

Computational Processing and Modeling of Intravascular Images Precisely Couple Arterial Morphology and Biomechanics

by

Max Louis Olender

M.S.E. Biomedical Engineering
University of Michigan, 2016

B.S.E. Mechanical Engineering
University of Michigan, 2015

SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING IN PARTIAL
FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

AT THE

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

FEBRUARY 2021

© 2021 Massachusetts Institute of Technology. All rights reserved.

Signature of Author:

Department of Mechanical Engineering
January 15, 2021

Certified by.....

Elazer R. Edelman
Edward J. Poitras Professor in Medical Engineering and Science
Thesis Supervisor

Certified by.....

Peter T. C. So
Professor of Mechanical and Biological Engineering
Thesis Committee Chair

Accepted by.....

Nicolas G. Hadjiconstantinou
Professor of Mechanical Engineering
Chairman, Committee on Graduate Students

Computational Processing and Modeling of Intravascular Images Precisely Couple Arterial Morphology and Biomechanics

by

Max Louis Olender

Submitted to the Department of Mechanical Engineering
on January 15, 2021 in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy in
Mechanical Engineering

ABSTRACT

Cardiovascular diseases, and coronary artery disease in particular, remain a persistent devastating and prevalent menace to health and wellbeing globally despite great strides in vascular biology and medicine. While biomechanical forces are known to play a driving role in the natural history of atherosclerosis, the nuanced yet profound impact of patient- and lesion-specific biomechanics in disease presentation, course, and treatment are not fully appreciated or accounted for in clinical practice. The incredible strides in melding image processing with artificial intelligence, computational modeling, and numerical methods is increasingly filling gaps in knowledge, especially at the intersection of pathological anatomy and biomechanical structural behavior.

We derived geometric and morphological structure, as well as constitutive material properties, from invasive intravascular image sequences to quantitatively assess and characterize the state of atherosclerotic arteries. Overcoming the challenge of limited penetration depth in the presence of signal-attenuating plaque, contextual information and spatial continuity was leveraged by a novel surface fitting method to fully delineate the mural conformation of the diseased vessel wall. Neural networks enriched with domain knowledge of vascular geometry and imaging classified pathological regions of interest within heterogeneous lesions. Construction of *in silico* computational models and *in vitro* phantom models facilitated the execution and validation of inverse methods to determine material constitutive mechanical properties non-destructively and in clinically amenable fashion. Strategic simplifying assumptions freed the approach from data acquisition limitations which inhibited previous methods of *in situ* mechanical characterization. Finally, to bridge the chasm between virtual and physical medicine and facilitate integration of these new capabilities into clinical practice, synthetic images were generated by an adversarial network trained in the familiar visual vernacular of vascular imaging.

Through the insights described in this thesis, greater information can be extracted, augmented, and made accessible from clinically-available imaging data. Approaches to more quantitatively and reliably assess, model, and convey biomechanical disease states may offer mechanistic insight into disease development, progression, and treatment response, ultimately leading to improved personalized patient care in an emerging era of computational cardiology.

Thesis Supervisor: Elazer R. Edelman, M.D., Ph.D.

Title: Director, Institute for Medical Engineering and Science & Edward J. Poitras Professor in Medical Engineering and Science

Acknowledgments

The work of this dissertation, and the range of experiences which permeate and extend beyond its pages, have been made possible and enriched beyond measure by the involvement of countless individuals. To each, not limited to those named here, I am eternally grateful.

I could not have been more fortunate than to have my advisor, Professor Elazer Edelman, to guide and join me through the journey of my PhD studies. His genuine and sincere care for all under his tutelage pervades every interaction and supersedes all professional considerations. I have benefited greatly not only from his magnanimity, but also his invaluable technical insight, advice, unrelenting support, and positive encouragement.

I am also incredibly grateful for the support and contributions of my thesis committee, comprised of a remarkable and inspiring collection of individuals. Professor Peter So, my committee chair, has asked insightful, challenging, and illuminating questions at each step, from my qualifying examination through the formalization and completion of my dissertation research. Professor Ellen Roche's practical and pragmatic perspective, informed by an incredible constellation of educational and professional experiences, shone through in her creative and innovative recommendations that elevated the quality of my work. Professor Guillermo (Gary) Tearney's expertise as a founding figure in intravascular imaging could not have been more relevant, and his deep familiarity with the field, historical insight, and strong vision as an active researcher were invaluable in determining the direction of my own research. His singular focus on translation and patient impact also helped to shape my perspective and direction. Finally, Professor Dimitrios Fotiadis, who graciously volunteered to serve as an external reviewer and examiner of my thesis, offered expertise and insight as a leader in the fields of biomedical and health informatics, co-author, colleague, and mentor that has made me a better researcher.

I have been indelibly impacted by the lasting friendships I have been privileged to form with those whose time in the Edelman Lab has intersected with my own. An account of even a fraction of these could match the length of this voluminous dissertation. Their dedication, commitment, and initiative was inspiring and motivating, but more importantly their companionship within and beyond the lab brought me immeasurable happiness. Fellow graduate students, both long-time comrades and visiting, have filled my mind with new perspectives and fond memories, and shared in the daily struggles, triumphs, and grind of research. From simple coffee and lunch breaks, to concerts and camping trips, the slog of graduate school was constantly studded with the joys of uplifting interactions.

There are a few individuals from the Lab who made enormous contributions to my development as a researcher and to this dissertation. Drs. Farhad Rikhtegar Nezami and Lambros Athanasiou, who initially convinced me to join the research group (for which I am indebted in perpetuity), continued to be central figures throughout my graduate experience. I feel incredibly fortunate to have had the opportunity to work and grown alongside both of them on a nearly daily basis, and to have benefited so much from their personal investment in my success. As technical experts and personal mentors, but foremost as trusted friends, they have guided me through and indelibly shaped my past 4+ years. Drs. José (Chema) de la Torre Hernández and Eyal Ben-Assa contributed

similarly in the most formative, early years of my graduate studies. Offering not just their expert medical expertise, clinical perspective, and patient tutoring, but also their warm, open-armed friendship, I am grateful for these deep relationships that have persisted over time and across oceans.

Beyond the laboratory, several communities have added dimension to my experience. My cohorts in the Mechanical Engineering, Health Sciences and Technology, and Graduate Education in Medical Sciences programs, as well as supporting administrators and faculty, have provided an accommodating and supportive academic home. The MIT Science Policy Initiative, Outing Club, MakerWorkshop, GradHillel, and Graduate Student Leadership Institute communities have offered opportunities and outlets to learn, grow, and find fulfillment.

Above all, though, I thank my family. Even when I lacked the patience to explain the details of my research, they were always convinced that my work was important and valuable, and offered their unwavering support. With absolute confidence and interest, they have listened with patience and offered encouragement, sharing in my successes and disappointments throughout my graduate studies, just as they have for all endeavors throughout my life.

Finally, I'd also like to acknowledge that I've been fortunate to receive generous support from various funding sources that have enabled me to pursue this work and enriching experiences throughout my graduate studies. I am grateful to have been supported by the MathWorks Engineering Fellowship, William Asbjornsen Albert Memorial Fellowship, and Whitaker Health Sciences Fund Fellowship, and to have received funding from the U.S. National Institutes of Health through a research grant to Dr. Edelman.

Table of Contents

Acknowledgments	3
Table of Contents	5
Introduction: Biomechanical Modeling & Assessment of Diseased Coronary Arteries	7
1.1. Disease of the Cardiovascular System.....	8
1.2. Imaging of Diseased Vessels: A Pillar of Clinical Practice.....	11
1.3. Arterial Modeling & Simulation: Promise & Challenges	14
1.4. Artificial Intelligence in Cardiology	19
1.5. Thesis Contributions & Organization.....	21
Extracting Vessel Geometry: Overcoming Current Limitations of Intravascular Imaging 25	
2.1. Diseased Vessel Geometry: Needs, Progress, & Challenges	26
2.2. Detecting Visible Segments of the Outer Border.....	36
2.3. Identifying the Complete Border from Visible Segments	44
2.4. Extracting or Tracking Geometry from Disparate Datasets	91
2.5. Implications & Future Directions.....	97
2.6. Conclusion.....	101
Determining & Characterizing Vessel Morphology	103
3.1. Diseased Vessel Morphology: Needs, Progress, & Challenges	104
3.2. A Flexible Approach for Intravascular Image Characterization: Fundamental Methodological Framework	112
3.3. Characterizing Tissue in Grayscale IVUS	119
3.4. Characterizing Tissue in Intravascular OCT	138
3.5. Segmenting Indwelling Devices in Intravascular OCT.....	144
3.6. Variations on the Framework	154
3.7. Importance of Domain Enrichment.....	167
3.8. Importance of Training Data.....	173
3.9. Implications & Future Directions.....	183
3.10. Conclusion.....	188

Computational Modeling & Simulation of Patient-Specific Vessels	189
4.1. Diseased Vessel Modeling: Needs, Progress, & Challenges.....	190
4.2. Constructing Detailed Image-Based Volumetric Models	199
4.3. Determining Patient-Specific Constitutive Parameters.....	206
4.4. Implications & Future Directions.....	212
4.5. Conclusion.....	215
Bridging the Chasm between Virtual & Physical Medicine with Synthetic Imaging	217
5.1. Synthetic Medical Imaging: Needs, Progress, & Challenges.....	218
5.2. Synthetic Image Generation.....	225
5.3. Applications in Cardiology	254
5.4. Implications & Future Directions.....	265
5.5. Conclusion.....	267
Conclusion: Biomechanical Assessment & Modeling of Diseased Coronary Arteries	269
6.1. Summary of Work & Contributions	270
6.2. A Vision and Roadmap for the Future of Cardiovascular Medicine	273
6.3. In Closing	278
References	279

Chapter 1

Introduction: Biomechanical Modeling & Assessment of Diseased Coronary Arteries

“The heart of living things is the foundation of their life, the sovereign of everything within them, the sun of their microcosm, that upon which all growth depends, from which all power proceeds.”¹

– William Harvey (1628)

Cardiovascular disease (CVD) is a monumental burden to society, accounting for over 30% of deaths globally [1]. In particular, atherosclerosis and coronary artery disease (CAD), characterized by plaque buildup in and hardening of artery walls [2], presents a significant and ubiquitous concern. Although tremendous progress has been made in combating and diminishing mortality of the disease [2], many questions regarding the mechanical progression of disease and patient response to intervention remain unanswered; mathematical and computational modeling are key to answering these biomechanical questions [3]. This thesis therefore melds novel quantitative analysis tools and models representing diseased coronary arteries to assess and characterize the state of atherosclerotic vessels, contributing insight at the intersection of pathological anatomy, imaging, and biomechanical structural behavior to enhance the diagnosis, prognosis, and treatment of CAD. In this chapter, major themes of this work, and scientific fundamentals of the fields it advances, are introduced, providing an overview of the disease, vascular imaging, and arterial modeling.

AIArtificial Intelligence
CADCoronary Artery Disease
CHDCoronary Heart Disease
CTComputed Tomography
CVDCardiovascular Disease
CFDComputational Fluid Dynamics
ESSEndothelial Shear Stress
FEAFinite Element Analysis
FSIFluid-Structure/Solid Interaction
IVUS ...Intravascular Ultrasound
MLMachine Learning
OFDI ..Optical Frequency Domain Imaging
OCTOptical Coherence Tomography
PCSPeak Cap Stress
PSSPlaque Structural Stress
RFRadiofrequency
TCFA ..Thin-Cap Fibroatheroma
WSSWall Shear Stress
3DThree-Dimensional

¹ “*Cor animalium, fundamentum est vitæ, princeps omnium, Microcosmi Sol, à quo omnis vegetatio dependet, vigor omnis & robur emanat.*”

1.1. Disease of the Cardiovascular System

The cardiovascular, or circulatory, system is the intricate and exquisitely complex network by which the body nourishes the immense mass of cells which constitute the living body [4]. The constituent vessels serve as a transportation system, delivering oxygen and nutrients while removing carbon dioxide and other waste products, as well as immune cells and antibodies, signaling factors, and other essential circulating molecules. Far from a passive system of pipes, the circulatory system is remarkably dynamic, responding to mechanical and chemical stimuli in order to actively regulate blood pressure, shunt blood to meet physical demands, maintain body temperature, and preserve core functions as needed, and even modulate vessel wall permeability to control transmural exchange [5], [6]. Circulating blood through this system—from prenatal development to death without respite—is the heart [4]. Like the vessels through which it drives blood, the heart is no monotonous pump, and in consort with the vessels, the heart controls the rate and volume of blood flow in response to the demands of the body by actively varying the strength and frequency of its beating [5], [6]. The dynamism, range, and responsiveness of this system, central to the sustenance of human life, is truly profound.

1.1.1. Impact & Pathology of Heart Disease

Given the complexity and importance of the system, it is no surprise that CVD, disease of the cardiovascular system, is both prevalent and dangerous. Indeed, CVD is the leading cause of death and responsible for 50% of all noncommunicable disease deaths—over 30% of all mortality worldwide [1]. Of those deaths, at least 57% can be conservatively and directly attributed to atherosclerosis in the form of coronary heart disease (CHD; 42.6%) or ischemic stroke (14.8%). In more concrete terms, CHD causes around 1 in every 7 deaths in the United States—an American has a heart attack (myocardial infarction) every 40 seconds [1]. Beyond this startling statistic, prevalence of subclinical atherosclerosis belies the true extent and impact of this pervasive menace [1], which is further compounded by the toll in diminished quality of life, financial burden, and emotional strain.

Atherosclerosis², named for the characteristic pathological presentation of arterial hardening, is a widespread chronic disease which can suddenly escalate in a devastating acute

² The terms “atherosclerosis” and “arteriosclerosis,” of similar etymology, are sometimes used interchangeably in the literature. However, the latter is formally sclerosis or scarring of the arteries *without* accompanying obvious lipid

health event. The development of atherosclerosis systemically affects arteries, the blood vessels which carry oxygenated blood away from the heart. These vessels, like nearly every tubular structure in the body, are comprised of three layers. The innermost layer, the (tunica) intima, is comprised of endothelial cells which lines the blood-filled lumen; it acts as an active barrier and serves important signaling functions. The (tunica) media, separated from the intima by the membranous internal elastic lamina, is comprised of smooth muscle and extracellular matrix, components which allow it to control vasomotion (and thereby regulate blood pressure and flow) and store potential energy during oscillatory periods of high pressure. Finally, the (tunica) adventitia or externa, separated from the media by the membranous external elastic lamina, supports the vessel with nerves, lymphatics, and blood vessels [5]–[7]. While the entire vessel structure can be modified and disrupted in atherosclerosis, the disease is considered “a disease of the intima” [8]. It is in this layer that endothelial dysfunction leads down a cascade of fatty deposition, inflammation, and dysmorphic remodeling, driven in large part by mechanical forces [5], [7], [9]. Composition of the resulting diseased intima can include calcifications, fatty lipid deposits or necrotic cores, fibrotic tissue, or numerous variations thereof. These depositions of so-called “plaque” result in a highly variable and heterogeneous mechanical structure.

