learning models. They have been proven to be universal function approximators in any finite dimensional space assuming that the network has enough parameters [53]. The trade-off between using a neural network versus a simpler machine learning model is that neural networks require much more data and computational resources to train. Another difference is that neural networks automatically extract salient features from the input data to perform the given task, unlike simpler machine learning models which require the user to identify or manipulate input parameters. The disadvantage of automated feature extraction is that it reduces interpretability of the resulting learned algorithm — the notorious black box. This is in contrast to simpler models, where the final algorithm can be analyzed to identify the most important input features.

There are several types of neural networks, each suited for different tasks. The simplest and classic neural network architecture is a multi-layer perceptron, or MLP [116]. This is a feed-forward network which is composed of an input layer, output layer, and one or more intermediate layers known as hidden layers 2.4. Each layer is composed of neurons, where weighted values from each neuron are linearly mapped to the next. Non-linear activation functions such as rectified linear units (ReLU), tanh, and sigmoid can be applied between layers to increase the complexity of the functions the neural network is able to approximate.

While MLPs can perform classification, regression, and other tasks, convolutional neural networks (CNNs) are more well-suited for image processing. The advantage of convolutions is that they can incorporate spatial relationships between features and extract features at many different scales. Features can be related to morphology such as edges or right vs left leaning curves, or more abstract such as eyes in a face. Other features are not necessarily interpretable to the human eye, but are mathematically useful for the network to approximate the desired function. Convolutions are also more computationally efficient than the linear or dense layers in a traditional MLP.

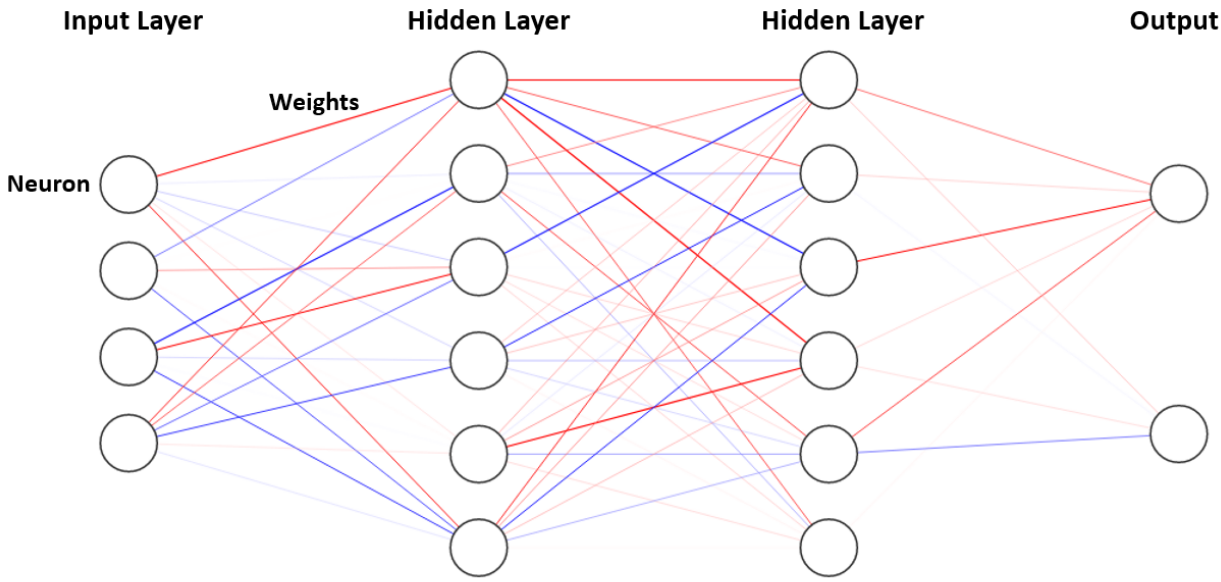


Figure 2.4: Schematic of Multi-layer perceptron

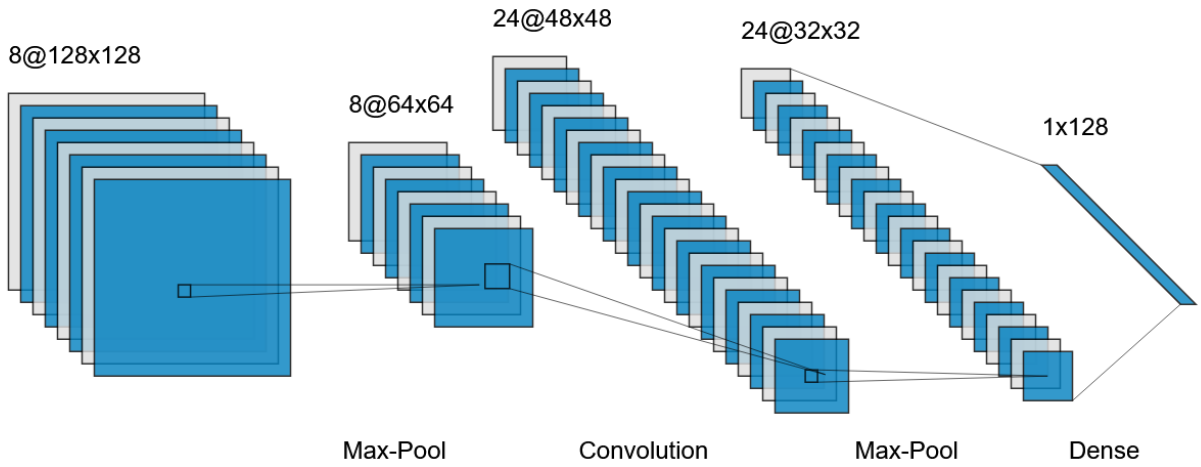


Figure 2.5: Schematic of convolutional neural network

More recently, there have been developments to make neural networks more interpretable. Attention units [144] are a state-of-the-art type mechanism whose purpose is to learn which input features are most important and the strength of the relationship between different features. The idea is that an attention unit can ”focus” on the salient information from the input and ignore irrelevant features. The strength of the network's attention provides a way to interpret which features

are most important to solve the given problem. The attention mechanism allows neural networks to perform even more sophisticated tasks than MLPs or convolutional neural networks; for example, combining natural language processing and image feature extraction as part of an image-captioning task [75]. Figure 2.6 demonstrates which parts of the image the network attends to in order to assign each word of the caption.



Figure 2.6: Visualization of attention weights for image captioning. The words in the caption are color-coded to correspond to the region of the image that the network associates with that word. Image reproduced from [75].

Attention mechanisms can also improve accuracy on more traditional tasks. For example, Attention U-Net [96] incorporated attention gates into the traditional U-net architecture[114] in order to improve pancreas segmentation from computed tomography (CT) images. The advantage of the attention mechanism was that it allowed the network to identify the important features at many scales and ignore parts of the CT that were not relevant, leading to increased Dice score and recall.

While attention units can be integrated into CNNs or other network architectures, they can also make up a purely attention-based network architecture known as a transformer. Transformers have shown an incredible capacity to learn complex tasks ranging from image classification, object detection, and multi-language translation [31, 18, 30]. Transformers require much larger datasets (on the order of 5-10k samples compared to a few hundred samples) and more computational

resources to train from scratch compared to other networks. They are also extremely sensitive to the learning rate at different stages of training, and require sophisticated learning rate schedules such as cosine annealing to be trained properly. Once trained, however, transformers are much more efficient at generalizing to new data or related tasks than other networks. For example, the Vision Transformer (ViT)[31] was trained to classify objects using the ImageNet benchmark dataset (21k images) and then fine-tuned on specialized datasets such as medical images and satellite images (1k images). ViT outperformed ResNet[50], a state-of-the-art convolutional network, on all datasets.

2.4 Projective Geometry and Stereovision

Projective geometry and stereovision are tools used to map the relationship between an object and its lower dimensional mappings, or projections. Given at least two images of a scene taken at different angles, we can use stereovision to compute the relative depth of different objects in the scene. Stereovision can also be used to reconstruct the position and shape of a 3D object given corresponding features on its projection images. Figure 2.7 shows a schematic of how a 3D point can be triangulated from two projection images taken at different angles. The position of the camera, or origin, for the left and right views are given by O_L and O_R. The ray of light $\overrightarrow{L}$ passing through O_L and the reference point of interest, x_L, must also pass through its corresponding 3D point X. The question is therefore where the 3D point lies along $\overrightarrow{L}$.

In order to triangulate the position of X, we must identify the point x_R that corresponds to the same feature on the image as x_L. Since we know the 3D point must be on $\overrightarrow{L}$, we know x_R must lie on the projection of $\overrightarrow{L}$ onto the right view. This projection is known as the epipolar line, and allows us to constrain the problem of finding the corresponding point x_R to 1D instead of searching the entire 2D image plane. The corresponding point x_R can be identified using image features. For example, if x_L corresponded to an eye in the left view, we look for the point on the epipolar line that intersects with the eye on the right view. Once x_L and x_R have been identified, we can triangulate the position of X as the intersection of $\overrightarrow{L}$ and $\overrightarrow{R}$, where $\overrightarrow{R}$ is the ray of light going through O_R and x_R.

3D reconstruction of an object from a limited set of projections is an ill-posed problem, as there are many possible shapes of the 3D object which could produce the same projections. For example, in Figure 2.7, we see many possible points X_i that could be the true 3D point corresponding to x_L depending on the choice of x_R. Identifying corresponding points based on image features is one way to limit the shape of the 3D object, but this can be a difficult task in practice. Identifying corresponding features on X-ray angiography images is a particularly challenging task, as these projection images are essentially shadows without features such as color, light reflection, or other details that enable feature matching in natural camera images. Another challenge is the number

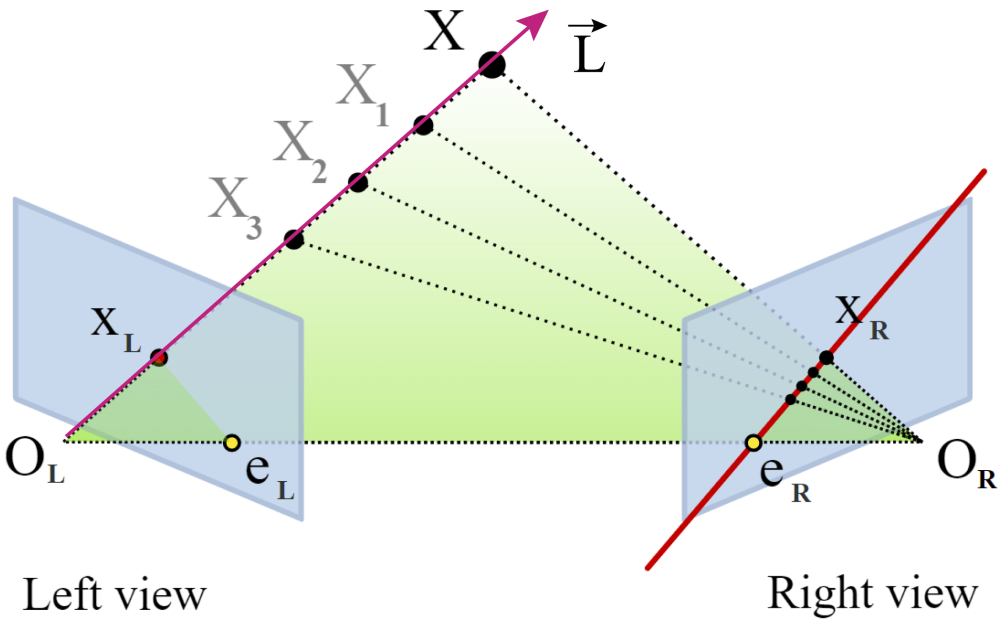
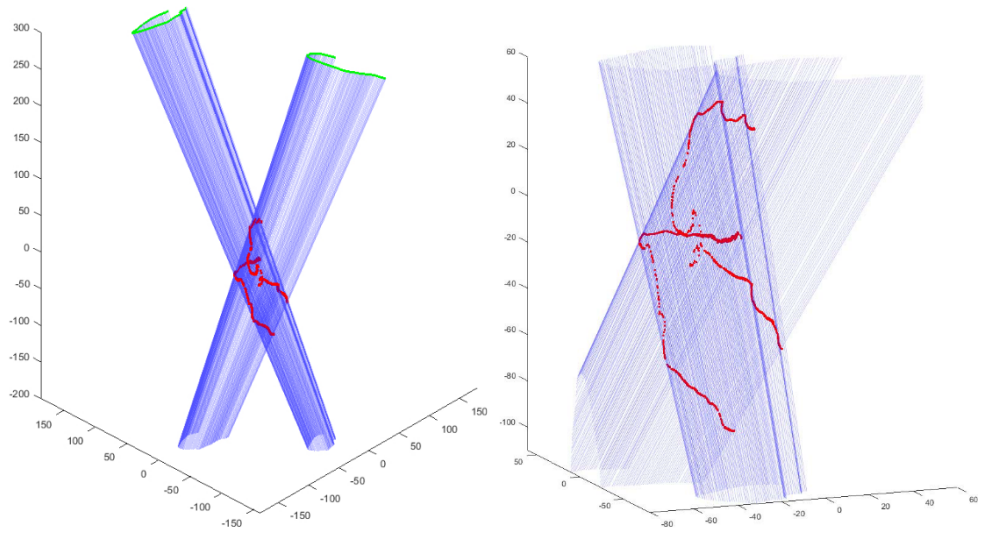


Figure 2.7: Diagram of 3D point triangulation using stereovision. The corresponding point x_R for a given reference point x_L must lie on the epipolar line e_R. Adapted from [95].

of projection images available. According to [134], we would ideally need 4 or more projections to reconstruct a unique 3D object, but in clinical practice only 2-3 angiograms are acquired of the right and left coronary trees.

To combat this, we use prior knowledge about the shape of coronary vessels to constrain the possible 3D shape. Instead of identifying corresponding points for the entire vessel, we focus on triangulating points on the 3D centerline. Once the 3D centerline has been determined, we can triangulate the position of vessel surface points on the lumen to determine the radius at several points along the centerline. Since we know vessels are tubular structures, we can re-create the vessel surface by lofting a NURBS surface along the 3D centerline using the contours from triangulated radius points. In this way, we eliminate the possibility of reconstructing arbitrary 3D shapes that match the projection images but are not tubular structures.

The issue of automatically identifying unique pairs of corresponding points on vessel centerlines remains a difficult task. As demonstrated by Figure 2.8, there can be many possible 3D centerline points found by triangulating rays that correspond to the 2D centerlines. As described in subsubsection 1.2.1.3, researchers have proposed ordering constraints, segment-wise correspondence, and other methods to uniquely identify the 3D centerline. Many of these methods still require users to identify corresponding branches or their endpoints in multi-branch vessel structures.



(a) Intersecting projection surfaces (b) Zoomed-in intersecting points

Figure 2.8: The intersections between two vessel surfaces from 2D images can produce several possible 3D centerlines. Reproduced from [10]

2.5 Hemodynamic Modeling in the Coronary Arteries

Computational modeling has been used to understand coronary hemodynamics for over 20 years. While early studies used idealized 3D models of coronary geometry, advances in medical image acquisition and image processing techniques have enabled the current paradigm of image-based patient-specific anatomical models [155].

In order to understand how hemodynamic quantities such as flow and pressure behave in a particular geometry, we use Computational Fluid Dynamics (CFD) analysis. In this approach, we are interested in solving the incompressible Navier-Stokes conservation of momentum and continuity equations, which characterize the behavior of a fluid:

$$\rho\frac{\partial v}{\partial t} + v \cdot \nabla v = -\nabla p + \mu\nabla^2 v + f \tag{2.1}$$

$$\nabla \cdot v = 0 \tag{2.2}$$

where v is the velocity vector in 3 dimensions, p is the pressure, ρ is the fluid density, μ is the dynamic viscosity, and f is the body force, which encompasses external forces such as gravity. Although the coronaries undergo large deformation and motion as the heart contracts, we typically exclude the effect of these external forces and gravity as they have a small effect on the behavior of the fluid [139, 108].

The Finite Element Method (FEM) is a numerical framework to solve partial differential equations such as the Navier-Stokes equations on a complex geometric domain. In this approach, the 3D geometry is discretized into smaller tetrahedral elements to form a mesh. The weak form, or integral form, of the Navier-Stokes equations are then solved locally on the nodes of the mesh. The size of the elements that form the mesh can have a large effect on the accuracy of the solution. For example, if mesh elements are too large, the simulation will not have enough resolution to resolve complex flow characteristics such as eddies. On the other hand, if the mesh elements are too small, the computational cost of running a CFD simulation can become prohibitive. It is therefore important to perform mesh-refinement studies, in which we ascertain the appropriate element size that achieves an accurate solution while minimizing computational cost.

Once the geometry has been discretized into a mesh, the next step is to specify the boundary conditions and underlying assumptions of the model. Many studies assume rigid walls in the coronaries and incorporate vessel compliance into lumped parameter models of the vasculature to reduce computational cost. Other researchers have used more complex computational frameworks such as the Arbitrary Lagrangian Eulerian method or immersed boundary method to model the interaction between a deformable vessel wall and the fluid. While non-rigid walls are more physiologically accurate, the small increase in accuracy when modeling the coronaries does not justify the much higher computational cost of such simulations [155].

The most common approach to model the behavior of the small arteries and microcirculation is to use reduced order models that are coupled to the outlets of the 3D domain. The complexity of the reduced order models can vary greatly. For example, Kim et al. [64] coupled a 0D heart model, which provided the inflow waveform, with lumped parameter models of the coronaries, which were sensitive to changes in left ventricular pressure. The heart model was also coupled with the aorta, and distal systemic resistance was modeled using 3-element Windkessel models. This sophisticated dynamic feedback loop enabled modeling of the unique flow waveform in the coronaries (see section 2.1) and enables understanding of complex mechanisms such as coronary autoregulation [4] or oscillatory shear index [79, 139].

When trying to understand the functional severity of CAD, however, the quantity we are most interested is the average pressure drop across the stenosis. Studies have shown that steady flow simulations are sufficiently accurate to estimate changes in pressure due to the presence of stenosis [81], and this assumption greatly simplifies the boundary conditions and reduces computational cost. Instead of a dynamic 0D heart model, we can simply impose a steady flow rate at the inlet of the coronary tree. The microvasculature and systemic resistance can be represented as simple resistors instead of a coronary lumped parameter model or Windkessel model, since compliance is not relevant in steady flow problems.

CHAPTER 3

AngioNet: A Convolutional Neural Network for Vessel Segmentation in X-ray Angiography

This chapter is adapted from a peer-reviewed publication of this work: https://doi.org/10.1038/s41598-021-97355-8 [55]

3.1 Introduction

In this work, we are interested in developing a method for automated coronary segmentation in X-ray angiography images. Fully-automated angiographic segmentation is a particularly challenging task due to the poor signal-to-noise ratio and overlapping structures such as the catheter and the patient's spine and rib cage [73]. Several filter-based or region-growing approaches [73, 52, 16, 36, 126, 149, 98, 85, 67, 22] have been developed for angiographic segmentation. The principal limitation of these methods is that they cannot separate overlapping objects such as catheters and bony structures from the vessels, requiring manual correction which can be time-consuming and subjective. To address these limitations, some have turned to CNNs for angiographic segmentation.

CNNs have been used for segmentation in numerous applications [109, 9, 104, 157, 99] . More recently, CNNs have been developed for segmentation of anatomical structures in medical images, including X-ray angiography [153, 33, 91, 122]. While these methods have demonstrated success in coronary segmentation, they still have several drawbacks as described in section 1.2. Namely, many of these methods are not robust to ignoring the catheter and other imaging artifacts, particularly in challenging cases such as low contrast images. Furthermore, other networks were designed for specific use-cases such as single-vessel segmentation [153] or place constraints during image acquisition such as minimal patient motion [33].

To address these shortcomings, we have developed a new CNN for angiographic segmentation: AngioNet, which combines an Angiographic Processing Network (APN) with a semantic segmentation network. The APN was trained to address several of the challenges specific to angiographic segmentation, including low contrast images and overlapping bony structures. AngioNet

uses Deeplabv3+ [21] as its backbone semantic segmentation network instead of U-Net[114] or other fully convolutional networks (FCNs), which are more commonly used for medical segmentation. In this paper, we explored the specific benefits of the APN – and the importance of using Deeplabv3+ as the backbone – by comparing segmentation accuracy in Deeplabv3+, U-Net, and Attention U-Net[96], trained both with and without the APN. Deeplabv3+ was chosen as the main backbone network since its deep architecture has greater expressive power [106, 76, 32] compared to the FCNs typically used for medical segmentation. Its ability to approximate more complex functions is what enables AngioNet perform well in challenging cases. We chose to investigate the effect of the APN on U-Net due to its extensive use in medical image segmentation. Attention-based networks have been shown to improve segmentation performance compared to pure CNNs by suppressing irrelevant features and learning feature interdependence [96, 37]; thus, Attention U-Net was chosen as an appropriate network backbone for comparison. Lastly, we performed clinical validation of segmentation accuracy by comparing AngioNet-derived vessel diameter against QCA-derived diameter.

The main contribution of this work is the introduction of an APN to various segmentation backbone networks, creating an end-to-end pipeline encompassing image pre-processing and segmentation for angiographic images of coronary arteries. The APN was shown to improve segmentation accuracy in all three tested backbone networks. Furthermore, networks containing the APN were better able to preserve the topology of the coronary tree compared to the backbone networks alone. Another contribution of our work is the comparison against a clinically validated standard for measuring vessel diameter, QCA. This validation study provides more insight into AngioNet's ability to accurately detect the vessel boundaries than traditional segmentation evaluation metrics. Our statistical analyses demonstrate that AngioNet and QCA are interchangeable methods of determining vessel diameter.

3.2 Methods

3.2.1 Datasets

Figure 3.1 summarizes the two patient datasets used in this work for neural network training and evaluation of performance. All data were collected in compliance with ethical guidelines.

3.2.1.1 UM Dataset

This dataset was composed of 462 de-identified angiograms acquired using a Siemens Artis Q Angiography system at the University of Michigan (UM) Hospital. The study protocol to access this data (HUM00084689) was reviewed by the Institutional Review Boards of the University

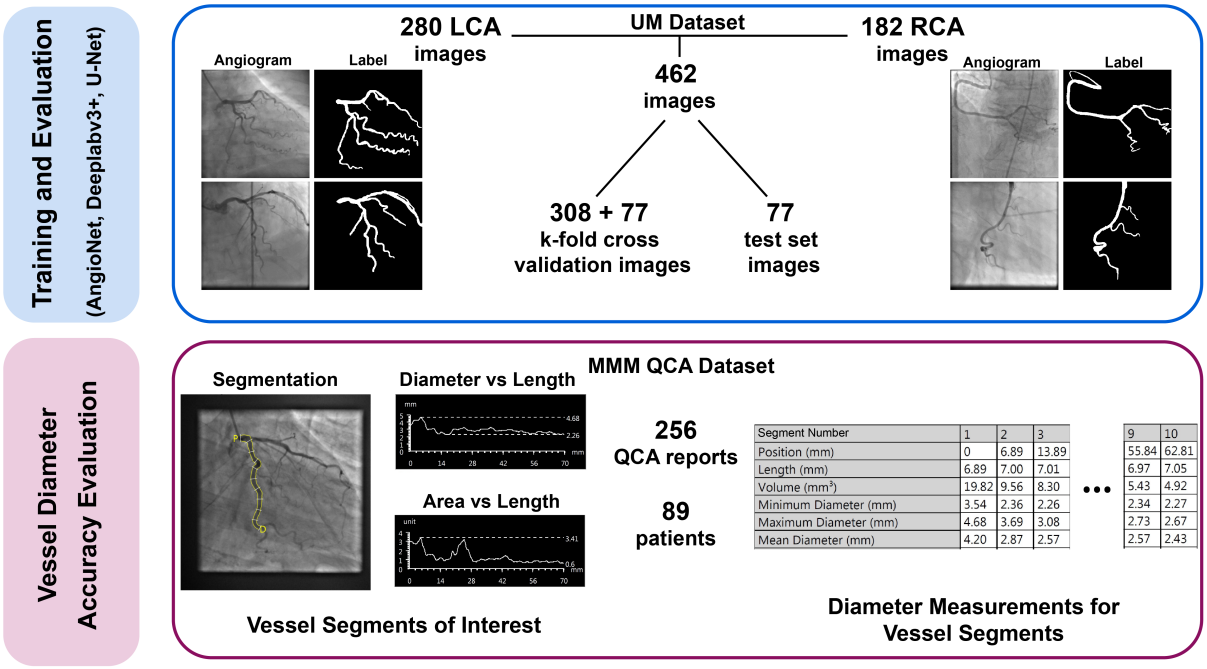


Figure 3.1: Diagram of datasets for CNN training and evaluation. AngioNet's performance was compared against state-of-the-art neural networks, all trained on the UM Dataset. The MMM QCA dataset was used to quantify segmentation diameter accuracy by comparing AngioNet's results against the diameters reported in QCA. Acronyms: left coronary artery (LCA); right coronary artery (RCA); Madras Medical Mission (MMM).

of Michigan Medical School (IRBMED). Since the data was collected retrospectively, IRBMED approved use without requiring informed consent. This dataset was composed of patients who were referred for a diagnostic coronary angiography procedure at the UM Hospital in 2017. Patients with pacemakers, implantable defibrillators, or prior bypass grafts were excluded, as these prior procedures introduce artifacts and additional vascular conduits. Furthermore, patients with diffuse stenosis were excluded as this is less common in arteries suitable for revascularization. In our sample of 161 patients, 14 had severe stenosis ($\geq$80% diameter reduction) and the remaining had mild to moderate stenosis. The dataset was composed of 280 images of the left coronary artery (LCA) and 182 images of the right coronary artery (RCA).

The data were equally split by patient into a 5-fold cross-validation set and test set to avoid having images from the same patient in both the training and test sets. Labels for all images were manually annotated to include vessels with a diameter greater than 1.5 mm (4 pixels) at their origin. Labels were created by selecting the vessels using Adobe Photoshop's Magic Wand tool followed by manual refinement of the vessel boundaries, and were reviewed by a board-certified cardiologist. The 5-fold cross-validation portion of the dataset was used for neural network training

and hyperparameter optimization, whereas the test set was used to evaluate segmentation accuracy. There is a great number of artifacts in XRA images, including borders from X-ray filters, rotation of the image frame, varying levels of contrast, and magnification during image acquisition. Data augmentation of the UM dataset was employed to account for this variability. Horizontal and vertical flips of the images were included to make the network segmentation invariant to image orientation. Random zoom up to 20%, rotation up to 10%, and shear up to 5% were used to account for variation in magnification and imaging angles. When zooming out, shearing, or rotating the image, a constant black fill was used to mimic images acquired using physical X-ray filters. The combination of the above data augmentations created a training dataset of over half a million images to improve network generalizability. Data augmentation was not applied to the test set. The augmented UM dataset was used for neural network training, and the test set was used to compare segmentation accuracy.

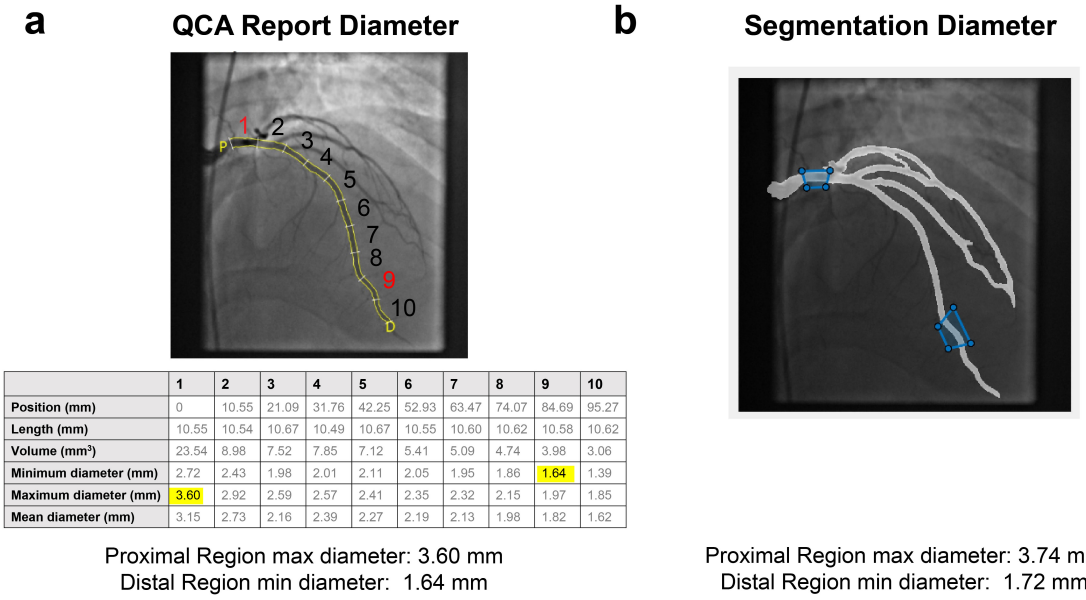
3.2.1.2 MMM QCA Dataset

The percent change in vessel diameter at the region of stenosis is a key determinant of whether a patient requires an intervention or not; therefore, the accuracy of AngioNet's segmented vessel diameters was assessed in addition to its overall segmentation accuracy. Although the main result of a QCA report is the overall percent change in vessel diameter, these reports also contain measurements of maximum, minimum, and mean diameter in 10 equal segments of the vessel of interest (Figure 3.1). These diameter measurements in the MMM QCA Dataset were used to evaluate the discrepancies between QCA and AngioNet.

The Madras Medical Mission (MMM) QCA dataset contained independently generated three-vessel QCA reports of 89 patients, encompassing 223 vessels in both the LCA and RCA. All patients presented with mild to moderate stenosis. The data were acquired from the Indian Cardiovascular Core Laboratory (ICRF) at the MMM, which serves as a core laboratory with experience in clinical trials and other studies and has expertise in QCA. The data provided by the MMM ICRF Cardiovascular Core Laboratory includes independent and detailed analysis of quantitative angiographic parameters (minimum lesion diameter, percent diameter stenosis, etc.) as per American College of Cardiology/American Heart Association standards, through established QCA software (CAAS-5.10.2, Pie Medical Corp). The study protocol for this data (Computer-Assisted Diagnosis of Coronary Angiography) was approved by the Institutional Ethics Committee of the Madras Medical Mission. This data was obtained using an unfunded Materials Transfer Agreement between UM and MMM. Since the data is completely anonymized and cannot be re-identified, it does not qualify as human subjects research according to OHRP guidelines.

To validate the accuracy of AngioNet's segmented vessel diameters, a MATLAB script was employed for user specification of the same vessel regions as those in the QCA report. Two regions

from the QCA report were sampled in each angiogram. The first was the most proximal region, containing the maximum vessel diameter, and the second was the region of stenosis (given in the QCA report), if present. If no stenosis was reported, the region containing minimum diameter was selected. A skeletonization algorithm [138] was used to identify the centerline and radius map of the selected vessel region. Using the output of the skeletonization algorithm, the script reported the maximum and minimum diameters at the selected regions and compared them against the diameters in the QCA report. Maximum and minimum vessel diameter were chosen rather than the diameters on either side of stenosis since the purpose of using the QCA reports was to systematically assess overall vessel diameter accuracy, not the percent diameter reduction. A diagram of the comparison between QCA and AngioNet diameters is shown in Figure 3.2.



	1	2	3	4	5	6	7	8	9	10
Position (mm)	0	10.55	21.09	31.76	42.25	52.93	63.47	74.07	84.69	95.27
Length (mm)	10.55	10.54	10.67	10.49	10.67	10.55	10.60	10.62	10.58	10.62
Volume (mm^3)	23.54	8.98	7.52	7.85	7.12	5.41	5.09	4.74	3.98	3.06
Minimum diameter (mm)	2.72	2.43	1.98	2.01	2.11	2.05	1.95	1.86	1.64	1.39
Maximum diameter (mm)	3.60	2.92	2.59	2.57	2.41	2.35	2.32	2.15	1.97	1.85
Mean diameter (mm)	3.15	2.73	2.16	2.39	2.27	2.19	2.13	1.98	1.82	1.62

Figure 3.2: a) Annotated QCA report, along with the corresponding diameters in the report table. Highlighted values correspond to maximum (proximal) and minimum (distal) diameters (segments 1 and 9, respectively). b) Schematic of how a MATLAB script was used to delineate regions in the neural network segmentation corresponding to the regions measured in the QCA report, along with the computed proximal and distal diameters.

3.2.2 CNN Design and Training

3.2.2.1 Design

AngioNet was created by combining Deeplab v3+ and an Angiographic Processing convolutional neural network (APN). A diagram of the network architecture is given in Figure 3.3. Each component of AngioNet, the APN and Deeplabv3+, was trained separately before fine-tuning the entire

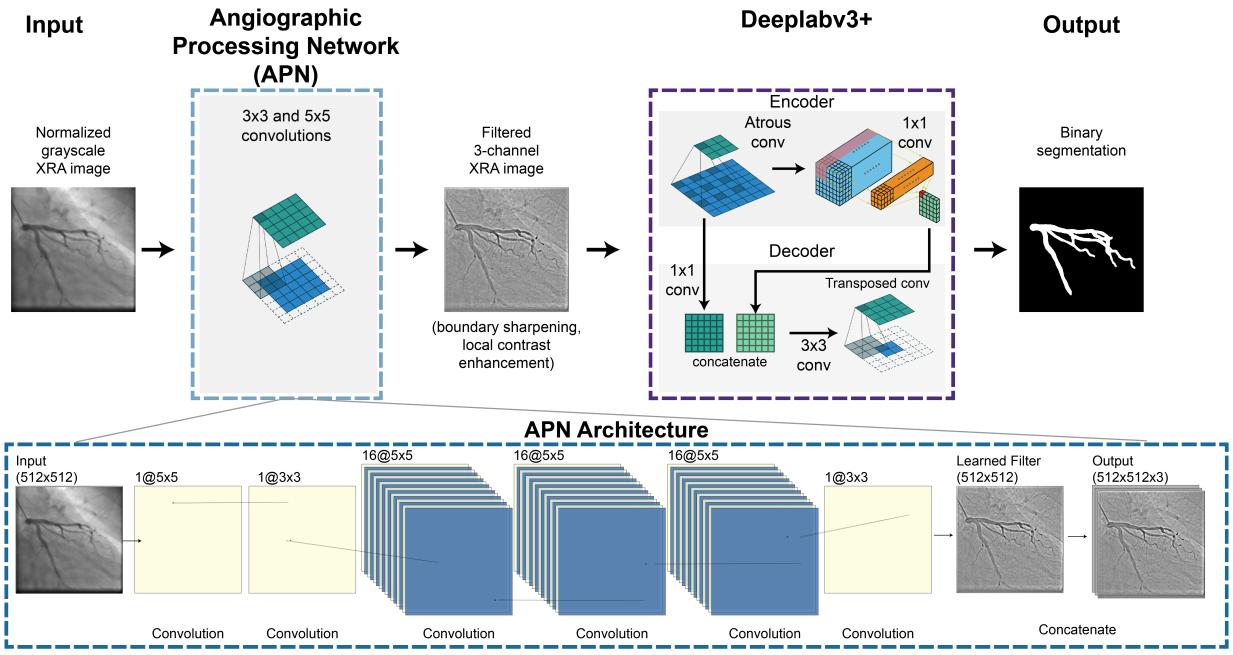


Figure 3.3: AngioNet Architecture Diagram. AngioNet is composed of an Angiographic Processing Network (APN) in tandem with Deeplabv3+. The APN is designed to improve local contrast and vessel boundary sharpness. The output of the APN, a single channel filtered image, is copied and concatenated to form a 3-channel image which is input into the backbone network.

network.

The purpose of the APN was to address some of the challenges specific to angiographic segmentation, namely poor contrast and the lack of clear vessel boundaries. The intuition behind the APN was that the best possible pre-processing filter and its parameters are unknown; we hypothesized that learning the best possible filter would lead to higher accuracy than manually sampling several filters. The APN was initially trained to mimic a combination of standard image processing filters instead of initializing with random weights, since it would later be fine-tuned with a pre-trained backbone network. A combination of unsharp mask filters was chosen as these can improve boundary sharpness and local contrast at the edges of the coronary vessels, making the segmentation task easier. Starting with an initialization of unsharp masking, the fine-tuning process was used to learn a new filter that was best suited for angiographic segmentation. The single-channel filtered image from the APN was then copied and concatenated to form a 3-channel image, as this was the expected input of Deeplabv3+, a network typically used to segment RGB camera images (see Concatenation and Output in the APN Architecture, Figure 3.3).

3.2.2.2 Training

The Deeplabv3+ CNN architecture was cloned from the official Tensorflow Deeplab GitHub repository, maintained by Liang-Chieh Chen and co-authors [21]. The network was initialized with pre-trained weights from their repository, as recommended by the authors for training on a new dataset. The input to this network were normalized angiographic coronary images, and the output was a binary segmentation. Training was conducted using four NVIDIA Tesla K80 GPUs on the American Heart Association Precision Medicine Platform (https://precision.heart.org/), hosted by Amazon Web Services. Hyperparameters such as batch size, learning rate, learning rate optimizer, and regularization were tuned. We observed that training with larger batch size led to better generalization to new data. A batch size of 16 was used as this was the largest batch size we could fit into memory using four GPUs. The Adam optimizer was chosen to adaptively adjust the learning rate, and L2 regularization was used to reduce the chance of over-fitting. The vessel pixels account for 15-19% of the total pixels in any given angiography image. Due to this class imbalance, it was important to encourage classification of vessel pixels over background using weighted cross-entropy loss.

The APN was initially trained to mimic the output of several unsharp mask filters applied in series (parameters: radius = 60, amount = 0.2 and radius = 2, amount = 1) as seen in Figure 3.4. This ensured the APN architecture was complex enough to learn the equivalent of multiple filters with sizes up to 121x121 using only 3x3 and 5x5 convolutions. The number of 3x3 versus 5x5 convolutions as well as the network width and depth were adjusted until the APN could reproduce the results of the serial unsharp mask filters. The inputs to the APN were the normalized single-channel images from the augmented UM Dataset, whereas the output was a filtered version of the image, copied to form a 3-channel image. The ground truth images for the APN were generated by applying several unsharp mask filters with various parameters to each normalized clinical image. The APN was composed of several 3x3 and 5x5 convolutional layers (Figure 3) and was trained to mimic the unsharp mask filters by minimizing the Mean Squared Error (MSE) loss between the prediction and ground truth images (MSE on the order of 1e-4). The APN design and training were carried out using TensorFlow 2.0, integrated with Keras.

Once the APN and Deeplabv3+ networks were individually trained, the two CNNs were combined to form AngioNet using the Keras functional Model API [154].The resulting network was trained with a low learning rate to fine-tune the combined model. Since neither the APN nor Deeplabv3+ were frozen during fine-tuning, both were able to adjust their weights to better complement each other: the APN learned a better filter than its unsharp mask initialization, and Deeplabv3+ learned the weights that could most accurately segment the vessel from the filtered image that was output by the APN. Network hyperparameters were once again tuned. The same process of pre-training, combining models, and fine-tuning was carried out with the APN and each

of the other network backbones, U-Net and Attention U-Net, to determine how much the backbone network contributes to segmentation performance. U-Net and Attention U-Net were not initialized with pre-trained weights as our dataset was adequately large to train these networks from random initialization. During all phases of training, batch normalization layers were frozen at their pre-trained values as we did not have a large enough dataset to retrain these layers. Furthermore, all hyperparameter optimization was performed on the 5-fold cross validation holdout set and accuracy was measured on the test set.