The consequences of plaque deposition and accumulation are variable, and evolve with the plaque morphology—composition, structure, and form—distribution, and ever-changing milieu—the microenvironment. While sub-clinical atherosclerosis is present to some extent in nearly every adult, the disease presents clinically when blood flow is inhibited and oxygen demand of downstream tissue exceeds the supply delivered by the blood, a state known as ischemia. This may be caused either by the protrusion of intima-bound plaque into the luminal space or rupture or erosion of the cap separating highly-thrombogenic plaque from the bloodstream, resulting in immediate clotting and luminal blockage [5]–[7]. Blockages starve distal (downstream) tissue of necessary oxygen; if flow of oxygen is not restored quickly, deprived tissue begins to undergo non-reversible injury and subsequently dies. When a blockage is located in the carotid artery or one of its distal branches, which supply blood to the brain, such a catastrophic event results in an ischemic stroke. When atherosclerosis is localized in the coronary artery, which supplies blood to

infiltration. This thesis primarily focuses on the former, of which intramural lipid accumulation is a hallmark, though is generally applicable in the latter pathology as well.

the heart itself, the disease is classified as coronary artery disease (CAD)³, and consequences of ischemia range from angina (ischemic chest pain) to deadly myocardial infarction (heart attack) [5], [6]. Atherosclerotic plaque can form elsewhere in the arterial vasculature (such as the renal, mesenteric, and peripheral arteries). These lesions can be devastating in their own right, though the consequences are typically less acute and severe as those caused by plaques in vessels supplying the brain and heart due to the indispensability and high oxygen requirements of their never-ceasing functions.

1.1.2. Diagnosis & Treatment of Atherosclerosis

The ischemic symptoms of atherosclerosis, and to some degree their cause, have been recognized for centuries. In the early 1500's, Leonardo da Vinci performed an autopsy and concluded that the cause of death was a "lack of sustenance...brought about by the continuous narrowing of the passage of the vessels...by thickening of the coats of these vessels" [10]. While da Vinci erred in some critical details—he, like his contemporaries, did not even grasp the circulatory nature of blood flow—his autopsy observations and hypotheses are forebears to modern understanding of disease that evolved over the following half-millennium [10]. By the late 1700's and early 1800's, medical experts were beginning to describe a dangerous disease that manifested as chest pain, particularly during physical exertion, which often concluded with the patient's death [11], [12]. However, no treatment was available; one doctor remarked at this time, "With respect to the treatment of this complaint, I have little or nothing to advance," recommending, quiet, warmth, alcohol consumption, and opium [12]. Even into the 1900's, while understanding was much improved, it was incomplete. Regarding "marked sclerosis of the arteries," one prominent doctor dismissively noted that "it is important not to over-estimate their seriousness. The outlook may depend much less upon the existence of these factors than on the sort of man in control of them," likening diseased arteries to a ship and the patient to a captain [13]. True diagnosis could be made only at postmortem autopsy, and prognosis was also unavailable; the same doctor noted that there was "no method by which we can reasonably tell whether a sufferer will be cut off the next minute or survive many years." While the importance

³ The terms "coronary artery disease" (CAD) and "coronary heart disease" (CHD) are often used interchangeably. Technically, however, CHD is the result of CAD, and occurs when disease of the blood vessels servicing the heart (coronary arteries) progresses to cause disease of the heart. No strict distinction is made between the two in this thesis.

of lifestyle changes were recognized, with doctors admonishing to “Go slowly” and “Eat less,” medical treatment was limited and clinicians sometimes resorted to palliative administration of chloroform [13]. The disease was still described at this time as “*meditatio mortis*,” translated as “contemplation of death” [13].

Comprehension and treatment of CAD has advanced rapidly in the age of modern medicine: facilitated by improving mechanistic understanding of the disease, overall mortality has dropped around five-fold in the past five decades, with in-hospital mortality declining even more so [1], [2]. Medical innovations propelling this trend, starting around the 1950’s, have included interventional devices (such as stents, scaffolds, and balloons), pharmaceutical agents and delivery methods, increasingly effective and decreasingly invasive (e.g. percutaneous) surgical procedures, and large coordinated risk-identification and education campaigns. Together, these advances have tapered the rising incidence and increased survivability of the disease.

1.2. Imaging of Diseased Vessels: A Pillar of Clinical Practice

Imaging has played an important direct and supporting role in lowering CVD mortality and growing knowledge of vascular physiology and pathology [2], [14]. While autopsy and histology studies helped to determine the cause of acute myocardial infarction, measure and quantify population distributions of normal and pathological vessel features, and classify plaque types, such approaches cannot, by their nature, be used to diagnose, assess, or guide treatment of current patients, or to monitor disease progression and response to treatment. For these purposes, *in vivo* imaging is required. *In vivo* imaging plays a role in all stages of disease diagnosis, prognosis, treatment, and monitoring [15], [16]. It is widely used to determine vessel geometry; such imaging can visualize macroscopic vessel and blood distribution patterns, identify and monitor stenoses (aberrant narrowings of blood vessels), measure vessel eccentricity, determine extent and direction of remodeling, and quantify plaque prevalence (i.e. plaque burden) [17]–[19]. Beyond geometry, plaque morphology can be assessed by determining plaque composition [18], [19] and biological features such as areas of inflammation can be identified and located [20]. The constellation of data and insights available from the imaging is beneficially used to identify plaques demonstrating high-risk features [14], [17], [21]–[24] and to guide, optimize, and assess outcomes of intervention (angioplasty and stent deployment), monitor patient reaction to, and recovery from, these interventions, and identify instances of subsequent device (e.g. stent or scaffold) failure [25]–[30].

An expanding arsenal of *in vivo* imaging modalities has been developed to service these clinical and research needs [31]. Noninvasive imaging techniques include multidetector computed tomography (CT), magnetic resonance imaging, nuclear imaging (including single-photon emission computed tomography and positron emission tomography), and contrast-enhanced ultrasound imaging. Invasive imaging techniques for arteries include angiography, intravascular ultrasound (IVUS), IVUS-radiofrequency (RF) analysis, palpography (intravascular elastography), optical coherence tomography (OCT), near-infrared spectroscopy, intravascular magnetic resonance imaging, angiography, and thermography [21]. Each imaging modality offers different benefits and strengths at different scales, and suffers from distinct shortcomings (Figure 1.1) [21], [26], [32]. For this reason, hybrid and multimodal technologies (combining modalities) are being explored to take advantage of distinct strengths and negate shortcomings [21], [26], [32]; such technologies, however, are not yet widely available. In fact, none of the non-invasive modalities are robustly validated and widely used, while among the invasive modalities only angiography, IVUS(-RF), and OCT are widely used in clinical routine; others are used only in limited and research contexts [21].

Intravascular OCT is a relatively new modality which offers exciting opportunities for clinical use and fresh challenges. Developed by researchers at MIT and Harvard Medical School [33]–[36], OCT is an invasive, catheter-based technique that utilizes near infrared light to construct cross-sectional slices of the artery and its internal (micro)structure [37]. The use of light instead of mechanical ultrasonic waves—used in the more established IVUS [19]—leads to remarkable resolution, particularly in frequency domain OCT/optical frequency domain imaging (OFDI), which is superior to the displaced time domain OCT [38]. Indeed, OCT resolution is typically 10–20 μm in the axial direction and 20–40 μm in the lateral direction [18], [39], [40], compared to 100–200 μm axial resolution and 200–300 μm lateral resolution in IVUS [39]. This difference in resolution is critical, as it uniquely allows *in vivo* measurement of particularly small features and identification of life-threatening thin-cap fibroatheromas (TCFAs), which are characteristically less than 65 μm in thickness [41], [42]. Furthermore, the light signal can better penetrate calcified tissue, which is opaque to IVUS' ultrasound signal and consequently obscures abluminal tissue (Figure 1.1C).




Figure 1.1. Various imaging modalities are associated with distinct benefits and limitations. The cross sectional profile of diseased coronary arteries can be visualized through (A) histology (von Kossa stain), (B) CT, (C) OCT/OFDI, and (D) IVUS. Each of the above illustrates a highly calcified lesion in a coronary artery at different scales and resolutions. While CT visualizes the cross-section of a large region, providing context and spatial information, and readily highlights calcified tissue, resolution is low and distinction of soft tissues is poor. OCT and IVUS provide highly detailed information with moderate tissue differentiation in smaller regions; the former has greater resolution but inferior penetration depth relative to the latter. Histology, with its arrays of specialized stains, provides the greatest resolution and tissue differentiation, but at sparse, discrete locations, and requires that the tissue be excised from the body—a critical restriction. (Histology image provided by CBSET.)

However, OCT has failed to overtake IVUS in global popularity due to several of its own shortcomings and challenges. Most importantly, the use of light instead of mechanical waves, while improving resolution, results in substantially lower penetration depth (0.1-2.5 mm in OCT, compared to 10 mm in IVUS [18], [39]). Though better able penetrate calcified tissue, light is

quickly attenuated by fatty tissue and lipid-rich lesions, obscuring abluminal tissue. The light signal also cannot penetrate the metal guidewire used to place the intravascular imaging catheters, resulting in a substantial reflection and shadow artifact (Figure 1.1B). Other artifacts sometimes also occur [43]. Furthermore, blood is a highly-attenuating medium for light, and so must be cleared from the field of view within the vessel. While older, slower forms of OCT (time domain) required proximal balloon occlusion, modern versions use non-occlusive flush of an optically transparent media (e.g. Lactated Ringer's or crystalloid/radiocontrast solution) alone; contrast may be administered through the guide catheter either by manual injection using a syringe or automatically with a power injector [18], [44]. While shown to be safe in typical quantities for those without contraindicative complications (e.g. renal insufficiency) [45], [46], this flush imposes additional challenges. Insufficient flush can result in obscured vessel wall tissue and artifacts, which significantly deteriorates image quality [18], [39]. Repeated acquisitions may be precluded due to the concomitant safety risks of additional media [47]. And flush administration can also impact the observed and measured geometry of the vessel by altering luminal pressure and inducing hypervolemic conditions [44], [48], [49].

Despite these shortcomings, OCT has been shown repeatedly to be not just reliable and safe for clinical use [45], [46], [50]–[54], but has demonstrated the capacity to offer insights and detailed information not offered by IVUS or other imaging modalities [50], [55]. OCT is regularly used in catheterization labs around the world to diagnose and prognosticate CAD [18], assess lesions and characterize plaque with excellent reproducibility [40], [56]–[58], guide coronary interventions such as stent deployment [45], [59], [60], and assess stents after deployment [37], [56], [61]. Ongoing work continues to demonstrate and improve the clinical utility of the modality through validation studies [62], [63], updated signal analysis and processing [31], [64]–[66], hardware modifications, and introduction of exogenous agents and fluorescence [66]. For a more extensive expert review of the intravascular applications of OCT, with comparisons to IVUS, see Prati *et al.* [37], [39].

1.3. Arterial Modeling & Simulation: Promise & Challenges

While imaging has led to greater understanding of basic pathology development and improved treatment in the clinic, it has also empowered researchers with the ability to develop patient-specific computational models to glean further insights. One vision for a burgeoning

paradigm in computational cardiology describes the mutual supplementation of imaging, modeling, and informatics in advancing understanding of atherosclerosis, novel treatment, and decision support systems [67]; this trend has already commenced. By generating virtual representations of physical biological systems—informed by imaging and patient data and adhering to fundamental governing principles—the system can be non-invasively interrogated in complex, comprehensive, and systematic ways that would be dangerous, expensive, or impossible in a patient.

In pursuit of the three ultimate goals—fundamental disease insight, novel and improved treatment options, and individualized patient care guidance—*in silico* (i.e. computer-based) models of three primary classes have become popular [68]. The first category of such models addresses fluid dynamics, primarily using computational fluid dynamics (CFD) [69], [70]; by simulating blood flow through a model vessel, researchers can approximate spatiotemporal pressure, speed, exerted forces (i.e. shear stress, including wall or endothelial shear stress), and gradient patterns in the vessel. The second category considers solid mechanics, primarily through finite element analysis (FEA) [71], [72]. By simulating vessels with applied blood pressure, heart wall deformation, balloon inflation (with or without stent deployment), or other loading conditions, researchers can approximate spatiotemporal stress distribution (circumferential, von Mises, principal, or other measure) throughout the vessel wall and device, if applicable. Finally, the third category addresses mass transfer—generally that of drug from a drug-eluting stent or balloon through a vessel wall or infiltration of molecules (such as lipoproteins) and/or cells (such as macrophages) from the blood stream into the vessel wall. Active and passive diffusion models are used to numerically predict the temporal and spatial distribution of the substance of interest. Advanced models intersect multiple categories: fluid-structure (or solid) interaction (FSI) models numerically compute the interplay between fluid dynamics and solid mechanics, allowing blood to deform vessel walls and walls to likewise exert pressure on the contained blood; multi-level (or multi-scale) modeling iteratively calculates mass transport, fluid dynamics, and/or solid mechanical state to numerically simulate the coupled system. Such multi-level approaches have utilized image, experimental, and blood exam data as inputs to reconstruct three-dimensional (3D) arteries, model blood flow and mass transport, simulate plaque growth and progression, and subsequently calculate clinical metrics (e.g. fractional flow reserve) and perform virtual treatment (e.g. stent implantation) [73].

Computational models arise from constitutive equations describing material and interface properties, initial and boundary conditions enforcing known or estimated states, driving forces applied to the system, and system geometry. By implementing the fundamental equations expressing conservation of mass and momentum within this context, a system of differential equations can be constructed and solved iteratively to calculate state parameters at various times and locations. For all of these model systems (Figure 1.2), geometries, constitution, constitutive models, initial and boundary conditions, and loading conditions can be, in whole or part, idealized, extracted from imaging data, determined experimentally (i.e. *ex vivo*), or derived by some combination of imaging or experimentation and simulation (e.g. inverse FEA). Various metrics may be calculated by these computational models. With CFD analyses, researchers can indirectly determine regions of disrupted flow, identify locations of elevated or depressed pressure, and calculate wall shear stress (WSS), which is referred to as endothelial shear stress (ESS) when occurring at the inner endothelial lining of a vessel. Associated metrics distilling spatial and temporal relationships are also commonly computed, including time-averaged WSS, area-weighted average WSS, axial WSS, WSS gradient, time-averaged WSS gradient, oscillatory shear index, modified oscillatory shear index, relative residence time, flow separation parameter, and other derived parameters. With FEA, researchers can locate points of elevated plaque structural stress (PSS) and calculate peak cap stress (PCS)—a value of interest because catastrophic mechanical failure of the cap results in immediate clotting and luminal blockage of the vessel. The metrics calculated by these models can provide insights in themselves, or may be coupled with clinical and/or animal model data to determine causal or mechanistic relationships.

Ongoing research in the arterial modeling field can be broadly divided into three partly-intersecting aims: improving modeling techniques, evaluating interventions, and seeking

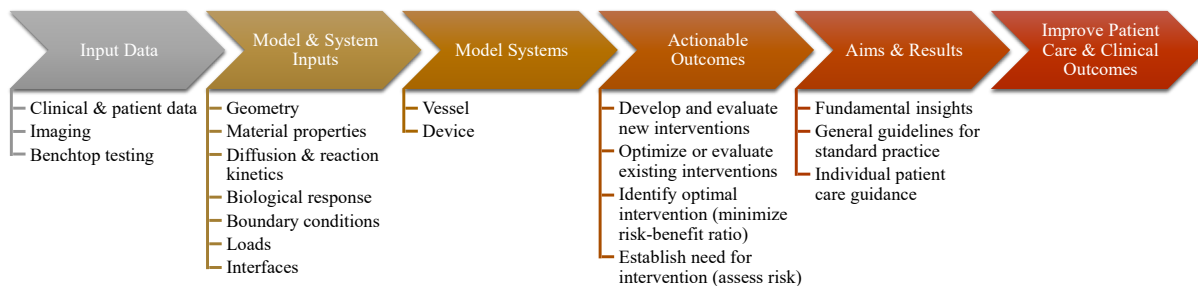


Figure 1.2. The validity of model outputs and the actionable outcomes derived therefrom ultimately relies on the validity of the input data. Input data is required to specify various model and system parameters, and can be derived from various sources.

(patho)physiological insights. Facilitated by improvements in computational capacity, modeling techniques have in some ways advanced substantially since some of the earliest, trailblazing work on vessel stress analysis [74], [75]. However, in many respects techniques are largely reminiscent of those from several decades ago, and much is still not understood regarding the selection of various parameters, inclusion of assorted features, and level of detail and resolution. Work is still ongoing to determine how to best model these complex systems, but findings continue to recommend implementation of ever-increasing complexity and specificity. Recent work has illustrated the importance of intima heterogeneity in accurately assessing stress (particularly PCS) [76], the high sensitivity of stent angioplasty simulation results to tissue constitutive, damage, and regrowth representation [77], [78], the need for multimodal deformations to accurately assess and predict stent performance [79], and the impact of anatomic and physiologic uncertainties on blood flow simulations [80]. Findings have also highlighted the need to incorporate wall compliance and displacement (through FSI, as opposed to CFD) to accurately compute WSS distribution and profile [81], the immense stress concentration factor that can be produced by micro-calcifications (and the dependence of calcification spacing, aspect ratio, and relative alignment) [82], and the importance of incorporating residual and initial stress/strain into vessel models for accurate stress calculation [72], [83]–[85], among other findings that seem to warrant models of escalating intricacy.

Meanwhile, more applied work has focused on intervention evaluation. Virtual execution of interventions *in silico* allows non-invasive assessment and/or prediction of both qualitative and quantitative procedure outcomes (such as lumen deformation and final shape of a stent following deployment and recoil, as well as stress and strain distribution in the vessel wall and stent or scaffold). Recent areas of investigation in this regard have included often-complex bifurcation stenting techniques [86], [87], the impact of plaque type, bifurcation angle, and procedure strategy on side branch compromise [88], the effect of predilation on stenting outcome (final lumen configuration and resulting tissue damage) [77], and approaches to optimize and select balloon and stent designs, materials, and deployment strategies [89]–[92]. Other work in this field seeks to project fatigue behavior and likely points of failure for stents and scaffolds, and determine the impact of calcifications on stent fatigue life [93].

Additional work has looked beyond devices and intervention to pursue pathophysiological insights. The most influential work in this field has been the use of CFD models to demonstrate

co-localization of low ESS and both plaque progression and post-implantation neointimal hyperplasia (often evaluated using OCT), suggesting new mechanistic pathways [94]–[97]. Less prevalent, but equally intriguing, is work dissecting the role of structural stress on plaque evolution; computational research has unveiled associations between elevated plaque structural stress (PSS) and increased atheroma volume and incidence of TCFA [98], as well as neointimal thickening after stenting [99]. More common work regarding structural stress doesn't assess plaque evolution, but rather seeks to identify areas susceptible to rupture or damage at a given time and the relative impact of various features (such as micro-calcifications, necrotic core, and overall plaque composition and geometry) on PCS to identify features of “vulnerable” plaques [100]–[102].

Idealized and image-based models of diseased coronary arteries are already used to better understand the pathological mechanisms of lesion progression and vulnerability [103]. There has long been interest in using such approaches to quantitatively segment patient populations, evaluate patient- and lesion-specific risk (or “vulnerability”) [104], and guide complex interventions [87]. As such, translation of patient-specific computational modeling, integrated with imaging and informatics, into the clinic has gained increasing attention and traction among researchers, clinicians, and regulators [67], [105]. Indeed, there have already been some limited successes in the field: a tool which uses computational fluid dynamics simulations—based on non-invasive imaging of coronary arteries—to calculate a critical hemodynamic measurement which drives clinical decision-making has been approved by the U.S. Food and Drug Administration [106]. However, fully realized concepts of “computational cardiology” remain encumbered by barriers to integration with physical medicine, stranding years of work in modeling and image analysis across a chasm from the clinical bedside. New approaches are needed to bridge the gap.

While all of the aforementioned models offer exciting prospects and opportunities to advance understanding and treatment of atherosclerosis, the vast majority of current models aren't reliable and lack the detailed input data necessary to robustly characterize and represent the systems and confidently address important questions. Despite the progress that has been made in refining and improving fidelity of models, uncertainty remains and the level of model complexity demanded to inspire confidence in results is ever-increasing and often prohibitive. Furthermore, since features such as micro-calcifications—which can induce stress concentration factors of at least 2 to 5 [82], [107]—are not visible in any routinely used clinical imaging modality, it is unclear

whether all relevant features needed to reliably estimate absolute (as opposed to relative) values of a given metric are actually represented in any generated model system. Implementing a model with sufficient but manageable complexity is indeed a challenge, in the determination of feature relevance, the attainment of input information needed to accurately represent these features, and the technical implementation of these features in a representative model.

Additional uncertainties arise from implementation of generic parameters in patient-specific models. While gross geometry—and to an extent constitution—can be replicated for individual patients, many parameters needed for modeling and simulation cannot be readily or practically determined *in vivo* on a per-patient basis. These parameters include boundary conditions (such as pressure and flow waveforms), initial and residual stress/strain, material mechanical properties, diffusion and reaction kinetics, and biological response profile. Many of these properties are determined *ex vivo* for a small number of samples and generalized to all patient models. However, elastic moduli and ultimate stresses of pseudo-homogeneous plaque components measured between—and even within—different studies have varied by multiple orders of magnitude; the various proposed constitutive models also predict drastically varying stress-strain behavior of the hyperelastic and highly anisotropic materials [101], [108], [109]. While some approaches to *in vivo* mechanical characterization have been proposed [110]–[112], none have been adopted for clinical use and a reliable method for identifying appropriate constitutive parameters remains an outstanding challenge. Coupled with a dearth of model validation, these challenges leave many *in silico* models unsubstantiated and do not allow reconciliation of their conclusions, diminishing prospects of delivering the oft-cited goal of providing real-time patient-specific modeling and clinical decision-making guidance [103], [113].

1.4. Artificial Intelligence in Cardiology

In addition to computational and image-based modeling, artificial intelligence (AI) is blossoming within the field of cardiology, buoyed by tentative optimism and high expectations among cardiologists and engineers alike [67], [114]–[120].

AI is the capacity of machines or computers to make decisions, perform higher-order analytical processes, or mimic other behavior typically attributed to natural human (or non-human animal) mental capacities. The concept is incredibly broad, and can range from a simple explicitly-encoded decision tree to a highly-complex self-trained network which can perform intricate tasks

like playing a game, driving a vehicle, or translating text into a different language. These latter examples, in which abilities are learned directly from a set of data without explicit programming—often because the tasks are too difficult for a programmer to provide a concrete, governing set of instructions or rules addressing all possible scenarios—fall within a subset of AI called machine learning (ML). ML methods often use a software structure modeled after the brain, called a neural network, to extract patterns from the data. Connections between sets of interconnected nodes, through which input data is transformed to an output value (or set of values), are optimized based on a set of training data to achieve the desired output results. When ML approaches implement this network structure, expanded to deeper levels, the process is called deep learning [121]. One specific type of deep learning neural network is effectively structured to handle data arrays, such as images and videos: convolutional neural networks. By leveraging prior structural knowledge of spatial locality and translation invariance—meaning that relevant information is typically located near each other and does not depend on its relative position in the array—the complexity of the network (and consequently the required volume of training data and the risk of over-fitting the model to the data) can be substantially reduced [121].

Though many fundamental concepts in AI are decades old, their potential has only begun to be realized. It was not until sufficient computational power recently became widely available to meet the needs of these demanding tools that the most powerful AI methods, including deep learning, became feasible. They are now widely applied to classification, prediction (or recommendation), and action selection (or planning) tasks in a wide array of fields and applications.

AI in healthcare, as in other fields, has experienced a resurgence in recent years. As for the wider field of medicine [122], the growing reliance on, and proliferation of, digitized data, supported by ever-increasing computational power, conjures an eagerness to harness and fully utilize this immense resource. Garnering particular interest are applications to collaboratively support clinicians in the amelioration of routine tasks [123]–[125] and processing of incomprehensibly vast troves of medical data [124]–[129], as well as translational research seeking fundamental insights from patterns in massive datasets [130]–[132]. Moving forward, AI to augment human intelligence and offer decision support offers perhaps the most challenging and beneficial future applications [133]–[135]. However, some have cautioned that the hype and tantalizing promise of AI has overshadowed typical scrutiny and vigilance of the field; ML-based

decision support systems are often not subject to sufficient critical assessment and may lead to unintended consequences through inappropriate use or application [136]–[138]. Furthermore, disconnects between AI specialists, clinicians, and medical infrastructure have impeded successful development and effective deployment of AI systems in clinical practice [139]–[141], and a lack of interpretability, explainability, and uncertainty-awareness among AI systems impedes establishment of trust necessary in such a high-stakes environment as healthcare [142], [143]. Still, however, with proper oversight and critical assessment, and collaboration between subject matter experts and stakeholders, such systems offer great promise in improving the delivery of healthcare, particularly in cardiovascular medicine.

Residing at the intersection of routine task automation, big data processing, and augmented decision support, deep learning for medical image analysis is a booming, promising specialty [144]–[147]. In cardiology, where imaging is a staple of clinical practice, opportunities to support clinicians in image interpretation are expected to have a particularly profound impact [148]–[152]. A niche corollary in medical image generation and synthesis has attracted some interest in the research community as well, but is yet to find a foothold in any clinical applications [153], [154]. However, recent exciting advances in ML addressing image processing and computer vision offer new possibilities to enable the continuum between physical and virtual medicine, the absence of which currently stifles translation of technical innovation into medical practice.

1.5. Thesis Contributions & Organization

Despite great strides in understanding and treating atherosclerosis, mortality and disease burden remains unacceptably high [1], [2]. While the diagnosis is fortunately no longer a sentence of *meditation mortis*, there remains limited actionable understanding or ability to differentiate those likely to experience an acute event, ineffectual treatment outcomes, or responsiveness to non-invasive interventions—a disappointing echo of the past century [13]. Tools of intravascular imaging, computational modeling, and AI promise to usher in a new era of computational cardiology [67] in which the field adopts more quantitative approaches leveraging the best available data to deliver patient-specific, personalized care. As noted, however, several challenges must first be overcome before this vision can be realized. The contribution of this thesis makes significant strides in overcoming several of these major obstacles. Using and advancing the aforementioned tools, the enclosed work develops new methods to determine arterial geometry,

extract vessel and plaque morphology and material properties, model and simulate patient-specific diseased vessels, and integrate this rich set of information into clinical practice and visual workflows (Figure 1.3), as described in the following chapters. The results are improved quantitative assessment and characterization of the state of atherosclerotic arteries to fill gaps in knowledge at the intersection of pathological anatomy, vascular imaging, and biomechanical structural behavior.




Figure 1.3. This thesis advances the confluence of imaging, modeling, and artificial intelligence. Processing intravascular images of diseased arteries provides information on lesion geometry and morphology, which can inform computational models and simulations. Models and simulations can conversely be supplemented with information extracted from imaging to elucidate lesion morphology and/or geometry. These intrinsically coupled systems can be integrated into clinical practice workflows through the generation of synthetic images. This thesis address and advances each of these synergistic elements.

Chapter 2 discusses a breakthrough method of extracting vessel geometry which overcomes current limitations of intravascular imaging. By employing *a priori* knowledge of arterial geometry and image acquisition procedure, and progressing from two-dimensional feature detection to 3D integration of adjacent and contextual information, existing barriers arising from limited penetration depth is largely surmounted to determine the geometry of diseased vessel wall geometry in high-resolution OCT [155]–[158]. Insights on the application and limitation of lumen-based image co-registration are also presented [159].

Chapter 3 describes the development and critical assessment of methods to determine and characterize vessel wall morphology. Aided by constraints imposed by vessel geometry and effective AI tools, plaque components of diseased vessels are labeled in OCT and IVUS images [160], [161]. The approach is extended as well to labeling intravascular devices (bioresorbable vascular scaffolds). The importance of enriching the method with domain knowledge and information on vessel geometry is quantitatively assessed, as is the impact of training data on overall performance [162].

Chapter 4 details the computational modeling and simulation of patient-specific vessels. Enabled by greater geometric and constitutive information extracted from clinically available intravascular imaging, improved patient-specific models of individual arteries are developed, offering the promise of more comprehensive assessment to guide clinical decision-making. The tradeoffs between two implemented modeling strategies are considered. An inverse approach utilizing the new models is also demonstrated and validated as a proof-of-concept for determining lesion-specific material properties.

Chapter 5 addresses a new way to bridge the chasm between virtual and physical medicine with synthetic imaging. Uniquely-applied methods of synthetic image generation are described and compared, and clinical applications of such images in cardiology are presented.

Chapter 6 concludes this thesis, discussing implications and future directions of the work to achieve a vision for a continuum of physical and virtual medicine enabling computational cardiology.

The work presented here has immediate impact through the insights revealed on geometric relationships, plaque visualization and appearance, impact domain knowledge enrichment, consequences of model structure, training, and evaluation metric decisions, and more. With further

refinement and validation, the achievements of this thesis will empower clinicians and researchers with enhanced information on diseased vessel constitution and biomechanical properties. With this information, which could previously only be derived through *in vitro* testing, coronary artery modeling and clinical assessment will allow clinicians to make more informed choices regarding their course of treatment and researchers to pursue new insights into the pathobiology of atherosclerosis and the impact of interventions.

Chapter 2

Extracting Vessel Geometry: Overcoming Current Limitations of Intravascular Imaging

The physical form of a blood vessel not only elucidates the progressive state of cardiovascular disease (CVD), but drives vessel functional capacity and, consequently, functional impact of the disease. Coronary artery disease (CAD) arises from insufficient delivery of blood to tissue of the heart due to atherosclerosis, which hardens and narrows the arteries [2]. The amount of plaque and extent of narrowing are, therefore, key metrics considered in the diagnosis, prognosis, and treatment of CAD—metrics inherent to the geometry of the vessel. Indeed, quantitative metrics of vessel geometry figure prominently in the expert consensus documents and clinical guidelines for intravascular imaging and coronary interventions [17]–[19], [25], [163]. However, such measures are typically made manually, presenting a host of challenges, and can be difficult in some modalities due to limited penetration depth. Such shortcomings severely restrict large-scale determination of vessel geometry, necessary not only for clinical assessment and clinical studies, but development of computational models of diseased arteries. In this chapter, a breakthrough method of extracting vessel geometry which overcomes current limitations of intravascular imaging [155]–[158] is presented, as well as a method and illuminating analysis of disparate geometric dataset alignment [159].

CAD	Coronary Artery Disease
CSA	Cross-Sectional Area
CT	Computed Tomography
CVD	Cardiovascular Disease
EEM ...	External Elastic Membrane
FPC	First Point of Correspondence
IEM	Internal Elastic Membrane
IVUS ...	Intravascular Ultrasound
OCT ...	Optical Coherence Tomography
SD	Standard Deviation
VH	Virtual Histology
2D	Two-Dimensional
3D	Three-Dimensional

2.1. Diseased Vessel Geometry: Needs, Progress, & Challenges

The geometry of a vessel can be distinctly altered by disease and profoundly impacts its function. It is therefore a critical consideration in clinical diagnosis, prognosis, and treatment-planning, as well as research. This importance has prompted efforts to offload the tedious work of annotation and measurement from human experts, but complications arising from the inherent limitations of imaging modalities have hindered progress in doing so.