3.3 Results

3.3.1 Learned Filters using the Angiographic Processing Network

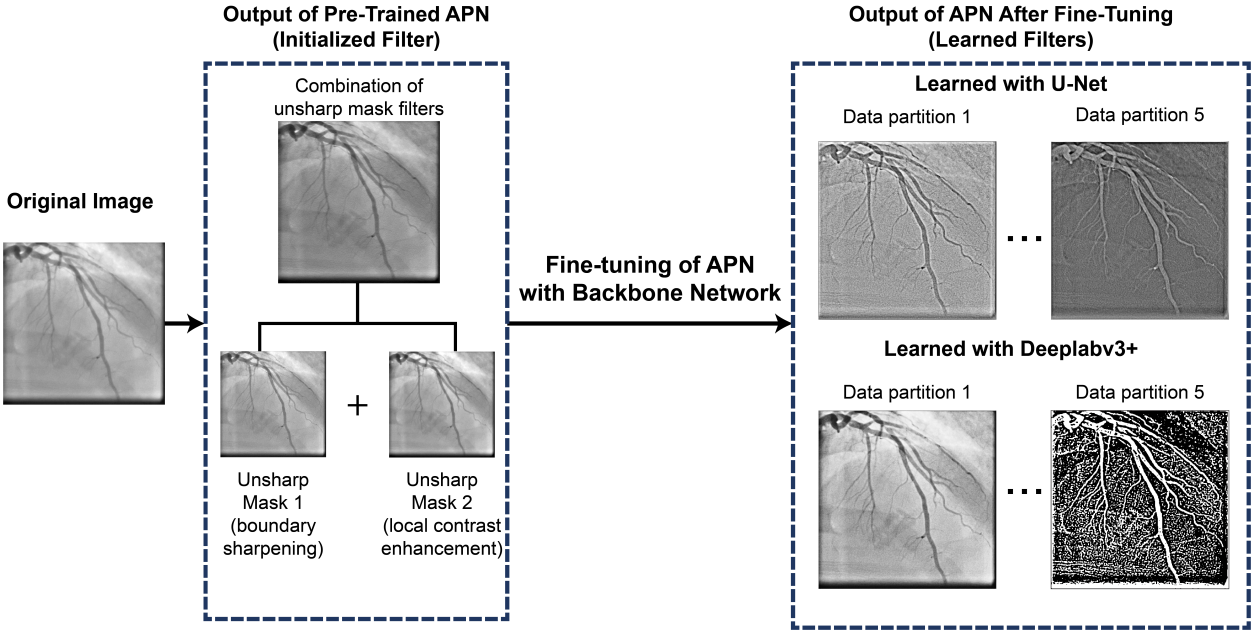


Figure 3.4: Examples of learned filters when the APN is fine-tuned together with Deeplabv3+ and U-Net. The APN was initialized with the combination of unsharp mask filters shown above, and learned new filters to aid segmentation. Each example image is the output of the APN after training with different data partitions.

Figure 3.4 contains examples of the filters that the APN learned when it was trained with Deeplabv3+ or U-Net, respectively. The images represent the output of the APN, and thus the input to the backbone network. Although the APN was initialized with the combination of unsharp mask filters shown in Figure 3.4, the network learns different filters that perform a combination of contrast-enhancement and boundary sharpening. The examples given are the results of training with different data partitions during k-fold cross-validation. The large variations in the learned

filters come from an inherent property of neural network training; since minimization of the neural network's loss function is a non-convex optimization problem [2], there are many combinations of network weights which will lead to similar values of the loss function, and consequently, similar overall accuracy. The effect of these varied learned filters on segmentation accuracy is described in the next section.

3.3.2 Segmentation Accuracy Metrics

Segmentation accuracy was measured using the Dice score, given by

$$Dice = \frac{2|Y \cap \hat{Y}|}{|Y| + |\hat{Y}|}. \tag{3.1}$$

Here, Y is the label image and $\hat{Y}$ is the neural network prediction, each of which is a binary image where vessel pixels have a value of 1 and background pixels have a value of 0. $|Y|$ denotes the number of vessel pixels (1s) in image Y, and $\cap$ represents a pixel-wise logical AND operation. Alternatively, the Dice score can be defined in terms of the true positives (TP), false positives (FP) and false negatives (FN) of the neural network prediction with respect to the label image, and is then given by

$$Dice = \frac{2TP}{2TP + FP + FN} \tag{3.2}$$

In addition to the standard Dice score, centerline Dice or clDice was also measured. Rather than measuring pixel-wise accuracy, clDice measures connectivity of tubular structures and can be used to determine how well the predicted image maintains the topology of the vessel tree in the label image [129]. clDice is measured by finding the centerlines of the prediction and label images, $cl_{\hat{Y}}$ and cl_Y. The proportion of $cl_{\hat{Y}}$ which lies in the label Y, $cl_{\hat{Y}}2Y$, and the proportion of cl_Y which lies in the prediction $\hat{Y}$, cl_Y2Y, are computed as analogs for precision and recall. clDice is then given by

$$clDice = 2 \times \frac{cl_{\hat{Y}}2Y \times cl_Y2\hat{Y}}{cl_{\hat{Y}}2Y + cl_Y2\hat{Y}} \tag{3.3}$$

We also report the Area under the Receiver-Operator Curve (AUC), which measures the ability of the network to separate classes, in this case, vessel and background pixels. An AUC of 0.5 indicates a model that is no better than random chance, whereas an AUC of 1 indicates a model that can perfectly discriminate between classes. Finally, we also report the pixel accuracy of the binary segmentation, defined as

$$PixelAccuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{3.4}$$

3.3.3 Comparison of AngioNet versus Current State-of-the-art Semantic Segmentation Neural Networks

The accuracy of AngioNet was validated using a 5-fold cross-validation study, in which the neural network was trained on 4 out of the 5 UM Dataset training partitions at a time, with the fifth partition reserved for validation and hyperparameter optimization (hold-out set). This process was repeated five times, holding out a different partition each time. The accuracy of the resulting five trained networks was measured on the sixth partition, the test set, which was never used for training. The mean k-fold accuracy and the accuracy when trained on all five training data partitions are summarized in Table 3.1. The network performs well on both LCA and RCA input images (Figure 3.5) and does not segment the catheter or other imaging artifacts despite uneven brightness, overlapping structures, and varying contrast.

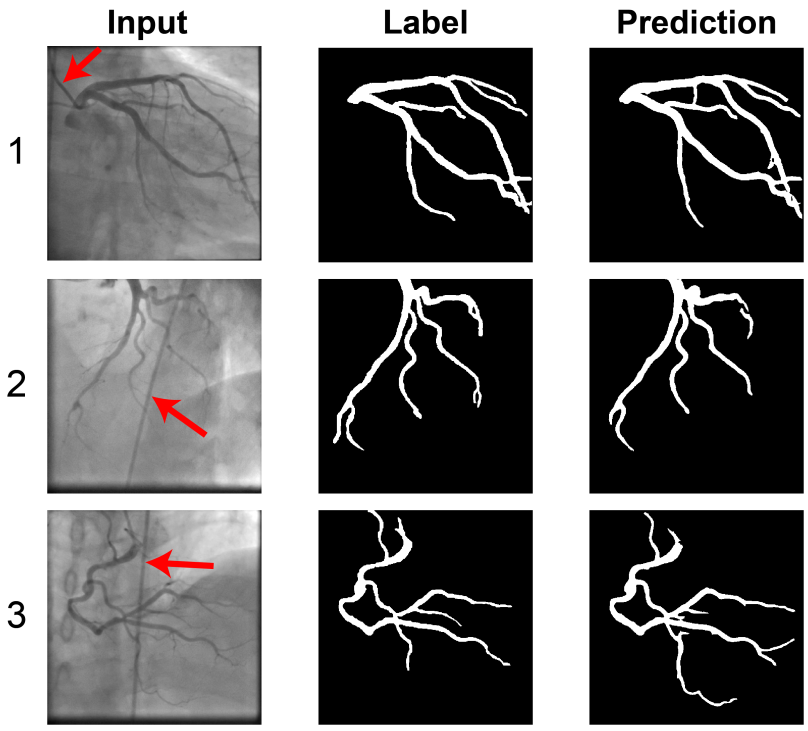


Figure 3.5: Examples of AngioNet segmentation on left coronary tree, taken at two different angles (1,2), and right coronary tree (3). AngioNet does not segment the catheter (red arrows), despite its similar diameter and pixel intensity as the vessels (2,3). It also ignores bony structures such as the spine in (3) and ribs in (1).

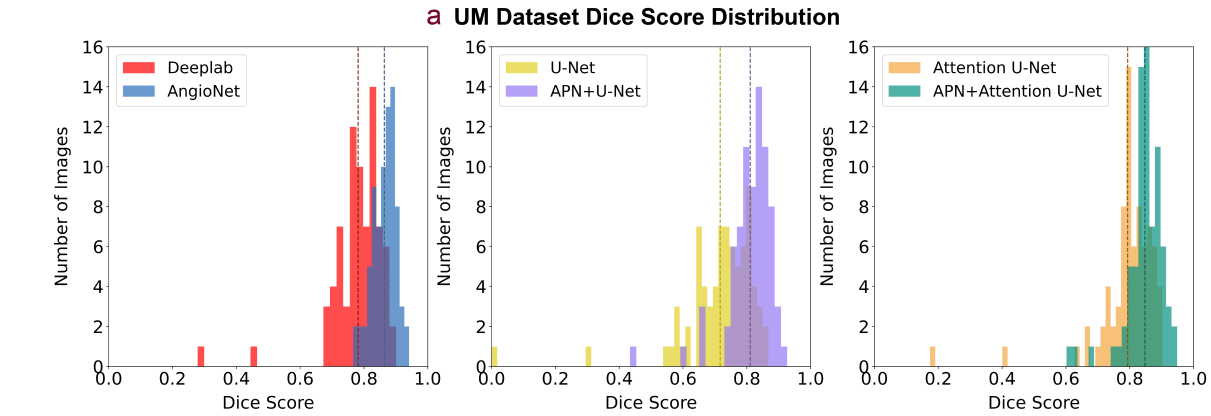
Table 3.1: Accuracy of AngioNet using K-Fold Cross Validation

	Dice Score	**Sensitivity**	**Specificity**	**AUC**	**Pixel Accuracy**
k-Fold Test	0.856±0.004	0.913±0.013	0.987±0.001	0.991±0.002	0.982±0.004
k-Fold Holdout	0.857±0.012	0.909±0.012	0.987±0.001	0.990±0.002	0.980±0.003
All Data	0.864	0.918	0.987	0.991	0.983

The Dice score distributions on the test set for AngioNet, Deeplabv3+, APN+U-Net, U-Net, APN+Attention U-Net, and Attention U-Net are shown in Figure 3.6. All networks were trained using the UM Dataset. AngioNet has the highest mean Dice score on the test set (0.864) when trained on all five partitions of the training data, compared to 0.815 for Deeplabv3+ alone, 0.811 for APN+U-Net, 0.717 for U-Net alone, 0.848 for APN+Attention U-Net, and 0.804 for Attention U-Net alone. On average, AngioNet has a 10% higher Dice score per image than Deeplabv3+ alone. APN+U-Net has a 14% higher Dice score than U-Net alone, while APN+Attention U-Net has a 10% higher dice score than Attention U-Net alone. A one-tailed paired Student's t-test (n=77) was performed to determine if the addition of the APN significantly improved the Dice score for all network backbones. The p-value between AngioNet and Deeplabv3+ was 5.76e-10, the p-value between APN+U-Net and U-Net was 2.63e-16, and the final p-value was 3.61e-11 between APN+Attention U-Net and Attention U-Net. All p-values were much less than the statistical significance threshold of 0.05, therefore we can conclude that there are statistically significant differences between the Dice score distributions with and without the APN. Furthermore, all three network backbones exhibit outliers with Dice score lower than 0.5, but adding the APN eliminates these outliers in all networks.

Although the Deeplabv3+ backbone had the highest Dice scores, the Attention U-Net backbone had the highest clDice scores. APN+Attention U-Net and Attention U-Net had clDice scores of 0.806 and 0.745 respectively, while AngioNet and Deeplabv3+ had clDice scores of 0.798 and 0.715. U-Net had the lowest clDice score of 0.715, and APN+U-Net had a clDice score of 0.781. In all three cases, the addition of the APN improved the clDice score by at least 5 points.

Compared to the other networks, AngioNet performs the best on the challenging cases shown in Figure 3.6b. The first row of Figure 3.6b shows segmentation performance on a low contrast angiography image. AngioNet can segment the major coronary vessels in this image, whereas Deeplabv3+ and Attention U-Net only segment one vessel and U-Net is unable to identify any vessels at all. APN+U-Net and APN+Attention U-Net can segment more vessels than either backbone alone, indicating once again that the addition of the APN improves segmentation performance on these low contrast images. In the second row, the networks including the APN can segment fainter and smaller diameter vessels than the backbone networks. Finally, we observe that AngioNet and Deeplabv3+ did not segment the catheter in the third column, although it is of similar diameter



	Test Dice	k-fold Test Dice	Test cl Dice
Deeplabv3+	0.8120	0.8153 ± 0.0057	0.7152
AngioNet	0.8640	0.8574 ± 0.0040	0.7977
U-Net	0.7166	0.8071 ± 0.0181	0.7153
APN+U-Net	0.8108	0.8093 ± 0.0059	0.7807
Attention U-Net	0.8035	0.8083±0.0380	0.7453
APN + Attention U-Net	0.8476	0.8314± 0.0086	0.8059

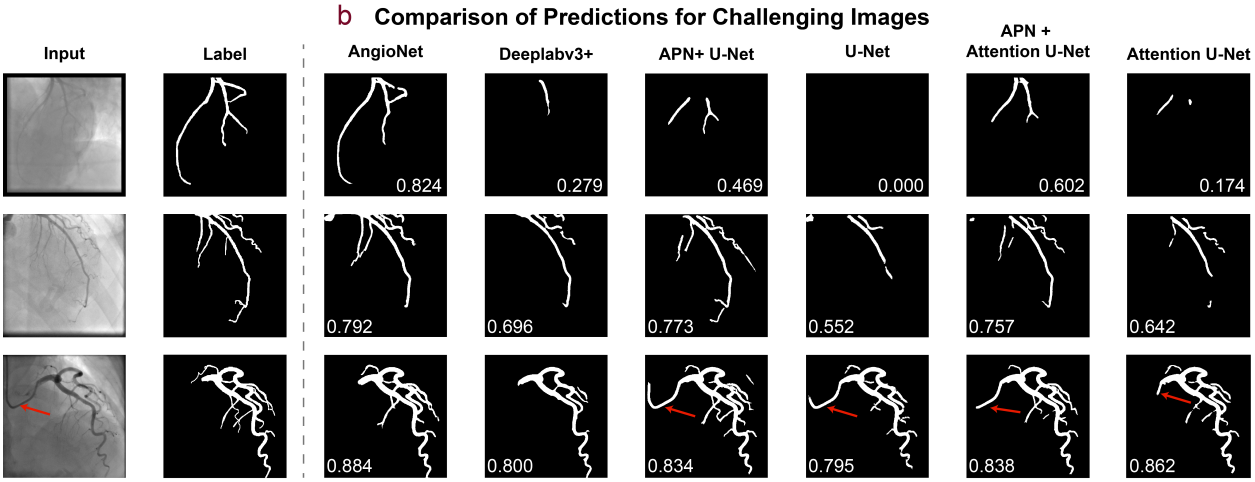


Figure 3.6: Summary of APN performance. All results are derived from the networks trained on all five partitions of the UM training set, unless otherwise noted as a k-fold result. a) Comparison of Dice score distribution on test set. AngioNet has the highest average Dice score, with scores ranging from 0.737 to 0.946. Adding the APN improves the Dice scores of all backbone networks (Deeplabv3+, U-Net, and Attention U-Net). Dashed lines correspond to the Test Dice in the table below. clDice is also summarized in the table, with the Attention U-Net backbone demonstrating highest topology preservation. b) Segmentation comparison on challenging images with low contrast, faint vessels, and a curved catheter. AngioNet can segment more vessels in these images without segmenting the catheter (red arrows). The Dice scores of all predicted images are indicated in the corner of each image.

and curvature as the coronary vessels. Conversely, the U-Net and Attention U-Net based networks included part of the catheter in their segmentations. Overall, AngioNet segmented the catheter in

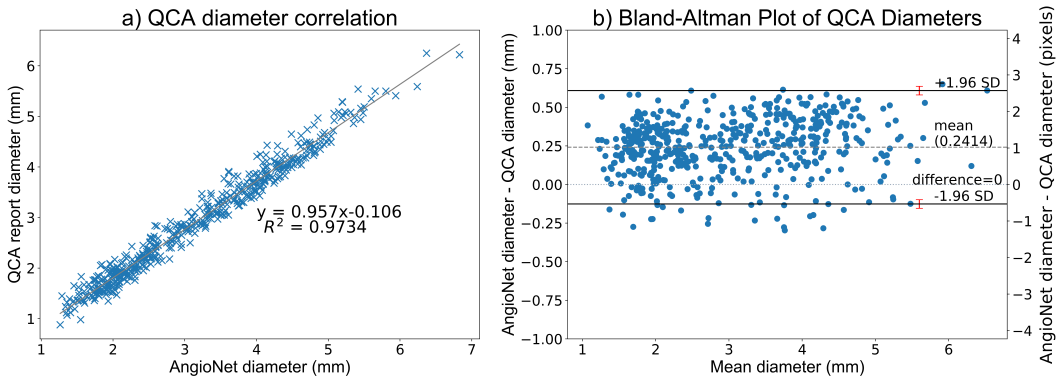


Figure 3.7: a) Correlation plot of QCA and AngioNet derived vessel diameters. b) The Bland-Altman plot demonstrates that AngioNet's segmentation and QCA are interchangeable methods to determine vessel diameter since more than 95% of points lie within the limits of agreement. The red error bars represent the 95% confidence interval containing the limits of agreement. The mean difference in diameter between methods is 0.24mm or 1.1 pixels.

2.6% of the images, where the catheter curved across the image and overlapped with the vessel. In contrast, Deeplabv3+ segmented the catheter in 6.4% of images. Both networks performed better than U-Net and APN+U-Net, which segmented the catheter in 19.5% and 9.1% of images respectively. The Deeplabv3+ backbone networks also outperformed the Attention U-Net networks, as Attention U-Net segmented the catheter in 11.7% of images and APN+Attention U-Net included in the catheter in 9.1% of images.

3.3.4 Evaluation of Vessel Diameter Accuracy versus QCA

Evaluation of vessel diameter accuracy was done using the MMM QCA dataset. Maximum and minimum vessel diameter were compared in 255 vessels including both the RCA and LCA. On average, the absolute error in vessel diameter between the AngioNet segmentation and QCA report was 0.272mm or 1.15 pixels.

The linear regression plot in Figure 3.7a shows that vessel diameter estimates of both methods (n=255) are linearly proportional and tightly clustered around the line of best fit, y=0.957x-0.106, Pearson's correlation coefficient, r=0.9866. The standardized difference [8], also known as Cohen's effect size [24], was used to determine the difference in means between the diameter distributions of AngioNet and QCA. The standardized difference can determine significant differences between two groups in clinical studies [8]. The standardized difference between the AngioNet and QCA diameter distributions is 0.215, suggesting small differences between the two methods.

Figure 3.7b is a Bland-Altman plot demonstrating the interchangeability of the QCA and AngioNet-derived diameters. The mean difference between both measures, $\bar{d}$, is 0.2414. The

magnitude of the diameter difference remains relatively constant for all mean diameter values, indicating that there is a systematic error and not a proportional error between the two measurements. The limits of agreement are defined as $\bar{d}$±1.96 SD, where SD is the standard deviation of the diameter differences. For both measurements to be considered interchangeable, 95% of data points must lie between these limits of agreement. In this plot, 96% of the 255 data points are within $\bar{d}$±1.96 SD. When including the 95% confidence interval of the limits of agreement as recommended by Bland and Altman [12], 97% of data points are within the range.

3.4 Discussion

In the following sections, we will demonstrate that AngioNet is comparable to state-of-the-art methods for angiographic segmentation. The key findings and clinical implications of our study are also described below.

3.4.1 Learned Filters using the Angiographic Processing Network

As seen in Figure 3.4, the APN learns many different preprocessing filters that improve segmentation performance based on the data partition used for training. Despite the variation in the learned filters, all learned filters exhibit both boundary sharpening and local contrast enhancement. Moreover, the overall segmentation accuracy of the various learned filters remains relatively constant as indicated by the small standard deviation from k-fold cross validation (0.004 for AngioNet, 0.006 for APN+U-Net, and 0.009 for APN+Attention U-Net, see table in Figure 3.6a). This demonstrates that the different combinations of learned weights all achieved similar local minima of the loss function, leading to similar Dice scores. Furthermore, the addition of the APN to U-Net and Attention U-Net led to a lower k-fold Dice score standard deviation (0.006 for APN+U-Net compared to 0.018 for U-Net, and 0.009 for APN+Attention U-Net compared to 0.038 for Attention U-Net). This could indicate that the networks incorporating the APN are more robust than the backbone networks alone.

To better understand the effect of using the APN for image preprocessing compared to selecting a particular pre-processing filter, we performed a follow-up experiment where Deeplabv3+ was trained on images pre-processed with the series of unsharp mask filters used to initialize the APN. These unsharp mask filters yielded the highest accuracy of all pre-processing filters we tested, including CLAHE for improved contrast, SVD-based background subtraction, and Gaussian blur for denoising. The average Dice score on the test set for unsharp Deeplabv3+ was 0.833, while the average Dice score for AngioNet was 0.864. A one-tailed paired Student's t-test (n=77) was used to determine if there were any significant differences between the two Dice score distributions.

The p-value of this test was 1.41e-11, which is much less than the threshold of 0.05. We therefore conclude that the APN's learned filter has a significant impact on overall segmentation accuracy compared to standard pre-processing filters. This suggests that the learned preprocessing filter implemented in this work is superior to manually selecting a particular contrast enhancement or boundary sharpening filter for preprocessing.

Our pre-training strategy for the APN and backbone networks allows the backbone network to gradually adjust its learned weights to better suit the filter learned by APN. Had the APN been randomly initialized, the filtered image may not be of high quality and could cause the quality of the backbone weights to degrade. Similarly, if only the APN was pre-trained and the backbone was randomly initialized, the much larger learning rate required to train the backbone network would cause the APN weights to degrade. Finally, initializing both the APN and backbone networks randomly can lead to the network getting trapped in local minima since the input to the backbone network may be of poor quality.

Another advantage of the pre-training and fine-tuning strategy is that it allows the user to design for the type of filter that is learned by the APN. By pre-training the APN to mimic a series of unsharp mask filters and applying a small learning rate during fine-tuning, we are able to prime the network to learn a filter that performs similar functions of boundary sharpening and local contrast enhancement. Indeed, despite the large variation in learned filters as seen in Figure 3.4, all the learned filters perform some variation of boundary sharpening and contrast enhancement at the edges of the vessel. This concept could be expanded to take advantage of all three channels of the filtered image that is input into the backbone network. Rather than training the APN to learn a single filter and concatenating it to form a three-channel image, we can pre-train the APN to mimic three different filters for various purposes such as edge detection, contrast enhancement (CLAHE), vesselness enhancement (Frangi filter), and texture analysis (Gabor filter). Learning several new filters for these purposes through the APN may provide the backbone network with richer information with which to segment the vessels.

3.4.2 Comparison of AngioNet to Current State-of-the-Art Semantic Segmentation Neural Networks

Several aspects of AngioNet's design contribute to its enhanced segmentation performance compared to existing state of the art networks. The APN successfully improves segmentation performance on low contrast images compared to previous state-of-the-art semantic segmentation networks (Figure 3.6a). The APN also enhances performance on smaller vessels, which have lower contrast than larger vessels because they contain less radio-opaque dye. Without the APN, Deeplabv3+, U-Net, and Attention U-Net are not equipped to identify these faint vessels and un-

derpredict the presence of small coronary branches. As seen in Figure 6a and Figure 6b, both Deeplabv3+ and AngioNet perform better than U-Net on angiographic segmentation. The addition of the APN to U-Net significantly increases the mean Dice score, facilitates segmentation of more vessels compared to U-Net alone, and greatly reduces the proportion of segmentation that include the catheter; yet APN+U-Net has some of the same drawbacks of U-Net such as disconnected vessels and more instances of the catheter being segmented compared to Deeplabv3+ and AngioNet. Although the Attention U-Net backbone and APN+Attention U-Net outperform U-Net and APN+U-Net, both Attention U-Net networks are still more susceptible to including the catheter (in 19.5% and 9.5% of test images) than Deeplabv3+ (6.4%) and AngioNet (2.6%).

The main benefit of the attention gates in the Attention U-Net backbone seems to be in suppressing background artifacts and preserving the connectivity of the vessels, which can be explained by the global nature of attention-based feature maps. This is demonstrated by Attention U-Net and APN+Attention U-Net having the highest clDice scores compared to the other network backbones. Deeplabv3+ and AngioNet still demonstrate higher mean Dice scores and lower standard deviations than the attention networks; the discrepancy between Dice and clDice scores could be because Deeplabv3+ and AngioNet are able to segment more vessels than the attention-based networks, but the branches are not necessarily connected. While vessel connectivity is an important feature of coronary segmentation, the drawback of catheter segmentation outweighs this benefit of the attention-based networks. It is difficult to avoid segmenting the catheter using filter-based methods, thus the main advantage of deep learning methods is their ability to avoid the catheter. AngioNet excels at this task compared to the other networks. Although U-Net and Attention U-Net have demonstrated great success in other binary segmentation applications [115, 1, 96], the presence of catheters and bony structures with similar dimensions and pixel intensity as the vessels of interest make this a particularly challenging segmentation task. Deeplabv3+ and AngioNet have a deeper, more complex architecture, which allows these networks to learn more features with which to identify the vessels in each image [106, 32, 76].

The effective receptive field size U-Net and Attention U-Net is 64x64 pixels whereas that of Deeplabv3+ is 128x128 pixels. A larger receptive field is associated with better pixel localization and segmentation accuracy, as well as classification of larger scale objects in an image [74, 147]. Deeplabv3+'s larger receptive field may explain why Deeplabv3+ and AngioNet are more successful in avoiding segmentation of the catheter, an object typically larger than U-Net or Attention U-Net's 64x64 pixel receptive field. The larger receptive field may also explain why Deeplabv3+ and AngioNet are better able to preserve the continuity of the coronary vessel tree and produce fewer broken or disconnected vessels than U-Net and APN+U-Net. The attention gates in Attention U-Net perform a similar function in terms of preserving connectivity, however, the attention mechanism is not able to suppress the features that define catheter. Thus, Deeplabv3+ was an

appropriate choice of network backbone for AngioNet.

AngioNet's strengths compared to previous networks include ignoring overlapping structures when segmenting the coronary vessels, smaller sensitivity to noise, and the ability to segment low contrast images. The ability to avoid overlapping bony structures or the catheter is especially important as this eliminates the need for manual correction of the vessel boundary, which is a major advantage over mechanistic segmentation approaches.

AngioNet's greatest limitation is that it overpredicts the vessel boundary in cases of severe ($>$85%) stenosis. The network performs well on mild and moderate stenoses, but it has learned to smooth the vessel boundary when the diameter sharply decreases to a single pixel. This is likely due to the low number of training examples containing severe stenosis: only 14 out of the 462 images in the entire UM Dataset contained severe stenosis, and two of these were in the test set. This drawback can be addressed by increasing the training data to encompass more examples of severe stenosis.

3.4.3 Evaluation of Vessel Diameter Accuracy

A significant clinical implication of our findings was the comparison between AngioNet and QCA. In Figure 7a, we observe that QCA and AngioNet results are clustered around the line of best fit, $y = 0.957x - 0.106$. Given that the slope of the line of best fit is nearly 1, the intercept is close to 0, and the Pearson's coefficient r is 0.9866, the line of best fit indicates strong agreement between these two methods of determining vessel diameter. The R^2 coefficient for the linear regression model implies that 97.34% of the variance in the data can be explained by the line of best fit. The standardized difference, or effect size, is a measure of how many pooled standard deviations separate the means of two distributions [8]). According to Cohen, an effect size of 0.2 is considered a small difference between both groups, 0.5 is a medium difference, and 0.8 is a large difference [24]. Given that the effect size between the QCA and AngioNet diameter distributions was 0.215 (91.5% overlap between the two distributions), we can conclude that the difference between QCA and AngioNet diameters are small. Furthermore, since the standardized difference indicated no large difference between QCA and AngioNet diameter estimations, these results suggest that both methods can be used interchangeably from a clinical perspective for the dataset examined.

The Bland-Altman plot in Figure 3.7b shows that the mean difference between QCA and AngioNet diameters is approximately 1.1 pixels. AngioNet under-predicts the vessel boundary by no more than 1 pixel and over-predicts by no more than 2.5 pixels. To put these values in context, the inter-operator variability for annotating the vessel boundary is 0.18±0.24mm or slightly above 1 pixel according to a study by Hernandez-Vela et al [51]. The 95% confidence intervals of the limits of agreement were taken into consideration when determining how many data points lie be-

tween the limits of agreement as recommended by Bland and Altman [12]. 97% of the data points lie within the range, which is greater than the 95% threshold. Given these results, and that the standardized difference test which produced no significant difference between the methods, one can conclude that QCA and AngioNet are interchangeable methods to determine vessel diameter. Given AngioNet's fully automated nature, the workload required for generating QCA due to human input could be substantially reduced. Although our direct comparison of AngioNet-derived diameters with QCA-derived diameters required user interaction, future work will focus on developing an automated algorithm for stenosis detection and measurement based on the outputs of AngioNet's segmentation.

3.5 Conclusion

In conclusion, AngioNet was designed to address the shortcomings of current state-of-the-art neural networks for X-ray angiographic segmentation. The APN was found to be a critical component to improve detection and segmentation of the coronary vessels, leading to 14%, 10%, and 10% improved Dice score compared to U-Net, Attention U-Net, or Deeplabv3+ alone. AngioNet demonstrated better Dice scores than all other networks, particularly on images with poor contrast or many small vessels. Although APN+Attention U-Net demonstrated the highest clDice score (0.806), the connectivity of AngioNet was comparable (0.798) and AngioNet was better able to avoid segmenting the catheter and other imaging artifacts. Furthermore, our statistical analysis of the vessel diameters determined by AngioNet and traditional QCA demonstrated that the two methods may be interchangeable which could have large implications for clinical workflows. Future work to improve performance will focus on increasing accuracy on severe stenosis cases and automating stenosis measurement. We also aim to increase the versatility of the APN by pre-training and learning several filters for edge detection, texture analysis, or vesselness enhancement in addition to our current contrast-enhancing implementation. Combining these learned filters into a multi-channel image may improve the semantic segmentation performance of AngioNet.

CHAPTER 4

Stenosis Characterization

4.1 Introduction

As discussed in Chapter 1, QCA is a software tool that can be used to quantify CAD severity from X-ray angiography images. The QCA workflow is composed of three major tasks: 1) Identification of the target vessel and its boundaries 2) Automatic stenosis localization within the target vessel and 3) Quantification of stenosis severity.

QCA is generally performed in patients with intermediate lesions, since these are the cases most affected by the subjectivity of visual estimation of disease severity [40]. Several studies have concluded that visual estimation over-predicts the presence of ischemic intermediate lesions [90, 44, 125], potentially leading to revascularization procedures in patients who did not have obstructive disease according to QCA. There is also large inter-operator variability in visual assessment of disease severity compared to the variability between QCA operators [39, 159, 132]. Quantitative metrics of disease severity are therefore a crucial tool to reduce the incidence of unnecessary revascularization procedures.

Despite the value of QCA in improving diagnostic outcomes, the workflow has several limitations. The initial steps of the QCA workflow require substantial operator input, particularly during manual selection of the vessel of interest and the correction of any errors in the automatic determination of the vessel boundaries. While lower than visual inspection alone, the inter-operator variability introduced by these manual steps is approximately 0.2mm for obstruction diameter and 6% for percent diameter stenosis[112]. Furthermore, QCA is typically performed at independent core laboratories with specially trained staff to reduce bias[40]. The high level of training required and limited access to independent core labs restricts the availability of QCA. Another drawback of the QCA workflow is the estimation of the reference diameter used to compute percent diameter stenosis. According to de Feyter et al.[29], the reference vessel diameter is prone to error, particularly if there is diffuse atherosclerosis or nonlinear tapering in the diseased portion of the vessel.

In recent years, there has been a push towards using neural networks to improve the automation and speed of stenosis severity quantification. Several neural network pipelines that automate the steps of QCA are presented in section 1.2. However, these pipelines did not *quantify* lesion diameter or percent diameter reduction, choosing instead to categorize severity based on threshold values and validate their results against visual estimates of stenosis severity.

In this work, we present a multi-stage workflow for 1) automated stenosis localization, 2) stenosis segmentation, and 3) stenosis severity quantification using neural networks and image processing algorithms. The major contribution of this work is an end-to-end pipeline to quantify stenosis severity in X-ray angiography images that has been validated against the clinical standard of QCA.

4.2 Methods

This study has two main endpoints. The first is to evaluate the accuracy of stenosis localization by our machine learning pipeline against a ground truth determined by 2 board-certified cardiologists. The second endpoint is to quantify stenosis severity, in which we compare both lesion diameter and percent diameter reduction from our pipeline against the ground truth of QCA. These endpoints were achieved using two datasets: X-ray angiography images from the University of Michigan (UM) were used for neural network training and algorithm development, while angiography images and their corresponding QCA reports from Madras Medical Mission Indian Cardiovascular Core Laboratory (ICRF) were used for external validation of our workflow Figure 4.1.

The UM dataset was composed of 984 angiogram images from 709 patients. These images were manually annotated and split into training and test sets which were used to train neural networks for the stenosis detection and segmentation tasks. The ICRF validation set was composed of 399 angiograms from 173 patients, each with a corresponding full-vessel QCA report. The Institutional Review Boards of the University of Michigan Medical School (IRBMED) and the Institutional Ethics Committee at the Madras Medical Mission approved the retrospective analysis of de-identified patient images (study protocol HUM00155491). Further details of the study population are presented in the results section.

Angiography images from the ICRF validation set were independently assessed by two board-certified cardiologists to determine whether they contained stenoses of greater than 50% diameter reduction. The cardiologists were blind to the location and severity of lesions measured by QCA. The cardiologists' consensus on $\geq$50% stenoses was used as the ground truth for our first endpoint, the stenosis localization task. For the second endpoint, we compared percent diameter reduction and lesion diameter against the measurements derived from QCA in vessels containing stenoses.

Our workflow to identify stenoses and quantify their severity is composed of three main stages: localization of stenoses in X-ray angiograms, stenosis segmentation, and evaluation of disease

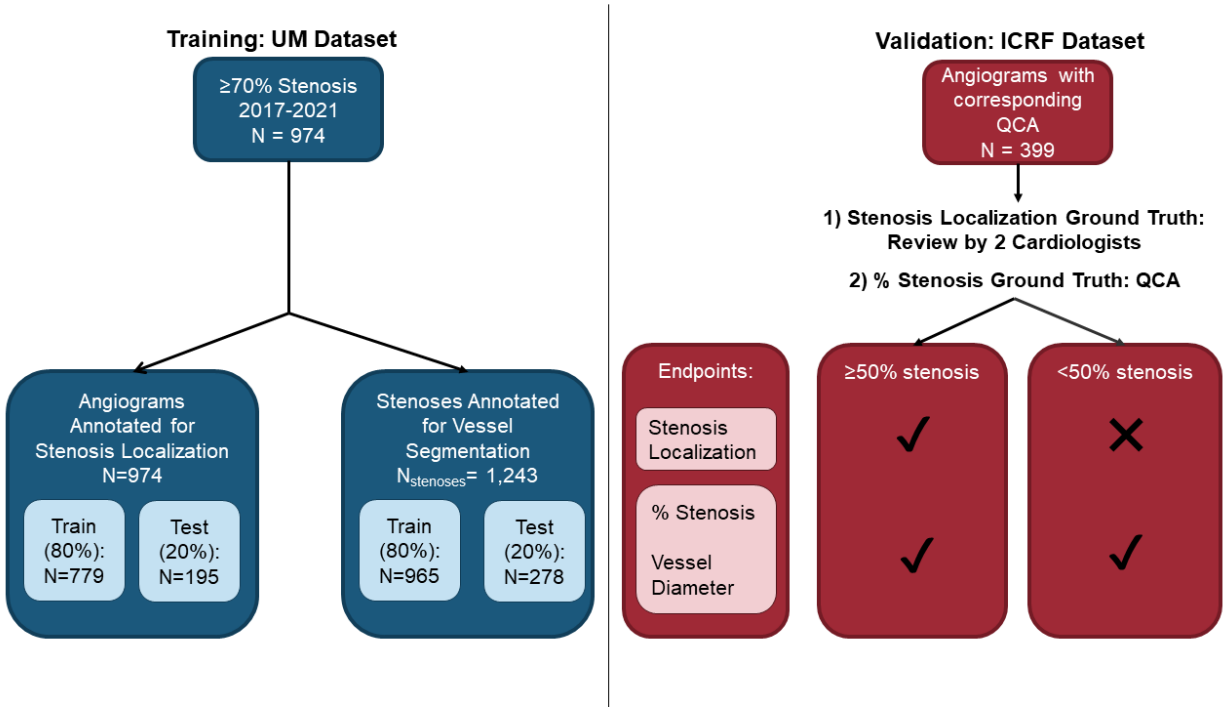


Figure 4.1: Diagram of training and validation datasets. Angiograms from UM were used to train neural networks for the stenosis localization and vessel segmentation tasks. Validation angiograms and their corresponding QCA reports were obtained from ICRF. The ground truth for stenosis localization was lesions of $\geq$50% diameter reduction according to two board-certified cardiologists, and the ground truth for vessel diameter and percent stenosis calculation was from QCA analysis of vessels containing $>$20% stenosis.

severity (Figure 4.2). Neural networks were used to perform the stenosis localization and segmentation tasks, while image processing algorithms were used to quantify lesion diameter and percent diameter reduction.

4.2.1 Stenosis Localization

The first stage of our stenosis severity quantification pipeline is a neural network which identifies stenoses in X-ray angiography images. The YOLOv5s network architecture [58], based on the work of Redmon and Farhadi [111], was chosen because of its near-real time inference speed and better performance at detecting small objects compared to other state of the art object detection networks [113].