2.1.1. The Need for Accurate Geometry

Long before vessels could be observed *in vivo*, the geometry of an artery has been known to vary with disease. Indeed, even before the circulatory system was known to be that—circulatory—Leonardo da Vinci made the astute pathological deduction on the basis of autopsy observations that certain CVD “is brought about by the continuous narrowing of the passage of...vessels by thickening of the coats of these vessels” [10]. (Here, “coats” refers to the layers of the vessel wall, which themselves have been imaginatively named “*tunicae*” due to their rough semblance to the tunic garment.) More recent detailed autopsy and histology studies have quantified the range of normal and diseased vessel geometry [8], [164]–[167]. Findings have shown that the plaque focally develops in the intima layer, driving most geometric changes, but also that the media layer (comprised of smooth muscle cells and connective tissue) is markedly thinner in diseased segments of atherosclerotic vessels than in healthy ones [164]. However, only measurement of the intima has been widely adopted as a strong signal of disease burden; standard quantitative metrics of the media alone are not noted in the standards for either [17]–[19], [25], [163].

To characterize a patient’s vessel and atherosclerotic lesion, if present, the dimensions of the vessel wall and plaque must be determined. With the advent of *in vivo* imaging techniques, including optical coherence tomography (OCT) and intravascular ultrasound (IVUS), recognition of the importance of vessel geometry led to several quantitative metrics being defined in the expert consensus documents and clinical guidelines for intravascular imaging and coronary interventions; these include, among others, (cross-sectional) area, diameter, and eccentricity of the lumen, plaque, or indwelling stent (when present), and additional metrics derived therefrom [18], [19]. Measurements are typically made in relation to the lumen, internal elastic membrane (IEM; equivalently the internal elastic lamina), and external elastic membrane (EEM; equivalently the

external elastic lamina). The IEM and EEM are fenestrated laminae separating the media layer from the intima and adventitia layers, respectively. Though too thin to be directly observed in currently-available clinical imaging modalities, these physical features are considered for practical purposes to reside at the visible border between layers. Of note, the media itself is differentiable from the adjacent layers in OCT but not IVUS—in which only the trailing edge of the media (EEM; adjacent to the adventitial layer) is visible [18], [19]. However, due to the limited penetration depth of OCT, the intima-media and media-adventitia borders (IEM and EEM, respectively) can be difficult to discern in diseased vessels, though studies have shown that experts can still derive accurate measurements of vessel area in arteries with highly-attenuating lipid-rich plaque [62], especially when the images are attenuation compensated and contrast enhanced [65]. Still, clinical measurements can be time-consuming and suffer from inter- and intra-observer variability, and are typically made at distributed, discrete points, limiting the attainable insights of holistic consideration. Furthermore, physicians without sufficient training, experience, or expertise are often unable to identify the complete outer border of vessels in OCT images, inhibiting important clinical measurement and limiting application in treatment decisions.

Aside from clinical practice, scientists and engineers seeking to computationally model arteries, particularly through structural mechanics or fluid dynamics analyses, are required to extract not just singular metrics of vessel geometry, but digitally recapitulate full two-dimensional (2D) or three-dimensional (3D) geometry as a cornerstone of the model. Furthermore, such geometric boundaries are needed to delineate the region of interest in images for any subsequent tissue characterization processes. Delineating vessel borders and characterizing healthy and diseased (plaque) wall constituents allows for the construction of models of patient-specific vessel geometry for research (and future clinical applications), as highlighted in Chapter 1. However, previous coronary artery anatomic reconstructions have been subject to limitations in application and/or fidelity arising from methodological approaches. Many modeling approaches have utilized either idealized geometries or relatively low-resolution image sources, such as intravascular ultrasound (IVUS), angiography, magnetic resonance imaging, or computed tomography (CT), resulting in necessarily simplified models. Others have relied on histology of excised tissue, which precludes predictive studies, serial monitoring, or application to current patients. Modeling approaches to date utilizing high-resolution OCT have been limited by assumptions or simplifications made during the reconstruction of the arterial wall necessitated by the absence of

geometric information; some use only the inner (lumen) border obtained from OCT (e.g. for computational fluid dynamics simulation without fluid-structure coupling) [70], [168]–[172], others have extrapolated radially from the inner (lumen) border by a fixed distance to approximate the outer (IEM or EEM) border [173], [174], and some have lofted circumferential cross-sections of indirectly-calculated diameters, defined only in plaque-free regions, to create external arterial wall surfaces [172].

The ability to reliably identify the full IEM and/or EEM in OCT images of diseased coronary arteries would provide great benefit by allowing high-resolution and holistic investigation of plaque and wall distribution and composition, along with macroscopic, continuous plaque burden measurements. Because automatic lumen segmentation and plaque characterization methods have been previously developed and validated for this modality [175], the major unresolved challenge is automatically identifying the outer border the vessel wall, which is often obscured by highly-attenuating plaque constituents which limit the penetration depth of OCT. Unfortunately, it is often difficult or impossible to identify the full IEM in a given frame. To generate the best possible estimate of the IEM within a frame, it is therefore necessary to use information from adjacent frames and *a priori* knowledge of vessel shape—strategies regularly employed by expert clinicians.

2.1.2. Progress in Automatic Geometry Extraction

In response to the clinical need for fast and consistent geometric quantification of vessels, several automated methods have been developed. In addition to progress in imaging methods and technology, methods for post-processing of IVUS and OCT images have brought about tools and capabilities that provide valuable insight to researchers and clinicians [175], [176]; algorithms now exist which can identify many key vascular structural and morphologic features with increasing accuracy, allowing for their measurement and quantification. With sufficient data, advanced algorithms can accurately detect lumen contours and identify stent struts [175], [177]–[179], and even rudimentary image processing algorithms can provide rough approximations of these features. The outer border can be readily segmented automatically in IVUS by various approaches [175], [180]. However, automatic segmentation of the outer border in OCT has been more elusive and challenging.

Automated methods for identifying the outer wall of a vessel in OCT are sparse, despite their potential value. Some approaches have sought to circumvent—and indeed leverage—the penetration depth limitations of OCT by using plaque free wall angle (i.e. visible arc) as a proxy for plaque burden, which is typically derived from the area delineated by the inner and outer borders of the vessel wall [181]. Others have directly delineated the outer wall, but only in non- or minimally-diseased vessels in which the entire outer border was clearly visible within the penetration depth-mediated field of view [182], [183]. Zahnd *et al.* were the first to extract multiple surfaces simultaneously in intravascular OCT, segmenting intima, media, and adventitia layers through a front propagation scheme [182]. However, the approach was only applied successfully in healthy regions of the vessel in which the layers were clearly visible; these regions neither present the most challenging features characteristic of intravascular OCT, nor do they represent segments of clinical relevance. A later graph-based approach introduced by Chen *et al.* was used to segment the lumen, intima, and media in cardiac allografts [183]. While applied to an interesting and clinically-relevant application of OCT image segmentation, the method analyzed vessel segments which represent relatively ideal conditions with clear visibility, and is unlikely to translate to more challenging cases such as advanced atherosclerosis. Curiously, these approaches were strikingly similar to OCT segmentation methods developed nearly a decade earlier for an entirely different application.

Unexpectedly, success in extracting geometry from OCT can be found in a different medical field: ophthalmology. Layer segmentation in OCT is not itself novel, though doing so in intravascular OCT is still a fledgling endeavor. When the revolutionary modality was first introduced, OCT was demonstrated in both coronary arteries and the peripapillary area of the retina [33]. Retinal OCT, which is non-invasive, has been widely adopted in ophthalmology; in this use case, penetration depth is not a significant issue, as the retina is generally less than 0.23 mm thick in its thickest region [184]. As such, layer segmentation has matured significantly in retinal OCT, and multiple methods and approaches have been developed to segment retinal layers in OCT images. Grayscale edge detection (with substantial pre-processing including specialized smoothing and de-noising filters) [185]–[187], active contours (energy-minimizing splines) [188], multi-step region growing [189], layered graph minimum-cost closed set determination (including feasibility constraints and regional information) [190], and coarse grained diffusion mapping using texture analysis to identify regional boundaries [191] have all been employed successfully to

segment layers in retinal OCT. Kafieh *et al.* reviewed and classified a number of other previous approaches that achieved varying success [191]. As engineers and scientists tackle the challenge of identifying and segmenting layers in intravascular OCT, it is prudent to learn from the preceding work in retinal OCT layer segmentation. Penetration depth aside, many complications are shared by both applications. Both retinal and intravascular OCT experience speckle noise, discontinuities in layer boundaries, intensity inhomogeneity, low image contrast, poor layer-to-layer image gradients, and feature shape irregularities. Unfortunately, these issues are more extreme in intravascular OCT. For example, layer boundary discontinuities in retinal OCT are caused by small intraretinal blood vessels, whereas large discontinuities in intravascular OCT are caused by shadows cast widely by a metal catheter guidewire, obscuring up to around 45° (one eighth) of the vessel wall. Nevertheless, challenge similarities suggest there may be much to be gleaned from retinal OCT layer segmentation approaches.

Despite all of the work to date, more progress must be made to enhance the information gleaned from intravascular images and to make this information more accessible to scientists and clinicians. Critically, current methods can only extract spatially superficial information from OCT images, providing incomplete information necessary for clinical assessment and biomechanical study.

2.1.3. Prior Interpolation and Surface Fitting Methods & Spring Model Precursors

Surface fitting and interpolation are common techniques for filling gaps in 3D information—approximating data in 3D when information is only available in certain areas of the solution space. It ultimately served a critical function in my solution to extract geometric information on diseased vessel walls from OCT images. Unfortunately, while many variations of interpolation and surface fitting exist, most are limited in flexibility, intuitiveness, ability to incorporate secondary information, and computational efficiency. Though existing approaches proved insufficient, a brief introduction to common approaches to interpolation and surface fitting is informative in understanding the shortcomings of prior approaches and the tremendous benefits of the method I developed for this purpose.

Interpolation methods are widely applied to construct a continuous function from irregularly spaced data, including in medical imaging [192]. In these schemes, functions are constructed that intersect known data points or positions, and utilize various methods to

approximate the values or positions lying between them [192], [193]. While some interpolation processes can theoretically be global, considering all known data points, nearly all are limited to local application to an arbitrary number of nearby points either by definition of the approach or computational limitations in considering large data sets. Furthermore, some popular approaches, including triangulation based methods, radial basis functions, and natural neighbor interpolation methods, require preprocessing steps, thereby further increasing computational complexity. Others, such as inverse distance weighted methods, artificially enforce radial symmetry around points, introduce artifacts, overweight clusters of points, or are highly sensitive to user-defined weight functions [192]. Fundamentally, smoothing is also limited in interpolation methods, as functions are constrained to pass through the available data points, which may themselves be non-smooth and could be unreliable.

Surface fitting overcomes some issues associated with interpolation techniques, but introduces alternative challenges [192]–[194]. Scattered data fitting techniques seek to find a reasonably approximating, ostensibly smooth surface that minimizes some error, such as cumulative difference between known data point values and surface values, to represent the underlying data trend. Such approximation is generally acknowledged to be preferable to interpolation when the available data may include noise, error, or inaccuracies, or are not reliable, as is the case with the application at hand. These techniques may be either global—such as polynomial least squares, discrete least-squares fitting by splines, discrete L_1 and L_∞ approximation, spline smoothing or smoothing splines, and continuous least squares—or local—such as patch methods and direct local methods [193]. One particularly popular local fitting technique is robust locally weighted regression and smoothing, in which polynomials are fit to local sets of data through weighted least squares such that the weighting provides robustness against deviant points [195]. The technique, initially introduced for application in 2D scatterplots, provides four mechanisms for customization: polynomial order, weighting function, smoothing parameter, and number of iterations [195]. However, the values and functions that can be selected are still highly constrained, and offer only modest flexibility [195]. While local techniques such as these require less computational capacity and provide other benefits, most do not produce globally smooth, or necessarily continuous, surfaces. Furthermore, some fitting techniques require regularly spaced data, may require specification of derivative data, or impose other limiting restrictions [193]. The shapes and contours that can be achieved with existing surface fitting

techniques are also fundamentally constrained—most surfaces must conform, locally or globally, to a polynomial function. This presents a vital limitation, particularly in applications to natural and complex systems which cannot be represented with sufficient fidelity by such simplified expressions.

Compounding the above issues is the restriction that neither standard interpolation nor surface fitting techniques can incorporate additional inputs from the rich information available in imaging data. Resulting approximations are consequently heavily biased when the input data points are not representative of, or equally-distributed from, the underlying data. In the present case of the vessel wall, we expect these techniques to underestimate unknown values of wall position because missing data points (non-visible wall segments) are inherently more likely to have values (radial depths) at the extremes of, or exceeding, those of the available data point values (visible wall segment depths).

The alternative approach developed in this thesis (Section 2.2.3) leverages key benefits of simple mechanical systems. For example, the simplicity and elegance of Hookean linear-elastic spring models lend themselves to unparalleled efficiency because they allow straightforward expression in matrix operations and the associated benefits in solving systems of equations. For this reason, deformable mass-spring tissue models have been implemented instead of finite element tissue models to realistically simulate interactive surgeries in real time [196], and a spring surface mesh has been implemented to simulate dynamic organ behavior corresponding to surgical operations with precision and speed [197]. While not applied to image processing, these applications exemplify the efficiency of Hookean linear-elastic springs models where large systems of equations must be constructed and solved to calculate three-dimensional response to external loads. Nevertheless, application of mechanical models to the field of image processing has been very limited. Bayesian statistics and statistical mechanics have been applied to probabilistic image restoration techniques and extended to additional image processing tasks, such as edge detection, image segmentation, motion detection, and image compression [198]. However, more dynamic systems like the one implemented in this work have not been explored and have no identified precedent, and are unique in their ability to extract geometric surfaces from incomplete spatial data.

2.1.4. Progress in Comparing Geometries between Image Datasets

Geometric information of interest is often distributed between multiple datasets. These datasets may be comprised of acquisitions from several complementary modalities, thereby providing a more complete description of vascular geometry to guide clinical decision-making [199]–[201], may be spatially sequential and partly overlapping acquisitions capturing a particularly long lesion (beyond the typical range captured by a single acquisition), or may be acquired at different time points to capture the longitudinal progression of disease or response to treatment. To enable the combination or comparison of these datasets, they must be properly aligned to match those segments which correspond to the same underlying physical structure(s). Alignment is typically achieved qualitatively through manual co-registration or matching using acquisition-specific identified fiduciary markers—clear and unique features or landmarks visible in both datasets, such as side-branches [18], [19], [202]. However, this process requires much of the time and expertise requirements, and is subject to much of the inter- and intra-observer variabilities, that plague manual geometry segmentation itself.

Several frameworks have been developed to address the challenges of dataset co-registration and matching, with ancillary tools for clinical assessment focused on algorithmic development for intravascular imaging. For circumstances in which multiple image acquisitions of the same vessel region are captured (by some combination of OCT, IVUS, and angiography), algorithms have been developed which are able to assist with co-registration and spatial alignment of these images, a laborious and time consuming task once reserved for trained experts.

Angiography is frequently aligned (or “fused”) with intravascular imaging data from OCT or IVUS, a capability included in user-friendly integrated systems [203]. However, the alignment process is only semi-automated, requiring a human operator to reconstruct vessel centerlines from the angiography and register the datasets by indicating a baseline position along the centerline corresponding to a selected axial position in OCT or IVUS (such that the remaining dataset is automatically aligned once the expert designates a single point of correspondence along the vessel axis) [204], [205].

Several approaches have also been proposed for the co-registration and spatial alignment of multiple IVUS acquisitions. Some methods have focused efforts on circumferentially aligning virtual histology (VH) IVUS, relying on manual co-registration in the longitudinal direction but

subsequently completing the alignment automatically based upon VH-IVUS-defined parameters [206]. Others have taken a step further to automatically align IVUS sequences longitudinally as well. Using automatically obtained features such as side-branch location, vessel area, lumen area, and areas of calcified and fibro-lipid plaques, non-rigid alignment of acquisitions through dynamic time warping was achieved and validated [207]. The most advanced method to date for aligning IVUS acquisitions achieves simultaneous registration of location and orientation using both feature-based and direct approaches [208]. Using perivascular tissue, plaque appearance, and plaque thickness, including direct comparison of image pixel intensity, pullback pairs are registered via 3D graph-based optimization [208]. While alignment of two IVUS acquisitions has been achieved, even the most successful and capable approaches are limited in their scope. In particular, grayscale intensity can't be readily compared directly between modalities to align multimodal image sets, and features relied upon to guide registration are not readily visible or extracted automatically in other modalities (e.g. intramural features of a diseased vessel in OCT). Therefore, inter-modality applications of these methods presents a serious challenge, despite this capability's desirability.