The input angiography images were standardized using z-score normalization to increase robustness of the neural network to images from different datasets, where the distribution of pixel intensities may not match that of the training dataset. The training dataset was composed of 984 standardized angiography images from the UM dataset and their corresponding annotations. An-

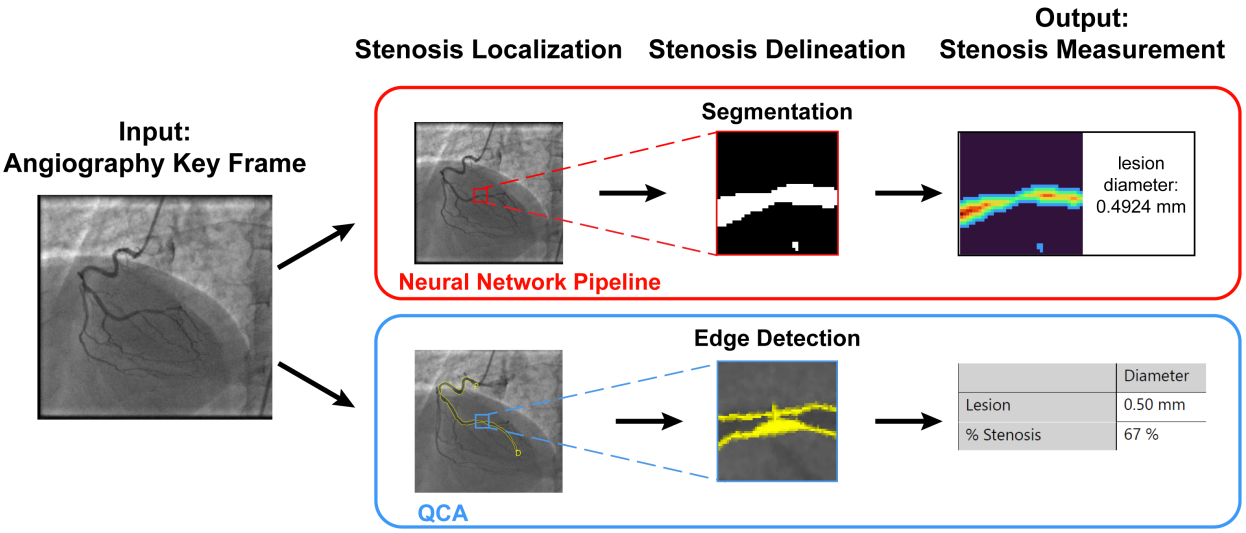


Figure 4.2: Diagram of the proposed neural network stenosis characterization pipeline (top) and the QCA workflow (bottom) to determine lesion diameter and stenosis severity.

notations were manually generated by two cardiologists using a script to click on the location of stenoses in each image and save the normalized coordinates of the resulting bounding box in a text file. All bounding boxes had dimensions of 40x40 pixels, which was empirically determined to encompass the region of stenosis in most images. 787 images were used for training the network and 197 images were used for validation, corresponding to an 80/20 train-validation split.

Transfer learning was employed due to the relatively small size of our training dataset. YOLOv5 was initialized with pre-trained weights from ImageNet and trained for 400 epochs on our dataset using an ADAM optimizer with a learning rate of 1e-3. The loss function was a combination of binary cross entropy loss and focal loss, which measured how well the network predicted bounding boxes in the correct locations and classified them as stenoses. The output of the stenosis detection network was the coordinates defining the 40x40 bounding boxes of all stenoses in an image. The bounding boxes were used to crop small patches of the original image for further analysis.

4.2.2 Stenosis Segmentation

In order to quantify diameter reduction, the region of the image containing the vessel must be delineated. In the QCA workflow, this is achieved using edge detection algorithms guided by user input. In order to automate this process, we trained a neural network for vessel segmentation.

Vessel segmentation was performed on image patches measuring 40x40 pixels, the same size as the bounding box generated by the stenosis localization network. The segmentation network is a variation of our previous work, AngioNet[55], which employs an Angiographic Processing Network (APN) to learn the best possible pre-processing filters that will improve vessel segmentation.

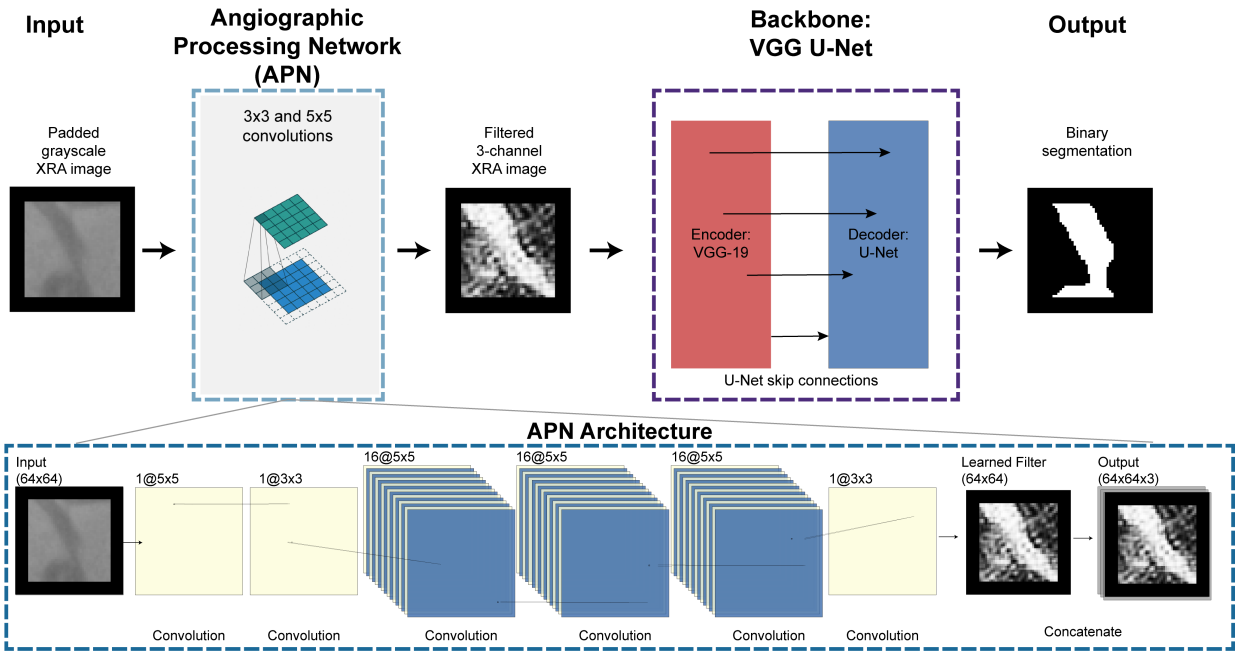


Figure 4.3: Architecture of our stenosis patch segmentation network, based on AngioNet [55]. The angiographic processing network (APN) learns filters that improve segmentation performance, and the VGG U-Net backbone performs the segmentation task.

In this work, the segmentation network was composed of the APN and a backbone network, U-Net [115] with a VGG-19[131] encoder Figure 4.3. The network was trained using 965 image patches containing stenoses from the UM dataset, while 278 image patches were used for validation. Our patch segmentation was trained for 400 epochs with a learning rate of 1e-4 and ADAM optimizer. To minimize overfitting, weight decay (L2) regularization was imposed with a coefficient of 1e-3 during the training process. The loss function was a weighted average of Tversky loss [120], which penalizes false positive vessel pixels, and centerline-Dice loss[129], which measures the accuracy of the vessel shape or morphology.

4.2.3 Stenosis Severity Quantification

Segmented image patches were evaluated using image processing algorithms to measure lesion diameter and quantify percent diameter reduction. A Euclidean distance transform was applied to the segmented binary image to create a distance map. A skeletonization algorithm [156] was then applied to identify the vessel centerline. The radii along the vessel segment were determined by sampling the Euclidean distance map at the locations of the centerline pixels. The pixel diameters were converted to millimeters via a calibration factor from the patient's QCA report, derived from the width of a 6F catheter. A 1-D signal processing algorithm was applied to the list of radii to

identify the peaks (diameters proximal and distal to the stenosis) and troughs (lesion diameter). Percent stenosis was calculated as

$$\% \; stenosis = 1 - \frac{lesion \; diameter}{0.5 \times proximal \; diameter \times distal \; diameter} \tag{4.1}$$

where the reference diameter was the average of the diameters proximal and distal to the stenosis. The absolute errors in percent stenosis and lesion diameter were calculated between the values computed using our workflow and the values recorded in the corresponding QCA report.

4.2.4 Statistics

The accuracy of stenosis localization was assessed using precision and recall. Stenosis severity characterization was broken up into two separate tasks: measuring lesion diameter and quantifying percent diameter reduction. Pearson's correlation coefficient was computed to determine the strength of the association between the stenosis diameters measured by QCA and our pipeline. The normality of the error distribution (QCA diameter – pipeline diameter) was verified using a Shapiro-Wilkes test for normality with an alpha level of 0.05. Since the normality assumption was met, a Bland-Altman analysis was performed to determine the interchangeability of both methods to measure lesion diameter and percent diameter reduction. All statistical analysis was performed using the python SciPy stats package (software version 1.7.3).

4.3 Results

4.3.1 Study Population

We now describe the study population characteristics for the UM dataset, used to train the neural networks in our pipeline, and the ICRF QCA dataset, which was used to validate our method's accuracy.

4.3.1.1 UM Training Dataset

Retrospective patient data from the University of Michigan were identified using the EMERSE[48] database. The study criteria for this dataset were patients who had a coronary angiography procedure performed at the University of Michigan hospital between 2017 and 2021, with coronary findings of at least one stenosis greater than 70% diameter reduction. Patients with bypass grafts or prior stent implantation were excluded from this study. All angiograms were de-identified before analysis in compliance with ethical guidelines.

4.3.1.2 ICRF QCA Dataset

The inclusion criteria for the ICRF QCA dataset were patients older than 18 with stable angina (acute coronary syndrome) symptoms for over 24 hours, with or without percutaneous intervention. Exclusion criteria included patients with a history of STEMI, ongoing arrhythmia, CABG, heart valve surgery, moderate to severe valve stenosis or regurgitation, tachycardia, atrial fibrillation, or cardiogenic shock. Angiographic exclusion criteria were lesions in the left main artery or at a bifurcation. Additionally, target vessels with in-stent restenosis, chronic total occlusion, severe diffuse disease, or collaterals were excluded from this study.

Patients were categorized according to the presence or absence of obstructive disease, defined as diameter reduction greater than 50% according to QCA. Out of 173 total patients, 51 patients had at least one vessel with obstructive disease. In total, 399 vessels were analyzed, of which 69 vessels had obstructive disease.

4.3.2 Stenosis Localization

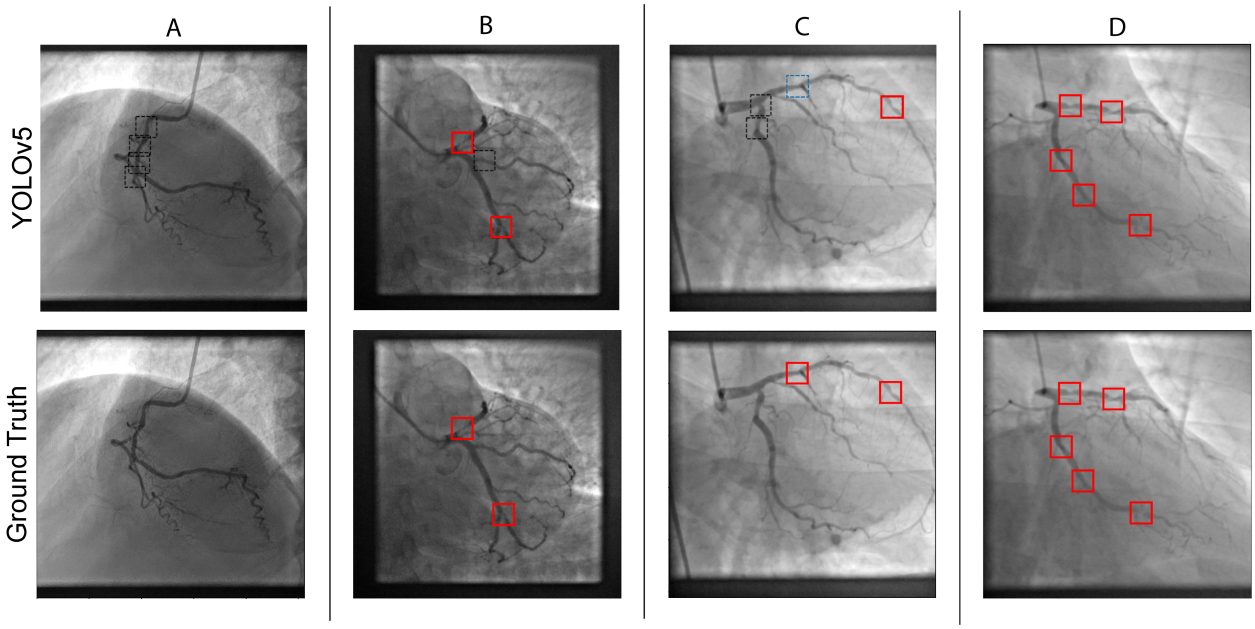


Figure 4.4: Examples of stenosis localization by the YOLOv5s network (top) and their corresponding ground truth as determined by 2 cardiologists (bottom). In these examples, correctly identified regions containing stenoses are given by the solid red bounding boxes. False positive neural network detections are shown using black dashed boxes, and false negative detections are shown using blue dashed boxes. Some of the false positive detections do contain stenoses; however, these were determined to be <50% stenoses by the cardiologists.

Figure 4.4 shows several examples of stenosis localization by the YOLOv5 neural network.

The black dashed boxes correspond to false positives, the blue dashed boxes are false negatives, and the red solid boxes indicate stenoses that were correctly identified by the neural network. In Figure 4.4A, the neural network identified several stenoses on the RCA which were all determined to be less than 50% diameter reduction by the cardiologists. Figure 4.4B and C are examples where the neural network correctly identified most of the lesions identified in the ground truth, although there are also some examples of false positives and one false negative. Finally, Figure 4.4D shows an example where the neural network correctly identifies all the stenoses in the ground truth without any false or missing detections.

To quantify localization accuracy, the precision and recall were computed with respect to the ground truth values determined by cardiologist review. Precision, or positive predictive value, and recall, or true positive rate, are defined as

$$Precision = \frac{TP}{TP + FP} \tag{4.2}$$

$$Recall = \frac{TP}{TP + FN} \tag{4.3}$$

where TP represents true positives, FP represents false positives, and FN represents false negatives.

The precision-recall curve in Figure 4.5 shows the trade-off that occurs between these two quantities. When the precision is high, a large proportion of the predicted stenoses are indeed stenoses according to the ground truth. This is achieved when the confidence threshold for predicted stenoses is high. The drawback of choosing a high confidence threshold is that the recall will be low due to a large number of false negative predictions. Conversely, when the confidence threshold for predicted stenoses is low, there will be a small number of false negatives and thus a high recall. In this case, the low confidence threshold will also lead to many false positives and consequently low precision. Our goal is to maximize both precision and recall, which occurs when the precision of our neural network is 0.63 and the recall is 0.67. The mean average precision across all values of recall was 0.65. Meanwhile, we would expect an untrained or random algorithm to predict stenoses with a precision of 0.5 no matter the recall value.

To put these values in context, Figure 4.6 shows the agreement between the two cardiologists who evaluated the ICRF angiograms for the presence of stenoses with greater than 50% diameter reduction. Compared to each other, the precision and recall of Cardiologist A compared to Cardiologist B was 0.61 and 0.65 respectively, and vice versa for Cardiologist B. Therefore, our neural network for stenosis localization is performing at a similar level as the human experts. This highlights the difficulty in identifying stenoses of a certain severity using visual estimation rather than quantitative tools.

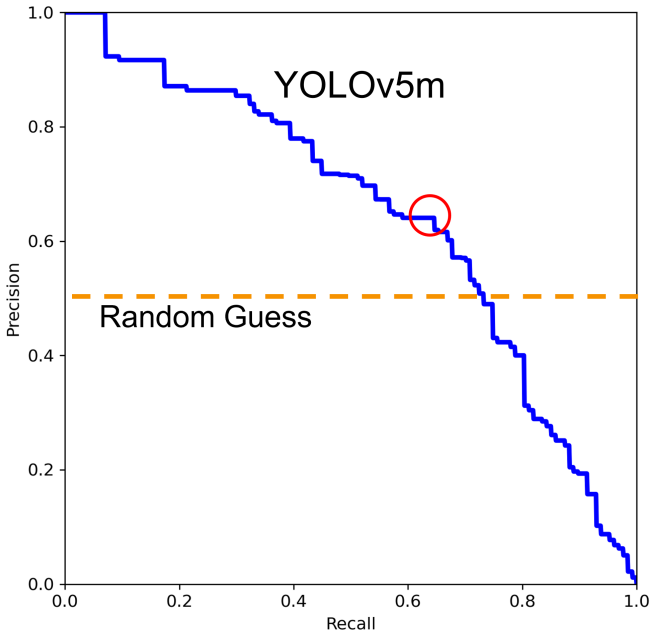


Figure 4.5: Precision-Recall curve for YOLOv5 stenosis localization. There is a trade-off that occurs between precision and recall, since minimizing false positives by increasing the confidence threshold for stenoses can lead to an increase in false negatives and vice versa. The mean average precision was 0.65 across all levels of recall, and the best performance occurs at a precision of 0.63 and recall of 0.68 (red circle).

4.3.3 Stenosis Severity Quantification

In this section, we are interested in comparing lesion diameter and percent diameter reduction from our pipeline with the measurements derived from QCA. Although QCA was performed on 399 vessels, not all of these analyzed vessels had stenoses. Cardiologists previously identified obstructive lesions in this set of angiograms for the stenosis localization task, however the lesions were found throughout the coronary tree while QCA was only performed in a single vessel for each angiogram. Furthermore, the visual estimate of the $\geq 50\%$ diameter reduction is not as robust as a quantitative measurement such as QCA, particularly due to the large inter-operator variability seen in Figure 4.6. Thus, we restrict our analysis to vessels that were analyzed using QCA and contained stenoses $\geq 20\%$ diameter reduction.

Localized image patches of the 308 stenoses in the ICRF QCA dataset were segmented and analyzed to determine vessel diameter at the lesion. Of these stenoses, 76 were measured to have diameter reduction $\geq 50\%$. Examples of vessel segmentation and computed lesion diameters are shown in Figure 4.7. The mean difference in vessel diameter was 0.206 ± 0.155, or approximately 1 pixel. There was no bias towards under- or over-estimating diameter, as shown in the relative error box plot in Figure 4.8A. Pearson's r was calculated to be 0.9098 (Figure 4.8B), indicating

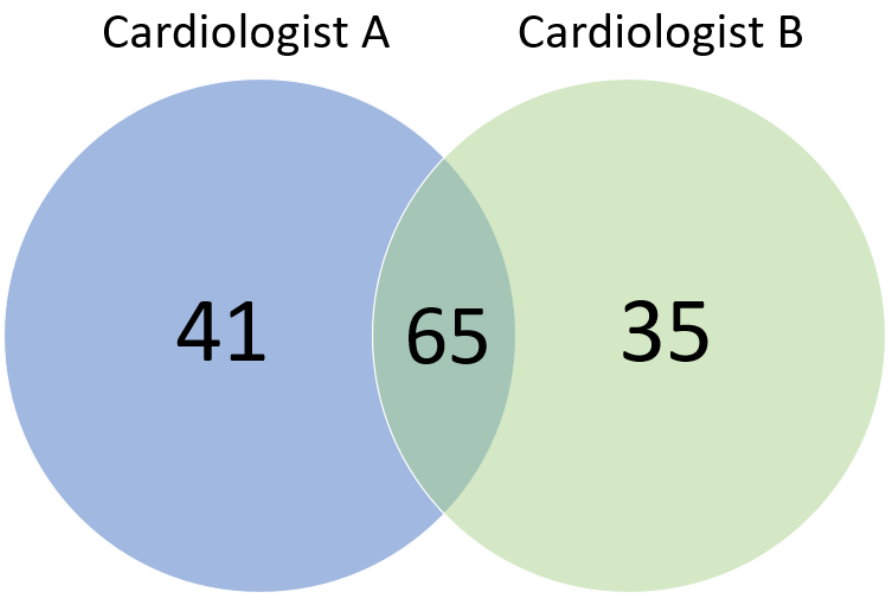


Figure 4.6: Variability in ground truth stenosis localization. A total of 141 stenoses $\geq$50% were localized by both cardiologists. Both cardiologists independently identified the same lesion 65 times, and disagreed on 76 lesions.

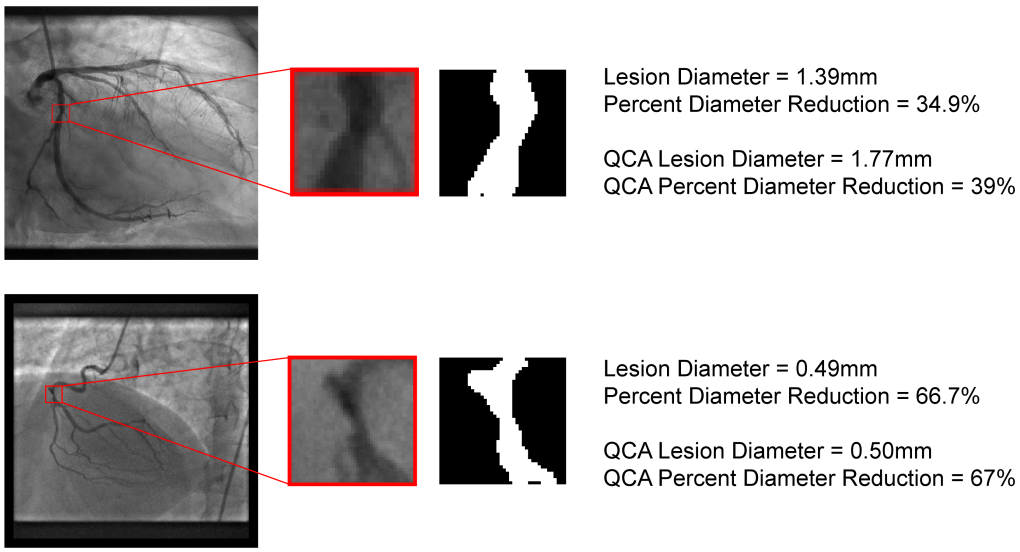


Figure 4.7: Examples of stenosis segmentation and their corresponding measurements from our pipeline and QCA.

strong positive correlation between lesion diameters derived from both methods.

A Bland-Altman plot was used to assess how well the vessel diameters from our stenosis charac-

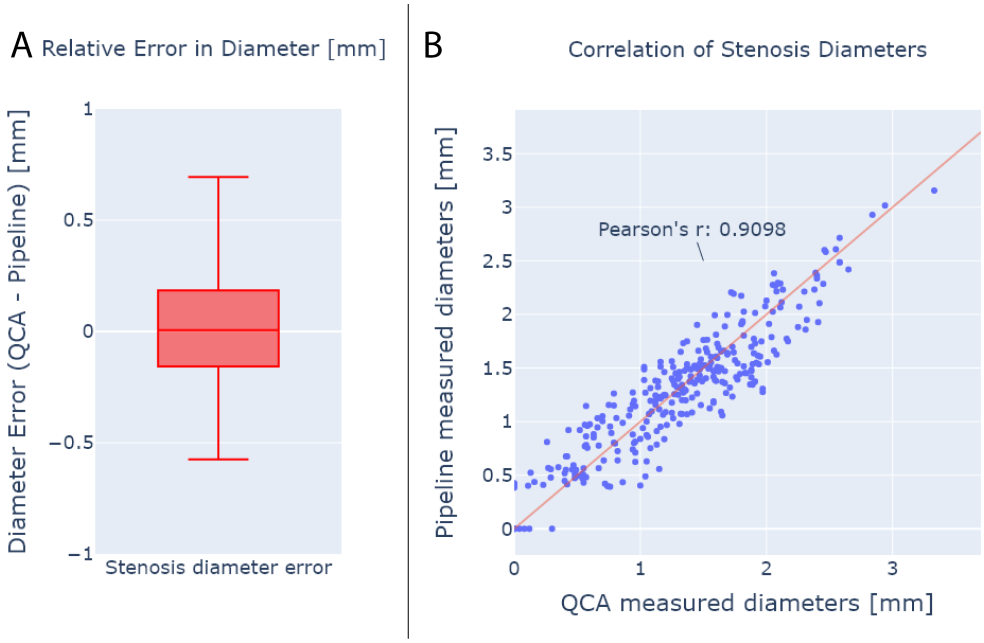


Figure 4.8: Agreement between lesion diameters measured by QCA and our pipeline. A) Box plot of relative error in millimeters (QCA – Pipeline diameters), with a median value of 0.006mm. B) Correlation plot of diameters measured using both methods and the line of best fit, with a Pearson's r value of 0.9098.

terization pipeline agree with those measured by QCA (Figure 4.9). Before generating the Bland-Altman plot, the normality of the error distribution was assessed. The p-value of the Shapiro-Wilkes test for normality was 0.205, indicating that we cannot reject the null hypothesis that the errors are normally distributed using an alpha level of 0.05.

The mean vessel diameter determined using both methods was plotted against the difference in diameter to generate a Bland-Altman plot. The mean error μ was 0.045 mm, indicating a slightly positive bias. The limits of agreement were set at 0.523 mm and -0.485 mm, which correspond to $\mu \pm 1.96\sigma$ where σ is the standard deviation of the error. 96.4% of the 308 data points were within the limits of agreement.

In addition to comparing the error in lesion diameter, we evaluated the agreement between percent diameter reduction as measured by QCA and our stenosis characterization pipeline. On average, the error in percent diameter reduction was 12.37$\pm$9.85% in the same 308 stenoses which were considered in our lesion diameter analysis. A Bland-Altman plot of the percent diameter reduction is shown in Figure 4.10. The limits of agreement are 34.3% and -20%, and 294 out of the 308 data points lie within these limits of agreement.

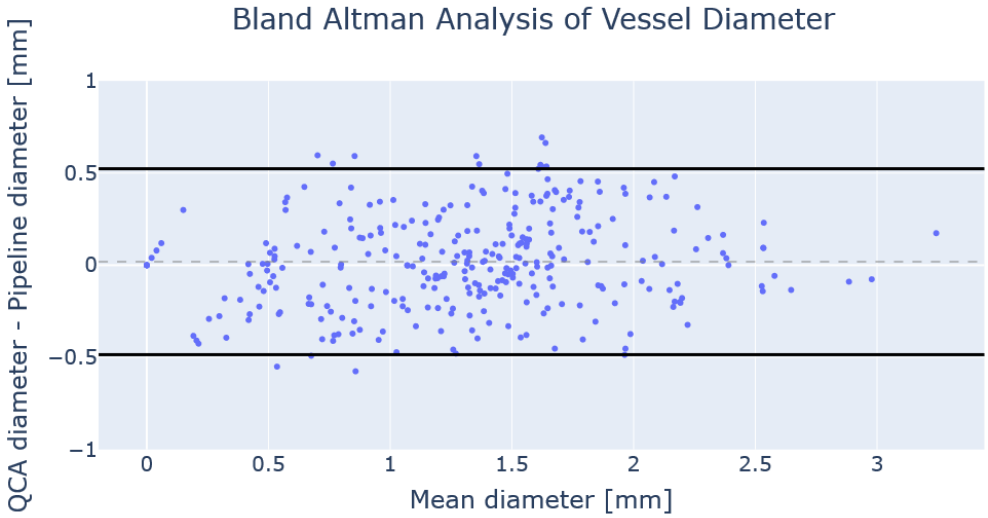


Figure 4.9: A Bland-Altman plot comparing vessel diameter measured by QCA and our stenosis characterization pipeline in vessels containing stenoses greater than 20% diameter reduction. The gray dashed line is the mean relative error, 0.045mm. The limits of agreement are at 0.52 and -0.49mm

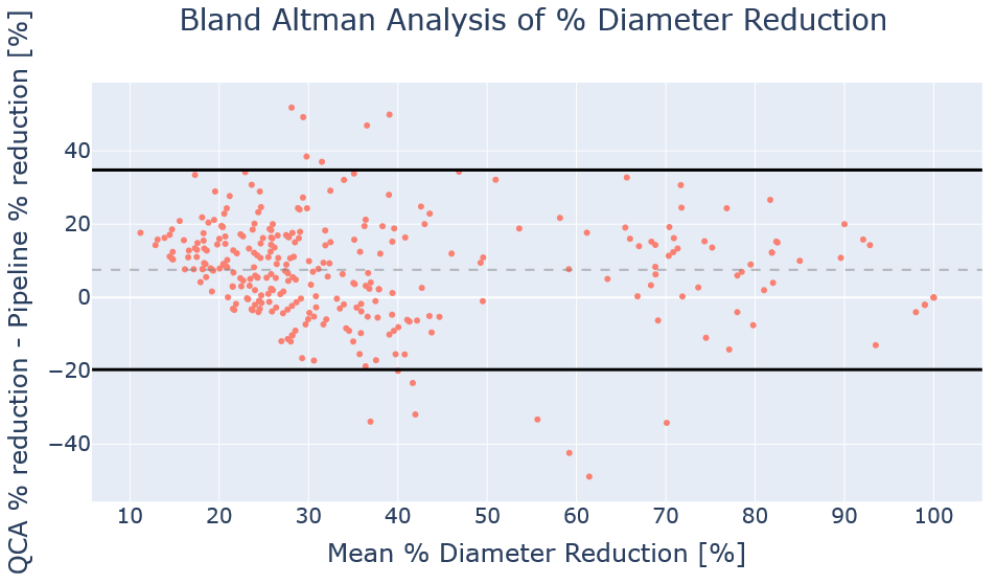


Figure 4.10: A Bland-Altman plot comparing percent diameter reduction measured by QCA and our stenosis characterization pipeline. The mean relative error of 7% is given by a gray dashed line. The limits of agreement are given by solid lines at 34.3% and -23%.

4.4 Discussion

4.4.1 Stenosis Localization

The major advantage of our stenosis localization network that it is not dependent on an operator identifying a vessel or region of interest in which to search for stenoses. Unlike QCA, which identifies stenoses as the location of minimum diameter within the user-specified region of the vessel, the neural network is sensitive to local changes in diameter and other learned features of stenoses. The fully automated neural network approach reduces subjectivity and decreases the time to localize stenoses in an angiogram.

As seen in Figure 4.4, the YOLOv5 neural network had a tendency to over-predict the presence of stenoses compared to the ground truth consensus of two cardiologists. This is mainly due to the fact that the neural network was trained to identify all stenoses in an image; therefore, the high number of false positive localizations is unsurprising since the ground truth only contains $\geq$50% stenoses. A closer look at Figure 4.4 shows that the false positive detections, shown in the black dashed boxes, do contain mild to moderate stenoses. Even if the network falsely localized extraneous regions of the angiogram, from the perspective of the diagnostic pipeline, the only drawback is the slight increase in computational time to segment and measure the additional possible stenoses. The increased time is on the order of less than a second per detection, which does not significantly slow down the diagnostic pipeline.

False negative detections are more dangerous, since a clinically significant lesion could be overlooked. We could reduce false negative detections and increase the recall by decreasing the confidence threshold for localizing stenoses, but this would lead to a large number of false positive detections. A more robust strategy to improve stenosis localization performance would be to increase the quality of the training data annotations by having more cardiologists review the images and retrain the network. This would push the network towards erring on the side of false positive rather than false negative detections. Another aspect of our evaluation which affects the precision and recall is the trustworthiness of the ground truth stenoses in the QCA validation set. Figure 4.6 demonstrates that two experts only agree on the location of obstructive stenoses with a precision of 0.61, so increasing the number of experts may lead to a better consensus on the ground truth locations of stenoses. Furthermore, since the neural network was trained to detect stenoses of any severity, a fairer comparison would be to include all levels of stenoses in the ground truth determined by the cardiologists as well. In this study, only stenoses $\geq$50% diameter reduction were included in the ground truth since identification of milder stenoses is more subjective. Gathering input from more cardiologists to determine the ground truth stenoses could again alleviate this issue.

4.4.2 Stenosis Severity Quantification

The box plot in Figure 4.8A demonstrates that the median error was close to zero (0.006mm), which reflects that the segmentation network does not have a biased tendency to either under- or over-predict the boundaries of the stenosis. The range of errors was between 0.69mm and -0.57mm, where half of the predicted diameters were within 0.18mm of the ground truth diameter. These findings are supported by the Bland-Altman plot in Figure 4.9, which has a mean close to zero and limits of agreement at at 0.52mm and -0.49mm. Since 297 of the 308 data points, or 96.4%, were within the limits of agreement, this indicates that our pipeline and QCA are interchangeable methods to determine vessel diameter. We must also consider whether the limits of agreement, or the discrepancy between both methods, is at an acceptable level. The limits of agreement are around 0.5mm, which means that 96.4% of the lesion diameters are within 2 pixels of the ground truth. This is a sufficiently small error for most stenoses, but could lead to large errors in percent diameter reduction in severe stenoses where the lesion diameter itself is only 2-3 pixels wide.

The error in percent stenosis is much larger than would be expected considering the small error in lesion diameter. This is evidenced by the Bland-Altman plot of percent diameter reduction in Figure 4.10. Our pipeline is biased towards under-predicting diameter reduction by 7% compared to QCA. The discrepancy between the two methods, between 34.3% and -23%, is large enough to be clinically important. The main reason for the large discrepancy is the difference in how our pipeline and QCA compute the reference diameter for the stenosis. Our method uses a simple linear interpolation between the diameters proximal and distal to the stenosis in order to approximate the healthy vessel diameter at the site of the lesion. Meanwhile, QCA uses a proprietary algorithm to draw the contour of the healthy vessel relative to the stenosed vessel in the angiogram. Since the algorithm QCA uses to determine the reference vessel diameter is unknown, it is difficult to assess its accuracy. A striking example is when QCA's interpolation of the healthy vessel contour leads to a negative percent stenosis in full-vessel QCA, which occurred in 6 out of the 399 cases. In addition, the edge detection algorithm used by QCA to delineate the vessels and approximate the reference diameter is influenced by corrections made by the operator, which introduces subjectivity. Since the true healthy lumen diameter at the stenosis location is unknown, both our pipeline and QCA are approximate methods of determining percent diameter reduction and we cannot know which is more accurate.

4.5 Conclusions and Future Directions

In this work, we have demonstrated that our fully automated pipeline for stenosis severity quantification is an interchangeable method with QCA to measure lesion diameter. The accuracy of

percent diameter reduction could be improved using a more sophisticated algorithm to estimate the reference diameter; however, the true reference diameter cannot be known since only the unobstructed portion of the lumen is visible in an angiogram. Another area for improvement is to reduce the occurrence of false negative stenosis detections by our localization neural network, which can be achieved by retraining the network on a more robust training dataset. The main benefit of this pipeline compared to QCA is that it does not require user input or trained operators to make corrections at different stages of stenosis characterization. This reduces the time for analysis from 20-30 minutes for QCA to 1-2 seconds per patient using our automated pipeline.

CHAPTER 5

3D Reconstruction of the Coronaries from 2D X-ray Angiography Images

As discussed in section 1.2, there have been many efforts to use blood flow modeling to derive computational estimates of clinically-measured functional metrics, including FFR and iFR, as well as new metrics such as QFR [35, 150, 19, 137, 71]. A key component required to perform hemodynamic computational modeling is a 3D anatomical model of the patient's coronary arteries. Since 2D X-ray angiography projections are the gold standard imaging modality for CAD, we require a method to reconstruct the 3D anatomy of the coronary arteries from these projection images.

5.1 3D Reconstruction using Projective Geometry

In order to understand the uses and limitations of projective geometry for coronary reconstruction from X-ray angiography images, several methods from recent papers were explored. Clinical angiograms introduce a lot of uncertainty to geometric reconstruction, including the angle of image acquisition, motion of the patient between images, or images taken at different phases of the cardiac cycle. In order to fairly evaluate the performance of these reconstruction methods, synthetic projections with known projection angles were used instead. This allows us to evaluate the best-case performance of these reconstruction methods and identify limitations of the methods rather than limitations due to noisy input data.

The ground truth geometry was a right coronary tree that was segmented from CCTA. Several virtual projection images were created at angles where most branches were clearly visible. Figure 5.1 shows the ground truth 3D geometry and the projection images which were used as the input for the following reconstruction methods:

1. Forward Projection
2. Back projection: Epipolar Matching
3. Point cloud method

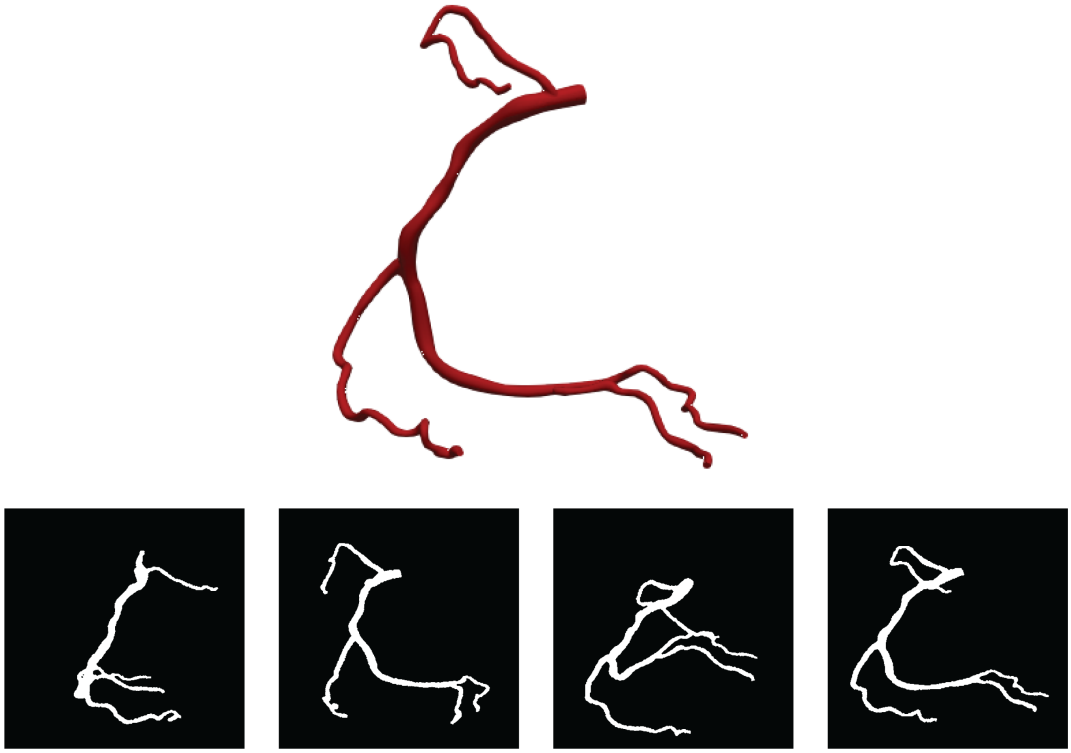


Figure 5.1: 3D right coronary tree from CCTA and its corresponding projection images. These images were used as the input for all projective geometry reconstruction methods.

The set up for the projection-based methods is shown in Figure 5.2. The position and angles of the virtual X-ray source and image planes were plotted in 3D space. The vessel coordinates on each binary projection image were then mapped onto the corresponding image plane in 3D space. Rays between the X-ray source and the 3D-mapped image coordinates were used to determine the position and shape of the 3D coronary tree.

5.1.1 Forward Projection

In forward projection, we create a virtual box around the imaging isocenter where we expect the 3D object to be. The box is sub-divided into N^3 voxels of a specified size. In this experiment, we chose to divide the virtual volume into a $150 \times 150 \times 150$ voxel box where each voxel occupied $0.5mm^3$.