To address the shortcomings of single-modality alignment approaches, methods have been developed to assist with the co-registration of multimodal image datasets of OCT and IVUS. Building upon the IVUS-IVUS alignment method of Alberti *et al.* [207], Molony *et al.* similarly apply dynamic time warping for longitudinal co-registration of IVUS and OCT, and supplement the algorithm with dynamic programming for circumferential co-registration [209]. Features used in the alignment include lumen area, lumen eccentricity, calcification arc angle, and catheter angle—a modified set of features typically identifiable in both modalities (cf. [207]). However, the framework requires prior segmentation of the images and identification of calcifications in both image sets. While some automated approaches exist for such segmentation and identification, the authors note that these methods are imperfect, and that their methods may be sensitive to resulting segmentation errors [209]. Another two-phase, coarse-to-fine registration process uses features such as lateral branch points and stent points, contour plots (vessel shape), and grayscale distribution [210]. However, while the approach obtains feature points through automatic detection, it requires manual pair matching of the feature points between IVUS and OCT image sets to complete the process [210]. Less comprehensive approaches have also been developed. A semi-automated method for matching OCT and IVUS images identifies frame-to-frame (2D)

matches on the basis of lumen shape characteristics, considering longitudinal position/order but not aligning the entire acquisition [211]. While offering some utility, each of these methods either offers incomplete capabilities or is not readily automated, and is typically very computationally demanding.

In addressing multimodal image datasets, additional steps are sometimes needed to fully align geometries in co-registered image frames and achieve transverse plane registration. Because images are acquired both at different times and through different mechanisms, visualizations of the same underlying physical segment sometimes do not appear in the same configuration or proportions. To address this issue, non-rigid image registration techniques [212] have been applied. A computerized methodology has been developed that “inflates” micro-CT and histology data to align in-plane with IVUS images; distance map-based spatial transformation through 2D interpolation was applied to manually paired frames [213]. Another method focuses on semi-automatic transverse plane registration of 2D OCT and IVUS image frames [214]. The two phase approach requires user selection of key image frames to bypass longitudinal registration in 3D, though subsequent 2D registration is automated. Features, including inner arterial wall, side branches, and calcification regions, are detected when present; such specific landmarks are required to perform the registration, but are not present in each frame. As such, users must also select key image frames with discernable features to which the algorithm should be applied [214].

In concert with automated tools to identify key features, the aforementioned ancillary tools allow for quantification of geometric changes, such as remodeling, over time or in response to treatment. However, with frequent inconsistency in imaging modality through longitudinal patient care, current methods leave much to be desired. In particular, there remains an outstanding need for a simple method that utilizes universally identifiable features which can be automatically extracted from the image datasets and robustly applied in the task of co-registration. This, of course, relies on the ability to accurately extract complete geometry to be registered and compared.

2.1.5. The Endeavor to Extract Vascular Geometry from OCT

Guided by the clinical and research needs, and drawing technical inspiration from previous segmentation approaches in retinal OCT and mechanical systems, I embarked on work to automatically detect and delineate both the inner and outer border of diseased coronary arteries in intravascular OCT images. Physiological understanding and constraints were brought to bear on

the problem. Plaque burden follows a continuity inside the vessel and does not change dramatically from frame to frame. Therefore, using *a priori* knowledge of the structural and geometric constraints of vessel walls and working in 3D space, rather than on a frame-by-frame basis, the limited visibility of the outer vessel wall were largely overcome. Indeed, the non-visible parts of the outer vessel border could be approximated in the 3D space to estimate the full arterial wall position.

Initial approaches leveraged basic filtering and edge detection operators to identify visible segments of the outer border and fit an ellipse in each frame to the segments found in locally adjacent frames [155], [156]. While increasingly sophisticated and 3D routines identified larger portions of the outer border, it became apparent that gains in this regard would be limited, and that improved methods of filling the space between the detected segments would be required. Existing interpolation and surface fitting techniques proved insufficient; they tend to underestimate the radial depth of “missing” segments (as the visible segments are overwhelmingly located proximal to non-visible segments), and provided little control over the level of smoothing (resulting in loss of features due to over-smoothing or sharp, non-physiological features due to errors in the visible segment detection and/or image artifacts). A surface fitting approach inspired by Newtonian physics was highly effective in fitting variably-sparse and questionably-reliable data [157], [158]. Intuitively applying an anisotropic linear elastic mesh with a force function allowed me to autonomously delineate outer borders of diseased vessels in OCT images. A relatively simple, lumen-based approach to automatic co-registration of vascular imaging data was developed and its applicability assessed [159], allowing for tracking changes in, or merging, such geometric information across disparate imaging datasets.

2.2. Detecting Visible Segments of the Outer Border

The first step in determining vessel geometry is identifying the visible features of the inner and outer border. If there is minimal plaque burden or other impediment obscuring the view of the outer wall or necessitating smoothing, this is, in fact, the only step needed. However, if segments of the vessel wall are obscured, these visible segments provide important context in determining the location of the non-visible segments.

Visible segments of the outer border were identified through image processing of the OCT images. These methods were implemented entirely in MATLAB (MathWorks, Natick, MA).

Approaches to detecting visible segments of the outer border began with basic 2D edge detection utilizing standard functions and modules (Section 2.2.1) [155], [156], and advanced in complexity to 3D edge detection employing specialized approaches tailored to the task of identifying the outer border of vessels in OCT images (Section 2.2.2) [157], [158].

2.2.1. Basic Two-Dimensional Edge Detection

Processing of the OCT images (Figure 2.1) began with artifact detection and image resizing. Each image was first converted to a binary format with a global threshold determined by Otsu's method [215]. The bright regions segmented by this method roughly correspond to proximal regions of the vessel wall, as well as image artifacts, while eliminated dark regions correspond to the lumen and abluminal regions in which the OCT signal has been attenuated. Properties of each connected and continuous region were measured to differentiate segments corresponding to the vessel wall from others (i.e. noise and artifacts). Because OCT artifacts universally appear in polar coordinates as bands at the top of the frame extending the width of the image (Figure 2.1A), thresholded regions that were highly oblong in shape were removed and designated as likely artifacts. The average distance from the edge of the frame to the designated artifacts was measured; the region within this distance of the edge was eliminated from further analysis. This distance was also used for scale determination, as artifacts reliably manifest at a known distance from the edge of the image (based on imaging catheter size). Regions with an area smaller than a set threshold were also eliminated as noise. An approximation of the inner (lumen) border was then made by tracing the innermost non-zero point across the width of the binary image (Figure 2.1B). Projecting the lumen 2 mm distally provided the extent of the analyzed region, since regions beyond this distance are generally deemed beyond the penetration depth of OCT, and therefore likely to contain no information of value. The image was then reduced in size to include only the analyzed region of interest between the artifact and the extent of the penetration depth (Figure 2.1B).

Processing continued with filtering and edge detection (Figure 2.1C). Bilateral filtering, which preserves edges while smoothing an image, was applied in 2D following the procedure of Tomasi and Manduchi [216]. Subsequent edge detection was performed using the Canny method, which is designed to be robust in identifying weak edges amongst noise [217]. Parameters for the filter and edge detection method were adjusted incrementally in an informal grid search to achieve optimal identification of the target features, determined qualitatively.

Because several edges in each image were identified, and not just those of the outer border, edge selection was next performed. Edge detection inevitably resulted in the identification of extraneous features beyond the IEM (Figure 2.1C). To eliminate those undesired edges while preserving the appropriate features, numerous physiologically-informed characteristics of each edge were assessed. The IEM corresponds to the location between the intima and media, layers




Figure 2.1. Sample OCT image of a minimally-disease vessel cross-section illustrating the progression of the basic 2D border detection algorithm. (A) The original unprocessed image, displayed—as acquired—in polar coordinates, showing the key physiological structures and image artifacts. (B) The image following inner (lumen) border detection (overlaid in green), artifact elimination, and image resizing. (C) The image following edge detection (edges overlaid in cyan). (D) The image, converted to its familiar Cartesian coordinate system, with all detected edges overlaid (red) and retained edges highlighted (cyan).

whose composition and structure cause differences in relative optical reflectivity (i.e. brightness) and entropy. To identify the location of the IEM, information on these features was therefore leveraged. Relative differences in average OCT signal intensity grayscale values and neighborhood entropy above and below each segment were calculated. Upper and lower thresholds were set on the basis of characteristic values recorded for several manually identified IEM edge segments; segments with values outside the allowable range for these characteristics were eliminated. The image was then converted from polar to Cartesian coordinates, resulting in the familiar circular image (Figure 2.1D).

As these appearance (grayscale intensity and entropy) criteria did not fully isolate the desired edges, geometric criteria were subsequently applied because vessels, even when diseased, are known to maintain roughly elliptical shapes. To each remaining segment was fit a circle and ellipse:

$$(x - x_0)^2 + (y - y_0)^2 = R \quad (1)$$

$$x - x_0 = a \cos t + h \sin t \quad (2)$$

$$y - y_0 = b \sin t + h \cos t \quad (3)$$

Here, x and y represent the coordinates of a point (pixel) in the image, x_0 and y_0 represent the center of the circle or ellipse, R is the radius of the circle, t is a value between 0 and 2π , a and b are parameters determining the eccentricity and size of the ellipse, and h is a parameter determining the axial rotation of the ellipse. Best-fit circles (Equation 1) and ellipses (Equations 2 and 3) were therefore respectively determined by selecting $x_0, y_0, R, a, b,$ and h to minimize the following losses (\mathcal{L}) over all points (x,y) of a given edge segment:

$$\mathcal{L}(x_0, y_0, R) = \sum_{\forall(x,y)} \left| \sqrt{(x - x_0)^2 + (y - y_0)^2} - R \right| \quad (4)$$

$$\mathcal{L}(x_0, y_0, a, b, h) = \sum_{\forall(x,y)} \sqrt{\left((x - x_0) - (a \cos t + h \sin t) \right)^2 + \left((y - y_0) - (b \sin t + h \cos t) \right)^2} \quad (5)$$

Values were loosely constrained to feasible ranges values relative to the image dimensions (e.g. such that the center of the shape, (x_0, y_0) fell within the image field of view), and initial guesses for each value were similarly provided to the minimization function as proportions of the image dimension (e.g. radius R equal to a quarter of the image width). The resulting values defined a

best-fit circle (with 3 degrees of freedom) and ellipse (with 5 degrees of freedom) for each edge segment, which was then used to further eliminate unwanted segments. Segments corresponding to best-fit circles with a radius (R_{circle}) significantly smaller than the radius of a circle with the same area as the lumen region (R_{lumen}) were eliminated:

$$R_{circle} > R_{lumen} = \sqrt{\frac{A_{lumen}}{\pi}} \quad (6)$$

Additionally, if the best-fit circle did not enclose most of the lumen region extrema points, its corresponding edge segment was eliminated. Finally, if the quality of the ellipse fit was poor, meaning the minimized sum was not sufficiently small relative to the total number of points contained in the segment, the corresponding segment was eliminated. Thus, only segments corresponding to sufficiently large, reasonably located, and well-fitting best-fit circles and ellipses were retained (Figure 2.1D). These segments were designated as the visible segments of the outer border, to be used later in determining the full vessel wall.

Refinement improved the quantity and accuracy of visible outer border detection, but the method carries inherent limitations. This approach uses only in-plane (2D) information, even though information from adjacent frames may provide greater insight on whether a given contour corresponds to that which is desired, a different feature, or an artifact (e.g. noise). Furthermore, and as a consequence of the 2D nature of the edge detection, there was no guarantee of continuity, and contours detected in subsequent frames were frequently discontinuous. Finally, the method was sensitive to catheter location within the lumen, as the relative position dictated the polar representation of the vessel wall (which in turn impacts filtering and other operations). Further work to advance the visible border detection addressed these shortcomings.

2.2.2. Advanced Three-Dimensional Edge Detection

As for basic 2D edge detection, advanced 3D edge detection began with lumen detection. Similar to the basic approach, initial detection of the lumen followed the procedure used by Athanasiou *et al.* [173]. In brief, image data were converted to binary format with Otsu's thresholding method [215], common artifacts (i.e. catheter sheath and guidewire backscattering) were purged, and the single most proximal point in each column (non-zero point of lowest radial distance) was retained. Simple 2D smoothing was performed on these points, and they were then connected through linear interpolation to points in laterally-adjacent columns with straight lines,

resulting in a continuous contour within each frame. However, while minimally smoothed, the resulting lumen surface was not necessarily continuous in three dimensions (longitudinally). To further smooth the lumen as a continuous 3D surface, an anisotropic linear-elastic mesh was fit to the points using the approach described later (Section 2.3.3), diverging from the previous, basic 2D approach. The result was relatively reliable and accurate; error in lumen border detection generally arose only occasionally in the presence of significant image acquisition artifact, such as insufficient blood flush or irregular catheter rotation, as well as side-branches. With the lumen accurately identified, outer border segment detection could proceed.

Identification of the (visible) outer border candidate edges followed lumen detection and drew upon image enhancement and edge detection approaches specifically adapted for this application. Undulation of the vessel wall is natural in the acquired polar image of the vessel wall due to the nominally non-circular shape of the vessel and non-concentricity of the imaging catheter within the vessel lumen. In a comparable situation, the rounded nature of the eye, and consequently retina, leads to large curvature of the acquired images. Therefore, an approach drawn from retinal OCT layer segmentation was utilized to flatten the image relative to the detected lumen contour to better align layers horizontally (Figure 2.2). This feature improved filtering efficacy, facilitated reduction of image sizes, and eased layer boundary identification through improved shape consistency [185], [190]. In each column of a frame, the lumen was shifted to the top of the image. Following this procedure, the outer border presents as a straight line in healthy arteries (arteries without plaque and a thin intima layer of consistent thickness). However, in diseased arteries, plaque buildup displaces the outer wall distally to generally unequal depths along the angular span of the wall (Figure 2.2).




Figure 2.2. OCT images of the arteries are flattened in polar coordinates relative to the lumen, improving filtering efficacy, size reduction, and layer boundary identification through improved shape consistency.

A compensation and contrast enhancement algorithm employing exponentiation was next applied to the flattened polar images using methods proposed by Girard *et al.* [218]. To enhance the image, we used a transformation in which contrast enhancement through exponentiation precedes image compensation:

$$HD\{I^n(r)\} = \frac{I^n(r)}{2 \int_r^\infty I^n(u) du} \quad (7)$$

where HD is a compensating transformation operator for a given exponentiated column of the image, I is pixel intensity, n is the exponential factor, and r is (propagating beam) radial depth [218]. As in the referenced study, an exponential factor of 2 was used, and the discretized implementation of the algorithm described by Girard *et al.* [218] was utilized.

Following this compensation and enhancement, which operated in just one dimension, frames were concatenated on each side with their adjacent frames to provide appropriate boundaries for subsequent multi-dimensional filtering. Because OCT data are acquired in a continuous helical sweep [18], rather than at discrete incremental positions as suggested by its normal display, these adjacent edges truly represent the boundaries of the frames and provide a continuation of the image data along the frame edges. The entire boundary-identification procedure was performed in this concatenated volume to avoid erroneously penalizing or disadvantaging segments along the edge of a frame, as criteria (i.e. size/area, major-to-minor axis length ratio, and orientation) would otherwise be skewed for sub-segments of boundaries that wrap to the opposite side of the polar image.