The next step is to map the positions of the X-ray source and image planes in 3D space relative to the isocenter. We can then trace rays from each X-ray source that pass through each of the N^3 voxels. When the rays through a voxel land on a vessel pixel in all projection images, we can classify that voxel as part of the 3D object. We use a soft constraint to determine if a ray lands on a

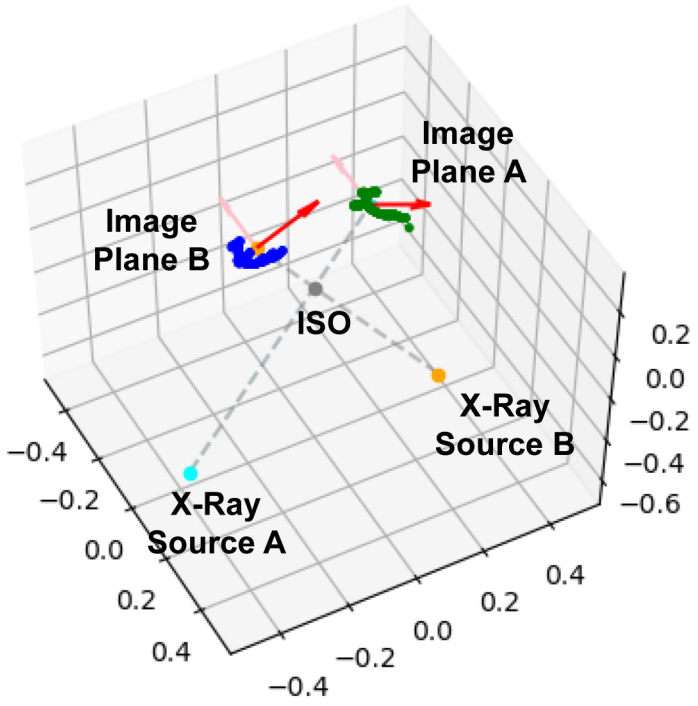


Figure 5.2: This figure demonstrates how 2D projection images were mapped into 3D space to mimic the geometry of image acquisition. The X-ray source and detector/image plane positions are given by their angle of rotation about the isocenter (ISO). The vessel pixels on each image can then be mapped into coordinates on the 3D image plane. The 3D object is assumed to be centered on the isocenter.

vessel pixel by computing the shortest distance between the point on the ray and the vessel pixels in an image. If the shortest distance is within a specified threshold given by ϵ, we conclude that the ray intersects with the vessel pixel.

The main advantage of this approach is that it is fully automated and does not require the user to identify corresponding features in the input images. Another advantage is that we are not making any assumptions about the shape of the vessels, for example radially symmetric contours, that could limit the accuracy of the reconstruction. With sufficient memory (34GB), the distance computation and voxel classification could be performed in a single step and greatly decrease runtime. Here we report the runtime on a laptop with 16GB RAM, which required looping through the N^3 voxels.

Figure 5.3 shows the results of 3D reconstruction when using two vs three input images. We can observe that using two input images provides a non-unique set of 3D points that could have generated those two projections. The true 3D geometry is within this set of points, but it is difficult to identify which branches are spurious and which could be part of a physiologically accurate coro-

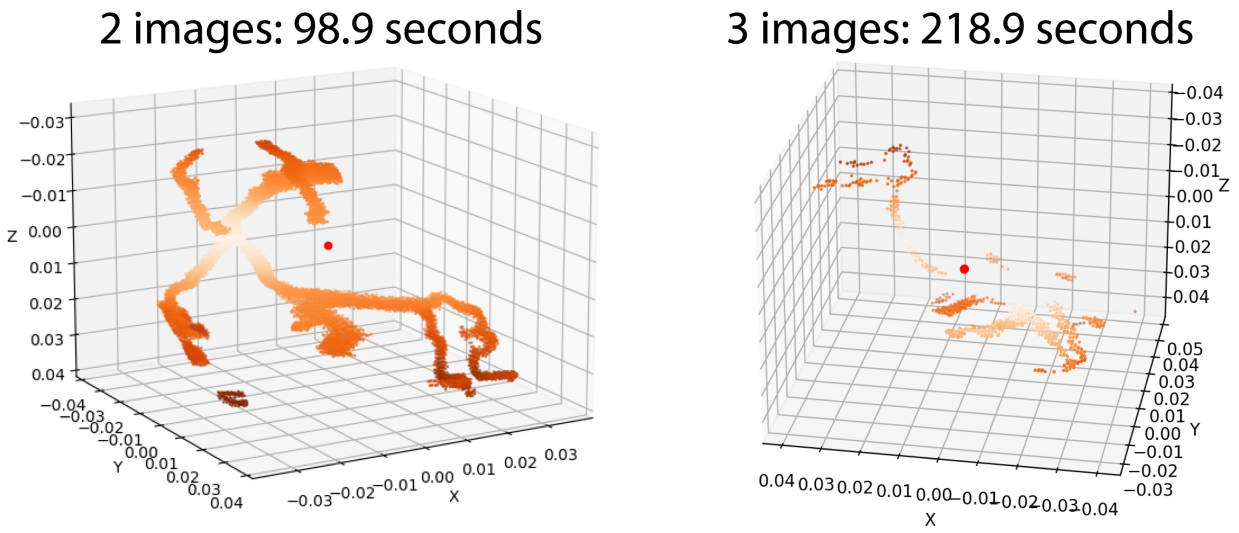


Figure 5.3: Forward projection results. When using 2 images for reconstruction, there is uncertainty about which of the many possible branches are correct. When 3 images are used, the reconstruction time increases by a factor of 2.2.

nary tree. When using three input images, the set of possible 3D points becomes more constrained, but we are still not able to fully reconstruct a unique coronary tree.

Using more input images would constrain the solution further, but even this method has its drawbacks. Due to the finite resolution of the voxels that we are projecting, it is possible that a ray passing through the center of a voxel could land just outside the boundary of the vessels in one or more images even if the voxel contained a point on the true 3D object. This would lead to a false negative classification of the 3D voxel, as seen when using 3 input images in Figure 5.3. Using a finer resolution of voxels would alleviate this issue; however, the computational cost would increase exponentially.

One way to reduce computational cost would be to only consider points on the vessel centerlines rather than all the vessel pixels in the projection images. Figure 5.4 shows an example of possible 3D centerline points constructed using forward projection onto 2 input images. Unlike the original forward projection approach, splitting up centerline and vessel surface allows us to ensure that the surface reconstruction is smooth and produces a tubular structure rather than an arbitrary shape. The main drawback of this approach is that it is difficult to split up the 3D centerline set into individual branches using clustering algorithms, and even more difficult to identify which branches are correct and which are projection artifacts.

While forward projection is fully automated, it is not a practical approach for 3D vessel reconstruction due to the computational cost and lack of unique 3D reconstructed geometries.

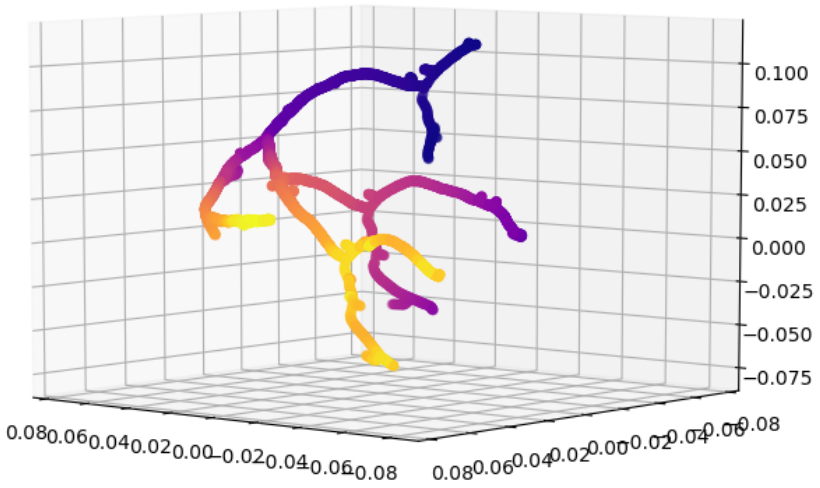


Figure 5.4: An example of forward projection of centerline points. There are several possible centerlines since we only have a limited number of projection images. The color of the points reflects their depth into the page, where the darkest points are the furthest away.

5.1.2 Epipolar Matching

An alternative to forward projection is back-projection, in which the epipolar constraint is used to determine corresponding points on the vessel centerlines. Pairs of corresponding points are then used to triangulate points on the 3D centerline as explained in section 2.4. This method requires more user input than forward projection, but the determination of corresponding point pairs ensures a unique 3D centerline. The difficulty lies in identifying which points correspond to each other. Once the 3D centerline has been determined, the vessel surface is generated by projecting the radius vectors of each 2D centerline point onto the plane normal to their corresponding 3D centerline point. The projected radius vectors can be used to create circular or elliptical contours, and the vessel surface can be created by lofting through the generated contours.

Here we present a semi-automated algorithm to determine centerline point correspondence. A skeletonization algorithm is first used to identify the 2D vessel centerlines in the input images. The skeleton is automatically broken up into segments by removing the bifurcation points, and the user identifies which segments correspond to the same vessel branch in all input images. By identifying which segments correspond to the same branch, we eliminate false correspondences between points on the same epipolar line that belong to different vessel branches. A nearest neighbor algorithm is then used to sort the centerline points from each image into an ordered list and identify the branch

endpoints.

In order to identify corresponding points, one image is chosen as the reference plane A and we consider every nth point on its vessel centerline as a reference point p_A^i. Corresponding points p_B^i are found at the intersections between the epipolar lines and the centerline in the second image plane B. When more than one intersection occurs, for example when the epipolar line crosses a C-shaped or tortuous vessel, we can use an ordering constraint to eliminate false correspondences. The intuition behind the ordering constraint is simple. If we choose reference points p_A^i in order along the centerline in A, the corresponding points p_B^i must follow the same order. Thus, we can automatically eliminate any intersections that occur behind or too far after the last identified point p_B^i. Figure 5.5 demonstrates the epipolar matching workflow.

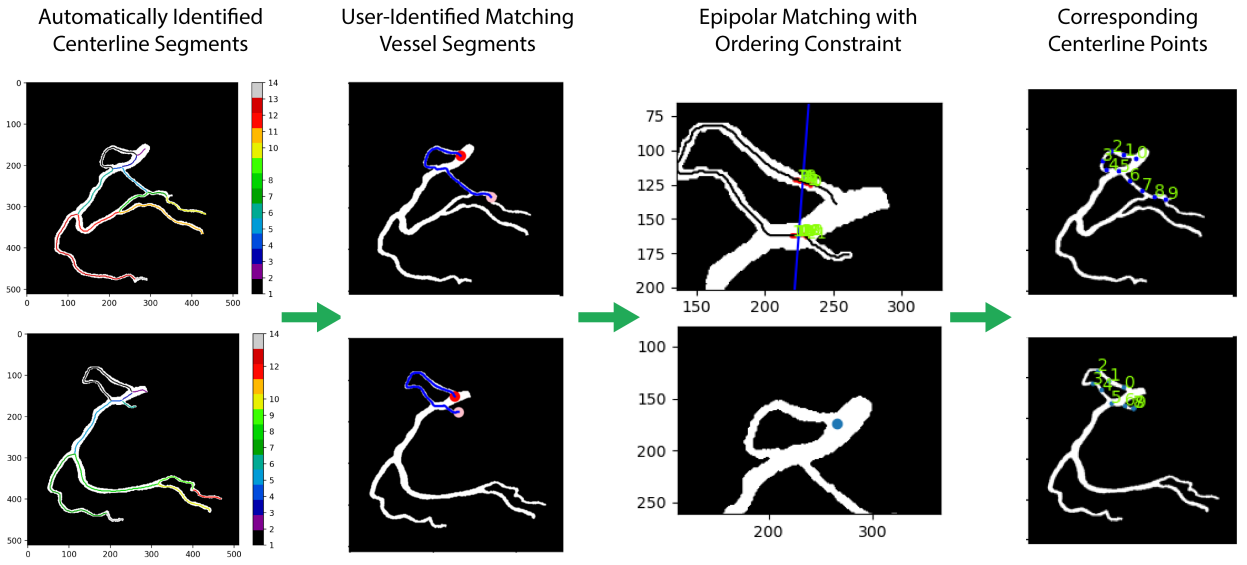


Figure 5.5: Epipolar matching method. The algorithm first automatically identifies vessel centerlines and breaks them up into segments. The user chooses which segments belong to the same vessel branch. Reference points on the branch from one image are considered sequentially, and matching points are identified using a combination of epipolar line intersections and an ordering constraint. Finally, an example set of corresponding point is shown.

One challenge of the ordering constraint is when the angle of the projection causes foreshortening or bending of the vessel in a way that the true correspondence point is indeed behind the last determined correspondence point or much farther away in one image than the other. If a mismatch due to the ordering constraint occurs, subsequent corresponding points may not be correctly identified. Foreshortening can cause other issues as well. Due to foreshortening in B, it is possible for many reference points to correspond to the same point on the second image. If foreshortening occurs in A, the epipolar line may not pass through the true corresponding point in B. We therefore include a 3-5 pixel tolerance around the intersections of the epipolar line and centerline

points to determine corresponding points. When the algorithm cannot automatically determine point correspondences due to these issues surrounding foreshortening, the user specifies the true corresponding point using their intuition and visual landmarks such as bifurcations.

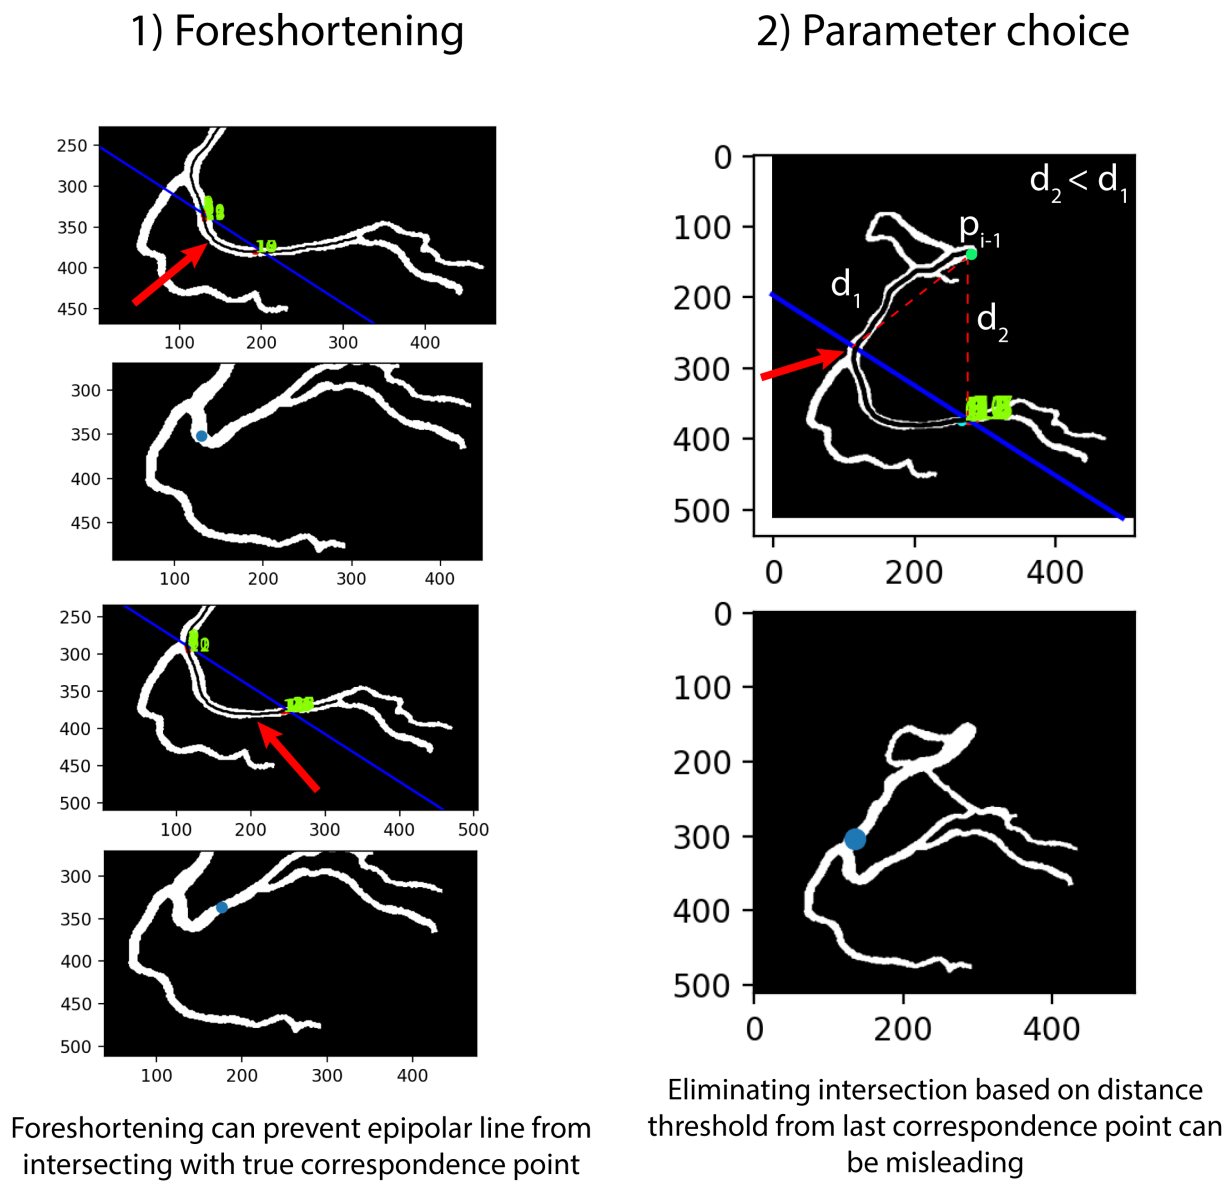


Figure 5.6: The two main challenges of this method are foreshortening in the input images and choice of algorithm parameters. Foreshortening in one or more images can lead to false correspondences using epipolar matching. Other parameters such as the number of centerline points to sample can affect accuracy in highly curved vessels. In this example, minimizing the distance between the previously identified point p_{i-1} and the two epipolar line intersections does not identify the correct intersection.

Compared to forward projection, the advantage of this method is its speed (30-60 seconds per branch, depending on the amount of user interaction required) and using prior knowledge about matching branches to avoid reconstructing spurious branches. The challenges of this method are how sensitive it is to the choice of parameters such as the spacing between reference points n, tolerance around intersection points, threshold for acceptable distance between consecutive points, and which image to use as the reference plane. The uncertainty due to foreshortening can make point matching difficult even with user input, which in itself is subjective.

5.1.3 Point Cloud Method with Ordering Constraint

In order to test a fully automated method to determine point correspondences for back-projection, we combined the work of two recent papers into one reconstruction algorithm. Like the epipolar matching algorithm described in the previous section, this algorithm first requires the user to identify corresponding branches in the segmented images.

In the point cloud approach, developed by Banerjee et al [10], all the centerline points of a vessel in one image were considered as possible correspondence points for a point in the other image. Since this method is fully automated, all centerline points could be sequentially considered as the reference point rather than sampling every nth point as in the semi-automated epipolar matching method. The error in X-ray intersection between all centerline point pairs were assessed and the top k correspondence points with minimum intersection error were considered. Following this, we eliminated false correspondences using a cost function described by Vukicevic et al [145]. The cost function was a weighted average of the distance between consecutive matching points (ordering constraint) and the reprojection error. Finally, the 3D centerline points that remained in the point cloud were ordered and fit with a b-spline to form a smoothly connected centerline. Ordering points in 3D was not a trivial task as a nearest neighbor algorithm alone was not sufficient. In order to simplify the problem, we projected centerline points onto a lower dimension using PCA and ordered the points in the lower dimensional space. The same ordering was then applied to the 3D point cloud.

This fully automated method of determining correspondence points was the fastest of the projection based methods, with a reconstruction time of 2.03 seconds per branch (not including manual specification of matching vessels). There were several challenges when implementing this approach. The first was determining a reasonable set of parameters, such as the values for X-ray intersection error threshold, k possible correspondence points, and relative weights of the ordering and reprojection error terms in the cost function. Determining possible 3D points via minimum ray intersection error implicitly encodes the same constraints as epipolar matching. Thus, if k is too small, it is possible that the points with minimum ray intersection error do not include the true

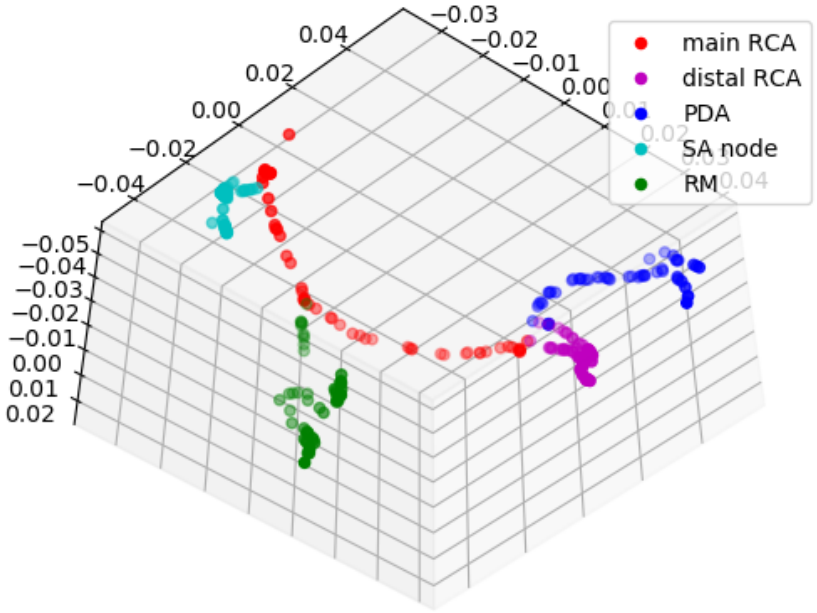


Figure 5.7: An example of a right coronary tree reconstructed using the point cloud approach. Depending on the relative orientation of different vessels in the tree, there can be gaps in the centerline due to foreshortening or regions of uncertainty with many possible 3D points corresponding to multiple epipolar line intersections.

correspondence point if there are multiple epipolar intersections. If k is too large, it becomes difficult to eliminate points based on the cost function. For example, the reprojection error in the cost function would be similar for all possible correspondence points near epipolar intersections. This method is also subject to many of the same limitations as the epipolar matching method described above. Vessel foreshortening in the input images can lead to missing correspondence points. The ordering constraint based on distance from the last determined correspondence point can also be misleading under certain conditions, particularly in C-shaped and tortuous vessels as shown in Figure 5.6.

By exploring several projection-based methods for 3D coronary reconstruction, we identified several challenges. The overarching theme is that the accuracy of these methods is heavily dependent on the input images. The angles of image acquisition affect vessel foreshortening and overlap, epipolar line intersections, and reprojection error. Even with perfect synthetic projections and known angles of image acquisition, these methods could not always reconstruct the 3D vessel centerlines from arbitrary input images. When using clinical data, the inherent uncertainty of im-

age acquisition parameters would limit the accuracy of these methods even further. Due to these limitations, we turned to a machine learning approach for 3D reconstruction.

5.2 A Machine Learning Approach for Coronary 3D Reconstruction from Uncalibrated X-ray Angiography Images

5.2.1 Introduction

Projective geometry has been the basis of 3D reconstruction methods for several decades; however, as observed in section 5.1 these methods have several limitations. Many projection-based methods rely on user input to identify corresponding vessels or features in the input images. They also either require precise calibration during image acquisition or algorithms to correct the recorded acquisition angles, since recorded values can have up to a 10% tolerance. Furthermore, projection-based methods are susceptible to how clearly visible all coronary branches are in each angiography image. Overlapping branches and foreshortening are two challenges which can introduce uncertainty into the reconstructed 3D geometry.

Machine learning has the potential to perform fully automated reconstruction without the need for calibration during imaging. A few groups have recently explored using a combination of neural networks and projective geometry for coronary reconstruction from X-ray angiography images. Bappy et al [11] developed a neural network for coronary segmentation and employed stereovision for model reconstruction. In their recent paper, Jiang et al. [56] used machine learning and projective geometry principles to segment and reconstruct coronary geometry from a combination of X-ray angiography and intravascular ultrasound (IVUS) images. To our knowledge, this work is the first purely machine learning approach for reconstruction of the coronaries using X-ray angiography alone.

While previous machine learning approaches for reconstruction encoded 3D objects as voxels, point clouds, or octrees, these methods can produce discontinuous or irregular surfaces which are unsuitable for downstream applications such as computational hemodynamic modeling. Furthermore, from a diagnostic perspective, the vessel radius and stenosis severity are the most important features of the 3D geometry. Therefore, an analytical representation of the 3D coronary geometry which explicitly includes vessel radius is required.

This work focuses on 3D reconstruction of the right coronary tree as the large anatomical variation in the left coronary arteries [15] makes reconstruction more challenging. The main contribution of this work is a multi-stage neural network which outputs an analytical representation of a 3D coronary tree given several binarized X-ray angiography projection images. By evaluating geometric reconstruction error, we demonstrate that the multistage approach is superior to a single

stage neural network (subsubsection 5.2.3.1) and traditional projection-based methods (subsubsection 5.2.3.2). We also use hemodynamic simulations to verify that reconstruction error is at an acceptable threshold to accurately model clinical quantities of interest (subsubsection 5.2.3.4).

5.2.2 Methods

In this section, we describe the training and evaluation methods of a neural network for 3D coronary reconstruction We first present a synthetic coronary tree generation algorithm that was used to create a training dataset. We then discuss our multi-stage neural network design and training strategy. Finally, we outline the assumptions and parameters of the hemodynamics simulations used to assess the accuracy of the neural network-derived 3D reconstructions.

5.2.2.1 Synthetic 3D Coronary Tree Generation

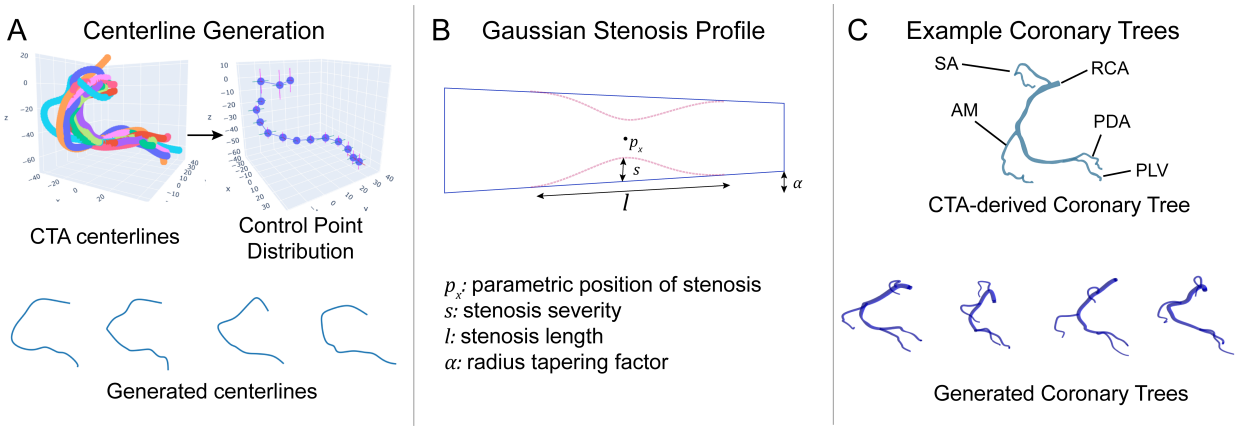


Figure 5.8: Diagrams of coronary tree generation. A: Centerlines derived from patient image data are aligned to determine the distribution of their control points. New control points are randomly sampled from this distribution to produce new centerlines. B: Schematic of how the gaussian stenosis profile is introduced in a tapering vessel C: Examples of a real patient coronary tree and synthetically generated coronary trees

In order to train a neural network, we require thousands of annotated examples; however, we only had access to a limited number of patient-specific 3D coronary geometries derived from CCTA. To address this, a synthetic vessel generator was developed to produce a sufficiently large training dataset. This generator encodes an analytical representation of each right coronary tree geometry in an $M \times N \times 4$ matrix. In this encoding, M is the number of branches in the vessel tree and N is the number of points on each vessel centerline. The final dimension corresponds to the (x, y, z) coordinates of each centerline point and its corresponding radius, r. This vessel matrix

served as the ground truth for the neural network, while projection images of the 3D coronary tree were the input for the network.

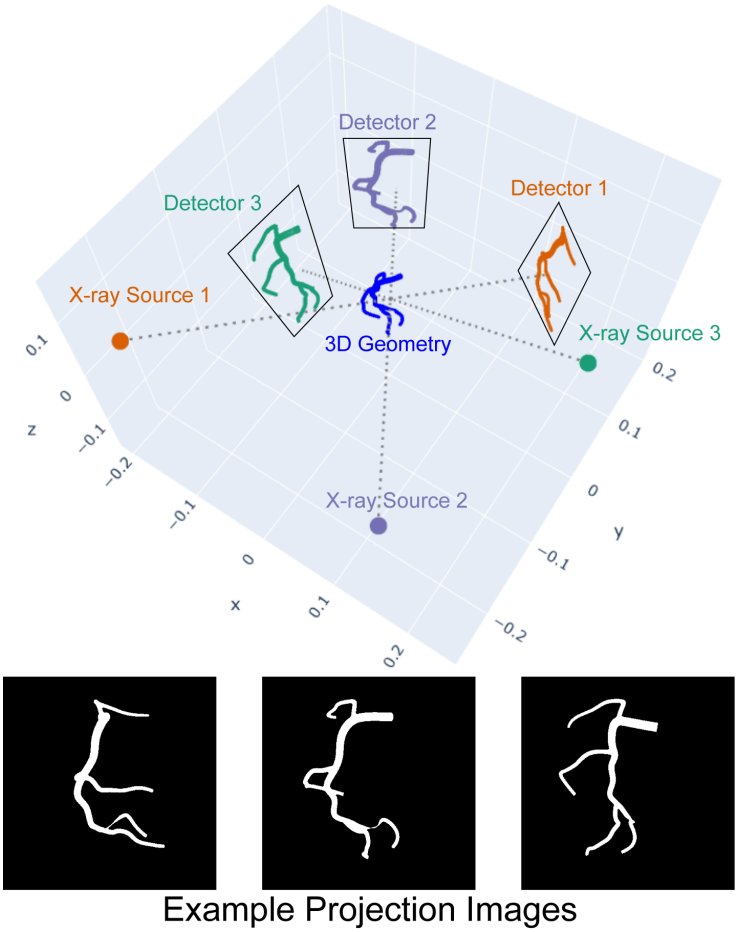


Figure 5.9: Schematic of Synthetic X-ray Projection Generation. Projections were modeled as a cone beam of X-rays from a point source hitting a detector at a series of sampled random positions.

Values derived from literature on coronary anatomy and clinical image data were used to inform the geometric modeling assumptions of the synthetic vessel tree generator. The details of generating synthetic coronary trees and their projection images are provided in Appendix A. Figure 5.8 provides an overview of vessel tree generation. In short, new vessel centerlines were randomly sampled from the distribution of patient-specific centerlines from CCTA. Linearly tapering radius was assigned to each branch, and stenoses with a gaussian profile were randomly introduced. Branches were combined into a tree and augmented with random rotation, shear, or warping. Examples of randomly generated projection images that simulate X-ray angiography acquisition are shown in Figure 5.9.

5.2.2.2 Neural Network Design and Training

The dataset for training our 3D reconstruction network was composed of 5,000 synthetically generated coronary trees. The input of the network was a set of binary projection images for each synthetic coronary tree, acquired at randomly sampled angles. A Euclidean distance transform was applied to the input projection images for mathematical smoothness. The transformed images were randomly shuffled and the first three were taken as the input for the neural network. Due to the random nature of the projection angles, 3 views were used to ensure that all branches could be clearly seen in at least 2 images. The output of the neural network was the $M \times N \times 4$ vessel matrix. The data were split into 4,500 coronary trees for training and 500 for validation (90-10 training split).

A multistage neural network was designed to output vessel centerlines and radii (Figure 5.10). The centerline and radius stages both employed a convolutional neural network backbone to learn the relevant features of coronary tree structure from the input images. In this work, we chose to use ResNet101 as the backbone network. In the centerline network, the final fully connected layer of the backbone is replaced by a 4-layer Multi-layer Perceptron (MLP) with ReLU activation and batch normalization between each layer. The output was a linear layer of dimension $M * N * 3$, which was reshaped into a vessel matrix before computing the loss. Meanwhile, the radius network replaced the final layer of the backbone with a separate 3-layer MLP for each branch. The backbone of the radius network was frozen and each vessel branch's MLP was trained separately in order to best capture the radii in the branches with stenoses. The output of each MLP was a radius vector of dimension N.

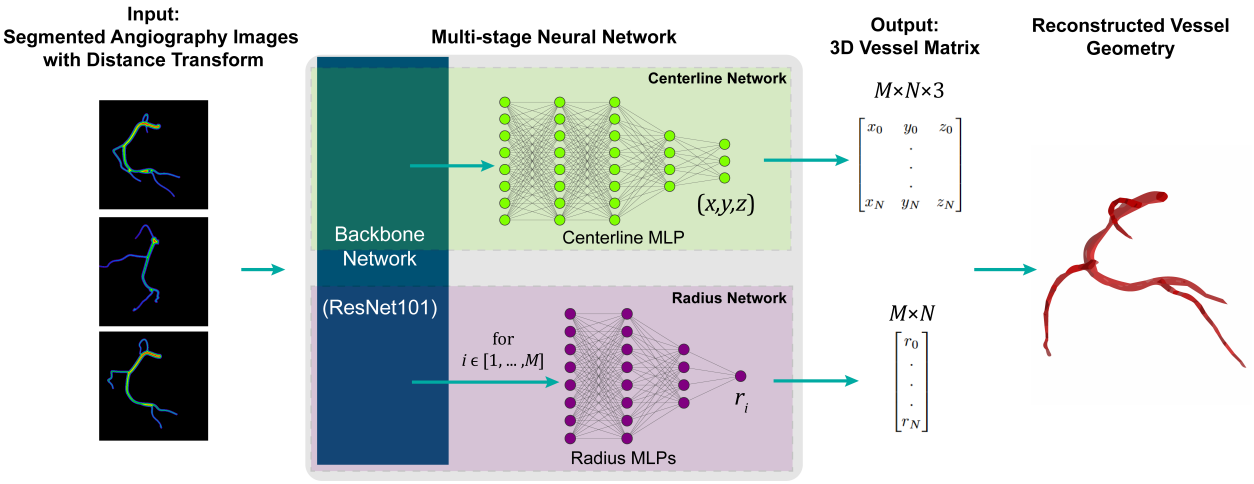


Figure 5.10: Neural network architecture. The input is binarized X-ray angiography images of the ground truth vessel geometry with a Euclidean distance transform applied for mathematical smoothness. The network is composed of a backbone network and multi-layer perceptrons which predict vessel centerlines and radii.

Both the radius and centerline stages were trained separately using mean squared error

$$MSE = \frac{1}{n}\sum_{i=0}^{n}(y_i - \hat{y}_i)^2 \tag{5.1}$$

as the loss function. The centerline network loss had an additional regularization term $\lambda \sum S_y - S_{\hat{y}}$ where λ is the regularization rate, S_y is the arclength of the ground truth branches, and $S_{\hat{y}}$ is the arclength of the predicted branches. This vessel length penalty term was included since the vessel length and radii are the most important characteristics to estimate pressure drop across a vessel, which the quantity of interest.

An ADAM optimizer with learning rate 5e-4 and weight decay (L2) regularization was used to train both stages. The batch size was 8 and networks were trained for 300 epochs. For the radius MLPs, the backbone was frozen after 300 epochs and each MLP was re-trained for an additional 50 epochs.

Training was carried out using 2 RTX2080Ti GPUs. The training time was approximately 12-14 hours per network, and inference time using a single GPU was 0.23 seconds per tree.

5.2.2.3 Hemodynamic Simulation Set-up

Computational fluid dynamics simulations were performed to compare pressure drop in ground truth and reconstructed coronary trees. The 3D geometry of each coronary tree was generated by importing the vessel centerlines and radius contours from our ground truth and predicted vessel matrices into CRIMSON [3] and lofting them to create solid 3D CAD models of each coronary tree. The 3D geometry was then discretized into an anisotropic finite element mesh composed of over 1 million elements, with smaller mesh elements in regions of stenosis. The 3D Navier Stokes equations were solved on the nodes of this mesh to derive velocity and pressure.

Since these coronary trees were synthetically generated, there were no clinical data to inform patient-specific boundary conditions; therefore, literature values [119] and allometric scaling laws [148, 143] were used instead. Identical boundary and initial conditions were applied to the ground truth and reconstructed coronary trees, namely:

1. 2x hyperemic steady flow at the inlet (0.2 L/min)
2. Resistances at the outlets, whose values were tuned using literature and allometric scaling
3. No slip at all walls
4. Initial pressure of 80mmHg

Although coronary flow is pulsatile in nature, we can model the blood flow as steady since we are interested in the average pressure drop across the coronaries under hyperemic conditions, or

conditions of maximum flow. As discussed in section 2.5, steady flow has previously been validated as sufficient to accurately estimate pressure drop across a stenosis [81]. Another assumption in our simulation is that the walls of the vessels are rigid and do not undergo any deformations. These simplifying assumptions greatly reduce computational cost. Simulations were run on 36 cores for 7-12 minutes.

5.2.3 Results

In this section, we present several methods to evaluate vessel reconstruction error. We first compared the reconstruction error of our multi-stage network against a single-stage neural network to demonstrate the advantages of the multi-stage approach. We then performed a head-to head comparison of our neural network reconstruction method and projection-based methods. Both of these comparisons were performed on single vessel geometries for simplicity. Next, we present an overview of centerline and radius reconstruction error in right coronary trees followed by an analysis of how geometric reconstruction error affects hemodynamics.

5.2.3.1 Comparison with Single-Stage Neural Network Reconstruction

The architecture of the multi-stage network was a backbone network for feature extraction followed by separate MLPs to learn the vessel centerlines and radii 5.10. The centerline MLP and radius MLPs were trained separately, and the results of both stages were concatenated to form the output vessel matrix. On the other hand, the single stage approach was composed of the backbone network and a single MLP that predicted the entire vessel matrix. The network was simultaneously trained to predict vessel centerlines and radii.

Evaluation of single- and multi-stage networks for vessel reconstruction was performed using a 10 synthetically generated RCA vessels. Stenoses between 20-90% were introduced into each of the 10 vessel geometries and the reconstruction error was assessed for each level of stenosis. An example of a vessel with 70% stenosis and its corresponding reconstructions are shown in Figure 5.11. On average, the multi-stage centerline error, expressed as root mean squared error (RMSE), was 0.83±0.29mm and the single-stage centerline RMSE was 1.73±0.42mm. The multi-stage centerline error was therefore 52% lower compared to the single-stage centerline RMSE (Figure 5.9B).

Mean absolute error (MAE) at the stenosis was used to evaluate radius reconstruction instead of RMSE in the whole vessel due to the importance of accurately predicting stenosis severity. The MAE across all levels of stenosis was 0.117±0.068mm for the multi-stage network and 0.927±0.436mm for the single stage network. The radius error was skewed for both networks, with larger error in radius as stenosis severity increases (Figure 5.9B). The single-stage network

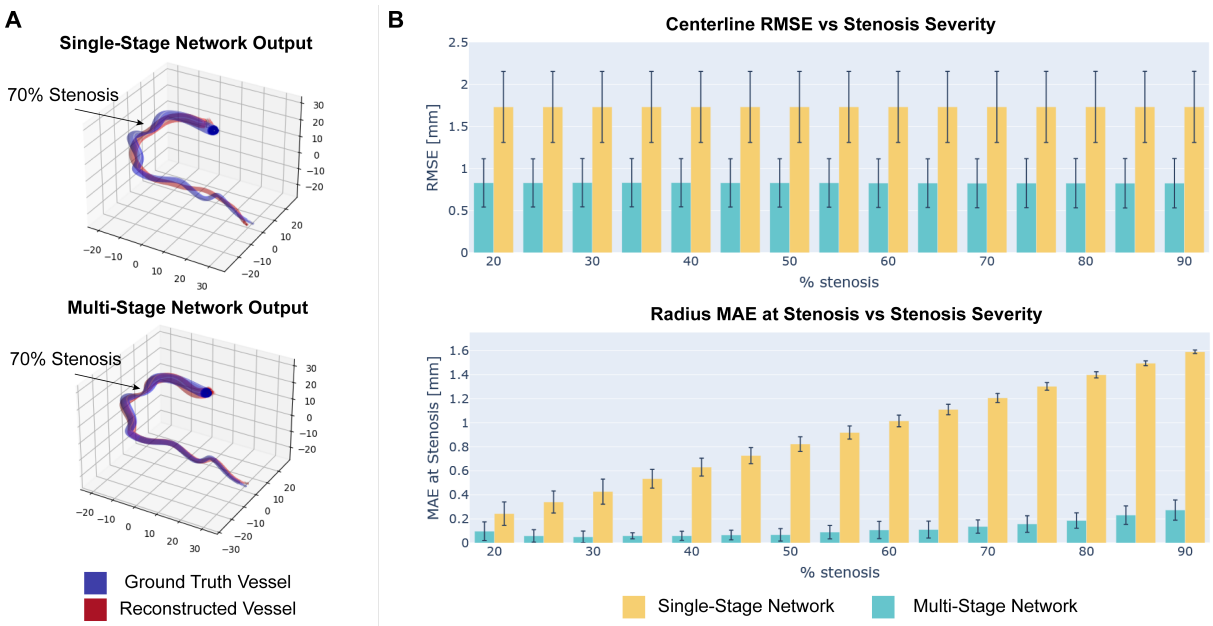


Figure 5.11: Single- and multi-stage neural network reconstruction results. 5.11A shows an example vessel with a 70% stenosis and the corresponding reconstructions. The single-stage network fails to predict the stenosis while the multi-stage network prediction does capture the stenosis and more closely follows the vessel centerline path. 5.11B shows the reconstruction error for the vessel centerlines (root mean squared error) and radii (mean absolute error at the stenosis) in 10 synthetically generated vessels, each with varying levels of stenosis from 20-90%. The multi-stage approach had a 52% lower centerline error compared to the single-stage approach across all levels of stenosis. The single-stage approach cannot predict the presence of stenosis, demonstrated by the increasing radius error for increasing stenosis severity.

was much more highly skewed because it failed to detect the stenosis at all, demonstrated by the example in Figure 5.9A. These results indicate that the advantage of a multi-stage network approach is increased accuracy in predicting vessel centerlines and radii, particularly the radius at the stenosis.

5.2.3.2 Comparison with Projection-based Methods of Reconstruction

We now compare the results of our neural network with projection-based methods in the simplest case of single vessel centerline reconstruction from 2 projection images. In this experiment, the projection-based approach was a combination of the algorithms proposed by Banerjee et al [10] and Vukicevic et al [145]. The point cloud approach was used to automatically identify up to 10 possible corresponding points for each centerline point in the reference image. The set of possible corresponding points was further refined using the reprojection error and ordering constraint cost function proposed in [145]. Matched points from both images were then back-projected and

interpolated with a b-spline to create the 3D centerline.

The comparison dataset was composed of 5 synthetic main RCA vessels, with projections taken at 15° intervals. Projection angles are given in spherical coordinates (θ,ϕ) where azimuthal angle $\theta \in [-45°, 45°]$ and elevation angle $\phi = 0°$ (Figure 5.12A). Each vessel was reconstructed using different pairs of projection images; the first image in every pair was fixed at $\theta = -45°$ and each of the subsequent projections was used as the second image of the pair. The reconstruction error as a function of the angle between projection images, $\Delta\theta$, was compared in both methods.

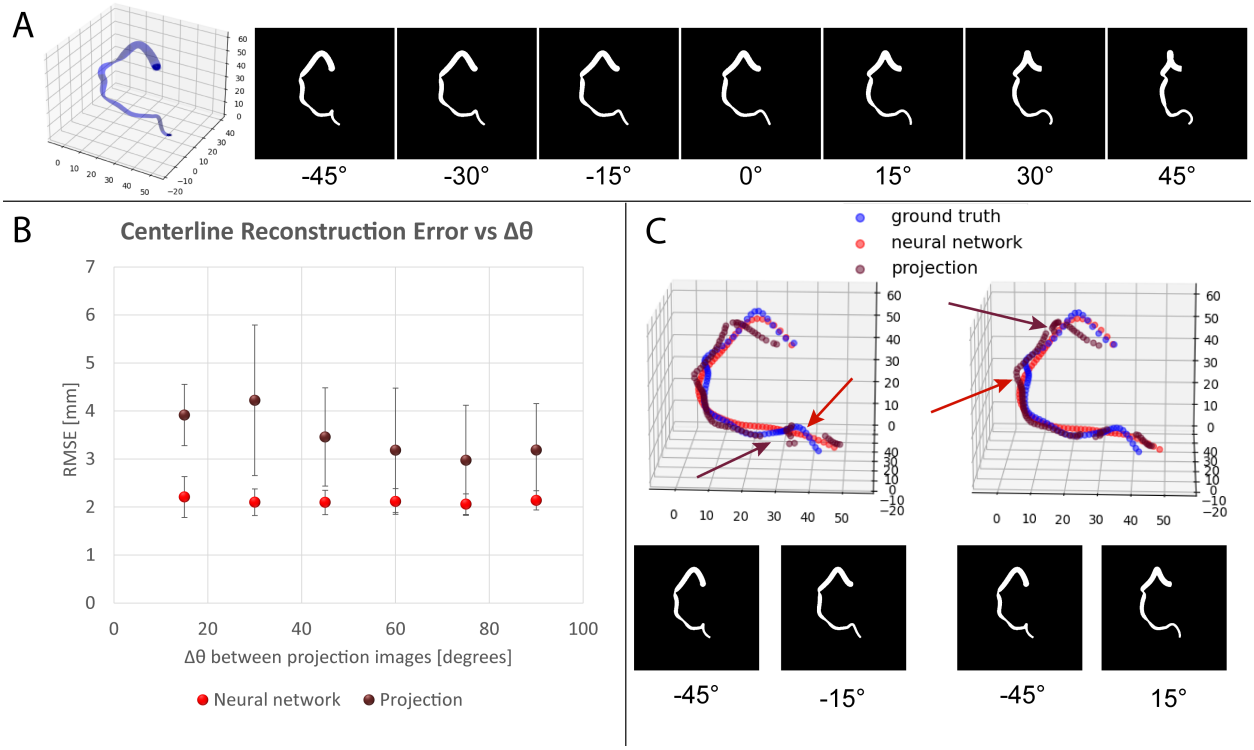


Figure 5.12: Comparison of neural network and projection-based reconstruction of a single vessel from different input images. A) Sample 3D RCA geometry and corresponding projection images. B) The neural network has on average 38% lower reconstruction error for all values of $\Delta\theta$, the angle between input projection images. C) Examples of the same 3D geometry reconstructed from different input images. Red arrows indicate regions where the neural network centerline does not follow the same bends as the ground truth while brown arrows indicate regions of uncertainty and gaps in the projection-based centerline.

As seen in Figure 5.12B, the centerline RMSE between 3D centerline points was lower in the neural network approach compared to the projection for all tested angles. The neural network approach also demonstrated lower variability: the RMSE of the neural network and projection methods were 2.12±0.05mm and 3.49±0.44mm respectively. Figure 5.12C shows two examples of the same vessel reconstructed using input images from different angles. The ground truth centerline is in blue, the neural network centerline is in red, and the projection-based centerline is in brown.

The neural network centerline followed the average curve of the ground truth centerline, but did not capture all the bends in the centerline path (red arrows). Meanwhile, the brown arrows indicate gaps or regions of uncertainty in the centerline reconstructed using projection-based methods.

5.2.3.3 Coronary Tree Reconstruction

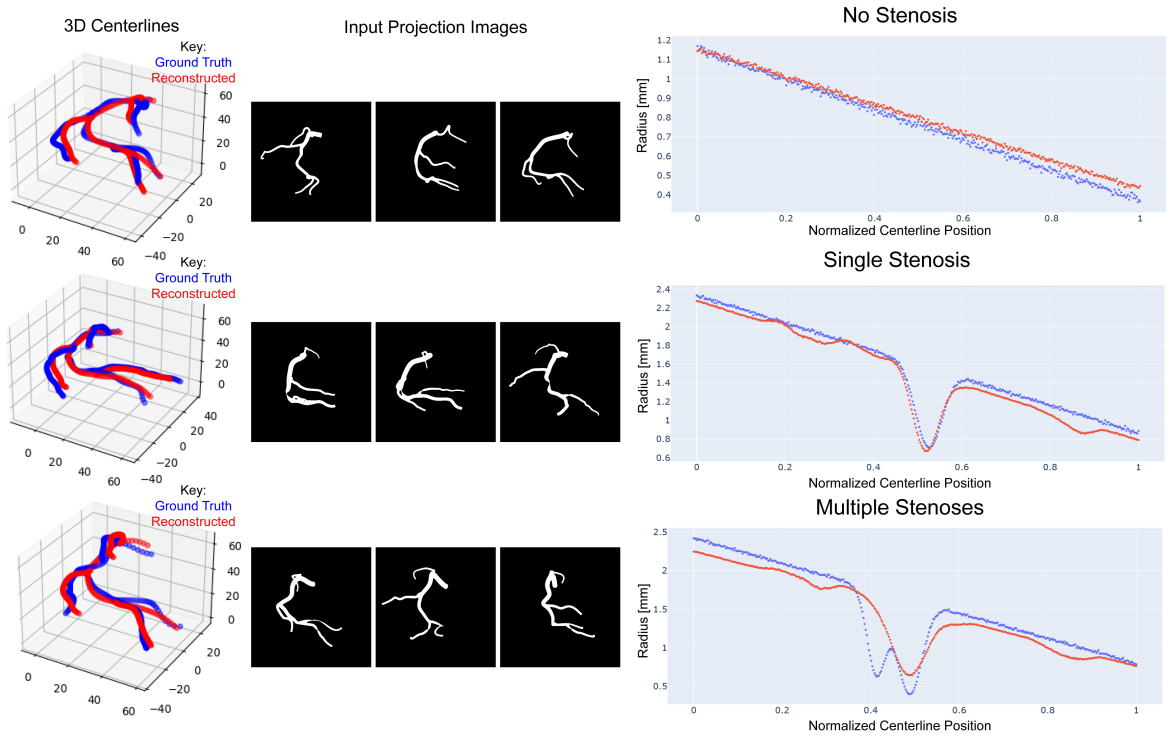


Figure 5.13: Left: Examples ground truth centerlines (blue), reconstructed centerlines (red), and their corresponding input images. As in the single vessel case, the neural network predictions follow the average path of all vessels but do not have the same tortuosity as the ground truth centerlines. Right: Examples of radius along the centerline in vessels with varying numbers of stenoses. When no stenosis is present, the neural network captures the linear tapering of the vessel well. The neural network is able to accurately predict the severity and location of a single stenosis, but in the multi-stenosis case, the neural network approximates serial stenoses as a single stenosis with an error of 0.3mm in stenosis severity.

In this section, we present neural network reconstruction accuracy of the right coronary tree. Figure 5.13 shows several examples of reconstructed coronary tree centerlines and their corresponding ground truth centerlines. We observed that the neural network learned the mean centerline path of each vessel in the tree; however, as in the single vessel case, the predicted vessels were not as tortuous as the ground truth. The RMSE between the ground truth and predicted centerline points in the validation set was 2.57±0.78mm. The mean absolute error in vessel length

was 8.83±4.81mm. This error in vessel length was 47% lower than the error of the same network trained without length regularization in the loss function.

We now consider the error in vessel radius along the centerline for all branches. The RMSE of the vessel radius was 0.16±0.07mm, which corresponds to sub-pixel resolution. The error in radius was higher in regions of stenosis (MAE = 0.27±0.18mm), particularly severe stenoses greater than 70% diameter reduction. This aligns with the trend seen in the single vessel case, where MAE at the stenosis was larger for more severe stenoses (Figure 5.11B). Although there was sometimes a shift in the location of the stenosis within the branch, the network predicted diameter reduction at the stenosis within approximately 1 pixel of the ground truth (Figure 5.13). Stenosis severity was less accurately predicted when serial stenoses were present, as the network approximated them as a single, longer stenosis.

5.2.3.4 Effect of reconstruction error on hemodynamics

Our measure of success for coronary reconstruction is whether the reconstruction error of our neural network low enough to accurately model pressure drop across a stenosis, which allows us to gauge disease severity. Hemodynamic simulations were performed in one example of a healthy coronary tree without stenosis, and one example with stenoses. The results are shown in Figure 5.14.

In the healthy coronary tree, the pressure drop ΔP across the RCA in the reconstructed geometry was 1.4mmHg lower than the pressure drop in the ground truth. In the stenotic coronary tree, the error in ΔP was 4.5mmHg. When considering only the pressure drop across the stenosis, the error was 3mmHg. In addition to the error in ΔP, the pressures in the reconstructed geometry were approximately 10mmHg lower than the pressures in the ground truth.

The larger pressure errors in the stenosis case compared to the healthy case can be attributed to the radius reconstruction error at the stenosis. The ground truth stenotic geometry had serial 51% and 57% stenoses in the mid-RCA. The neural network reconstruction approximated this as a longer 53% stenosis, corresponding to a radius error of 0.07mm. The ratio of pressures on either side of the stenosis, or FFR, was 0.886 in the ground truth and 0.895 in the reconstructed vessel. Despite the 10mmHg difference in inlet pressure, there was only a 1% error in the clinical quantity of interest.

5.2.4 Discussion

5.2.4.1 Comparison with Single-Stage Neural Network Reconstruction

The results of the comparison between single- and multi-stage neural networks demonstrate that separating the tasks of centerline and radius reconstruction greatly improves reconstruction accu-

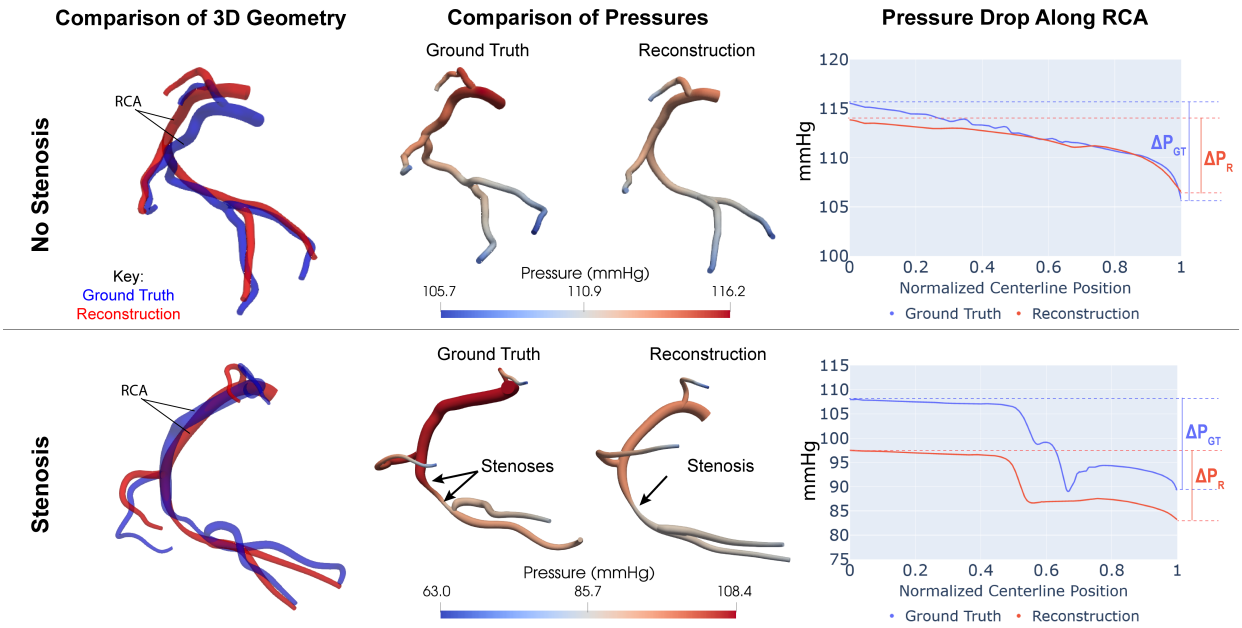


Figure 5.14: Comparison of pressures in ground truth and reconstructed coronary trees. When no stenosis is present, the error in pressure drop ΔP along the main RCA vessel is small (1.4mmHg) despite the error in geometric reconstruction. In the stenosis case, sub-pixel errors in stenosis severity (0.07mm) can lead to a larger error in pressure drop (4.5mmHg) along the vessel.

racy, particularly at the stenosis. The single-stage network was able to accurately predict the trend of linearly decreasing radius along the vessel but could not predict stenoses. The stenosis radii made up less than 2% of the $N \times 4$ vessel matrix. It is therefore unsurprising that the stenoses had such a small contribution to the loss function that the network did not learn to predict them. In the multi-stage approach, the stenosis radii had a much larger contribution to the loss function since the mean squared error of the radii was optimized separately from the centerlines. The network was therefore able to learn and accurately predict stenosis severity. This observation motivated going one step further for the coronary trees and using a separate MLP to learn the radii for each vessel. This was necessary since not all vessel branches contained stenoses; if the network was trained to predict the radii of all branches at the same time, the stenoses would once again make up less than 2% of the $M \times N$ output matrix. The challenge of predicting stenoses is that they are the most important diagnostic feature, but only make up a small proportion of the vessel matrix. The multi-stage approach described in this work amplifies the stenoses and improves prediction accuracy.

5.2.4.2 Comparison with Projection-based Methods of Reconstruction

In this experiment, we observed that the neural network RMSE was on average 39% lower than the projection method RMSE across all values of $\Delta\theta$, even though the neural network was not trained on images at these projection angles. Together with the lower standard deviation, this could indicate that the neural network is more robust to different characteristics of the input images. For example, each of the 5 RCA vessels may have had more foreshortening or overlap in projections taken at different angles. While the projection method reconstructions are less accurate when the input images have high levels of foreshortening or overlap, the neural network could learn universal characteristics of the RCA vessel shape to overcome these challenges and maintain similar accuracy for all input images. The theory that the neural network is more robust to foreshortening and overlap in the input images is supported by the qualitative examples in Figure 5.9. The brown arrows, which point to gaps or regions of uncertainty in the projection method centerline, correspond to regions on the input image where the vessel is highly foreshortened or bending at an angle that obscures other parts of the vessel. The red neural network centerline interpolates the 3D centerline in these regions, resulting in a continuous centerline.

The main limitation of the neural network compared to projection methods is that it does not capture all the bends in the vessel centerline path despite accurately following the average path. This could be a consequence of the network learning the features of a typical 3D RCA shape. While these features ensure that the network can interpolate the 3D centerline even when the input images are challenging, they also may cause the network to predict similar centerline paths for all inputs. Meanwhile, the projection-based method can accurately capture the details specific to each centerline path using the features of the input images. The neural network was trained using mean squared error loss, which only accounts for point-wise differences between the ground truth and predicted centerline. A loss function that incorporates more global shape information or reprojection error could improve the centerline path accuracy.

5.2.4.3 Coronary Tree Reconstruction

As seen in our single vessel reconstruction experiment, the neural network predictions once again captured the average path of all branches in the coronary tree but did not match their tortuosity. This issue is analogous to a common problem in curve-fitting regression tasks: when the network is not large enough to capture the complexity of a dataset, the predicted curves match the average behavior of the data. We hypothesized that increasing the complexity of the neural network by increasing the size of the MLP may alleviate this issue. We tested networks with 128, 256, 512, 1024, 2058, and 4096 neurons in each of 1-4 hidden layers. From this scaling study, we observed 1% decrease in RMSE between the 4096 $\times$ 4 network and 128$\times$1 network. Therefore, the com-

putational overhead of increasing the MLP size did not justify the resulting decrease in error. As discussed in subsubsection 5.2.4.2, a loss function which incorporates global shape may be a more effective strategy to improve centerline accuracy.

Another challenge of coronary tree reconstruction compared to single vessel reconstruction was accurately predicting radius at the stenosis. As described in subsubsection 5.2.4.1, employing a multi-stage approach and training a separate MLP for each vessel branch enabled the network to learn how to predict stenoses. Despite this, the MAE at the stenosis was higher for coronary trees (0.27mm) compared to a single vessel (0.12mm). This may be due to the random projection angles, as overlapping branches may have obscured the stenosis in some of the input images.

5.2.4.4 Effect of reconstruction error on hemodynamics

Our hemodynamics simulations demonstrate that the geometric reconstruction error of our neural network is small enough to accurately model pressure drop across the vessel of interest. Although the neural network centerlines were not as tortuous as the ground truth, this did not greatly affect the pressure since the error in vessel length was sufficiently small. The centerline path accuracy would be more critical to estimate local hemodynamic parameters such as wall shear stress. In the future, improvements to the neural network centerline could enable modeling and estimation of more hemodynamic quantities.

The error in radius had a much larger impact on the pressure drop, particularly error in stenosis radius. As seen in the stenosis example, the neural network approximated two serial stenoses as a single stenosis with their average severity. Despite only a 0.07mm difference in stenosis radius, the difference in pressure drop across the stenosis was 3mmHg. Since FFR is computed using ratios of pressures, this error in pressure drop was relatively small and only led to a 1% error in FFR. We believe that accurately capturing the vessel lengths and radii at the stenosis, as we have done in this work, contributed to the accurate estimation of FFR. At higher levels of stenosis, however, small errors in radius can lead to much higher differences in pressure drop.

This highlights the importance of accurately determining vessel radius at the stenosis. The main limitation of improving radius estimation is the resolution of the X-ray angiography image. In an angiogram, the difference between a 70% and 75% stenosis cannot be distinguished; the projections of both levels of stenosis are identical due to the pixel resolution. It is therefore difficult to differentiate between stenoses at a higher resolution even with a deep neural network.

5.2.4.5 Future Work

In addition to improving the accuracy of the neural network centerline predictions, future work will focus on increasing the scope of input images to match real-world data more closely. In terms of

synthetic data generation, we will expand our generator to incorporate anomalies and variations in the number and types of branches in the right and left coronary trees. Besides synthetic projections, this neural network can easily be applied to clinical X-ray angiograms that have been processed with a segmentation algorithm to create binary input images. AngioNet, described in chapter 3, can facilitate this.

In the catheterization lab, it is common practice to move the patient table during image acquisition to follow the flow of injected dye. Since the motion is controlled via joystick and is not often recorded, these images are extremely challenging to use as inputs for projection-based reconstruction methods because they require precise knowledge of image acquisition geometry. Neural networks could potentially reconstruct the coronary tree despite unknown motion of the 3D vessel tree between projections.

In conclusion, this work is a proof-of-concept that neural networks can be used for coronary reconstruction from uncalibrated X-ray angiography images, with sub-pixel resolution of vessel radius. We have demonstrated that reconstruction error is at an acceptable level to accurately model hemodynamic quantities such as pressure drop across a vessel.

CHAPTER 6

Relative Flow Extraction from Angiographic Time Series

6.1 Introduction

In Chapters 3, 4, and 5 we presented several methods that can be combined to produce a 3D geometric model of a patient's coronary anatomy. In order to perform hemodynamic simulations in these geometries and assess functional stenosis severity, the next step is to determine patient-specific boundary conditions.

During cardiac catheterization, the patient's mean arterial pressure and cardiac output are measured, and these can be used to set the inlet pressure or flow rate. The outlet boundary conditions are more difficult to determine, as it is not practical to measure pressure or flow in each of the branches in the coronary tree. Some researchers have proposed using transthoracic Doppler echocardiography to non-invasively determine coronary flow in the major branches [54, 84]. This approach has been validated against invasive guidewire flow measurements and widely cited as a promising method to determine coronary flow and CFR. One drawback of this method is that certain vessels are difficult to access with the Doppler probe due to their position on the heart. This can limit the accuracy or even prevent measurement of flow in certain portions of some vessels. Another drawback is that Doppler echocardiography is not a part of the standard clinical workflow and has still not been widely adopted despite the success of a landmark study in 2008 [84].

We are interested in a method that can extract boundary conditions from an X-ray angiography acquisition without the need for additional measurements or procedures. To achieve this, we aim to exploit the implicit information about coronary flow from an angiographic time series. A series of X-ray angiography images can provide information about the relative flow in each branch of the coronary tree, particularly when a flow limiting stenosis is present. The TIMI (Thrombolysis in Myocardial Infarction) frame counting method [42, 66] is a technique used by cardiologists to classify coronary flow into one of 4 categories:

1. TIMI grade 0: complete occlusion
2. TIMI grade 1: severe epicardial disease and/or microvascular disease, where flow occurs past an obstruction but the distal vasculature is not perfused
3. TIMI grade 2: flow occurs past the obstruction with partial or slow perfusion of the distal microvasculature
4. TIMI grade 3: healthy coronary flow

Vessels are classified into one of these categories by counting the number of frames it takes for dye to travel between standardized vessel landmarks. Each category is assigned a range of frame counts, normalized by acquisition frame rate and the relative length of each vessel. While TIMI-frame counting can standardize coronary flow assessment and provide insight into differences in relative flow due to disease, it does not quantify the flow rate and cannot be used to inform boundary conditions. Still, the TIMI frame counting technique has inspired several computational algorithms that quantify mean coronary flow from X-ray angiography series.

A review of algorithms used to estimate coronary flow from the movement of contrast agent is given in [130], which covers densinometric algorithms developed prior to 2000. In these approaches, the bolus of contrast agent was tracked by measuring the density of contrast along the vessel identifying the time of peak intensity as the bolus crosses a specified landmark. The time-density curve could be integrated to identify the "plug" flow of the contrast agent at different landmarks. Variations of this approach tracked total or partial opacification to identify bolus position, used cross correlation of time-density curves to identify bolus transport times, or measured instantaneous axial velocity instead of flow. The limitations of these early algorithms were that they couldn't account for the pulsatile nature of blood flow or irregular/asymmetric vessel shapes.

In 2012, Molloi and colleagues [88] developed a densinometric approach to estimate coronary flow. Volume was calculated using an equation relating integrated contrast density and mass in the region of interest, injected contrast agent concentration, and a correction factor for magnification of vessels at different depths. Change in volume over consecutive frames as used to determine coronary flow. The advantage of this method was that it does not make assumptions about coronary vessel shape. Improvements in flat-panel detector design also removed a lot of the distortion and signal decay which negatively affected early densinometric methods. Like previous methods, pulsatile flow and motion compensation were not accounted for. Another drawback was that vessel phantoms of known diameter were required to calibrate and determine correction factors for vessel volume.

More recently, groups have developed more sophisticated algorithms that use motion compensation algorithms and image segmentation to track the progression of dye through the coronaries. In their approach, Sokolov et al. [133] used edge-detection filters to segment only the vessel containing a stenosis via user specification of its centerline. Cross-correlation was used to align frames

from different parts of the cardiac cycle to compensate for coronary motion during the cardiac cycle. The dye front was tracked along the centerline to determine the path length S traveled between co-registered frames, and axial velocity was estimated as $\frac{\Delta S}{\Delta t}$. Meanwhile, Khanmohammadi et al. [62] segmented all vessels in the coronary tree using unsharp mask filtering followed by the Frangi vesselness filter[36], where filter parameters were iteratively tuned for each vessel. Non-rigid registration based on affine transformation and B-spline free-form deformation were used to co-register consecutive frames from different parts of the cardiac cycle. This method assumed constant radius for each section of a vessel, where a section was defined as the vessel segment before a bifurcation. Using this assumption, velocity could be estimated using the equation $v = \frac{\sum L_i}{T_i}$ where L_i is the length of each vessel section and T_i is the transit time per section. A plot of contrast agent area over time was fit with a 7-degree polynomial and used to identify which frames of the series were optimal for tracking the filling of the coronaries.

We aim to build on the segmentation-based approach for coronary flow estimation in order to measure relative flow in each branch, which can be used to inform boundary conditions for a hemodynamic simulation. Our goal is to extract relative flow in each coronary branch rather than the true flow since estimating the 3D distance the dye has traveled from 2D projection images is challenging. Previously published algorithms demonstrated that angiographic flow extraction algorithms estimated slightly different values of flow than doppler-measured flow, especially when flow extraction was repeated using different projections, but the general trends were preserved.

Here, we present a relative flow estimation algorithm that tracks change in segmented vessel area over time. The goal of this chapter is to demonstrate the feasibility of this algorithm by validating its performance against data derived from a ground truth hemodynamics simulation.

6.2 Methods

6.2.1 Synthetic Angiography Series Generation

Clinical X-ray angiography series have many characteristics that make flow estimation difficult, including nonrigid deformation and motion of the coronaries. We can test the feasibility of our flow estimation algorithm in a simpler case by generating synthetic angiography series instead. These synthetic angiography series have the advantages of a stationary coronary tree, thus doing away with the need for non-rigid registration algorithms, and a higher frame rate than clinical angiograms, which enables more accurate estimation of how far the dye has traveled between frames.

The first step in creating synthetic angiography data was to set up a simulation of dye transport in a patient-specific coronary tree. In this case, a right coronary tree was segmented from CCTA,

then blended and meshed using CRIMSON [4]. The flow initial and boundary conditions of the simulation were as follows:

1. steady 0.1 L/min inflow
2. outlet resistances based on allometric scaling laws and literature values [143, 119]
3. no slip at all walls
4. initial pressure of 80mmHg

Once the flow simulation reached steady state, at approximately 0.5s of physical time, a scalar transport problem was introduced. The concentration of dye at the inlet was fixed at 200 mol/mm^3, and dye was allowed to wash out of the 3D domain at the outlets. The walls of the vessel were considered to be impermeable to prevent the loss of dye via flux through the walls. The diffusion coefficient was 2.03e-3 mm^2/s. This was an advection-dominated transport problem, so the contribution of diffusion to the dye bolus transport was minimal. This dye transport simulation allowed us to model the flow of dye until the entire coronary tree was filled.

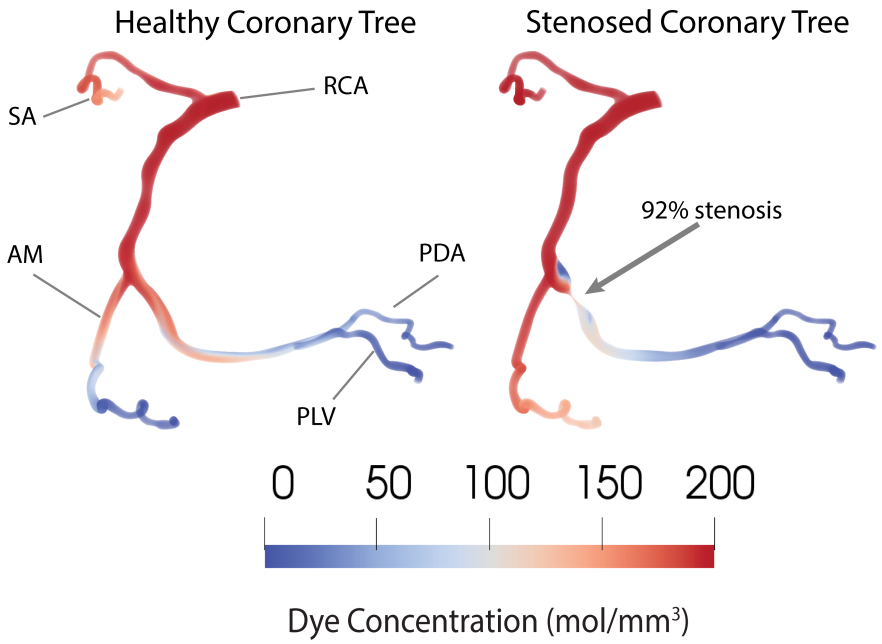


Figure 6.1: A snapshot of dye transport in a healthy and stenosed coronary tree, where dye concentration is highest in the red regions. At the same point in time, the dye has traveled farther down the RCA in the healthy coronary tree compared to the diseased case due to the flow limiting severe stenosis in the mid RCA.

Figure 6.1 shows a snapshot of dye transport in the same coronary tree under two different conditions: one healthy tree, and one with a 92% flow limiting stenosis in the mid-RCA. At the same time step, we observe that the dye has traveled farther along the RCA in the healthy tree

compared to the tree with stenosis. In the diseased coronary tree, the contrast agent has progressed farther in the proximal branches, the SA and AM, due to the flow-limiting stenosis.

Once the simulation of dye transport was complete, the concentration of dye was thresholded at 190 mol/mm^3 to identify how far the bolus of dye had traveled in each time step. Thresholded points on the RCA geometry were projected to create synthetic binary angiography images of dye transport (Figure 6.2). Projections were acquired every 0.02s of physical time for a total of 35 frames in which the coronaries were filling with contrast agent. The angles of the projections were chosen such that all the outlet branches (SA, AM, PDA, and PLV) were clearly visible with minimal foreshortening.

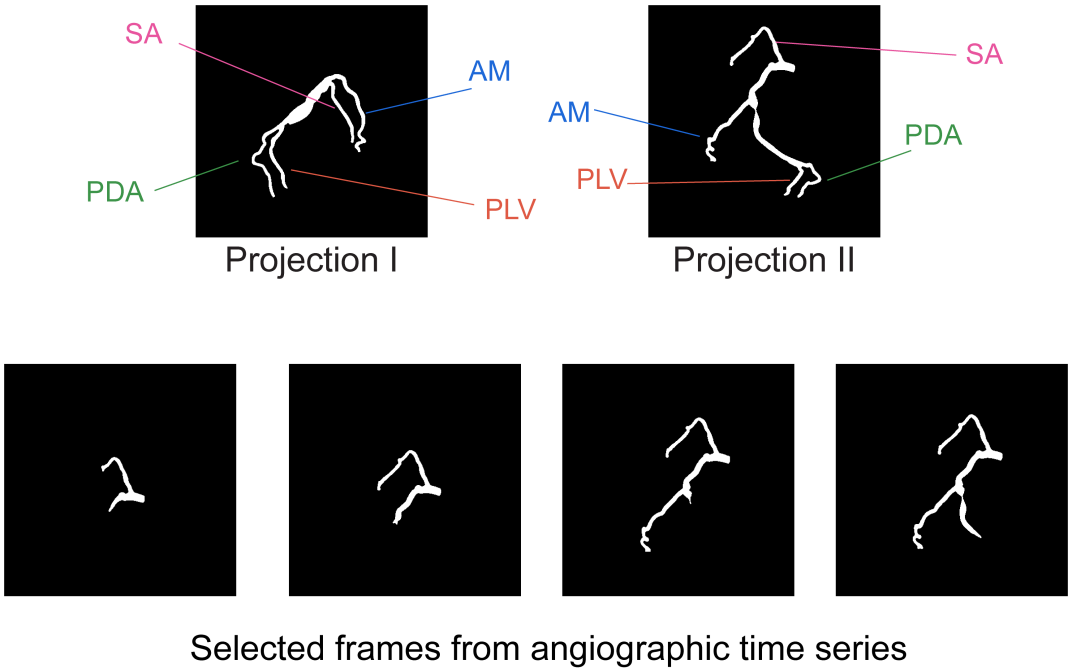


Figure 6.2: Examples of synthetic X-ray angiography projections acquired at different angles by projecting regions containing dye concentration above 190 mol/mm^3. Selected frames from the time series of projection angle II are shown below.

6.2.2 Relative Flow Estimation Algorithm

An overview of our flow extraction algorithm is given in Figure 6.3. The first step is to automatically skeletonize the series of binary images to identify vessel centerlines, followed by user-specification of the vessel of interest in all frames. This step is shown in more detail in Figure 5.5. A Euclidean distance transform is applied to the images and the values at the centerline points are sampled to determine the vessel radius in pixels. The area of a single vessel can then be approximated as the sum of diameters along the centerlines.

The change in area between consecutive frames, $dA(t)/dt$, was determined by taking the derivative of vessel area over time $A(t)$. The derivative was calculated using the first differences approximation. We could then plot $dA(t)/dt$ over time and divide these values by mean vessel radius to get the mean axial velocity over time. Next, we multiplied the axial velocity by mean cross sectional area to calculate the flow rate at each time step. In this example, the mean flow for each vessel was calculated using the time-averaged axial velocity since we our inlet flow was steady; however, this method can also be used to track the flow rate over time under pulsatile conditions. Finally, we divided the flow computed at each outlet by the inlet flow to calculate the relative flow distribution. This process was repeated for both series of angiographic projections and the relative flow was averaged across both projections. The purpose of taking the average over multiple projections was to reduce miscalculations in vessel length — and thus axial velocity — due to foreshortening in any given projection.

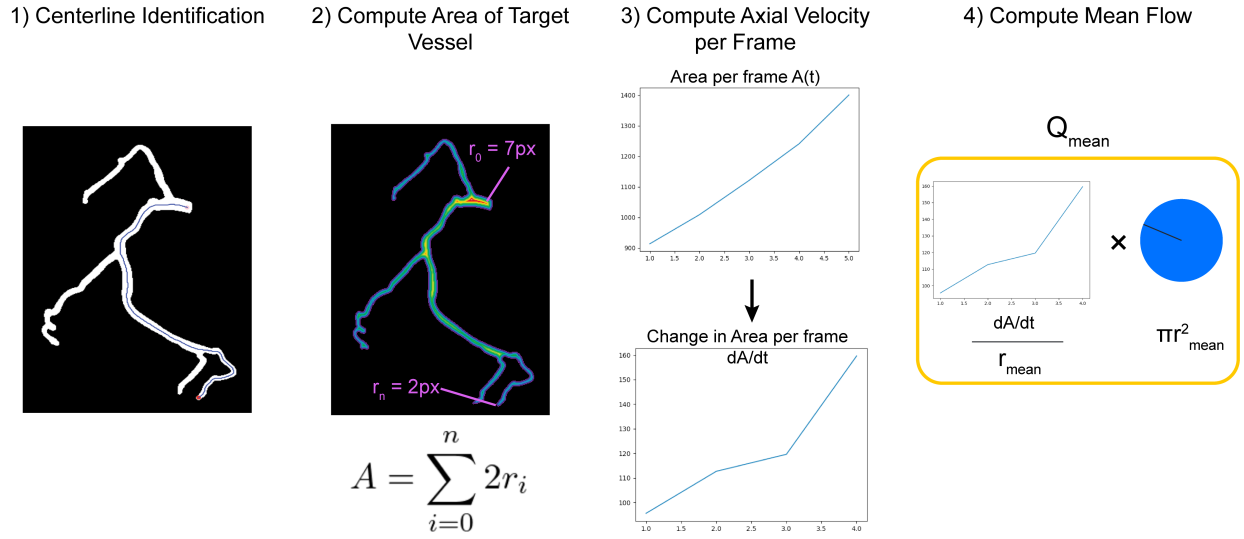


Figure 6.3: Overview of flow extraction algorithm. The centerline of the vessel of interest is first specified. Radii along the centerline are used to compute the area of the target vessel. This process is repeated for all frames, and a plot of area per frame is generated. The first differences method is used to find the derivative of the graph, which can be normalized by mean vessel radius to get axial velocity. The velocity is then multiplied by cross sectional area to compute the flow.