Speckle noise was magnified during the transformation, so simple median filtering was performed to de-noise the resulting compensated and unattenuated form of the images [219]. As the flattened polar images were arranged consecutively in 3D space (with pixels functioning as voxels), filtering was performed with a 3D kernel to leverage information from adjacent frames and minimize the impact of local artifacts.

Sobel-Feldman edge detection kernels [220] detected grayscale gradients in the radial (vertical) direction. Skewed vertical distribution kernels—specially designed to account for the nominal 200 μm thickness of the media layer [8], [164]—were subsequently convolved with the result. Characteristic radial light-dark-light intensity (grayscale value) transitions indicative of the media layer [18], the inner surface of which was considered here as the outer border of the vessel, were thereby identified as candidate outer border segments (Figure 2.3). Segments of the outer




Figure 2.3. Identifying the edge contours corresponding to the outer border. (*Top Left*) The original flattened image on which the operations were performed; substantial fatty plaque (*) partly obscures abluminal structures of the vessel wall. (*Right*) Convolution results calculating the weighted mean Sobel output above (*Top*) and below (*Bottom*) each pixel. Red indicates a positive value (light-to-dark transition); yellow indicates a negative value (dark-to-light). (*Bottom Left*) Heat map illustrating the product of the proximal and distal convolutions meeting the imposed signage requirements (light-to-dark transition above and dark-to-light transition below a given pixel).

border were then identified from the matrix of candidate segments through a 3D region-growing scheme that extracted segments of prominent, well-defined transition patterns. The 2D properties of resulting regions' intersections with each given plane were calculated to help eliminate non-feasible segments—regions with small major-to-minor axis length ratios (i.e. too round, rather than elongated) or orientations offset substantially from horizontal were removed.

As the standard Sobel-Feldman and vertical distribution kernels did not capture edges oriented at an angle as well as those that are horizontal, off-axis angles were substantially disadvantaged in the global region-growing. To overcome this shortcoming and capture those angled border segments—typically located at the transition to diseased regions of the vessel—the same edge identification procedures described above were repeated with Sobel-Feldman kernels and distribution kernels oriented at both 45° and -45° . Subsequent orientation feasibility criteria were modified accordingly. Edges detected with the standard and off-axis procedures were

combined to form a 3D binary representation of the detected segments of the outer border in the flattened 3D image volume.

Segments were identified by the presented approach in just over half of all radial projections in a diseased dataset (described in Section 2.3.3). However, some segments visible to the trained eye were not identified. Of greater concern were the falsely identified segments that did not correspond to the actual outer border. These segments usually arose at the adventitia-periadventitia transition or, more problematically, within heterogeneous plaques where the selected grayscale pattern occurred by happenstance (e.g. when thin bands of calcified tissue were present). To help address these shortcomings, additional strategies for segment detection improvement could include explicit incorporation of penetration depth, the external limit of which is not currently considered by the algorithm. Inclusion of this factor could be used to delimit a region of interest within a proximal subset of the image, thereby eliminating opportunities for false detection beyond this region and reducing computational operations in regions of the image in which no viable information is likely to be present. Still, however, this advanced 3D edge detection approach proved sufficient in the available dataset and for the application at hand, and represents a technical feat.

Advanced 3D edge detection provided more accurate and reliable designation of visible outer border segments than the basic 2D edge detection. By using out-of-plane (3D) information, information from adjacent frames could be leveraged which provided greater insight on whether a given contour corresponded to that which was desired, a different feature, or an artifact (e.g. noise). Furthermore, and as a consequence of the 3D nature of the edge detection, continuity was far more common, and contours detected in subsequent frames were frequently continuous. Finally, the method was far less sensitive to catheter location with the lumen, as the flattening step largely negated sensitivities in the polar representation of the vessel wall. Still, fundamental limitations of the modality's relatively low penetration depth—with strong dependence on tissue optical properties—does not allow for complete visualization of the entire vessel wall in diseased cases. To extract the full geometry, non-visible segments of the border still have to be identified.

2.3. Identifying the Complete Border from Visible Segments

As noted earlier, if there is minimal plaque burden or other impediment obscuring the view of the outer border of the vessel wall or necessitating smoothing, identifying the visible features

of the inner and outer border is the only step needed to extract vessel geometry. However, if segments of the vessel wall are obscured, visible feature identification is just the first step, and these visible segments provide important context in determining the location of the non-visible segments. In such cases, physiological understanding and constraints can be brought to bear on the problem. Plaque burden follows a continuity inside the vessel and does not change dramatically from frame to frame. By employing *a priori* knowledge of the structural and geometric constraints of vessel walls and image acquisition procedure, the limited visibility of the outer vessel wall can be largely overcome. Indeed, the non-visible parts of the outer vessel border can be approximated in the 3D space to estimate the full arterial wall position, thereby determining the geometry of diseased vessel wall geometry.

Complete outer border contours were determined through strategic fitting of the detected visible segments (Section 2.2). These methods were implemented entirely in MATLAB (MathWorks, Natick, MA). Approaches to fitting visible segments to estimate the full outer border (Figure 2.4) began with simple 2D ellipse fitting to visible segments identified in the given frame (Section 2.3.1) [155] and progressed to frame-by-frame shape fitting to visible segments identified in both the given and adjacent frames, thereby beginning to integrate contextual information and improve longitudinal continuity (Section 2.3.2) [156]. However, findings from these initial approaches made clear that 3D surface fitting is required for this task, and subsequent approaches implemented this by using a unique mechanical approach for smooth surface fitting (Section 2.3.3) [157]. Leveraging the flexibility and power of this new approach, further work extended the method to simultaneously fit coupled surfaces to delineate both the IEM and EEM (Section 2.3.4) [158].

2.3.1. Simple Two-Dimensional Shape Fitting

Perhaps the simplest approach to estimate the outer border of a vessel when not fully visible is to fit a shape, known to roughly approximate that of the vessel, to the visible segments in a given frame. In one implementation (Figure 2.5), an ellipse was fit to those visible segments identified with basic 2D edge detection (Section 2.2.1) [155]. With the edge segments selected down to reasonable candidates for belonging to the IEM (Figure 2.1), a best-fit circle and ellipse was determined, by the same procedure as outlined in Section 2.2.1, for all remaining points (by minimizing Equations 4 and 5 for the shapes defined by Equations 1–3). Ellipses were found to

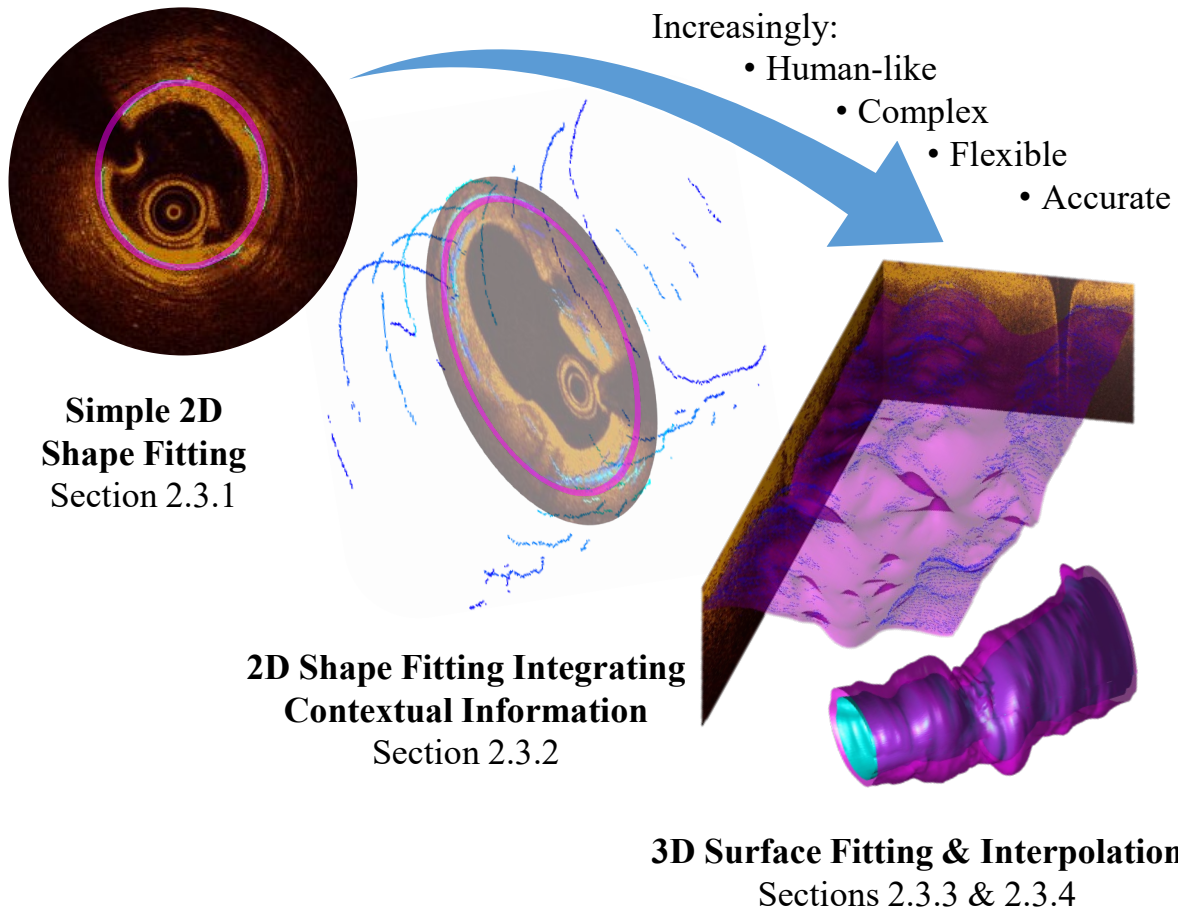


Figure 2.4. Methods to determine the complete outer border on the basis of partial, visible segments progressed from simple 2D shape fitting to 3D surface fitting and interpolation, offering increasingly human-like, complex, flexible, and accurate approaches.

produce better results than the circles, so were used exclusively thereafter. The area enclosed by the ellipse (A) was then determined computationally and calculated independently:

$$A = ab\pi \quad (8)$$

Validation of this approach was carried out in anonymized OCT and IVUS images of the same vessel region in a single patient (a 77-year-old male with non-ST elevated acute coronary syndrome). First, the IVUS images, which are considered the clinical gold standard for *in situ* spatial measurement of the vessel wall, were manually co-registered with the OCT frames by experienced experts. The meaningful region of the OCT pullback was identified and all OCT frames from this region with corresponding co-registered frames in the IVUS pullback were identified ($N = 37$). An expert interventional cardiologist, blind to the algorithm results, traced the outer wall of the vessel in each corresponding IVUS frame, and the enclosed area was measured. This allowed direct comparison between the cross-sectional area (CSA) estimated in OCT images




Figure 2.5. Sample OCT image of minimally-diseased vessel cross-section illustrating the outcomes of simple 2D shape fitting following basic 2D edge detection. These images correspond to the same frame shown in Figure 2.1, and display the lumen delineation (green), detected edges corresponding to the outer border (cyan), and best-fit ellipse (magenta) corresponding to the final outer border estimate for the frame. (*Inset*) The original (unlabeled) image, displayed in Cartesian coordinates, shows the distinctive three concentric layers of the vessel wall (visible in most of the cross-section).

by the simple 2D shape fitting algorithm with those determined in IVUS images by an experienced interventional cardiologist (Figure 2.6). Overall results were calculated and reported as mean \pm standard deviation (SD).

The algorithm successfully identified segments of the IEM in most analyzed frames, though visible segments were not always identified and extraneous edges were not always




Figure 2.6. Co-registered IVUS and OCT images used in the validation of the simple 2D ellipse fitting method applied in tandem with basic 2D edge detection. (*Left*) A human expert (interventional cardiologist) manually annotated the outer border of the vessel wall in IVUS (green), considered the clinical gold standard for this task. (*Right*) The automated computer algorithm annotated the outer border of the vessel wall in OCT by fitting an ellipse to the visible segments of the wall automatically identified in the frame (magenta).

excluded. In one frame ($N = 1$), no visible outer border segment was detected through the automated method, and was discarded. OCT images overlaid with the best-fit ellipse of the IEM are shown in Figures 2.5 and 2.6.

The IEM CSA estimates offered by the algorithm differed from the gold standard measurements made by experienced expert clinicians in IVUS images by $1.0 \pm 1.1 \text{ mm}^2$ (median: 0.7 mm^2). Area estimates derived from the OCT images by the algorithm were, on average, 0.6 mm^2 (10%) smaller than those derived from the IVUS images by an expert operator (Figure 2.7)⁴. Furthermore, despite some promise and success, CSA annotated by the human expert in IVUS and the simple algorithm in OCT did not achieve statistically significant correlation.

Discrepancies between the IEM CSA estimates and gold standard measurements (Figure 2.7) were generally dominated by systematic factors which skewed the automated estimates to artificially low values relative to the IVUS measurements. For example, the frame displayed in Figure 2.6, which shows qualitatively accurate identification of the IEM, still results in an area estimation error of 10% (relative to IVUS “gold standard”). Some underestimation relative to

⁴ Note that these results differ from those reported in [155]; the reported mismatch in areas prompted a review of the imaging systems, which identified a miscalibration. Correcting for this scaling issue yielded the results reported here.




Figure 2.7. Vessel CSA estimates made by in-frame shape fitting deviate from the expert annotations. (*Left*) The vessel area estimates derived automatically from OCT and measured from human annotations in IVUS differ, with the latter providing consistently larger values (0.6 mm^2 , or 10%, larger on average). (*Right*) The absolute difference in area largely falls around the 1.0 mm^2 mean (solid line) and within the 95% confidence interval (dashed line; $\text{SD} = 1.1 \text{ mm}^2$). A few poorly-estimated outliers skew the average difference from the 0.7 mm^2 median.

IVUS is expected—while the algorithm aimed to identify and fit visible segments of the IEM, only the (abluminal) EEM is visible in IVUS, and furthermore studies have found IVUS measurements to overestimate area measurements by 10-15% [40]. However, other significant factors identified here contributed to CSA estimate discrepancy and error. False and missed edge identification contributed to error, and dominated those estimates experiencing the largest errors. These situations, in which the algorithm did not perform well, are often those in which the IEM border is difficult to identify visually, even by expert clinicians. The missed edge identification corresponded mostly to the minimum lumen area sites with the largest plaque burden and IEM areas. Additionally, unexpected shadowing or significant variation in signal intensity of similar tissue within a single frame resulted in sub-optimal performance, as did major or unusual imaging artifacts.

Continued development of geometry-extraction algorithms therefore sought to identify and address the cause of the systematic discrepancy and overcome the edge detection obstacles in order to improve the accuracy of the algorithm. A major need arising from this work was to extend the 2D approach to 3D, enabling improvements in accuracy by allowing estimations to leverage well-defined and identified segments of the IEM in adjacent frames, i.e. contextual information.

2.3.2. Two-Dimensional Shape Fitting Integrating Contextual Information

Building upon previous work applying simple shape fitting to visible edges in 2D, 2D shape fitting was applied in a manner which integrated contextual information. Recognizing that to estimate the IEM within a frame, it is necessary to use information from adjacent frames and *a priori* knowledge of vessel shape, a somewhat more complex algorithm was developed that leverages these approaches to estimate the IEM in OCT frames autonomously [156]. As for the method described in Section 2.3.1, a shape known to roughly approximate that of a vessel was fit to visible segments of the outer border to estimate the full outer border of a vessel when not fully visible. This time, however, an ellipse was fit to the visible segments within a given frame *as well as* those identified in adjacent frames [156].

In each frame, an ellipse was fit to detected IEM segments in the current and adjacent frames to approximate an annotation of the feature. As before, visible segments were identified with basic 2D edge detection (Section 2.2.1; Figure 2.1) [155]. The ellipse was fit by optimizing its geometric parameters (Equations 2 and 3) to minimize the weighted sum of distances between the ellipse and the points constituting the detected segments (Figure 2.8):

$$\sum_{\forall(x,y,z)} w(x,y,z) \sqrt{\left((x-x_0) - (a \cos t + h \sin t)\right)^2 + \left((y-y_0) - (b \sin t + h \cos t)\right)^2} \quad (9)$$

$$w(x,y,z) = w_{\text{uniqueness}}(x,y) \times w_{\text{distance}}(|z-z_0|) \quad (10)$$

Note the distinction from Equation 5; Equation 9 includes an error weighting factor w (Equation 10) for each detected point (x, y, z) and considers points not located at the longitudinal position (z) of the considered plane (z_0) . Weights were based on distance from the current frame (current frame weighted most heavily) and location uniqueness (less common location along arc weighted more heavily). The area enclosed by each ellipse (A) was calculated (Equation 8).

The algorithm was validated using experts' CSA measurements in both OCT and gold standard IVUS images (Figure 2.9). OCT and IVUS pullbacks from two patients were considered for this study (77- and 61-year-old males experiencing non-ST elevated acute coronary syndrome). The pullbacks from the two modalities were co-registered, and coincident frames in useable regions were annotated ($N = 215$). Useable regions were designated as those unobstructed by unflushed blood or other significant artifacts. Two experienced interventional cardiologists annotated the IEM in the OCT frames; one also annotated the EEM in the corresponding IVUS

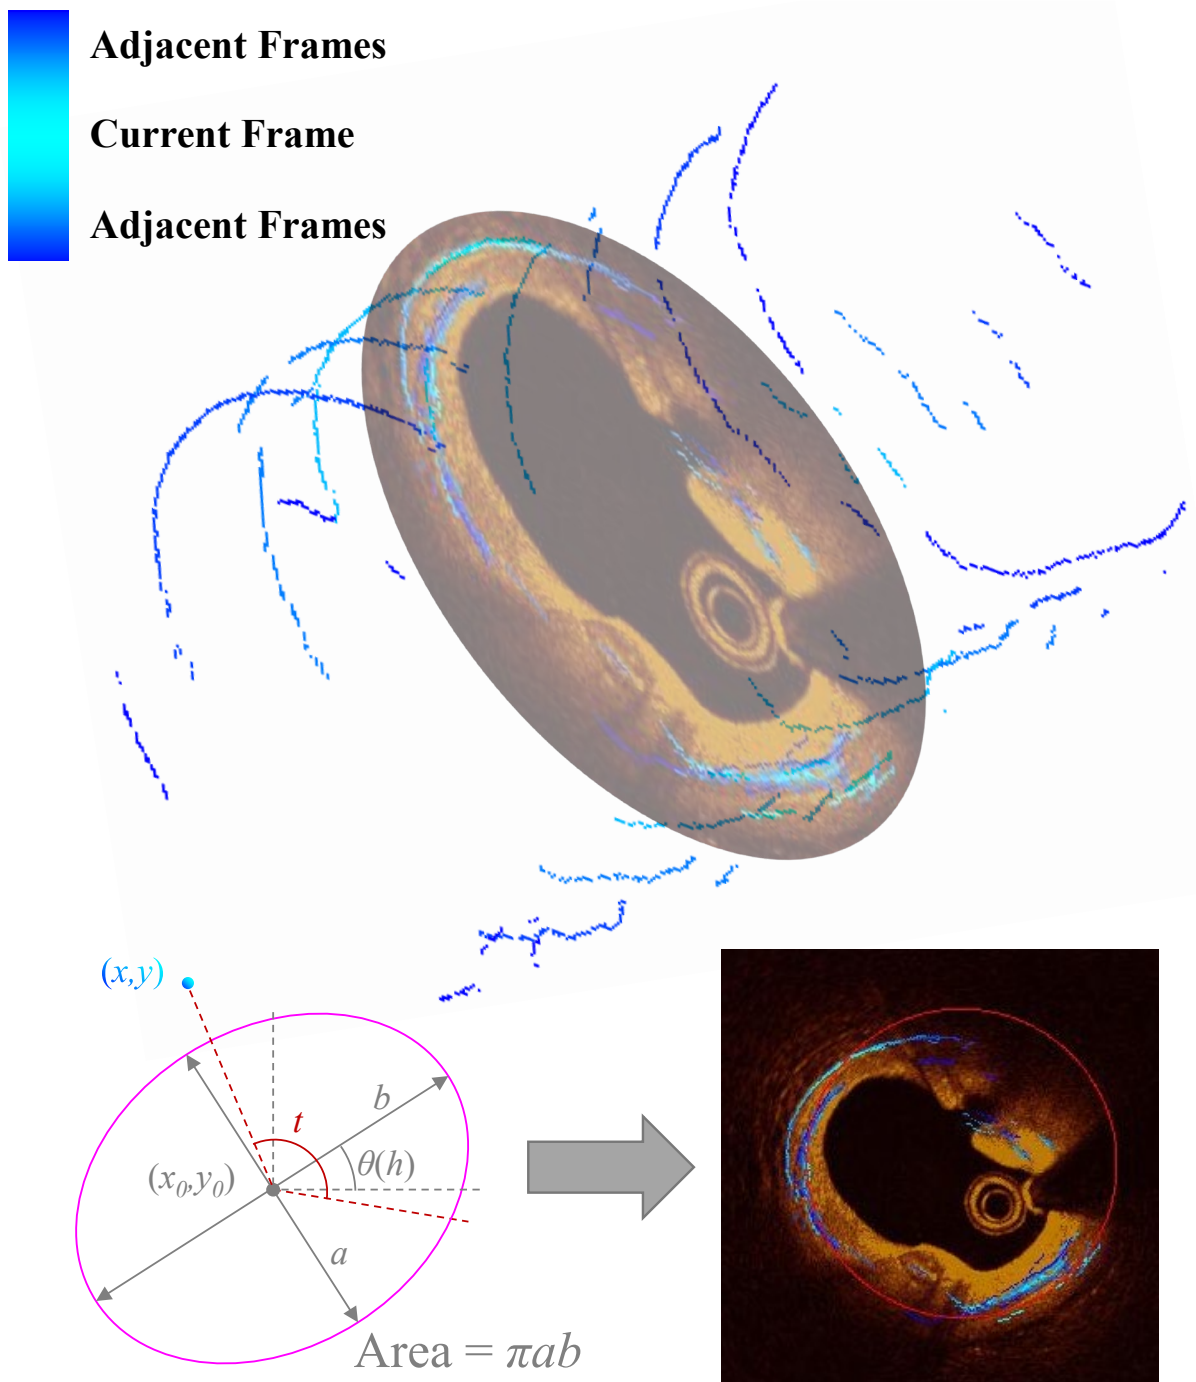


Figure 2.8. An ellipse was fit to visible segments of the IEM detected in the current and adjacent frames. While shape fitting was performed in 2D, contextual information from adjacent frames was integrated by projecting the visible edges detected in adjacent frames into the current one; acknowledging that frame-to-frame variation is expected, error contribution of a point was degraded with longitudinal distance from the frame under consideration. Parameters for an ellipse described by Equations 2 and 3 were optimized to minimize the resulting cumulative error (Equation 9). Area was computed for the resulting ellipse (Equation 8).

frames (Figure 2.9). Furthermore, to assess the impact of vessel wall visibility, OCT frames were designated by one of the experts as “Readable” or “Unreadable” depending on whether the IEM

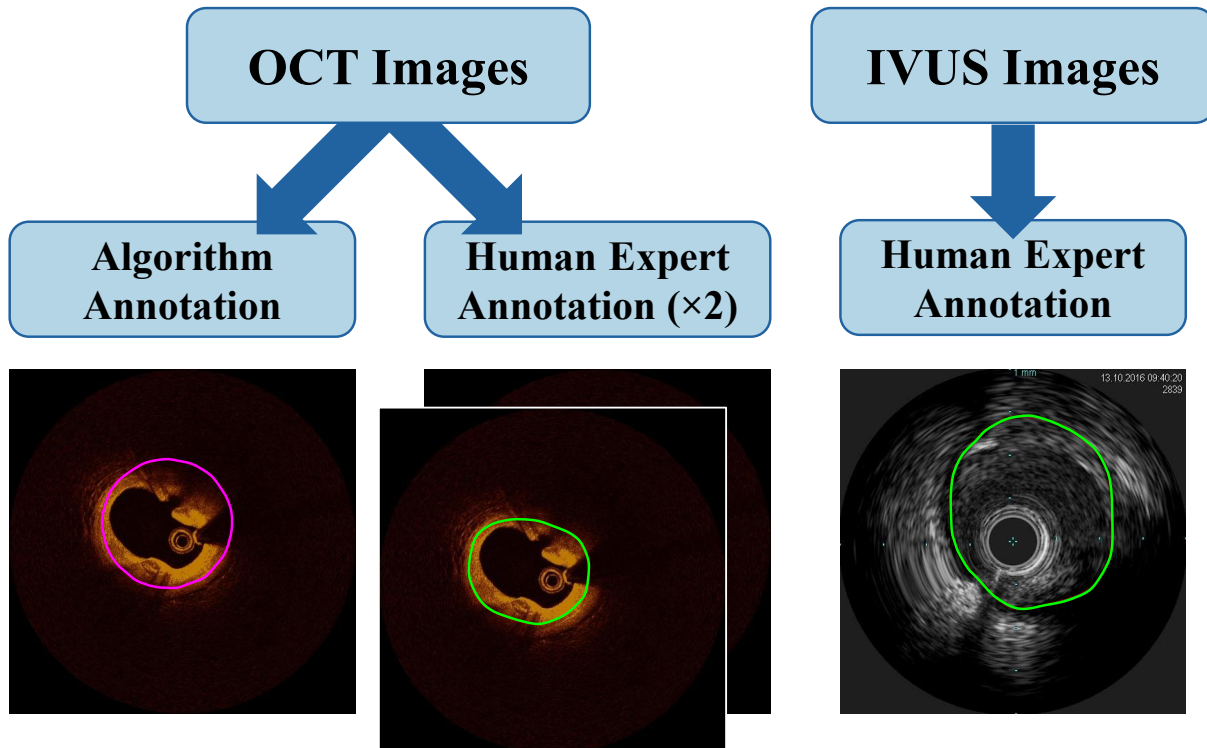


Figure 2.9. The method employing 2D shape fitting and integrating contextual information to delineate the vessel wall in OCT was validated through comparison with human expert annotations in both OCT and IVUS images. The diagram shows the annotations compared and analyzed. (*Left*) The automated computer algorithm annotated the outer border of the vessel wall in OCT by fitting an ellipse to the visible segments of the wall automatically identified in the current and adjacent frames. (*Center*) Two human experts (interventional cardiologists) manually annotated the outer border of the vessel wall in the same OCT frame. (*Right*) A human expert (interventional cardiologist) manually annotated the outer border of the vessel wall in IVUS, considered the clinical gold standard for this task.

was visible in at least 3 quadrants ($>270^\circ$ visible). IVUS, which has greater penetration depth and is well validated for identification of lumen vessel outer wall, was treated as the gold standard.

The computer algorithm performed better than that which did not consider contextual information from adjacent frames (Section 2.3.1), and achieved strong statistical significance, but did not perform as well as the interventional cardiologists (Table 2.1). Inner (lumen) border detection was excellent in emulating human experts (Figure 2.10), but performance was less reliable for the outer border (Figure 2.11). The average differences (magnitude) between the CSA of the vessel determined by algorithm in OCT and experts in OCT and IVUS are presented in Table 2.1, along with correlation (p) values between each set. CSAs of OCT annotations—both those generated by the algorithm and the experts—better matched IVUS annotation CSAs in frames identified as “Readable” (showing visible IEM in 3+ quadrants; Figure 2.11; Table 2.1). As for the previous approach (not integrating contextual information), the computer algorithm

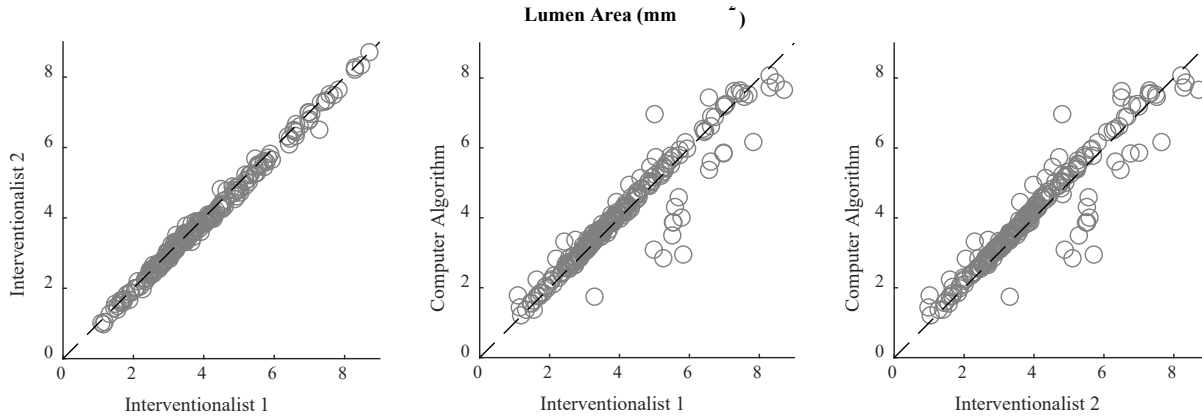


Figure 2.10. Direct comparison and Bland-Altman plots of lumen annotation CSA (mm^2) from OCT. (*Left*) Very little inter-observer variation was observed in annotation of the lumen in OCT images. (*Center & Right*) The simple lumen detection algorithm (delineating the inner border of the vessel wall) was very accurate and closely reflected human experts, with a handful of outlier cases in which it performed relatively poorly. Black dashed lines indicate a 1-to-1 ratio, which reflects complete agreement between the measurements.

tended to underestimate CSA, particularly in “Unreadable” frames with less visibly identifiable IEM (Figure 2.11, right). Notably, however, unlike the strictly 2D approach (Section 2.3.1), there were no frames which had to be discarded due to a complete absence of visible segments. For OCT frames in which the IEM is largely visible (“Readable”), the presented computer algorithm provided similar measurements of cross-sectional area when compared to expert annotation in OCT, especially in smaller vessels with greater IEM visibility (Table 2.2). The two expert annotations in OCT varied by $1.29 \pm 1.18 \text{mm}^2$ (13.8 \pm 11.2%; $R = 0.82$, $p < 10^{-53}$) in all frames and by $1.09 \pm 1.22 \text{mm}^2$ (11.0 \pm 9.0%; $R = 0.93$, $p < 10^{-33}$) in “Readable” frames. As expected, “Readable” and “Unreadable” frames differed significantly in IEM CSA and plaque burden (Table 2.2).