6.3 Results and Discussion

Figure 6.4 shows the calculated flow from each projection, color coded by vessel branch. The solid lines are the flows calculated in the healthy geometry, while the dashed lines are the flows calculated in the coronary tree with stenosis. The rate of filling for individual branches in the

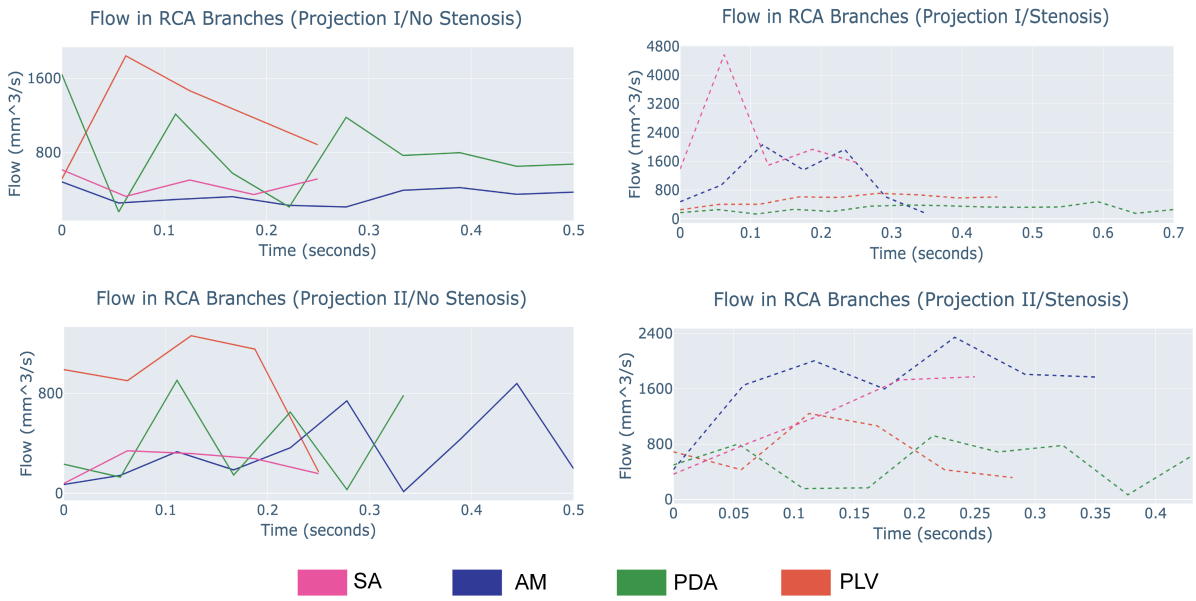


Figure 6.4: The computed flow between frames is shown for each vessel branch in the healthy coronary tree (solid lines) and stenosed coronary tree (dashed lines). For each coronary tree, the flow computed using each of the 2 projections is shown. Although the values of flow in both projections differ, the average values of flow in each branch are similar.

tree ranged from 5 to 10 frames, depending on their relative length in the projection image. As observed by Sokolov [133], the calculated values of flow in each branch are not consistent across different projections; we also observe that the calculated flow is not constant although the inflow waveform was steady. These observations both stem from the fact that performing measurements on a 2D projection cannot accurately capture path length throughout a tortuous 3D vessel.

Figure 5.8 demonstrates how the calculated relative flow in each branch compares to the ground truth simulation. The ground truth simulation results are depicted using solid bars. As expected, when comparing the relative flow in the healthy case (blue) and stenosed case (red), more flow is diverted to the branches proximal to the stenosis in the diseased coronary tree. Specifically, the SA and AM branches experience 17.1% and 17.7% increases in flow in the stenotic tree. Consequently, the distal branches experience a reduction in flow compared to the healthy coronary tree, namely 15.1% in the PDA and 19.8% in the PLV. We observe the same trend in the computed results, shown as dashed bars. Based on our calculated relative flow in each branch, the SA and AM branches experienced a 26.9% and 17.1% increase in flow, while the flow in the PLV and PDA decreased by 19.6% and 22.4% respectively. Our algorithm could not recover the exact changes in flow per branch between the healthy and stenosed case; however, trends were preserved and the algorithm could distinguish between patients with and without epicardial disease, at least when the stenosis was severe enough to be flow-limiting. On average, the error $\hat{Q}_{simulation} - \hat{Q}_{computed}$ was 5.15%

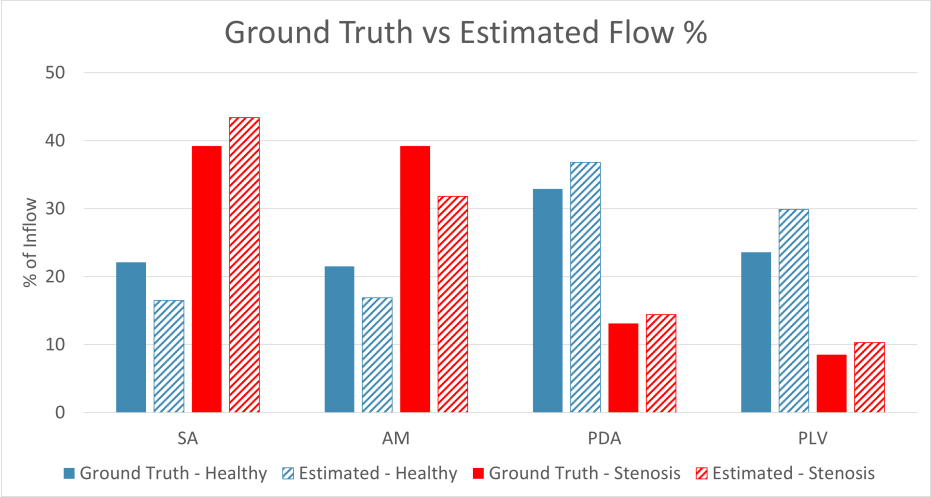


Figure 6.5: Comparison of ground truth flow splits and computed flow splits in the healthy (blue) and stenotic (red) coronary trees. The average difference in % flow across all branches was 5.15% in the healthy tree and 3.68% in the diseased tree.

for the healthy case and 3.68% for the stenosed case, where $\hat{Q}$ represents relative flow expressed as a percentage of inlet flow.

The accuracy of this method could be improved by increasing the number of angiography series to get a better sense of the 3D path length. Even under the simplest possible conditions of steady flow, the complexity of the geometry led to variability in flow rate between consecutive frames when averaging the results of only two projections. This issue would be exacerbated further when increasing the complexity of the problem to include more realistic pulsatile flow. Tracking the movement of contrast agent in 3D would lead to more accurate estimates of coronary flow. Saito et al. [118] developed a method for this in the early 1990s, but their 3D QCA method was limited to a single vessel and required a lot of user interaction. The 3D coronary centerline was reconstructed using back-projection of manually-chosen correspondence points, and vessel contours were assumed to be ellipses whose major and minor axes corresponded to the radius in each angiography view. The traveling dye front from 10 consecutive frames was then projected onto the 3D centerline to measure the distance traveled between frames. The accuracy of flow measurements was limited by the simplifying assumptions of the vessel contours; the radii in each branch would only correspond to the major and minor axes if the two projection angles were 90° apart. This method was also slow and required a lot of manual user interaction, which makes expanding the method to all branches in the coronary tree infeasible. The machine learning 3D reconstruction algorithm presented in section 5.2 is a possible alternative that could automate the process of 3D reconstruction. The network would first have to be retrained to reconstruct partial coronary trees from time series input images. If the centerline reconstruction was improved to capture the true tortuosity and

path length of the 3D coronary vessels, this method could be integrated with the flow extraction algorithm presented here to more accurately estimate flow rate over time without greatly increasing the computational time.

6.4 Future Work: Application to Clinical Angiograms

In the example above, the projection images from the time series were automatically binary images because we thresholded the dye concentration prior to creating the projection. In a real clinical angiogram, the vessels would need to be segmented prior to being input into our flow extraction algorithm. AngioNet, explained further in chapter 3, can be used for this task. An alternative approach would be to exploit the grayscale pixel intensity values of the dye by masking the original image with the AngioNet segmentation rather than using the binary image directly as the input for this algorithm.

The biggest challenge in applying this flow extraction algorithm to real clinical angiograms is developing a suitable algorithm for motion compensation. Since the coronaries undergo nonrigid deformation throughout the cardiac cycle, it is challenging to determine how far the dye has traveled between frames since the geometry itself is changing over time. Even projecting the dye onto reconstructed 3D centerlines would not alleviate this problem as it would be difficult to map the last known position of dye onto subsequently deformed 3D coronary trees. It is therefore essential to determine a method to co-register frames and determine the path length that the dye has traveled. The b-spline free form deformation algorithm used by [62] is promising, but computationally expensive. Most of the recent developments in non-rigid registration for medical images has been for applications in 3D imaging modalities such as CT or MRI, although some methods for 2D registration are also available [61].

CHAPTER 7

Hemodynamic Studies

7.1 Effect of Reconstruction Error on FFR Computation for Varying Levels of Stenosis

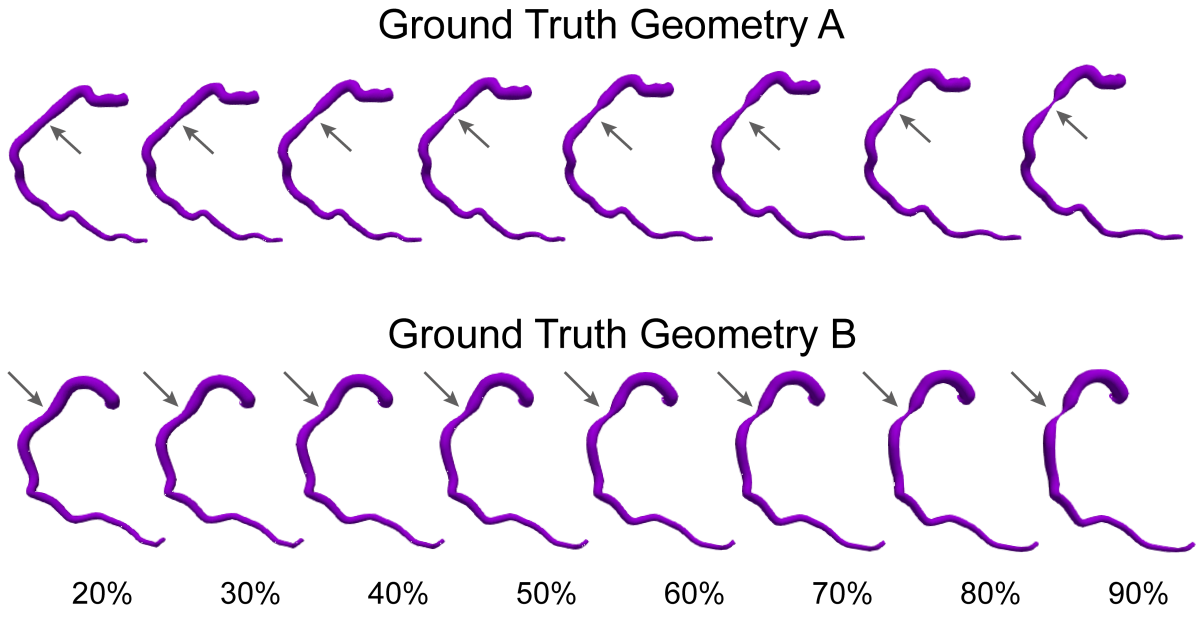


Figure 7.1: Ground truth RCA vessel geometries with 20-90% stenoses (gray arrows). Geometry A had the lowest reconstruction error out of the 10 geometries that were reconstructed in subsubsection 5.2.3.1 while Geometry B had the maximum reconstruction error. FFR computed in each of these geometries was compared to FFR in their reconstructions.

In Section 5.2.3.1 of Chapter 5, we measured the reconstruction error for vessel radii and centerlines in 10 RCA vessel geometries with varying levels of stenoses. In this section, we are interested in understanding how these errors in reconstruction, particularly the error stenosis severity, affects hemodynamic quantities such as FFR. To do so, we performed CFD simulations in 2 of the 10

RCA geometries: one with the best overall radius reconstruction error and one with the worst overall radius reconstruction error. Figure 7.1 shows the two ground truth RCA vessel geometries with 20-90% diameter reduction. Geometry A had the lowest error in radius, with a mean error at the stenosis of 0.03mm, while Geometry B had the highest mean error of 0.20mm at the stenosis.

7.1.1 Methods

Binary projections of the ground truth geometries were input into the 3D reconstruction network as in Section 5.2.3.1 to generate the reconstructed vessels, represented as a matrix. The vessel matrix was converted into a 3D model suitable for simulations as follows: Approximately 18-30 centerline points and their corresponding radii were evenly sampled along each vessel and imported into CRIMSON [3]. Care was taken to ensure that the point of minimum vessel diameter at the stenosis was included in the sampled points. The vessels were lofted and turned into solid models using CRIMSON. The vessels were then subdivided into tetrahedral elements to create a 3D mesh. The maximum mesh element size was taken to be 10% of the minimum distal vessel diameter, or 0.2mm. Curvature refinement was applied to increase the number of elements in the stenosis. The number of mesh elements ranged from approximately 800,000 elements in the mildest stenosis cases to 2 million elements in the 80% and 90% stenosis geometries. A mesh refinement study was performed in these severe stenosis geometries to confirm that the mesh was sufficiently fine and that meshes with smaller elements did not change the solution.

A total of 32 CFD simulations were performed: pairs of ground truth and reconstructed geometries for 8 levels of stenosis (20, 30, 40, 50, 60, 70, 80, and 90% diameter reduction) in each of the two RCA geometries. In order to perform a fair comparison of the pressure and FFR in all geometries, identical boundary conditions were applied in all 32 simulations. The steady inflow value was calculated by multiplying literature values for mean velocity in the distal RCA and the average cross-sectional area of the vessel. The mean velocity from the literature was measured using an invasive Doppler guidewire under hyperemic conditions [70]. A resistance of 7.5 g/mm$^4 \cdot$ s was applied at the outlet, and the no slip condition was enforced at the walls of the vessel.

The 3D Navier-Stokes equations were solved for each pair of ground truth and reconstructed geometries. The pressure along the centerline of each vessel was extracted to determine the pressures proximal and distal to the stenosis. The ratio of these pressures was used to compute FFR.

7.1.2 Results

Figure 7.2 shows an example of vessel reconstruction and pressures in the ground truth and reconstructed geometries for the 70% stenosis case. The top half of the figure overlays the ground truth and reconstructed geometries, with an inset showing the reconstruction error at the stenosis. In this

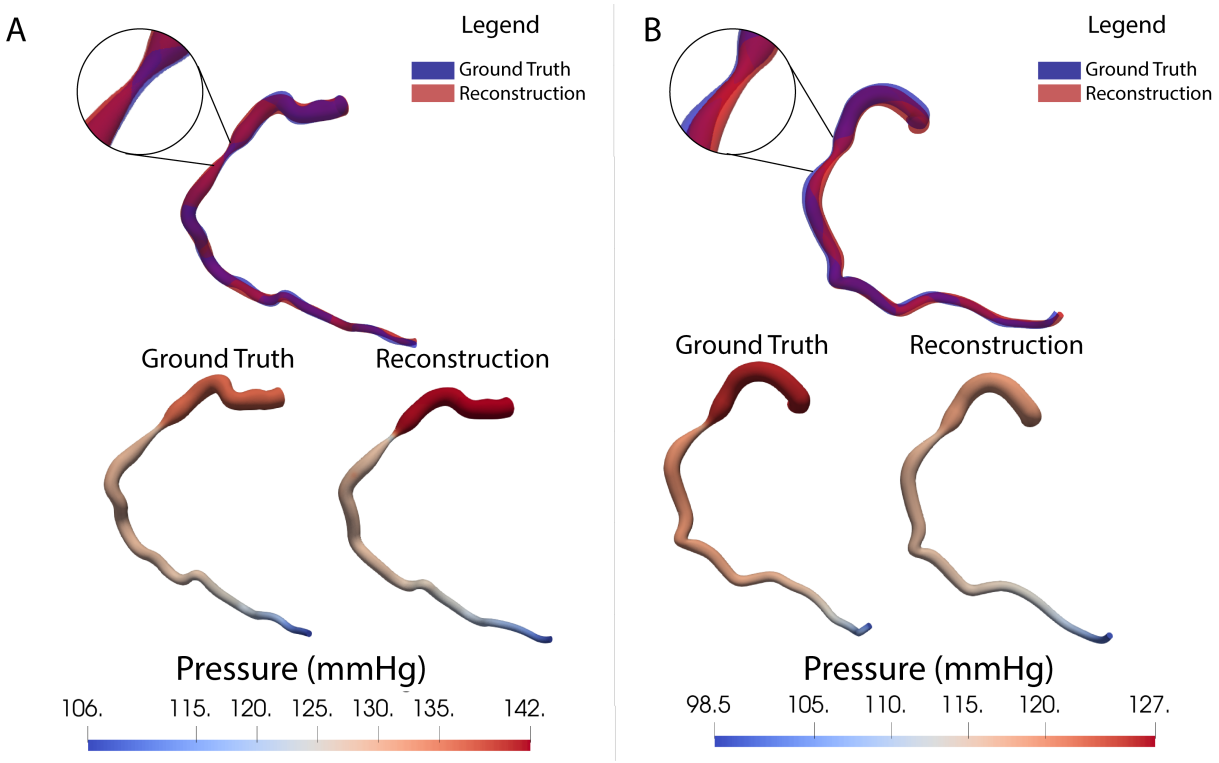


Figure 7.2: Comparison of ground truth and reconstructed geometries with 70% stenosis. The reconstruction in A underestimated stenosis diameter by 0.007mm, leading to slightly higher pressure at the inlet. Meanwhile, reconstruction B overestimated stenosis diameter by 0.203mm and underestimates pressure at the inlet.

case, the error in stenosis radius was -0.007mm for geometry A and 0.203mm for geometry B. The reconstructed centerlines closely match the ground truth, although the reconstructed centerlines are once again smoother than the original geometries. The centerline root mean squared errors were 0.699mm and 0.643mm in geometry A and B respectively. The bottom half of the figure compares the pressures calculated in the ground truth and reconstructed vessels for both geometries. The difference in inlet pressures between the ground truth and reconstructed vessel for geometry A was 5.7mmHg or 4%. Geometry B had a slightly larger difference in inlet pressure (7.3mmHg, or 5.7%).

The pressures along the centerline in the 50%, 70%, and 90% stenosis cases are shown in Figure 7.3. In most cases, the neural network underestimates stenosis severity, which leads to decreased resistance in the 3D domain of the vessel and consequently decreased pressure along the centerline compared to the ground truth. The pressure along the centerline is more sensitive to the error in radius for higher levels of stenosis. For the 50% stenosis case, the best-case error in inlet pressure is 2mmHg and the worst-case error is 4.4mmHg. Meanwhile, for 90% stenosis, a 0.11mm error in stenosis radius leads to a 54.1mmHg error in inlet pressure (geometry A) and a 0.39mm

error in stenosis radius leads to 79.5mmHg error in inlet pressure (geometry B).

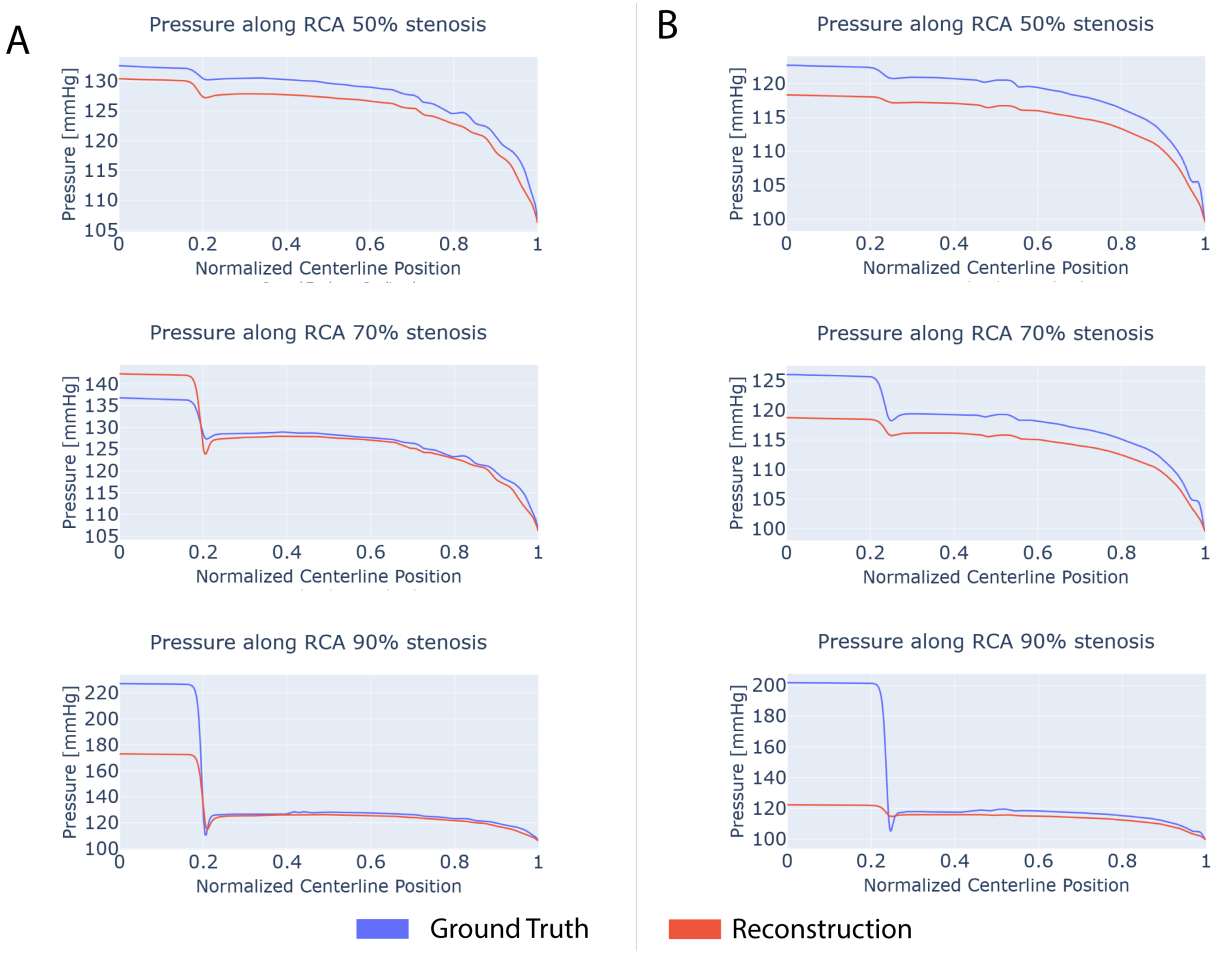


Figure 7.3: Comparison of ground truth vs reconstruction pressure along the centerline in cases with 50,70, and 90% stenosis.

Although Figure 7.3 demonstrates the error in pressure along the entire centerline, we are more interested in the error in FFR, or the ratio of pressures proximal and distal to the stenosis. Figure 7.4 plots the error in stenosis radius and the computed FFR for each stenosis level. In the radius error plot, the orange points summarize the mean and standard deviation of the errors in all 10 single vessel RCA geometries that were reconstructed in Section 5.2.3.1. The green dots are the radius error for geometry A, the geometry with minimum radius reconstruction error, while the purple dots are the radius error for geometry B which had the maximum radius error. In the graph of FFR, the green lines correspond to FFR computed in the ground truth geometry (solid line) and the reconstructed geometry (dashed lines). The FFR values computed in geometry B are given in purple. The gray dashed line is the FFR threshold value of 0.8; values above 0.8 indicate functionally insignificant lesions while values below 0.8 indicate a significant flow-limiting lesion.

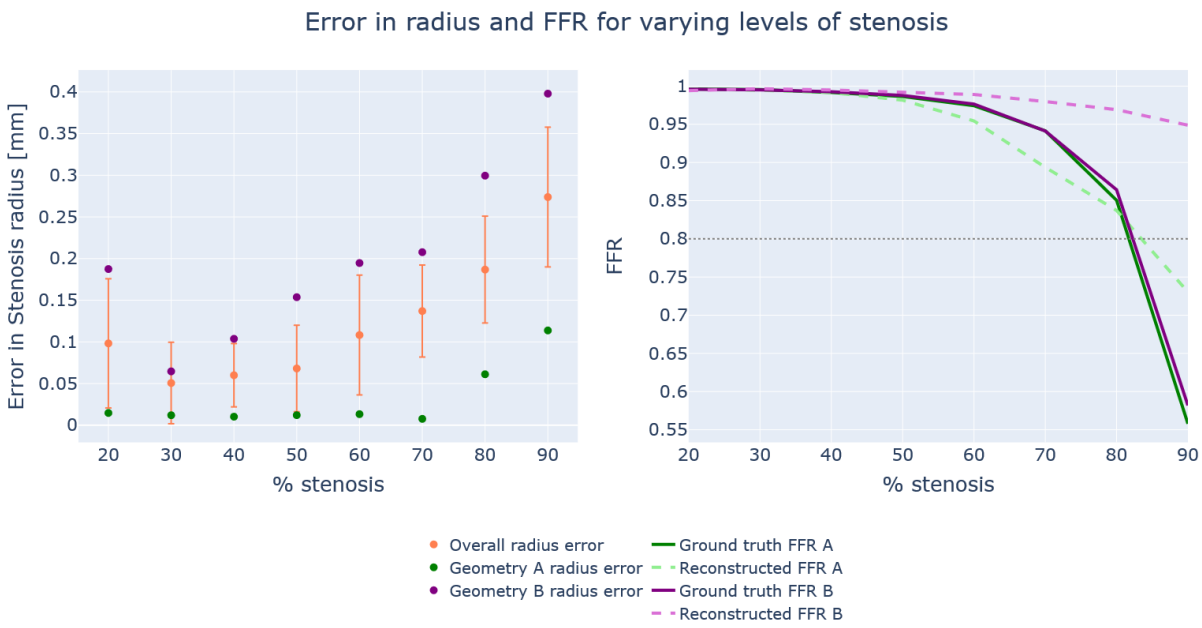


Figure 7.4: Comparison of stenosis radius error and error in FFR for varying levels of stenosis. In the radius graph (left), the orange series summarizes the mean and standard deviation of the set of 10 geometries from which geometry A and B were selected. Geometry A has the minimum error, in green, and Geometry B has the maximum error, in purple. The FFR graph (right) demonstrates the error in FFR in the ground truth (solid lines) and reconstructed (dashed lines) geometries A and B. Error in FFR increases sharply as stenosis severity increases past 50%.

	Geometry A		Geometry B	
% stenosis	Ground Truth	Reconstruction	Ground Truth	Reconstruction
20	0.996	0.996	0.996	0.994
30	0.995	0.995	0.995	0.997
40	0.992	0.991	0.993	0.995
50	0.986	0.982	0.988	0.992
60	0.974	0.954	0.976	0.989
70	0.941	0.894	0.941	0.980
80	0.850	0.837	0.864	0.969
90	0.558	0.730	0.582	0.949

Table 7.1: FFR values in ground truth and reconstructed geometries for varying levels of stenosis

FFR computed for each stenosis level are summarized in Table 7.1. In both geometries, the FFR from the reconstructed geometry is almost identical to the FFR from the ground truth for stenosis severity $\leq$50%. As stenosis severity increases, the discrepancy between reconstructed and ground truth FFR also increases. In the best-case scenario with minimum radius reconstruction error, we see that the green dashed line remains on the same side of the threshold value of 0.8 for all levels of

stenosis. In the worst-case scenario, the purple dashed line diverges from the solid purple line for 80% and 90% stenoses. The 90% stenosis case is falsely determined to be FFR >0.8 although the true FFR value is 0.58. The percent error in FFR ranged from <1% for stenoses ≤70% diameter reduction to 5.3% and 47% error for the 80% and 90% diameter reduction stenoses.

7.1.3 Discussion

From Figure 7.3 and Figure 7.4, we can see that the error in reconstructed radius increases with stenosis severity. This may be due to the difficulty in identifying vessel radius from projection images when the diameter of the stenosis is only a few pixels wide. Although they are the most challenging case for reconstruction, the severe stenoses are also the cases where pressure is most sensitive to small errors in vessel radius. In the best-case scenario, the radius errors are small enough such that the FFR in the reconstructed geometry is on the correct side of the 0.8 threshold value for all levels of stenosis. Even in the worst-case scenario, the FFR values closely agree until stenosis severity increases to be >70%. These findings are not however generalizable to all vessels since they are highly dependent on the boundary conditions chosen for this study.

The main limitation of this study is the fact that we are only considering single vessels and prescribing constant flow at the inlet for all levels of stenosis severity. In reality, the stenotic vessel would be part of a coronary tree and flow would be diverted to other branches in the presence of a flow-limiting lesion. Since we are modeling the single-vessel case, the amount of flow entering the vessel had to be consistent to make a fair comparison. The inflow rate was determined from average literature values; however, patients have a wide range of values for cardiac output and CFR. If the inflow rate was chosen to be higher, it is possible that FFR < 0.8 could have been achieved for moderate stenoses instead of only the 90% stenosis case. This may also have led to greater differences in the pressures and FFR computed in the reconstructed vessels. One disadvantage of this is that increasing the inflow would also have led to unreasonable pressure values in the 80% and 90% stenosis cases; the pressures were already outside of the normal physiological range with the flow rate chosen for this study.

These findings demonstrate the importance of refining reconstruction accuracy of the stenosis radius. One strategy may be to incorporate information from the stenosis detection and measurement algorithms presented in Chapter 4. Although it is difficult to measure the true vessel diameter from 2D projections, the average diameter reduction in all of the input images could be used to correct the predicted stenosis radii from the neural network. One disadvantage of this approach is that estimating diameter reduction using automated image processing algorithms is also prone to error as discussed in Section 4.4.2. From the results of the geometry with minimum reconstruction error, we know that an average error of 0.03mm at the stenosis is acceptable to accurately compute

FFR. Ideally, the 3D reconstruction neural network could be retrained to predict vessel diameter within this threshold. The challenges are the pixel resolution of the input images (0.2-0.3mm/px) and the increased difficulty of predicting the 3D geometry of the full coronary tree rather than a single vessel.

7.2 FFR Verification in Clinical Case

In this section, we compare FFR calculated using the methods developed in this dissertation with clinically measured FFR in the right coronary tree of one patient. This entails 1) using AngioNet (Chapter 3) to segment several angiographic frames, 2) using YOLOv5 and patch AngioNet to detect and refine the segmentation at the stenoses (Chapter 4), 3) reconstructing the 3D geometry of the patient using our 3D reconstruction neural network (Chapter 5), and 4) calculate FFR from a patient-specific simulation. The flow extraction algorithm from Chapter 6 was not used to inform the boundary conditions of the simulation since the method needs to be developed further before it can be applied to clinical angiograms.

7.2.1 Vessel Segmentation

The first step was to identify frames from the same part of the cardiac cycle in 3 angiography series of the patient and segment the coronary arteries. Cross-correlation was used to identify the number of image frames in one cardiac cycle. By identifying number of frames between the reference frame in each angiography series and the frame with the highest cross-correlation value, the length of the cardiac cycle was determined to be 8 frames. This corresponded to a heart rate of 64 beats per minute and closely matched the heart rate of 66 beats per minute reported in the patient's procedure details. In order to co-register the cardiac cycle across all 3 angiography series, peak systole was manually identified as the frame at which the coronaries were closest to the top of the image. The coronaries travel downward as the heart muscle relaxes in diastole, therefore this was a reliable landmark that could be identified in all series. The peak systole reference frames and length of the cardiac cycle were used to find the best set of co-registered frames from each series. Since each angiography series contains multiple cardiac cycles to maximize the visibility of the vessels at the same point in the cardiac cycle.

These angiography frames were then segmented using AngioNet, shown in Figure 7.5. All the major branches of the coronaries are included in the segmentation besides the SA branch. Some of the segmented smaller and tortuous branches are not continuous. None of the segmentation results include the catheter, even when it intersects vessels of similar diameter and pixel intensity in the first image.

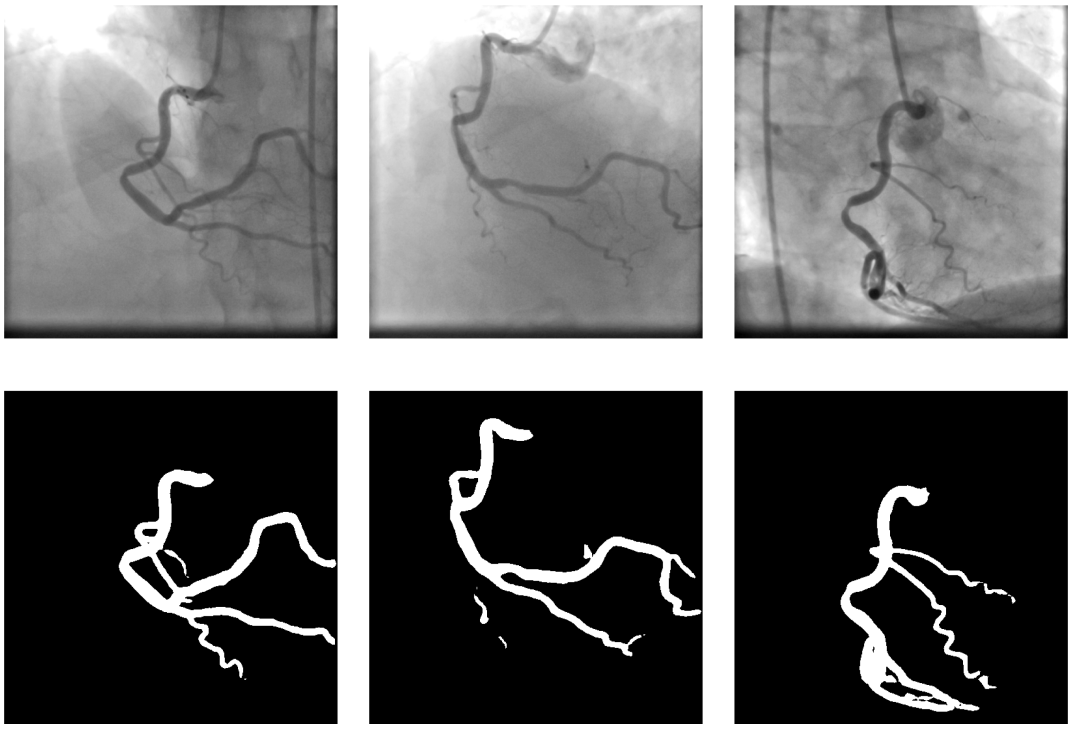


Figure 7.5: Input angiography frames (top) and their corresponding segmentations (bottom). The main branches of the coronaries are included in the segmentation images, besides the SA branch, and the catheter is excluded.

7.2.2 Stenosis Detection and Segmentation Refinement

The next step was to refine the segmentation at the stenoses. The coronary findings for this patient reported a 60% mid-RCA stenosis and 30% PLV stenosis. Both stenoses were detected in all three selected frames by YOLOv5. Although there were a few false positive detections, the refined segmentations by patch AngioNet in the extra regions could still be used to improve the original AngioNet segmentation. By measuring the vessel diameter in the segmented patches, we determined that the RCA contained a 61% diameter reduction stenosis while the PLV contained a 38% diameter reduction stenosis.

Figure 7.6 shows the detected stenoses and their corresponding segmentations overlaid on the input angiogram frames. The patch-level stenoses were used to replace the corresponding sections of the original AngioNet segmentation, shown in the bottom row of the figure. The red boxes enclose the regions of the AngioNet segmentation that were replaced by the patch-level segmentation. The blue arrows indicate regions of mismatched vessel diameter at the interface between the original and patch-level segmentation. They also indicate gaps or missing vessels in the patch level segmentation that were present in the original segmentation.

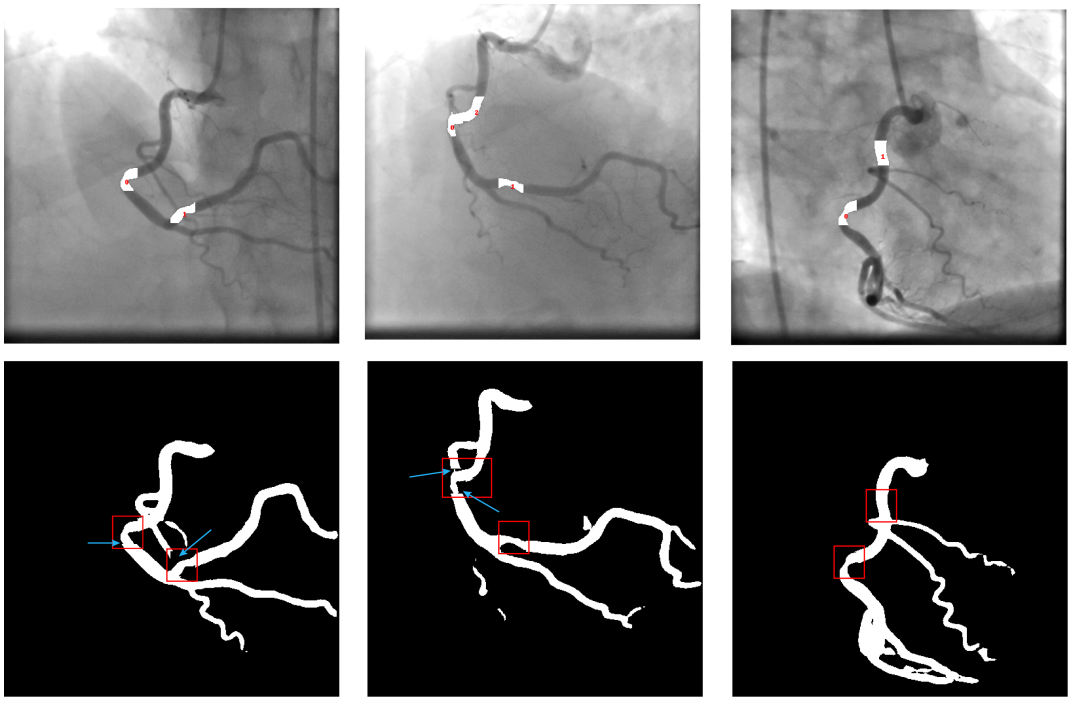


Figure 7.6: Stenosis segmentation and refinement. The top row shows the regions of the image that were localized as stenoses and their segmentation overlaid on the input image. The bottom row shows the refined segmentation image where the results of the patch segmentation replace the region in the original image segmentation. Blue arrows point to areas of diameter mismatch, gaps, or missing vessels.

7.2.3 3D Reconstruction

The corrected segmentation images were used as input to the multi-stage 3D reconstruction neural network. There were several aspects of the clinical angiograms that differed from the synthetic binary projections used to train the network. Firstly, the segmented images did not contain the SA branch and some included other side branches that were not used in synthetic data generation. The relative position of the patient in each angiography series is unknown, as the patient table was moved throughout each acquisition to follow the flow of dye. Meanwhile, the synthetic data that was used to train the network did not have any translation of the 3D object between projections. Another difference is that some of the branches are cut off at the edges of the image, unlike the synthetic data where the entire coronary tree was visible. Finally, the distances between the X-ray source and the isocenter or image plane were fixed at 0.75m and 1.2m respectively in the synthetic data but were variable in the clinical data. The source to image plane distance varied between 1.05 and 1.1m in the clinically acquired angiograms, while the source to isocenter distance varied between 0.693m and 0.764m.

The reconstructed RCA geometry is shown in Figure 7.7. Since the network was trained to

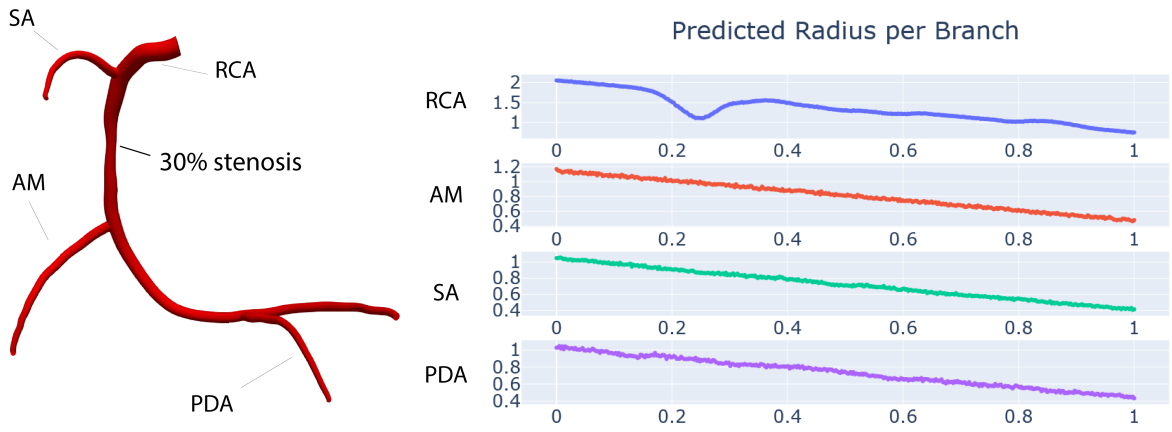


Figure 7.7: Reconstructed coronary tree and the predicted radii along each branch. The network identified a stenosis in the RCA and did not predict a stenosis in the PLV (distal half of the RCA vessel in this graph).

always predict 4 branches, it includes the SA branch even though it was missing in the input segmentation images. In the same vein, the reconstruction is missing some of the side branches that were present in the segmentations. The neural network did predict a stenosis in the mid-RCA, however it was only a 30% stenosis and occurred proximal to the AM branch while the true geometry had a 60% stenosis distal to the AM branch. The stenosis in the PLV was not detected.

7.2.4 FFR Computation

In order to perform a hemodynamic simulation to compute FFR, we must first determine the patient-specific boundary conditions. To do so, we manually segmented and meshed the patient's coronary geometry from CCTA and adjusted the boundary conditions in this ground truth geometry to match measurements acquired during cardiac catheterization. The patient's procedure notes contained invasively measured aortic pressure at baseline conditions, 142/87 mmHg. They also contained non-invasive cuff pressure recorded at baseline and hyperemic conditions: 160/91 and 151/78 mmHg. Since FFR was measured under hyperemic conditions, we estimated the hyperemic aortic pressure using the ratio of non-invasive blood pressures at baseline and hyperemia. The mean hyperemic aortic pressure was calculated as the weighted average of the systolic and diastolic pressures. The resulting mean pressure, 97mmHg, was used to inform the expected pressure at the inlet of the coronary tree.

Since no information about cardiac output or flow was indicated in the patient's chart, the baseline flow was initially estimated as 2% of the average adult cardiac output, 5 L/min. Assuming hyperemia led to a 2.5 times increase in coronary flow, the inlet flow was initially set to be 0.25

L/min. Resistances were applied to the outlets using allometric scaling laws [143]. The values were chosen such that the equivalent resistance would lead to the desired inlet pressure, using the relation $P_{\text{inlet}} = Q_{\text{total}} R_{\text{total}}$ where R_{total} represents the resistances assigned to the outlets as well as the inherent resistance of the 3D domain. The flow and equivalent resistance were then iteratively adjusted until the pressure drop across the stenosis yielded an FFR value of 0.88 to match the patient's recorded coronary findings. The final inflow that produced an FFR of 0.88 was 0.1 L/min.

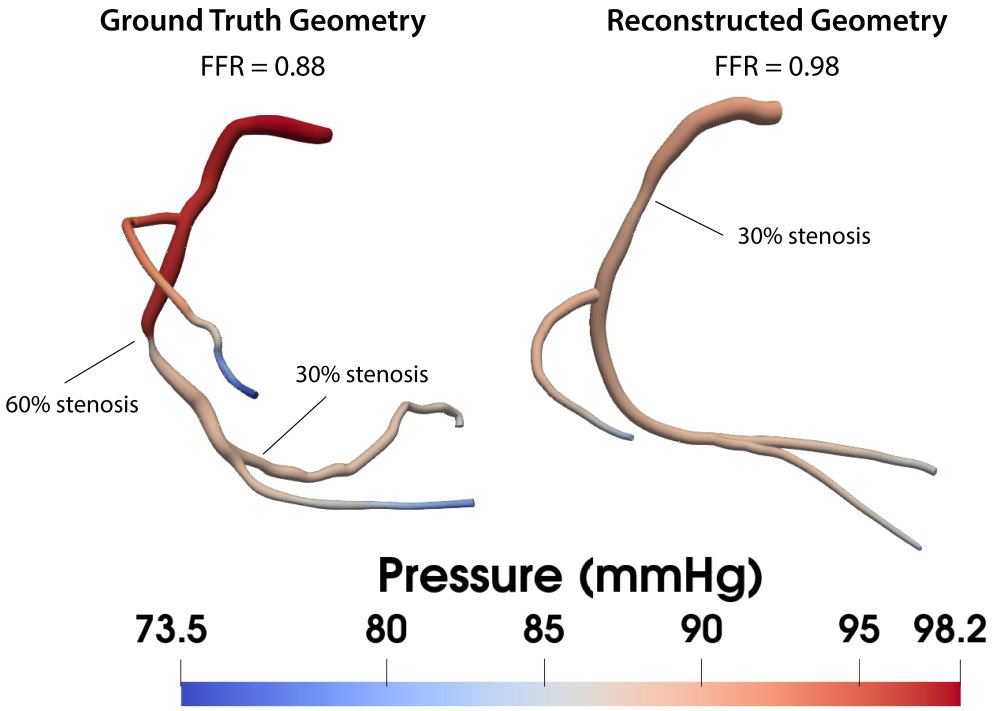


Figure 7.8: FFR computation in the ground truth geometry from CCTA compared with the reconstructed geometry. The reconstruction underestimates stenosis severity and has the RCA stenosis in the wrong position relative to the AM branch, leading to a lower pressure drop across the stenosis and higher FFR.

Once the correct inflow and resistance values were determined from the ground truth geometry, the same boundary conditions were applied to the reconstructed geometry. The SA branch was omitted from the reconstructed geometry to facilitate this. The results are shown in Figure 7.8. The inlet pressure in the ground truth was 98.2mmHg while the inlet pressure in the reconstruction was 90mmHg. The pressures on either side of the ground truth mid-RCA stenosis were 96.8Hg and 85.5mmHg, therefore FFR was 0.883. Meanwhile, the pressures on either side of the reconstructed 30% stenosis were 88.1mmHg and 86.3mmHg, leading to an FFR of 0.980.

7.2.5 Discussion

Accurate vessel segmentation is the most vital part of the FFR computation pipeline, since inaccurate segmentation introduces error into all subsequent steps. Although AngioNet is able to segment the branches of the coronary arteries, in general it has a tendency to over-estimate vessel diameter. This can lead to under-estimation of stenosis severity, which is why the stenosis detection and patch segmentation pipeline was employed to refine the segmentation at the stenoses. The challenge lies in integrating the segmentation of the entire image and the patch. As seen in Figure 7.6, there is often a mismatch in vessel diameter between the AngioNet and patch segmentations. This boundary can be smoothed using top hat filters or other morphological operations; however, it is difficult to know which segmentation contains the correct vessel diameter. Another challenge is that the segmented patches are sometimes missing side branches that AngioNet did include. This could be because of the small size of the patches: 40x40 pixels in each patch compared to 512x512 pixels in the entire image. Without the context of the larger image, faint side branches can be perceived as background noise in the patch. The advantage of AngioNet is its ability to recognize faint vessels using global context, while the advantage of patch AngioNet is to accurately capture vessel diameter locally. Perhaps increasing the size of the patches and retraining the patch segmentation network would create a better balance between global context and local diameter. The segmentation results of an intermediate patch network could perhaps be more easily integrated into the full image segmentation.

The error in vessel segmentation may have compounded the error in 3D reconstruction; however, other aspects of the clinical input images may have caused most of the reconstruction error. As outlined in subsection 7.2.3, there were several differences between the synthetic images used to train the 3D reconstruction network and these clinical angiograms. The patient was moved between image acquisitions, which could have led to difficulty with the neural network identifying matching features such as the stenosis in all images. It is possible that the translation of the 3D object caused the network to spatially integrate vessel diameter from different parts of the vessel. This could explain why the reconstructed coronary tree contained a 30% stenosis rather than 60% stenosis and why the stenosis was in the wrong position. The addition of an auxiliary training task such as classifying branches may help the network learn to identify which branches and features correspond to each other in all images. Another difference was the relative distances between the isocenter, imaging plane, and X-ray detector in the clinical images. This could have led to scaling issues since these distances define the magnification of the object in each image. Furthermore, the PDA and PLV vessels were not clearly visible or cut off in several of the input images. This could explain why the reconstructed branch centerlines are shorter than their CCTA counterparts. Finally, the network was trained with a strong assumption about the structure of the coronary tree: 4 branches, expected to be within certain positions relative to the main RCA vessel. This patient's

coronaries did not match the expected tree structure, and the network still predicted an SA branch even though it was not present in the input images. More flexibility could be integrated into the tree structures used to train the 3D reconstruction network to account for anatomical variations. Despite these challenges, it is encouraging that the reconstruction still resembles a coronary tree, and the general shape of the main RCA matches the ground truth from CCTA.

We now discuss the results of the hemodynamic simulation in the reconstructed coronary tree. The ground truth geometry from CCTA was used to tune the boundary conditions rather than the reconstruction since we only had a limited number of clinical measurements: aortic pressure and FFR in the mid-RCA. We have more parameters than patient-specific measurements, therefore we could tune the inflow and resistances in the reconstructed geometry to artificially match the clinically measured pressures. Since the FFR value depends on stenosis severity, it was imperative that the geometry used to tune the inflow and outlet resistances had the same diameter reduction as the vessel in which FFR was measured during cardiac catheterization. We are confident that the coronary geometry from CCTA closely matches the true coronary tree in which the cardiac catheterization measurements were made.

Through this tuning process, we identified that an inflow of 0.1 L/min led to the appropriate pressure drop across the stenosis and FFR value. This inflow rate may seem too low for hyperemic conditions since it is equal to the baseline flow when the cardiac output is assumed to be 5 L/min. However, cardiac output can be as low as 2.5 L/min in some patients. Moreover, if microvascular disease were present, CFR may be as low as 1.5 leading to small differences between baseline and hyperemic flow. Given the wide range of coronary inflow and the lack of flow measurements for this specific patient, it is indeed possible that the hyperemic flow rate was 0.1 L/min.

When the tuned boundary conditions were applied to the reconstructed geometry, the inlet pressure was lower than the clinically measured value of 97mmHg. The reconstructed geometry underestimated the RCA stenosis severity and did not include the stenosis in the PLV; furthermore, several of the branches were shorter and less tortuous than their ground truth counterparts. These contributed to a decreased resistance in the reconstructed 3D domain compared to the ground truth, and therefore an underestimation of inlet pressure. The decreased inlet pressure and underestimation of stenosis severity both contributed to the overestimation of FFR.

In future clinical applications where a ground truth geometry from CCTA is not available to tune the boundary conditions, the relative flow extraction algorithm could be applied to determine patient-specific outlet resistances rather than relying on allometric scaling laws. Although this algorithm can provide accurate estimates of relative flow in each branch, it cannot accurately estimate the true inflow value. For prospective analysis, cardiac output or CFR measurements in the clinic could be acquired and used to inform the boundary conditions of the simulation.

CHAPTER 8

Discussion and Future Directions

8.1 Discussion

In this work, we have developed a collection of automated algorithms to assess anatomical and functional severity of CAD. The purpose of these algorithms is to provide cardiologists with quantitative information about disease severity that can assist in diagnostic decisions. While visual estimation of disease severity is the most common method of diagnosis in practice, it can be subjective and often over-estimates disease severity. This in turn can lead to unnecessary revascularization procedures. We know that quantifying stenosis severity using QCA or FFR can lead to better patient outcomes; the automated approaches presented in this work can provide the same information as QCA and FFR without the associated cost, time, and variability of these operator-dependent measurements. Since the only input for these algorithms is the X-ray angiography images, they can be integrated into the clinical diagnostic workflow with minimal changes to current clinical practice.

The first algorithm we developed was AngioNet, a neural network for vessel segmentation. This algorithm was the foundation for several other methods developed in this thesis, since the 3D reconstruction and flow extraction methods used binary segmentation images as their input. Compared to other neural networks for coronary segmentation, AngioNet was able to consistently ignore the catheter and other imaging artifacts even in challenging images where these structures overlapped with the vessels. The key innovation in this work was the development of the Angiographic Processing Network to learn the best possible set of pre-processing filters that improved segmentation accuracy. We demonstrated that the APN led to better segmentation accuracy than commonly used preprocessing filters such as unsharp masking or CLAHE, particularly in images or small vessels with low contrast. Even in the much simpler 40x40 pixel image patches containing stenoses, the addition of the APN to the backbone network increased the Dice score by over 5%.

While the APN developed for AngioNet improved segmentation performance in patch AngioNet, there were some lessons learned from training the patch-level network that led to improve-

ments in the original AngioNet as well. Patch AngioNet was trained and validated using data from the University of Michigan (UM) hospital, and its performance was measured using angiograms from ICRF. We initially observed that the segmentation performance was much poorer on the ICRF angiograms compared to the validation images from UM. We hypothesized that the angiograms acquired from the two different hospitals may have had different pixel intensity distributions. Instead of normalizing the pixel values to be between [0,1] or [-1, 1], we performed z-score standardization to homogenize the images from different hospitals. This led to a marked improvement in segmentation accuracy on the ICRF images. When the original AngioNet was retrained with z-score normalization, we once again noted an improvement on segmentation performance on the ICRF angiograms. These findings highlight how sensitive neural networks can be to implicit features of different datasets. The UM and ICRF angiograms appeared visually similar, yet there were underlying mathematical differences that led to poor generalizability when the datasets were not standardized.

The lesson that a neural network can be sensitive to invisible features or learn unintended patterns was also apparent when training the 3D reconstruction neural network. When determining the feasibility of using a neural network for 3D reconstruction, we developed a simple dataset consisting of cylinders with variable diameters and their projections. All the cylinders were set to have a 50% linear reduction in diameter, which the neural network was able to accurately capture in the validation set. When the network was tested on images of cylinders with a 30% or 70% tapering instead, the predictions broke down. The network had in fact learned to correctly predict the maximum diameter and then assumed a linear decrease of 50% irrespective of the pixel width of the cylinder in the input images. This problem was solved by introducing a random level of tapering to each cylinder in the training dataset. Although a human may interpret this as the network learning to recover the radius of the cylinder from the pixels in the image, this assumption is not necessarily the case. For example, when stenoses were randomly introduced into the cylinders, the network correctly predicted linear decreases in vessel diameter but did not correctly identify the stenosis diameter. This was resolved by adopting a multi-stage neural network approach; however, the underlying assumption of linear tapering remained in the training data. There are countless other embedded features and underlying assumptions in any synthetic training dataset, and not all of them can be easily identified and addressed like the constant radius tapering. It is therefore important to understand the limitations of a synthetic training dataset and how the assumptions of data generation can cause a network to learn unintended features that would not generalize to clinical input data.

One surprising aspect of training the 3D reconstruction neural network was that changing the backbone network to be a transformer instead of a CNN did not improve performance. We hypothesized that since attention mechanisms are meant to suppress irrelevant features and learn the

strength of the association between different features, they would help the network better identify the corresponding centerlines and radii for each branch in the coronary tree. Furthermore, transformers were initially designed to predict sequences of words in a sentence, where the likelihood of one word depends on the previous one. This was analogous to a series of centerline points or radii, where the spatially sequential relationship between points could be used to make better predictions. Despite these theoretical advantages, in practice, the multi-stage network with a ViT backbone had almost identical performance as the Resnet101 backbone presented in Chapter 5. Transformers are notoriously difficult and resource intensive to train; perhaps increasing the training data and tuning the training parameters could lead to better performance in the future.

Another surprise in this work was the ability for a neural network to learn to localize stenoses. Object detection networks are notoriously unsuccessful at detecting "small" objects, defined in the benchmark MS COCO dataset as objects whose dimensions are less than 32x32 pixels. Using our empirically determined 40x40 pixel bounding box size, the stenosis dimensions were of a similar size to the small object class. An added challenge is that stenoses are a small part of a larger whole, rather than a distinct object such as a car or a fruit in a natural camera image. Finally, object detection networks often require on the order of 3-5k images to train, and we attempted to train the YOLOv5 network with less than 1k annotated angiograms. Given these limitations, the stenosis localization network performed better than expected. The main hurdle was to create a robust training dataset, since manual identification of stenoses in hundreds of images can be both subjective and error prone. In many cases, the location of a stenosis in an image is obvious. In cases with diffuse disease or coronary aneurysms, it can be more difficult to identify the boundaries of a stenosis. Cardiologists typically use a time series of angiograms to localize stenoses in practice, therefore annotating single angiography frames proved somewhat challenging.

The final algorithm we developed in this work was to extract relative flow in the branches of the coronary tree. In the simplest possible case, where the input was a synthetic angiographic series with steady flow and no coronary motion, the algorithm could approximate the relative flow in each branch within 5%. We would expect the error to increase for more complex cases, such as simulated pulsatile flow and eventually real clinical angiograms. A major source of error is the foreshortening of different vessel segments due to the projection angle. As discussed in section 6.3, a possible solution would be to use the 3D reconstruction network to reconstruct the 3D path of the dye for each set of frames. In subsection 7.2.1, we used strategies such as cross correlation and the motion of the coronaries to co-register frames from different series. 3D reconstruction from the co-registered frames was still challenging since some vessels were not clearly visible in all images. Another challenge would be to reconstruct the deforming 3D coronary tree from subsequent frames of the angiographic series, since different vessel segments may be obscured over time due to the motion of the coronary tree. An alternative could be to use the principles of the energy-based

approach for 3D reconstruction to translate deformation forces in 2D images onto the 3D structure rather than independently reconstructing the coronary tree from successive frames.

Although this work did not present any new methods to model coronary flow and pressure, the hemodynamic studies led to some interesting insights. One finding of interest was that the pressure in the coronaries is less sensitive to errors in vessel radius for lower levels of stenoses. The best-case radius error for the 90% stenosis case and the worst case error for the 40% stenosis case were both approximately 0.1mm. Meanwhile, the error in FFR for the 90% case was 0.173 compared to 0.002 the 40% case. Therefore, the results of FFR computation in a reconstructed geometry are most trustworthy for mild to moderate stenoses. These are the cases where functional insight is most useful: $>$80% diameter reduction stenosis is an extreme case and is most likely flow limiting, while it is less intuitive to identify if a series of mild to moderate stenoses may be flow limiting. Another observation was that the accuracy of FFR in this stenosis severity experiment was determined relative to the 0.8 threshold value, however this threshold may not be applicable in all cases. In patients with low inflow such as the one modeled in the second half of Chapter 7, the FFR value would be high simply because the flow is not large enough to cause a large pressure drop across the stenosis. If the patient were experiencing chest pain or shortness of breath, these would be indications of functional severity that contradict the high FFR value. Meanwhile, the same stenosis in a patient with higher flow, perhaps due to a higher CFR, could lead to an FFR value less than 0.8. The threshold value was determined empirically from clinical studies in patients without severe microvascular dysfunction or low cardiac output; however, many patients with epicardial disease also experience microvascular dysfunction. Thus, the relationship between CFR and the appropriate FFR threshold value to determine a flow limiting lesion must be better understood.

8.2 Future Directions

In subsection 7.2.5 and section 8.1, we identified several areas for improvement when applying the methods developed in this thesis to clinical angiograms. Since accurate vessel segmentation is imperative for the subsequent algorithms, we discussed several strategies to improve segmentation at the image and patch levels. These include standardizing the input images using z-scores rather than simply normalizing them for greater generalizability and increasing the size of the patches for stenosis segmentation, which would incorporate more spatial context and improve segmentation of faint side branches. Since one of the most important features of vessel segmentation is vessel diameter, the AngioNet and Patch AngioNet could perhaps be trained to learn the Euclidean distance map rather than a binary image; this distance map explicitly encodes vessel diameter and could be easily thresholded to create binary segmentation images for downstream applications.

The stenosis localization network could be improved using a more robust training dataset, anno-

tated by several cardiologists to reduce subjectivity. Another strategy would be to train the network to identify stenoses in time series of angiograms as in [97], which more closely matches how cardiologists perform this task. By identifying the same stenoses across all frames in a time series, we could reduce the incidence of false positives – stenoses that are only localized in a few sporadic frames – and false negatives – stenoses that may not be apparent in certain frames could still be detected in others.

Improvements in stenosis localization and diameter measurement could also be used to improve the results of machine learning-based 3D reconstruction. As described in subsection 7.2.5, the measured stenosis in 2D could be used to update the neural network radius predictions if the stenosis is overlooked. Another finding was that there is more anatomical variation in clinical cases compared to synthetically generated trees, so more flexibility in the number and types of branches needs to be introduced into the data encoding for the network. The challenge of creating a suitable training dataset is that there are a relatively small number of ground truth patient geometries that can be acquired from medical images, while synthetic data inherently contains embedded assumptions and patterns. A hybrid dataset consisting of both clinical cases and synthetic data may lead to better reconstruction. Data augmentation could also be expanded from affine transforms to nonrigid deformation prevent the network from learning patterns in the data generation. Another area to explore is changing the encoding of the 3D geometry to enhance the accuracy of stenosis reconstruction. For example, rather than implicitly encoding stenoses via a series of N radii along the centerline, the radii could be encoded as a list containing the proximal and distal diameter of each vessel, and the parametric position and severity of any stenoses. This sparse encoding explicitly encodes the stenoses, and would simplify the radius MLP architecture in our multi-stage network to be a single MLP instead of one MLP per branch. Finally, as discussed in subsubsection 5.2.4.2 and subsubsection 5.2.4.1, a loss function which incorporates the relationship between points rather than simply measuring the error of predicted points independently could improve performance. We already experimented with encoding the vessel centerlines and radii in terms of the initial point and the first derivatives; however, this approach had limited success. A more sophisticated loss function that incorporates prior knowledge of coronary geometry such as curvature or tortuosity may be more successful.

The future directions for the flow extraction algorithm include validating its accuracy using simulated angiograms of pulsatile flow and incorporating nonrigid registration algorithms so that flow in clinical angiograms can be measured. Another improvement mentioned in section 8.1 would be to integrate the 3D reconstruction neural network with the flow extraction algorithm to reduce error due to foreshortening.

In order for these algorithms to be implemented in the clinical workflow, it is imperative that they can produce results in a timely manner on standard hardware. The automated algorithms for

vessel segmentation, stenosis detection, and reconstruction can produce results in a matter of seconds; however, the relative flow extraction algorithm and CFD simulations still require substantial user input and time to run. The flow extraction algorithm must be automated before it is practical to include in a clinical workflow. The auxiliary task classifying types of branches for the 3D reconstruction network could also assist in this method, since identifying corresponding branch centerlines is the main bottleneck of this approach. Finally, the CFD simulations presented in this work were computationally intensive and required at least 36 cores to run within 5-20 minutes. Reduced order models for flow could provide more rapid results and make clinical integration more feasible.

8.3 Summary of Ongoing Challenges and Potential Solutions

Here we summarize the future directions presented in section 8.2 as action items for each of the algorithms developed in this thesis.

8.3.1 Vessel Segmentation

1. Challenge: Improve vessel diameter accuracy
   (a) AngioNet and Patch AngioNet: Learn distance map instead of binary segmentation image. This can be an auxiliary task or replace the segmentation objective.
   (b) Patch AngioNet: increase the size of the patches (64x64 or 100x100 pixels instead of 40x40) to incorporate more context for which pixels are vessels vs noise
2. Challenge: Integrate AngioNet and Patch AngioNet segmentation
   (a) Use morphological operations and filters to smooth segmentation interfaces

8.3.2 Stenosis Localization

1. Challenge: Improve ground truth annotations
   (a) Incorporate annotations from more cardiologists to reduce subjectivity
   (b) Train cardiologists on synthetic angiograms of phantoms with known diameter reduction to improve their expertise
2. Challenge: Reduce false negative detections
   (a) Train network on angiographic series instead of single frames

8.3.3 3D Reconstruction

1. Challenge: Increase anatomical variation in training dataset
   (a) Include more possible vessel branches in the synthetic tree generator, and a [0] or [None] encoding to indicate missing branches
   (b) Use a combination of synthetic data and real patient data with affine and nonrigid transformations to augment the dataset
2. Challenge: Improve stenosis reconstruction
   (a) Change radius encoding to be a list of proximal and distal vessel diameters, and parametric position and severity of stenoses instead of N radii along the centerline. This would require set-matching before the loss function since the number of stenoses and therefore output size is variable. If a maximum number of stenoses is set, 0 padding could be used to keep the output size fixed.
3. Challenge: Improve centerline path reconstruction
   (a) Change loss function, regularization term, and/or data encoding to incorporate global relationship between points (curvature, 1st or 2nd order derivatives, branch point parametric position)

8.3.4 Relative Flow Extraction

1. Challenge: Measure relative flow in pulsatile conditions
   (a) Run scalar simulations of dye transport to inject dye at different parts of the cardiac cycle and see how this affects recovered flow. Use regression as needed to adjust the recovered flow based on cardiac cycle
   (b) Develop motion compensation algorithms using rigid and nonrigid transformations to track dye path length in clinical angiograms
2. Challenge: Improve flow extraction accuracy
   (a) Use 3D reconstruction network to reconstruct 3D geometry and project dye onto it at each time step to determine true 3D path length
3. Challenge: Automate tracking flow in each branch
   (a) Train a neural network or develop an image processing algorithm to automatically localize each branch across a time series

APPENDIX A

Synthetic Vessel Generation Parameters

In this Appendix, we include the parameters and methods to create 3D synthetic coronary trees and their projections, initially described in Chapter 5.

A.1 Synthetic Coronary Generation

Our goal is to create a synthetic vessel generator that can produce thousands of unique right coronary trees. This generator encodes an analytical representation of each right coronary tree geometry in an $M \times N \times 4$ matrix. In this encoding, M is the number of branches in the vessel tree and N is the number of points on each vessel centerline. The final dimension corresponds to the (x, y, z) coordinates of each centerline point and its corresponding radius, r.

Values derived from literature on coronary anatomy and clinical image data were used to inform the geometric modeling assumptions of the synthetic vessel tree generator. The clinical image data were additionally used to create a set of 3D vessel centerlines that formed the basis of spline-based synthetic centerline generation.

A.1.1 Geometric Modeling Assumptions

The right coronary tree is composed of the right main coronary artery and its branches. In our model we include:

1. Right main coronary artery (RCA)
2. Sino-atrial node branch (SA)
3. Acute Marginal branch (AM)
4. Posterior Descending artery (PDA)

Branch	Length [mm]	Maximum Diameter [mm]	Parametric Position
RCA (+PLV)	120-140mm	4-5.5mm	N/A
SA	35-50mm	1.8-2.2mm	0.03-0.14
AM	48.5-78mm	2-2.4mm	0.22-0.4
PDA	40-60mm	1.9-2.5mm	0.59-0.72

Table A.1: Dimensions of RCA tree branches for synthetic vessel generation

5. Posterior Ventricular artery (PLV): This branch is included implicitly as part of the RCA in our model, as the RCA bifurcates into the PDA and PLV

Clinical image data were analyzed to determine the dimensions, orientations, and positions of each vessel branch. CCTA images from 10 patients were acquired from the UM hospital. All data were de-identified and collected retrospectively according to ethical guidelines, and the study protocol was approved by University of Michigan IRBMED (HUM00155491).

In each CTA image volume, 3D vessel centerlines were delineated for each branch of the coronary tree using the CRIMSON software package [3]. These centerlines were used to calculate vessel lengths and positions of each branch in the coronary tree. The relative position of a given branch with respect to the main right coronary was defined as the parametric position of its bifurcation point along the RCA.

We now list the anatomical assumptions that underpin our model:

1. The dimensions and relative position of each branch are informed by a combination of literature values [146] and the 10 CTA-derived patient geometries described above:

2. Vessel radius
   (a) We assume that the vessels have circular cross-sections (radially symmetric)
   (b) The radius decreases linearly along the length of the vessels. The value of tapering for each branch is randomly sampled from a uniform distribution
      i. 60-70% along the main RCA, as the distal diameter is typically 1.5-2mm [146]
      ii. 30-50% in the side branches
3. Stenosis characteristics
   (a) Following the work of Lee and Fung [68], we assume the stenosis has a gaussian profile
   (b) Stenosis length ranges from 6-22mm [82]
   (c) Local diameter reduction, or stenosis severity, ranges from 20-90% which aligns with clinical standards for mild, moderate, and severe stenoses [117]

(d) Although stenoses can present as concentric or eccentric lesions, we assume concentric stenoses in this model

(e) Stenoses are introduced in the RCA, PDA, and PLV branches, as stenoses in other branches are not typically treated with interventions: 0-2 stenoses on the RCA+PLV and 0-1 stenoses on PDA

(f) We assume stenoses do not occur at bifurcation points

A.2 Vessel Matrix Generation

As described in A.1, vessel centerlines for 10 patients were derived from their CTA image volumes. B-splines representing the centerlines of each branch were aligned using rigid registration. The mean μ and standard deviation σ of the b-spline control points were computed; this created a distribution $\mu \pm 1.5\sigma$ from which to uniformly sample new sets of control points. Random shear and sine-based perturbation of control points were introduced as a form of data augmentation to reduce the likelihood of the neural network learning to predict an underlying mean vessel. New vessel centerlines were generated by interpolating b-splines through randomly sampled control points (Figure 5.8A).

Once all branch centerlines were generated, they were assembled into a tree. The relative positions of each branch were randomly sampled from the ranges in Table A.1. Branches were rotated by up to 10 degrees to increase the variation in the tree structures. The resulting centerline coordinates were concatenated into an $M \times N \times 3$ matrix containing N centerline points for each of the M branches.

The final step of creating the vessel tree matrix was to generate the radii of each branch. The maximum radius for each branch was randomly sampled from the values given in Table A.1. The radius decreased linearly by a factor of α along the centerline (Assumption 2b) to produce a radius vector of length N. Stenoses were then introduced in the synthetically generated vessel (Figure 5.8B). Stenosis position p_x, length l, and severity s were sampled from uniform distributions (Assumption 3). The profile of the stenosis was computed as

$$f(x) = s * \exp\left(\frac{-1}{2}\left(\frac{x - p_x}{\sigma}\right)^2\right) \tag{A.1}$$

such that the maximum value of the function $f(x) = s$ (stenosis severity) and $2\sigma = l/2$. The gaussian stenosis profile was incorporated into the new radius vector as follows:

$$stenosis_vector(x) = radius_vector(1 - f(x)) \tag{A.2}$$

The resulting radius vector for each branch was concatenated with the centerline matrix to produce the $M \times N \times 4$ vessel matrix. This vessel matrix serves as the ground truth for our 3D reconstruction neural network. The matrix can be converted to a surface representation by creating circular contours at planes perpendicular to each of the N centerline points. The matrix therefore encodes a tubular, tapering vessel surface with regions of local narrowing. Examples of generated right coronary trees are shown in Figure 5.8C.

A.3 Synthetic X-ray Angiography Image Generation

Synthetic angiography projections were created by projecting the 3D vessel surface onto several image planes. In order to closely match the acquisition of X-ray angiography images in the clinic, image acquisition was modeled as cone-beam projections from a point X-ray source. The coronary tree was re-centered such that its centroid was at the origin, or isocenter. The distance between the X-ray source and isocenter was fixed at 0.75m and the distance between the source and image plane was 1.2m. The angle of each imaging plane was given by an azimuthal angle θ and elevation angle ϕ. The imaging angles were sampled from 20° windows around commonly used clinical views and other views where all branches of the coronary tree can be visualized clearly. Binary projections were generated by mapping the X-rays from the point source that pass through the 3D vessel surface onto each imaging plane (Figure 5.9.)

BIBLIOGRAPHY

[1] Joel Akeret, Chihway Chang, Aurelien Lucchi, and Alexandre Refregier. Radio Frequency Interference Mitigation using Deep Convolutional Neural Networks. *Astronomy and Computing*, 18:35–39, September 2017.

[2] Zeyuan Allen-Zhu and Elad Hazan. Variance Reduction for Faster Non-Convex Optimization. In *Convex Optimization*, volume 48, page 9, New York, 2016. JMLR: W&CP.

[3] Christopher J. Arthurs, Rostislav Khlebnikov, Alex Melville, Marija Marčan, Alberto Gomez, Desmond Dillon-Murphy, Federica Cuomo, Miguel Silva Vieira, Jonas Schollenberger, Sabrina R. Lynch, Christopher Tossas-Betancourt, Kritika Iyer, Sara Hopper, Elizabeth Livingston, Pouya Youssefi, Alia Noorani, Sabrina Ben Ahmed, Foeke J. H. Nauta, Theodorus M. J. van Bakel, Yunus Ahmed, Petrus A. J. van Bakel, Jonathan Mynard, Paolo Di Achille, Hamid Gharahi, Kevin D. Lau, Vasilina Filonova, Miquel Aguirre, Nitesh Nama, Nan Xiao, Seungik Baek, Krishna Garikipati, Onkar Sahni, David Nordsletten, and C. Alberto Figueroa. CRIMSON: An open-source software framework for cardiovascular integrated modelling and simulation. *PLOS Computational Biology*, 17(5):e1008881, May 2021.

[4] Christopher J. Arthurs, Kevin D. Lau, Kaleab N. Asrress, Simon R. Redwood, and C. Alberto Figueroa. A mathematical model of coronary blood flow control: Simulation of patient-specific three-dimensional hemodynamics during exercise. *Heart and Circulatory Physiology*, 310(9), May 2016.

[5] American Heart Association. Heart Disease and Stroke Statistics—2022 Update: A Report From the American Heart Association. http://www.ahajournals.org/doi/epub/10.1161/CIR.0000000000001052.

[6] Benjamin Au, Uri Shaham, Sanket Dhruva, Georgios Bouras, Ecaterina Cristea, Alexandra Lansky MD, Andreas Coppi, Fred Warner, Shu-Xia Li, and Harlan Krumholz. Automated Characterization of Stenosis in Invasive Coronary Angiography Images with Convolutional Neural Networks, July 2018.

[7] Benjamin Au, Uri Shaham, Sanket Dhruva, Georgios Bouras, Ecaterina Cristea, Alexandra Lansky MD, Andreas Coppi, Fred Warner, Shu-Xia Li, and Harlan Krumholz. Automated characterization of stenosis in invasive coronary angiography images with convolutional neural networks, 2018.

[8] Peter C. Austin. Balance Diagnostics for Comparing the Distribution of Baseline Covariates between Treatment Groups in Propensity-score Matched Samples. *Statistics in Medicine*, 28(25):3083–3107, November 2009.

[9] Vijay Badrinarayanan, Alex Kendall, and Roberto Cipolla. SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:2481–2495, November 2017.

[10] A. Banerjee, F. Galassi, E. Zacur, G. L. De Maria, R. P. Choudhury, and V. Grau. Point-Cloud Method for Automated 3D Coronary Tree Reconstruction From Multiple Non-Simultaneous Angiographic Projections. *IEEE Transactions on Medical Imaging*, 39(4):1278–1290, April 2020.

[11] D. M. Bappy, Ayoung Hong, Eunpyo Choi, Jong-Oh Park, and Chang-Sei Kim. Automated three-dimensional vessel reconstruction based on deep segmentation and bi-plane angiographic projections. *Computerized Medical Imaging and Graphics*, 92:101956, September 2021.

[12] J Martin Bland and Douglasg. Altman. Statistical Methods For Assessing Agreement Between Two Methods Of Clinical Measurement. *The Lancet*, 327(8476):307–310, 1986.

[13] C. Blondel, G. Malandain, R. Vaillant, and N. Ayache. Reconstruction of coronary arteries from a single rotational X-ray projection sequence. *IEEE Transactions on Medical Imaging*, 25(5):653–663, May 2006.

[14] Jerrold T Bushberg, Anthony J Seibert, Edwin M Leidholdt, and John M Boone. *The essential physics of medical imaging; 4th ed.* Lippincott Williams & Wilkins, Philadelphia, PA, 2021.

[15] Filippo Cademartiri, Ludovico La Grutta, Roberto Malagò, Filippo Alberghina, Willem B. Meijboom, Francesca Pugliese, Erica Maffei, Anselmo Alessandro Palumbo, Annachiara Aldrovandi, Michele Fusaro, Valerio Brambilla, Paolo Coruzzi, Massimo Midiri, Nico R. A. Mollet, and Gabriel P. Krestin. Prevalence of anatomical variants and coronary anomalies in 543 consecutive patients studied with 64-slice CT coronary angiography. *European Radiology*, 18(4):781–791, April 2008.

[16] C. Canero and P. Radeva. Vesselness Enhancement Diffusion. *Pattern Recognition Letters*, 24(16):3141–3151, December 2003.

[17] Ruben Cardenes, Alexey Novikov, Julian Gunn, Rod Hose, and Alejandro F. Frangi. 3D reconstruction of coronary arteries from rotational X-ray angiography. In *2012 9th IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 618–621, May 2012.

[18] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. *CoRR*, abs/2005.12872, 2020.

[19] Jason M. Carson, Carl Roobottom, Robin Alcock, and Perumal Nithiarasu. Computational instantaneous wave-free ratio (IFR) for patient-specific coronary artery stenoses using 1D network models. *International Journal for Numerical Methods in Biomedical Engineering*, 35(11):e3255, 2019.

[20] Fernando Cervantes-Sanchez, Ivan Cruz-Aceves, Arturo Hernandez-Aguirre, Martha Alicia Hernandez-Gonzalez, and Sergio Eduardo Solorio-Meza. Automatic Segmentation of Coronary Arteries in X-ray Angiograms using Multiscale Analysis and Artificial Neural Networks. *Applied Sciences*, 9(24):5507, December 2019.

[21] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking Atrous Convolution for Semantic Image Segmentation. *arXiv*, 2017.

[22] S.J. Chen and J.D. Carroll. 3-D Reconstruction of Coronary Arterial Tree to Optimize Angiographic Visualization. *IEEE Transactions on Medical Imaging*, 19(4):318–336, April 2000.

[23] Adriaan Coenen, Marisa M. Lubbers, Akira Kurata, Atsushi Kono, Admir Dedic, Raluca G. Chelu, Marcel L. Dijkshoorn, Robert Jan M. van Geuns, Max Schoebinger, Lucian Itu, Puneet Sharma, and Koen Nieman. Coronary CT angiography derived fractional flow reserve: Methodology and evaluation of a point of care algorithm. *Journal of Cardiovascular Computed Tomography*, 10(2), 2016.

[24] Jacob Cohen. *Statistical Power Analysis for the Behavioral Sciences*. L. Erlbaum Associates, 1988.

[25] Chao Cong, Yoko Kato, Henrique D. Vasconcellos, Mohammad R. Ostovaneh, Joao A. C. Lima, and Bharath Ambale-Venkatesh. Deep learning-based end-to-end automated stenosis classification and localization on catheter coronary angiography, March 2021.

[26] Weijian Cong, Jian Yang, Danni Ai, Yang Chen, Yue Liu, and Yongtian Wang. Quantitative Analysis of Deformable Model-Based 3-D Reconstruction of Coronary Artery From Multiple Angiograms. *IEEE Transactions on Biomedical Engineering*, 62(8):2079–2090, August 2015.

[27] B De Bruyne, F Hersbach, N H Pijls, J Bartunek, J W Bech, G R Heyndrickx, K L Gould, and W Wijns. Abnormal epicardial coronary resistance in patients with diffuse atherosclerosis but "Normal" coronary angiography. *Circulation*, 104(20):2401–6, November 2001.

[28] Bernard De Bruyne, Nico H. J. Pijls, Guy R. Heyndrickx, Dominique Hodeige, Richard Kirkeeide, and K. Lance Gould. Pressure-Derived Fractional Flow Reserve to Assess Serial Epicardial Stenoses. *Circulation*, 101(15):1840–1847, April 2000.

[29] P J de Feyter, P W Serruys, M J Davies, P Richardson, J Lubsen, and M F Oliver. Quantitative coronary angiography to measure progression and regression of coronary atherosclerosis. Value, limitations, and implications for clinical trials. *Circulation*, 84(1):412–423, July 1991.

[30] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding, 2018.

[31] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020.

[32] Ronen Eldan and Ohad Shamir. The Power of Depth for Feedforward Neural Networks. *Journal of Machine Learning Research*, 49(June):907–940, 2016.

[33] Jingfan Fan, Jian Yang, Yachen Wang, Siyuan Yang, Danni Ai, Yong Huang, Hong Song, Aimin Hao, and Yongtian Wang. Multichannel Fully Convolutional Network for Coronary Artery Segmentation in X-ray Angiograms. *IEEE Access*, 6:44635–44643, August 2018.

[34] William F Fearon. Physiologic approach for coronary intervention. *The Korean journal of internal medicine*, 28(1):1–7, January 2013.

[35] William F. Fearon, Stephan Achenbach, Thomas Engstrom, Abid Assali, Richard Shlofmitz, Allen Jeremias, Stephane Fournier, Ajay J. Kirtane, Ran Kornowski, Gabriel Greenberg, Rami Jubeh, Daniel M. Kolansky, Thomas McAndrew, Ovidiu Dressler, Akiko Maehara, Mitsuaki Matsumura, Martin B. Leon, Bernard De Bruyne, and For the FAST-FFR Study Investigators. Accuracy of Fractional Flow Reserve Derived From Coronary Angiography. *Circulation*, January 2019.

[36] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale Vessel Enhancemenet Filtering. In William M. Wells, Alan and Colchester, and Scott and Delp, editors, *Medical Image Computing and Computer-Assisted Intervention — MICCAI'98*. Springer Berlin Heidelberg, 1998.

[37] Jun Fu, Jing Liu, Haijie Tian, and Yong Li. Dual Attention Network for Scene Segmentation. *arxiv*, page 9, 2018.

[38] Francesca Galassi, Mohammad Alkhalil, Regent Lee, Philip Martindale, Rajesh K. Kharbanda, Keith M. Channon, Vicente Grau, and Robin P. Choudhury. 3D reconstruction of coronary arteries from 2D angiographic projections using nonuniform rational basis splines (NURBS) for accurate modelling of coronary stenoses. *PLoS ONE*, 2018.

[39] Jo Etta Galbraith, Marvin L. Murphy, and Neil de Soyza. Coronary Angiogram Interpretation: Interobserver Variability. *JAMA*, 240(19):2053–2056, November 1978.

[40] Paolo Garrone, Giuseppe Biondi-Zoccai, Ilaria Salvetti, Noemi Sina, Imad Sheiban, Peter R. STELla, and Pierfrancesco Agostoni. Quantitative Coronary Angiography in the Current Era: Principles and Applications. *Journal of Interventional Cardiology*, 22(6):527–536, 2009.

[41] Arash Ghorbanniahassankiadeh, David S. Marks, and John F. LaDisa, Jr. Correlation of Computational Instantaneous Wave-Free Ratio With Fractional Flow Reserve for Intermediate Multivessel Coronary Disease. *Journal of Biomechanical Engineering*, 143(5), March 2021.

[42] C. Michael Gibson, Christopher P. Cannon, William L. Daley, J. Theodore Dodge, Barbara Alexander, Susan J. Marble, Carolyn H. McCabe, Lori Raymond, Terry Fortin, W. Kenneth Poole, and Eugene Braunwald. TIMI Frame Count. *Circulation*, 93(5):879–888, March 1996.

[43] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation, 2013.

[44] R K Goldberg, N S Kleiman, S T Minor, J Abukhalil, and A E Raizner. Comparison of quantitative coronary angiography to visual estimates of lesion severity pre and post PTCA. *American heart journal*, 119(1):178–184, January 1990.

[45] K L Gould, K O Kelley, and E L Bolson. Experimental validation of quantitative coronary arteriography for determining pressure-flow characteristics of coronary stenosis. *Circulation*, 66(5):930–937, November 1982.

[46] K. L. Gould, K. Lipscomb, and G. W. Hamilton. Physiologic basis for assessing critical coronary stenosis. Instantaneous flow response and regional distribution during coronary hyperemia as measures of coronary flow reserve. *The American Journal of Cardiology*, 33(1):87–94, January 1974.

[47] William Grossman. *Cardiac catheterization and angiography*. Lea and Febiger, Philadelphia, PA, 1986.

[48] David A. Hanauer, Qiaozhu Mei, James Law, Ritu Khanna, and Kai Zheng. Supporting information retrieval from electronic health records: A report of University of Michigan's nine-year experience in developing and using the Electronic Medical Record Search Engine (EMERSE). *Journal of Biomedical Informatics*, 55:290–300, June 2015.

[49] Bashar Hannawi, Wilson W. Lam, Suwei Wang, and George A. Younis. Current Use of Fractional Flow Reserve: A Nationwide Survey. *Texas Heart Institute Journal*, 41(6):579–584, December 2014.

[50] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

[51] A Hernandez-Vela, C Gatta, S Escalera, L Igual, V Martin-Yuste, M Sabate, and P Radeva. Accurate Coronary Centerline Extraction, Caliber Estimation, and Catheter Detection in Angiographies. *IEEE Transactions on Information Technology in Biomedicine*, 16(6):1332–1340, 2012.

[52] David M. Herrington, Maria Siebes, and Gary D. Walford. Sources of Error in Quantitative Coronary Angiography. *Catheterization and Cardiovascular Diagnosis*, 29(4):314–321, August 1993.

[53] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[54] Takeshi Hozumi, Kiyoshi Yoshida, Takashi Akasaka, Yoshio Asami, Yumiko Ogata, Tsutomu Takagi, Shuichiro Kaji, Takahiro Kawamoto, Yoshiaki Ueda, and Shigefumi Morioka. Noninvasive assessment of coronary flow velocity and coronary flow velocity reserve in the left anterior descending coronary artery by Doppler echocardiography. *Journal of the American College of Cardiology*, 32(5):1251–1259, November 1998.

[55] Kritika Iyer, Cyrus P. Najarian, Aya A. Fattah, Christopher J. Arthurs, S. M. Reza Soroushmehr, Vijayakumar Subban, Mullasari A. Sankardas, Raj R. Nadakuditi, Brahmajee K. Nallamothu, and C. Alberto Figueroa. AngioNet: A convolutional neural network for vessel segmentation in X-ray angiography. *Scientific Reports*, 11(1):18066, September 2021.

[56] Jun Jiang, Li Feng, Changling Li, Yongqing Xia, Jingsong He, Xiaochang Leng, Liang Dong, Xinyang Hu, Jian'an Wang, and Jianping Xiang. Fractional flow reserve for coronary stenosis assessment derived from fusion of intravascular ultrasound and X-ray angiography. *Quantitative Imaging in Medicine and Surgery*, 11(11):4543–4555, November 2021.

[57] Kyungmin Jo, Jihoon Kweon, Young-Hak Kim, and Jaesoon Choi. Segmentation of the Main Vessel of the Left Anterior Descending Artery Using Selective Feature Mapping in Coronary Angiography. *IEEE Access*, 7:919–930, 2019.

[58] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, TaoXie, Jiacong Fang, imyhxy, Kalen Michael, Lorna, Abhiram V, Diego Montes, Jebastin Nadar, Laughing, tkianai, yxNONG, Piotr Skalski, Zhiqiang Wang, Adam Hogan, Cristi Fati, Lorenzo Mammana, AlexWang1900, Deep Patel, Ding Yiwei, Felix You, Jan Hajek, Laurentiu Diaconu, and Mai Thanh Minh. Ultralytics/yolov5: V6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference. Zenodo, February 2022.

[59] Ghassan S. Kassab. Morphometry of Coronary Vasculature. In Ghassan S. Kassab, editor, *Coronary Circulation: Anatomy, Mechanical Properties, and Biomechanics*, pages 29–104. Springer International Publishing, Cham, 2019.

[60] Asma Kerkeni, Asma Benabdallah, Antoine Manzanera, and Mohamed Hedi Bedoui. A coronary artery segmentation method based on multiscale analysis and region growing. *Computerized Medical Imaging and Graphics*, 48:49–61, March 2016.

[61] András P. Keszei, Benjamin Berkels, and Thomas M. Deserno. Survey of Non-Rigid Registration Tools in Medicine. *Journal of Digital Imaging*, 30(1):102–116, February 2017.

[62] Mahdieh Khanmohammadi, Kjersti Engan, Charlotte Sæland, Trygve Eftestøl, and Alf I. Larsen. Automatic Estimation of Coronary Blood Flow Velocity Step 1 for Developing a Tool to Diagnose Patients With Micro-Vascular Angina Pectoris. *Frontiers in Cardiovascular Medicine*, 6:1, January 2019.

[63] Nancy Khav, Abdul Rahman Ihdayhid, and Brian Ko. CT-Derived Fractional Flow Reserve (CT-FFR) in the Evaluation of Coronary Artery Disease. *Heart, Lung and Circulation*, 29(11):1621–1632, November 2020.

[64] H. J. Kim, Irene E. Vignon-Clementel, J. S. Coogan, C. Alberto Figueroa, K. E. Jansen, and C. A. Taylor. Patient-specific modeling of blood flow and pressure in human coronary arteries. *Annals of Biomedical Engineering*, 38(10):3195–3209, 2010.

[65] Ran Kornowski, Ifat Lavi, Mariano Pellicano, Panagiotis Xaplanteris, Hana Vaknin-Assa, Abid Assali, Orna Valtzer, Yonit Lotringer, and Bernard De Bruyne. Fractional Flow Reserve Derived From Routine Coronary Angiograms. *Journal of the American College of Cardiology*, 68(20):2235–2237, November 2016.

[66] Vijayalakshmi Kunadian, Caitlin Harrigan, Cafer Zorkun, Alexandra M. Palmer, Katherine J. Ogando, Leah H. Biller, Erin E. Lord, Scott P. Williams, Michelle E. Lew, Lauren N. Ciaglo, Jacqueline L. Buros, Susan J. Marble, William J. Gibson, and C. Michael Gibson. Use of the TIMI frame count in the assessment of coronary artery blood flow and microvascular function over the past 15 years. *Journal of Thrombosis and Thrombolysis*, 27(3):316, April 2008.

[67] Daniel S D Lara, Alexandre W C Faria, Arnaldo De A Araújo, and D Menotti. A Novel Hybrid Method for the Segmentation of the Coronary Artery Tree in 2D Angiograms. *International Journal of Computer Science & Information Technology (IJCSIT)*, 5(3), 2013.

[68] Jen-Shih Lee and Yuan-Cheng Fung. Flow in Locally Constricted Tubes at Low Reynolds Numbers. *Journal of Applied Mechanics*, March 1970.

[69] Jonathon Leipsic, Tae-Hyun Yang, Angus Thompson, Bo-Kwon Koo, G. B. John Mancini, Carolyn Taylor, Matthew J. Budoff, Hyung-Bok Park, Daniel S. Berman, and James K. Min. CT Angiography (CTA) and Diagnostic Performance of Noninvasive Fractional Flow Reserve: Results From the Determination of Fractional Flow Reserve by Anatomic CTA (DeFACTO) Study. *American Journal of Roentgenology*, 202(5):989–994, May 2014.

[70] Harald Lethen, Hans P. Tries, Stefan Kersting, and Heinz Lambertz. Validation of noninvasive assessment of coronary flow velocity reserve in the right coronary artery: A comparison of transthoracic echocardiographic results with intracoronary Doppler flow wire measurements. *European Heart Journal*, 24(17):1567–1575, 09 2003.

[71] Jianping Li, Yanjun Gong, Weimin Wang, Qing Yang, Bin Liu, Yuan Lu, Yawei Xu, Yunlong Huo, Tieci Yi, Jian Liu, Yongle Li, Shaopeng Xu, Lei Zhao, Ziad A Ali, and Yong Huo. Accuracy of computational pressure-fluid dynamics applied to coronary angiography to derive fractional flow reserve: FLASH FFR. *Cardiovascular Research*, 116(7):1349–1356, June 2020.

[72] Rui Liao, Duong Luc, Yiyong Sun, and Klaus Kirchberg. 3-D reconstruction of the coronary artery tree from multiple views of a rotational X-ray angiography. *The International Journal of Cardiovascular Imaging*, 26(7):733–749, October 2010.

[73] Chih Yang Lin and Yu Tai Ching. Extraction of Coronary Arterial Tree Using Cine X-ray Angiograms. *Biomedical Engineering - Applications, Basis and Communications*, 17(3):111–120, June 2005.

[74] Yongge Liu, Jianzhuang Yu, and Yahong Han. Understanding the effective receptive field in semantic image segmentation. *Multimedia Tools and Applications*, 77(17):22159–22171, September 2018.

[75] Jiasen Lu, Caiming Xiong, Devi Parikh, and Richard Socher. Knowing when to look: Adaptive attention via A visual sentinel for image captioning. *CoRR*, abs/1612.01887, 2016.

[76] Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The Expressive Power of Neural Networks: A View from the Width. In *31st Conference on Neural Information Processing Systems*, 2017.

[77] Guangkun Ma, Jinzhu Yang, and Hong Zhao. A coronary artery segmentation method based on region growing with variable sector search area. *Technology and Health Care*, 28(S1):463–472, January 2020.

[78] Hua Ma, Ayla Hoogendoorn, Evelyn Regar, Wiro J. Niessen, and Theo van Walsum. Automatic online layer separation for vessel enhancement in X-ray angiograms for percutaneous coronary interventions. *Medical Image Analysis*, 39:145–161, July 2017.

[79] Zbigniew Malota, Jan Glowacki, Wojciech Sadowski, and Marcin Kostur. Numerical analysis of the impact of flow rate, heart rate, vessel geometry, and degree of stenosis on coronary hemodynamic indices. *BMC Cardiovascular Disorders*, 18(1):132, June 2018.

[80] Thomas H. Marwick, Iksung Cho, Bríain ó Hartaigh, and James K. Min. Finding the gatekeeper to the cardiac catheterization laboratory. *Journal of the American College of Cardiology*, 65(25):2747–2756, 2015.

[81] R E Mates, R L Gupta, A C Bell, and F J Klocke. Fluid dynamics of coronary artery stenosis. *Circulation Research*, 42(1):152–162, January 1978.

[82] Laura Mauri, A. James O'Malley, Donald E. Cutlip, Kalon K. L. Ho, Jeffrey J. Popma, Manish S. Chauhan, Donald S. Baim, David J. Cohen, and Richard E. Kuntz. Effects of stent length and lesion length on coronary restenosis. *The American Journal of Cardiology*, 93(11):1340–1346, A5, June 2004.

[83] Johns Hopkins Medicine. Anatomy and Function of the Coronary Arteries. https://www.hopkinsmedicine.org/health/conditions-and-diseases/anatomy-and-function-of-the-coronary-arteries, November 2020.

[84] Patrick Meimoun and Christophe Tribouilloy. Non-invasive assessment of coronary flow and coronary flow reserve by transthoracic Doppler echocardiography: A magic tool for the real world. *European Journal of Echocardiography*, 9(4):449–457, July 2008.

[85] Faten M'Hiri, Luc Duong, Christian Desrosiers, and Mohamed Cheriet. Vesselwalker: Coronary Arteries Segmentation Using Random Walks and Hessian-based Vesselness Filter. In *Proceedings - International Symposium on Biomedical Imaging*, pages 918–921, 2013.

[86] James K. Min, Jonathon Leipsic, Michael J. Pencina, Daniel S. Berman, Bon-Kwon Koo, Carlos van Mieghem, Andrejs Erglis, Fay Y. Lin, Allison M. Dunning, Patricia Apruzzese, Matthew J. Budoff, Jason H. Cole, Farouc A. Jaffer, Martin B. Leon, Jennifer Malpeso, G. B. John Mancini, Seung-Jung Park, Robert S. Schwartz, Leslee J. Shaw, and Laura Mauri. Diagnostic Accuracy of Fractional Flow Reserve From Anatomic CT Angiography. *Jama*, 308(12):1237, 2012.

[87] Bhavik N. Modi, Sethuraman Sankaran, Hyun Jin Kim, Howard Ellis, Campbell Rogers, Charles A. Taylor, Ronak Rajani, and Divaka Perera. Predicting the Physiological Effect of Revascularization in Serially Diseased Coronary Arteries. *Circulation: Cardiovascular Interventions*, 12(2):e007577, February 2019.

[88] Sabee Molloi, David Chalyan, Huy Le, and Jerry T. Wong. Estimation of coronary artery hyperemic blood flow based on arterial lumen volume using angiographic images. *The International Journal of Cardiovascular Imaging*, 28(1):1–11, January 2012.

[89] Jong Hak Moon, Da Young Lee, Won Chul Cha, Myung Jin Chung, Kyu-Sung Lee, Baek Hwan Cho, and Jin Ho Choi. Automatic stenosis recognition from coronary angiography using convolutional neural networks. *Computer Methods and Programs in Biomedicine*, 198:105819, January 2021.

[90] Brahmajee K. Nallamothu, John A. Spertus, Alexandra J. Lansky, David J. Cohen, Philip G. Jones, Faraz Kureshi, Gregory J. Dehmer, Joseph P. Drozda, Mary Norine Walsh, John E. Brush, Gerald C. Koenig, Thad F. Waites, D. Scott Gantt, George Kichura, Richard A. Chazal, Peter K. O'Brien, C. Michael Valentine, John S. Rumsfeld, Johan H. C. Reiber, Joann G. Elmore, Richard A. Krumholz, W. Douglas Weaver, and Harlan M. Krumholz. Comparison of clinical interpretation with visual assessment and quantitative coronary angiography in patients undergoing percutaneous coronary intervention in contemporary practice: The Assessing Angiography (A2) project. *Circulation*, 127(17):1793–1800, April 2013.

[91] E. Nasr-Esfahani, N. Karimi, M. H. Jafari, S. M.R. Soroushmehr, S. Samavi, B. K. Nallamothu, and K. Najarian. Segmentation of Vessels in Angiograms using Convolutional Neural Networks. *Biomedical Signal Processing and Control*, 40:240–251, February 2018.

[92] DICOM NEMA. Dicom ps3.17 2022b - explanatory information.

[93] Wilmer W. Nichols, Michael F. O'Rourke, Charalambos. Vlachopoulos, and Donald A. McDonald. *McDonald's Blood Flow in Arteries : Theoretical, Experimental and Clinical Principles*. Hodder Arnold, London, 2011.

[94] Koen Nieman, Benno J. Rensing, Robert-Jan M. van Geuns, Arie Munne, Jurgen M.R. Ligthart, Peter M.T. Pattynama, Gabriel P. Krestin, Patrick W. Serruys, and Pim J. de Feyter. Usefulness of multislice computed tomography for detecting obstructive coronary artery disease. *The American Journal of Cardiology*, 89(8):913–918, 2002.

[95] Arne Nordmann. Epipolar geometry, 2007.

[96] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, Ben Glocker, and Daniel Rueckert. Attention u-net: Learning where to look for the pancreas, 2018.

[97] Kun Pang, Danni Ai, Huihui Fang, Jingfan Fan, Hong Song, and Jian Yang. Stenosis-DetNet: Sequence consistency-based stenosis detection for X-ray coronary angiography. *Computerized Medical Imaging and Graphics*, 89:101900, April 2021.

[98] Thrasyvoulos Pappas and Jae S. Lim. A New Method for Estimation of Coronary Artery Dimensions in Angiograms. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 36(9):1501–1513, 1988.

[99] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A Deep Neural Network Architecture for Real-Time Semantic Segmentation. *arXiv*, June 2016.

[100] Mariano Pellicano, Ifat Lavi, Bernard De Bruyne, Hana Vaknin-Assa, Abid Assali, Orna Valtzer, Yonit Lotringer, Giora Weisz, Yaron Almagor, Panagiotis Xaplanteris, Ajay Kirtane, Pablo Codner, Martin B Leon, and Ran Kornowski. Validation Study of Image-Based Fractional Flow Reserve During Coronary Angiography. *Circ Cardiovasc Interv*, 2017.

[101] N. H. Pijls, B. De Bruyne, K. Peels, P. H. Van Der Voort, H. J. Bonnier, J. J. Bartunek J Koolen, and J. J. Koolen. Measurement of fractional flow reserve to assess the functional severity of coronary-artery stenoses. *The New England Journal of Medicine*, 334(26):1703–1708, June 1996.

[102] Nico H. J. Pijls and Pim A. L. Tonino. The Crux of Maximum Hyperemia. *JACC: Cardiovascular Interventions*, 4(10):1093–1095, October 2011.

[103] Nico H.J. Pijls, William F. Fearon, Pim A.L. Tonino, Uwe Siebert, Fumiaki Ikeno, Bernhard Bornschein, Marcel van't Veer, Volker Klauss, Ganesh Manoharan, Thomas Engstrøm, Keith G. Oldroyd, Peter N. Ver Lee, Philip A. MacCarthy, and Bernard De Bruyne. Fractional Flow Reserve Versus Angiography for Guiding Percutaneous Coronary Intervention in Patients With Multivessel Coronary Artery Disease: 2-Year Follow-Up of the FAME (Fractional Flow Reserve Versus Angiography for Multivessel Evaluation) Study. *Journal of the American College of Cardiology*, 56(3):177–184, July 2010.

[104] Tobias Pohlen, Alexander Hermans, Markus Mathias, and Bastian Leibe. Full-Resolution Residual Networks for Semantic Segmentation in Street Scenes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[105] Binjie Qin, Mingxin Jin, Dongdong Hao, Yisong Lv, Qiegen Liu, Yueqi Zhu, Song Ding, Jun Zhao, and Baowei Fei. Accurate vessel extraction via tensor completion of background layer in X-ray coronary angiograms. *Pattern recognition*, 87:38–54, March 2019.

[106] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl Dickstein. On the Expressive Power of Deep Neural Networks. *34th International Conference on Machine Learning, ICML 2017*, 6:4351–4374, 2017.

[107] Udaya Ralapanawa and Ramiah Sivakanesan. Epidemiology and the Magnitude of Coronary Artery Disease and Acute Coronary Syndrome: A Narrative Review. *Journal of Epidemiology and Global Health*, 11(2):169–177, June 2021.

[108] S. D. Ramaswamy, S. C. Vigmostad, A. Wahle, Y. G. Lai, M. E. Olszewski, K. C. Braddy, T. M. H. Brennan, J. D. Rossen, M. Sonka, and K. B. Chandran. Fluid Dynamic Analysis in a Human Left Anterior Descending Coronary Artery with Arterial Motion. *Annals of Biomedical Engineering*, 32(12):1628–1641, December 2004.

[109] Waseem Rawat and Zenghui Wang. Deep Convolutional Neural Networks for Image Classification: A Comprehensive Review. *Neural Computation*, 29(9):2352–2449, September 2017.

[110] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection, 2015.

[111] Joseph Redmon and Ali Farhadi. YOLOv3: An Incremental Improvement. *arXiv:1804.02767 [cs]*, April 2018.

[112] Johan H.C. Reiber. An overview of coronary quantitation techniques as of 1989. In Johan H. C. Reiber and Patrick W. Serruys, editors, *Quantitative Coronary Arteriography*, pages 55–132. Kluwer Academic Publishers, Dordrecht, 1997.

[113] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks, 2015.

[114] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

[115] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *MICCAI 2015: Medical Image Computing and Computer-Assisted Intervention*, pages 234–241, May 2015.

[116] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, pages 65–386, 1958.

[117] John A. Rumberger. Coronary Artery Disease: A Continuum, Not a Threshold. *Mayo Clinic Proceedings*, 92(3):323–326, March 2017.

[118] T. Saito, M. Misaki, K. Shirato, and T. Takishima. Three-dimensional quantitative coronary angiography. *IEEE Transactions on Biomedical Engineering*, 37(8):768–777, August 1990.

[119] Shingo Sakamoto, Saeko Takahashi, Ahmet U. Coskun, Michail I. Papafaklis, Akihiko Takahashi, Shigeru Saito, Peter H. Stone, and Charles L. Feldman. Relation of Distribution of Coronary Blood Flow Volume to Coronary Artery Dominance. *The American Journal of Cardiology*, 111(10):1420–1424, May 2013.

[120] Seyed Sadegh Mohseni Salehi, Deniz Erdogmus, and Ali Gholipour. Tversky Loss Function for Image Segmentation Using 3D Fully Convolutional Deep Networks. In Qian Wang, Yinghuan Shi, Heung-Il Suk, and Kenji Suzuki, editors, *Machine Learning in Medical Imaging*, Lecture Notes in Computer Science, pages 379–387, Cham, 2017. Springer International Publishing.

[121] P.M. Samuel and T. Veeramalai. VSSC Net: Vessel Specific Skip chain Convolutional Network for blood vessel segmentation. *Computer Methods and Programs in Biomedicine*, 198, 2021.

[122] Fabian Sanchis-Gomar, Carme Perez-Quilis, Roman Leischik, and Alejandro Lucia. Epidemiology of coronary heart disease and acute coronary syndrome. *Annals of translational medicine*, 4(13):256, July 2016.

[123] R. H. Selzer, C. Hagerty, S. P. Azen, M. Siebes, P. Lee, A. Shircore, and D. H. Blankenhorn. Precision and reproducibility of quantitative coronary angiography with applications to controlled clinical trials. A sampling study. https://www.jci.org/articles/view/113913/scanned-page/520, February 1989.

[124] Sayan Sen, Javier Escaned, Iqbal S. Malik, Ghada W. Mikhail, Rodney A. Foale, Rafael Mila, Jason Tarkin, Ricardo Petraco, Christopher Broyd, Richard Jabbour, Amarjit Sethi, Christopher S. Baker, Micheal Bellamy, Mahmud Al-Bustami, David Hackett, Masood Khan, David Lefroy, Kim H. Parker, Alun D. Hughes, Darrel P. Francis, Carlo Di Mario, Jamil Mayet, and Justin E. Davies. Development and Validation of a New Adenosine-Independent Index of Stenosis Severity From Coronary Wave–Intensity Analysis: Results of the ADVISE (ADenosine Vasodilator Independent Stenosis Evaluation) Study. *Journal of the American College of Cardiology*, 59(15):1392–1402, April 2012.

[125] Rohan Shah, Eric Yow, William Schuyler Jones, Louis P. Kohl, Andrzej S. Kosinski, Udo Hoffmann, Kerry L. Lee, Christopher B. Fordyce, Daniel B. Mark, Alicia Lowe, Pamela S. Douglas, and Manesh R. Patel. Comparison of visual assessment of coronary stenosis with independent quantitative coronary angiography: Findings from the Prospective Multicenter Imaging Study for Evaluation of Chest Pain (PROMISE) trial. *American Heart Journal*, 184:1–9, February 2017.

[126] Yanfeng Shang, Rudi Deklerck, Edgard Nyssen, Aneta Markova, Johan De Mey, Xin Yang, and Kun Sun. Vascular Active Contour for Vessel Tree Segmentation. *IEEE Transactions on Biomedical Engineering*, 58(4):1023–1032, April 2011.

[127] Shashank, Mahua Bhattacharya, and G.K. Sharma. Optimized Coronary Artery Segmentation Using Frangi Filter and Anisotropic Diffusion Filtering. In *2013 International Symposium on Computational and Business Intelligence*, pages 261–264, August 2013.

[128] Seung Yeon Shin, Soochahn Lee, Il Dong Yun, and Kyoung Mu Lee. Deep Vessel Segmentation By Learning Graphical Connectivity. *arXiv:1806.02279 [cs]*, June 2018.

[129] Suprosanna Shit, Johannes C. Paetzold, Anjany Sekuboyina, Ivan Ezhov, Alexander Unger, Andrey Zhylka, Josien P. W. Pluim, Ulrich Bauer, and Bjoern H. Menze. clDice - a novel

topology-preserving loss function for tubular structure segmentation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2021.

[130] Simon D. Shpilfoygel, Robert A. Close, Daniel J. Valentino, and Gary R. Duckwiler. X-ray videodensitometric methods for blood flow and velocity measurement: A critical review of literature. *Medical Physics*, 27(9):2008–2023, 2000.

[131] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv:1409.1556 [cs]*, April 2015.

[132] Per A. Sirnes, Yngvar Myreng, Per Mølstad, and Svein Golf. Reproducibility of quantitative coronary analysis. *The International Journal of Cardiac Imaging*, 12(3):197–203, September 1996.

[133] S. Yu. Sokolov, S. O. Volchkov, I. S. Bessonov, V. V. Chestukhin, G. V. Kurlyandskaya, and F. A. Blyakhman. Estimation of Blood Flow Velocity in Coronary Arteries Based on the Movement of Radiopaque Agent. *Pattern Recognition and Image Analysis*, 29(4):750–762, October 2019.

[134] Richard Szeliski. *Computer Vision*. Springer London, London, 2011.

[135] Irina Andra Tache and Dimitrios Glotsos. Vessel Segmentation and Stenosis Quantification from Coronary X-Ray Angiograms. In Ruidan Su and Han Liu, editors, *Medical Imaging and Computer-Aided Diagnosis*, Lecture Notes in Electrical Engineering, pages 27–34, Singapore, 2020. Springer.

[136] Chun Xiang Tang, Yi Ning Wang, Fan Zhou, U. Joseph Schoepf, Marly van Assen, Robert E. Stroud, Jian Hua Li, Xiao Lei Zhang, Meng Jie Lu, Chang Sheng Zhou, Dai Min Zhang, Yan Yi, Jing Yan, Guang Ming Lu, Lei Xu, and Long Jiang Zhang. Diagnostic performance of fractional flow reserve derived from coronary CT angiography for detection of lesion-specific ischemia: A multi-center study and meta-analysis. *European Journal of Radiology*, 116:90–97, July 2019.

[137] C. A. Taylor, Timothy A. Fonte, and James K. Min. Computational fluid dynamics applied to cardiac computed tomography for noninvasive quantification of fractional flow reserve: Scientific basis. *Journal of the American College of Cardiology*, 61(22):2233–2241, 2013.

[138] Alexandru Telea and Jarke J Van Wijk. An Augmented Fast Marching Method for Computing Skeletons and Centerlines. In D Ebert, P Bruent, and I Navazo, editors, *EUROGRAPHICS - IEEE TCVG Symposium on Visualization*, Barcelona, 2002.

[139] R Torii, J Keegan, N B Wood, A W Dowsey, A D Hughes, G-Z Yang, D N Firmin, S A Mcg Thom, and X Y Xu. The effect of dynamic vessel motion on haemodynamic parameters in the right coronary artery: A combined MR and CFD study. *The British Journal of Radiology*, 82(special_issue_1):S24–S32, January 2009.

[140] Shengxian Tu, Emanuele Barbato, Zsolt Köszegi, Junqing Yang, Zhonghua Sun, Niels R. Holm, Balázs Tar, Yingguang Li, Dan Rusinaru, William Wijns, and Johan H. C. Reiber.

Fractional Flow Reserve Calculation From 3-Dimensional Quantitative Coronary Angiography and TIMI Frame Count: A Fast Computer Model to Quantify the Functional Significance of Moderately Obstructed Coronary Arteries. *JACC: Cardiovascular Interventions*, 7(7):768–777, July 2014.

[141] Shengxian Tu, Daixin Ding, Yunxiao Chang, Chunming Li, William Wijns, and Bo Xu. Diagnostic accuracy of quantitative flow ratio for assessment of coronary stenosis significance from a single angiographic view: A novel method based on bifurcation fractal law. *Catheterization and Cardiovascular Interventions*, 97(S2):1040–1047, 2021.

[142] Shengxian Tu, Jelmer Westra, Junqing Yang, Clemens von Birgelen, Angela Ferrara, Mariano Pellicano, Holger Nef, Matteo Tebaldi, Yoshinobu Murasato, Alexandra Lansky, Emanuele Barbato, Liefke C. van der Heijden, Johan H.C. Reiber, Niels R. Holm, and William Wijns. Diagnostic Accuracy of Fast Computational Approaches to Derive Fractional Flow Reserve From Diagnostic Coronary Angiography: The International Multicenter FAVOR Pilot Study. *JACC: Cardiovascular Interventions*, 9(19):2024–2035, October 2016.

[143] Alina G. van der Giessen, Harald C. Groen, Pierre-André Doriot, Pim J. de Feyter, Antonius F. W. van der Steen, Frans N. van de Vosse, Jolanda J. Wentzel, and Frank J. H. Gijsen. The influence of boundary conditions on wall shear stress distribution in patients specific coronary trees. *Journal of Biomechanics*, 44(6):1089–1095, April 2011.

[144] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.

[145] Arso M. Vukicevic, Serkan Çimen, Nikola Jagic, Gordana Jovicic, Alejandro F. Frangi, and Nenad Filipovic. Three-dimensional reconstruction and NURBS-based structured meshing of coronary arteries from the conventional X-ray angiography projection images. *Scientific Reports*, 8(1), December 2018.

[146] Bruce F. Waller, Charles M. Orr, John D. Slack, Cass A. Pinkerton, James Van Tassel, and Thomas Peters. Anatomy, histology, and pathology of coronary arteries: A review relevant to new interventional and imaging techniques—Part I. *Clinical Cardiology*, 15(6):451–457, 1992.

[147] Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding Convolution for Semantic Segmentation. In *Proceedings - 2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018*, volume 2018-Janua, pages 1451–1460. Institute of Electrical and Electronics Engineers Inc., May 2018.

[148] Geoffrey B. West, James H. Brown, and Brian J. Enquist. A General Model for the Origin of Allometric Scaling Laws in Biology. *Science*, 276(5309):122–126, April 1997.

[149] Shaoyan Xia, Haogang Zhu, Xiaoli Liu, Ming Gong, Xiaoyong Huang, Lei Xu, Hongjia Zhang, and Jialong Guo. Vessel Segmentation of X-Ray Coronary Angiographic Image Sequence. *IEEE Transactions on Biomedical Engineering*, 67(5):1338–1348, May 2020.

[150] Bo Xu, Shengxian Tu, Shubin Qiao, Xinkai Qu, Yundai Chen, Junqing Yang, Lijun Guo, Zhongwei Sun, Zehang Li, Feng Tian, Weiyi Fang, Jiyan Chen, Wei Li, Changdong Guan, Niels R. Holm, William Wijns, and Shengshou Hu. Diagnostic Accuracy of Angiography-Based Quantitative Flow Ratio Measurements for Online Assessment of Coronary Stenosis. *Journal of the American College of Cardiology*, 70(25):3077–3087, 2017.

[151] Chenyang Xu and Jerry L. Prince. Generalized gradient vector flow external forces for active contours. *Signal Processing*, 71(2):131–139, December 1998.

[152] Jian Yang, Weijian Cong, Yang Chen, Jingfan Fan, Yue Liu, and Yongtian Wang. External force back-projective composition and globally deformable optimization for 3-D coronary artery reconstruction. *Physics in Medicine and Biology*, 59(4):975–1003, February 2014.

[153] Su Yang, Jihoon Kweon, Jae Hyung Roh, Jae Hwan Lee, Heejun Kang, Lae Jeong Park, Dong Jun Kim, Hyeonkyeong Yang, Jaehee Hur, Do Yoon Kang, Pil Hyung Lee, Jung Min Ahn, Soo Jin Kang, Duk Woo Park, Seung Whan Lee, Young Hak Kim, Cheol Whan Lee, Seong Wook Park, and Seung Jung Park. Deep learning segmentation of major vessels in X-ray coronary angiography. *Scientific Reports*, 9(1):1–11, December 2019.

[154] Emil Zakirov. Keras-deeplab-v3-plus. https://github.com/bonlime/keras-deeplab-v3-plus, 2019.

[155] Jun-Mei Zhang, Liang Zhong, Boyang Su, Min Wan, Jinq Shya Yap, Jasmine P. L. Tham, Leok Poh Chua, Dhanjoo N. Ghista, and Ru San Tan. Perspective on CFD studies of coronary artery disease lesions and hemodynamics: A review. *International Journal for Numerical Methods in Biomedical Engineering*, 30(6):659–680, 2014.

[156] Tongjie Y. Zhang and Ching Y. Suen. A fast parallel algorithm for thinning digital patterns. *Communications of the ACM*, 27(3):236–239, 1984.

[157] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid Scene Parsing Network. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6230–6239, December 2017.

[158] Sun Zheng, Tian Meiying, and Sun Jian. Sequential reconstruction of vessel skeletons from X-ray coronary angiographic sequences. *Computerized Medical Imaging and Graphics*, 34(5):333–345, July 2010.

[159] L M Zir, S W Miller, R E Dinsmore, J P Gilbert, and J W Harthorne. Interobserver variability in coronary angiography. *Circulation*, 53(4):627–632, April 1976.