Table 2.1. Average difference (magnitude) between vessel CSA obtained by different means

	OCT – Algorithm		OCT – Experts	
	All	“Readable”	All	“Readable”
IVUS – Expert	$2.51 \pm 1.63 \text{mm}^2$ 32.1 \pm 20.3% $R = 0.72$ $p < 10^{-34}$	$1.58 \pm 1.10 \text{mm}^2$ 19.9 \pm 11.3% $R = 0.88$ $p < 10^{-25}$	$1.40 \pm 1.17 \text{mm}^2$ 15.5 \pm 12.0% $R = 0.79$ $p < 10^{-91}$	$1.28 \pm 1.04 \text{mm}^2$ 14.8 \pm 9.1% $R = 0.86$ $p < 10^{-46}$
OCT – Algorithm	-	-	$2.02 \pm 1.71 \text{mm}^2$ 25.2 \pm 20.1% $R = 0.78$ $p < 10^{-89}$	$1.06 \pm 1.20 \text{mm}^2$ 12.1 \pm 11.4% $R = 0.91$ $p < 10^{-60}$

Table 2.2. Comparison of average vessel CSA and plaque burden—as annotated by human expert in IVUS—in frames deemed “Readable” and “Unreadable” by human expert in OCT

	“Readable” (N = 78)	“Unreadable” (N = 137)	p-value*
Vessel CSA [mm^2]	8.76	9.49	0.03
Plaque Burden	0.38	0.43	$< 10^{-4}$

*Two-tailed t-test

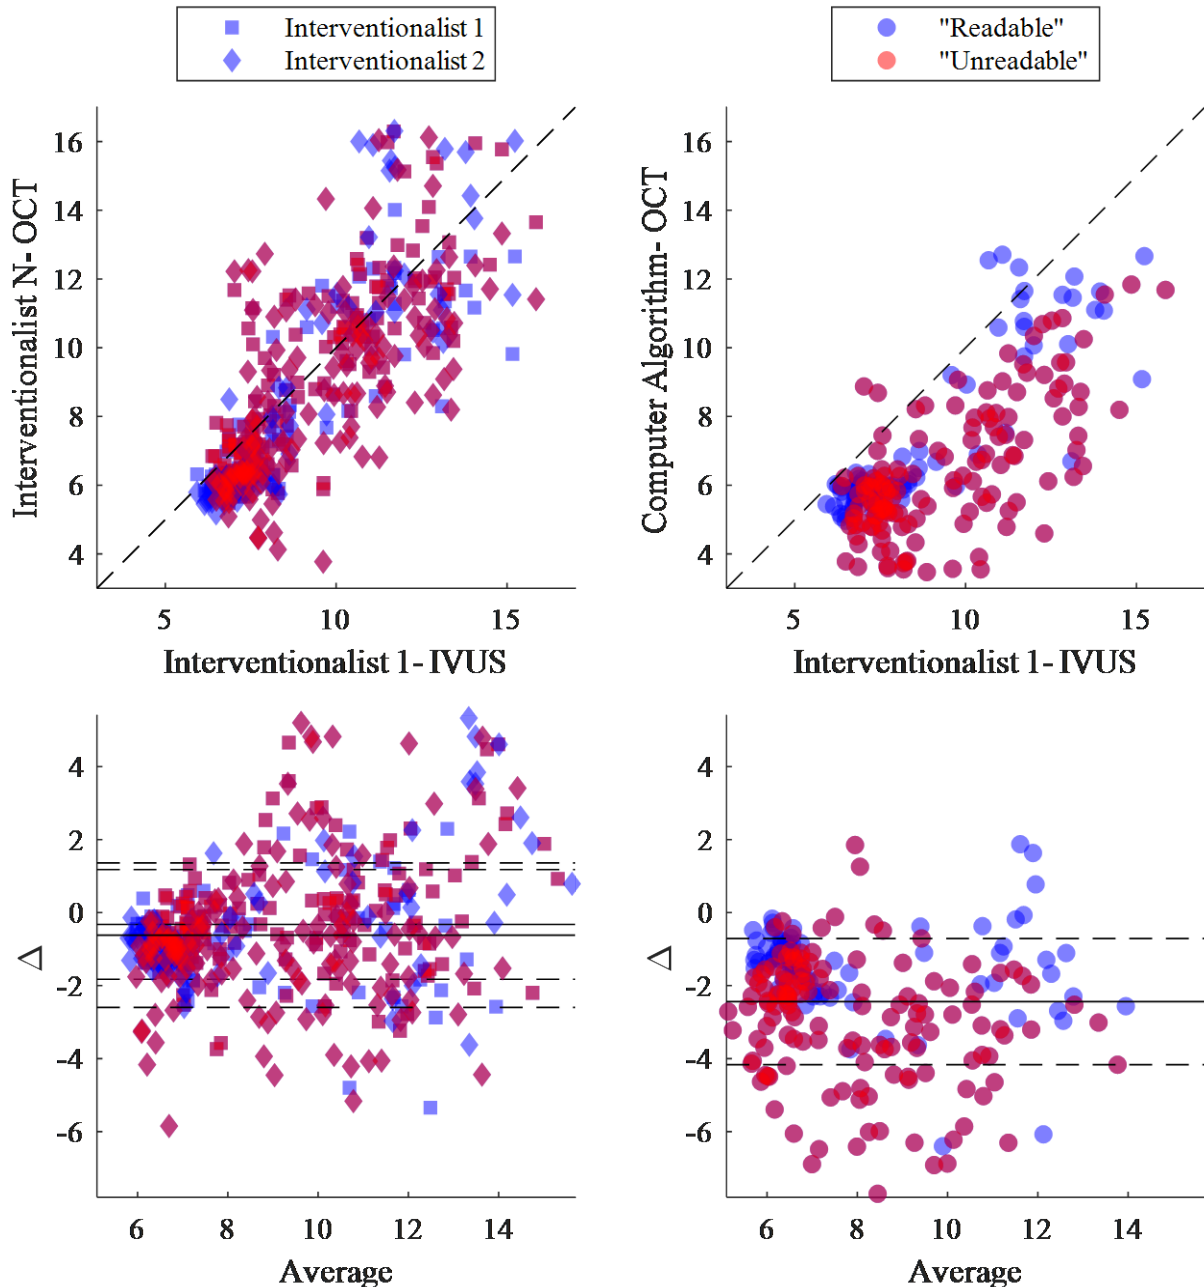


Figure 2.11. Direct comparison and Bland-Altman plots of IEM annotation CSA (mm^2). (Left) CSAs determined by two human experts in OCT were smaller than those determined by a human expert in IVUS by 0.32mm^2 and 0.62mm^2 (overall average absolute difference of $1.40 \pm 1.17\text{mm}^2$). (Right) CSAs determined by the algorithm in OCT were smaller than those determined by a human expert in IVUS by 2.44mm^2 (overall average absolute difference of $2.51 \pm 1.63\text{mm}^2$). While human performance was superior to the algorithm, both sets achieved strong statistical correlation, and performance distinction diminished substantially in “readable” frames (blue markers).

This method leveraged information from adjacent frames and *a priori* knowledge of vessel shape, enabling automated annotation and estimation of IEM CSA with strong statistical significance and producing similar CSA measurements in “Readable” frames compared to

interventional cardiologist [156]. This represented a considerable improvement over the previous method which relied solely on information extracted from the present frame. It remained, however, strongly affected by visibility of the IEM. Comparison between annotations made in OCT and IVUS revealed, as expected, that OCT readability was correlated with vessel CSA and (more so) plaque burden.

Integrating contextual information was an important step towards successfully and accurately extracting vessel wall geometry, and the results presented here unveiled important insights on achieving this—including how it is accomplished by humans. While human performance was, expectedly, better in “Readable” frames, it was surprisingly robust in “Unreadable” frames—far more so than the computer algorithm. It therefore became apparent that further features of human experience in identifying the outer border the vessel wall were not being captured by the computational algorithm. There also remained several obvious simplifying assumptions to be addressed. Error persisted in part due to the assumption that the IEM can be described by an ellipse. While this is a generally reasonable assumption justified by years of image review by a contributing clinician-researcher, it is clearly imperfect. Subsequent work explored more complex shapes to fit to the IEM (e.g. egg shapes), but longitudinal continuity could not be enforced with a 2D shape-fitting approach. Other approaches were needed to connect identified segments in a way which more accurately reflected the shape and continuity of the IEM, as well as the experience of humans in assimilating contextual information from the entire OCT acquisition.

2.3.3. Three-Dimensional Surface Fitting & Interpolation

The 2D shape-fitting approaches advanced the state-of-the-art, but failed to match the performance of human experts in delineating the geometry of diseased vessels imaged by OCT. Geometric limitations imposed by the available shapes limited flexibility in handling complex phenotypes, and even when contextual information was drawn from adjacent frames, longitudinal continuity could not be enforced with such 2D shape-fitting approaches. To relax geometric restrictions and impose such continuity—concepts intuitively grasped and implemented by human experts—3D approaches were needed to connect visible segments of a vessel’s outer border. Success was achieved through development of a promising new smoothing surface fitting method that overcame the challenges of existing interpolation and surface fitting methods, and enabled the

delineation of the entire inner and outer borders of vessels in OCT images [157]. Building upon the previous work in 2D shape-fitting [155], [156], the new method, which was inspired by intuitive linear-elastic springs common in classical (Newtonian) mechanics, was both efficient in 3D space and conspicuously effective. As such, the novelty of this work derived not only from its application—in autonomous delineation of the full outer border of a diseased coronary artery using OCT imaging alone—but also in the approach used to reconstruct and smooth surfaces using incomplete information [157].

The novel 3D surface fitting method, employed in concert with the advanced 3D edge detection (Section 2.2.2), followed a multi-step process (Figure 2.12):

- (1) the inner (lumen) border of the vessel was detected in each frame and smoothed as a 3D surface;
- (2) the image was flattened relative to the lumen (Figure 2.2) and visible edges of the outer wall were identified using image contrast enhancement, 3D filtering, application-specific edge detection kernels, and region-growing (Section 2.2.2; Figure 2.3);
- (3) an anisotropic linear-elastic mesh was fit to the identified edges with applied forces to generate a continuous, smooth surface along the outer border; and

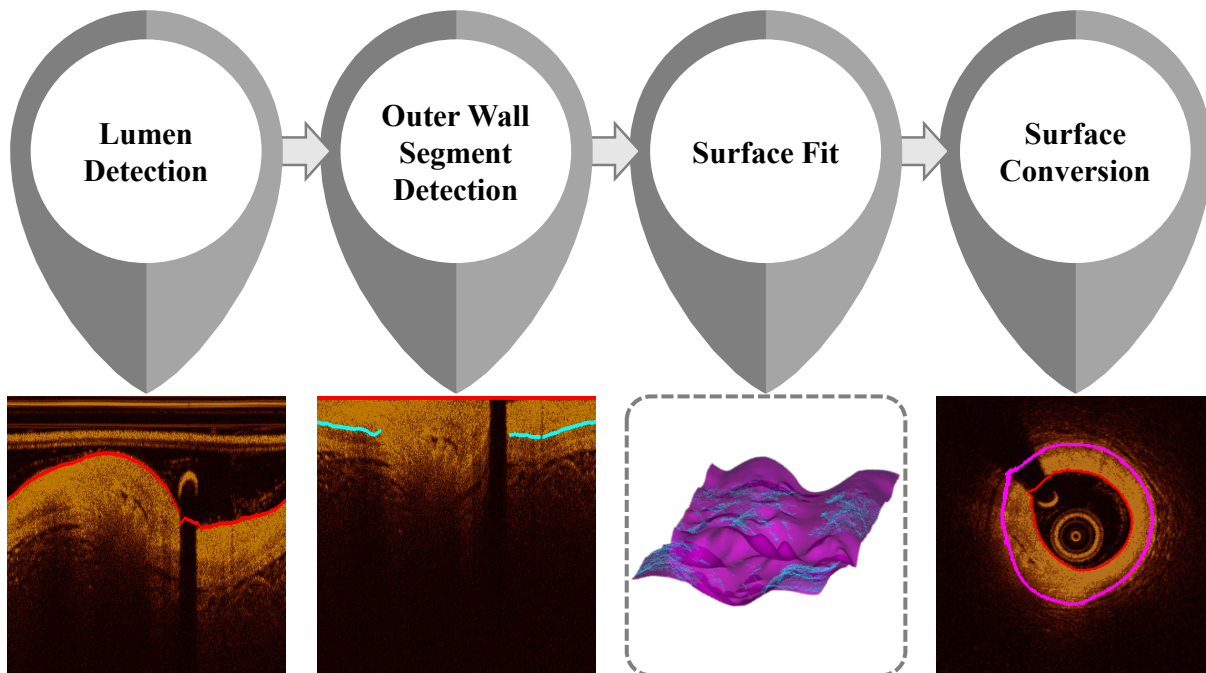


Figure 2.12. Major methodological procedures of the novel 3D surface fitting method included lumen detection and smoothing, visible outer border segment detection, surface fit to wall segments, and surface reconstitution with conversion from cylindrical to Cartesian coordinates. Initial lumen (inner border) contours and visible segments of the outer border were determined by advanced 3D edge detection (Section 2.2.2).

- (4) the results were finally reconstituted to the pre-flattened frame of reference and converted to a Cartesian coordinate system.

Note that the surface fitting method was used twice: first to smooth the inner border, and second to both complete *and* smooth the outer border. In the first case, no information was missing (nominal inner border was available for the full vessel from the initial detection procedure), but the given information was a bit noisy and not perfectly reliable. In the latter case, only incomplete, irregularly-spaced, and sometimes unreliable information on the outer border location was available from the initial detection procedure.

In providing a general description of the spring mesh surface construction, a finite set of m known points $\mathbf{P}_n \in \mathbb{R}^3$ and tolerances $\epsilon_n \in \mathbb{R}_+$ is defined, as for any surface fitting approach [194]. The desired surface $\mathcal{S} \subset \mathbb{R}^3$ approximates the positions of points such that

$$\|\mathcal{S}_{\mathbf{P}_n} - \mathbf{P}_n\| \leq \epsilon_n, n \in \{1, \dots, m\}. \quad (11)$$

While this general framework can accommodate continuous positions, known points are here restricted to discretized, indexed location to facilitate application to image processing. Specifically, here the framework is implemented in a 3D discretized space with q sequential planes $j \in \{1, \dots, q\}$ each containing p columns $i \in \{1, \dots, p\}$. Each point of the surface resides at some depth k , such that its position can be assigned an index of (i, j, k) .

The novel surface fitting approach is a global method implementing a virtual anisotropic linear-elastic mesh of interconnected nodes and springs (Figure 2.13). Each column of the discretized 3D space, assigned a distinct (i, j) index, contains a single node constrained to that column but free to translate in the \hat{k} -direction. This node represents the fit surface's position at the given column index, S_{ij} . Springs join adjacent nodes (in the \hat{i} - and \hat{j} -directions). If there exists a known point \mathbf{P}_n in a column (P_{ij}) , a position-fixed anchor may be placed at (i, j, k) , where index k is depth of the known point (p_{ij}) . A base layer (L_{ij}) of variable height (l_{ij}) may also be present, and represents a nominal, average, or expected position. Springs oriented in the \hat{k} -direction join nodes to either an anchor or the base layer (if no anchor is present in the given column). Several classes of springs may be present (Figure 2.13), e.g. those adjoining nodes:

- (1) within the plane j , from column i to $i+1$ and $i-1$ (Class 1);
- (2) between adjacent planes, from plane j to $j+1$ and $j-1$ (Class 2);

- (3) to anchor points (Class 3); and
- (4) to the base layer (Class 4).

Each class of springs may be of different stiffness, K . Note that any tolerance condition (Equation 11) is guaranteed to be met if anchored springs (Class 3) are infinitely stiff, i.e. $K_3 = \infty$. This assignment would enforce that a surface node position is absolute and fixed to coincide with an anchor present in its column, thereby guaranteeing no deviation from observed points while compromising smoothness. The stiffness of each spring could also be, though is not here, assigned individually (for example, such that K_3^{ij} is proportional to confidence in the anchor point position p_{ij} , or other approaches akin to a weighting function for smoothing weighted regression [195]).

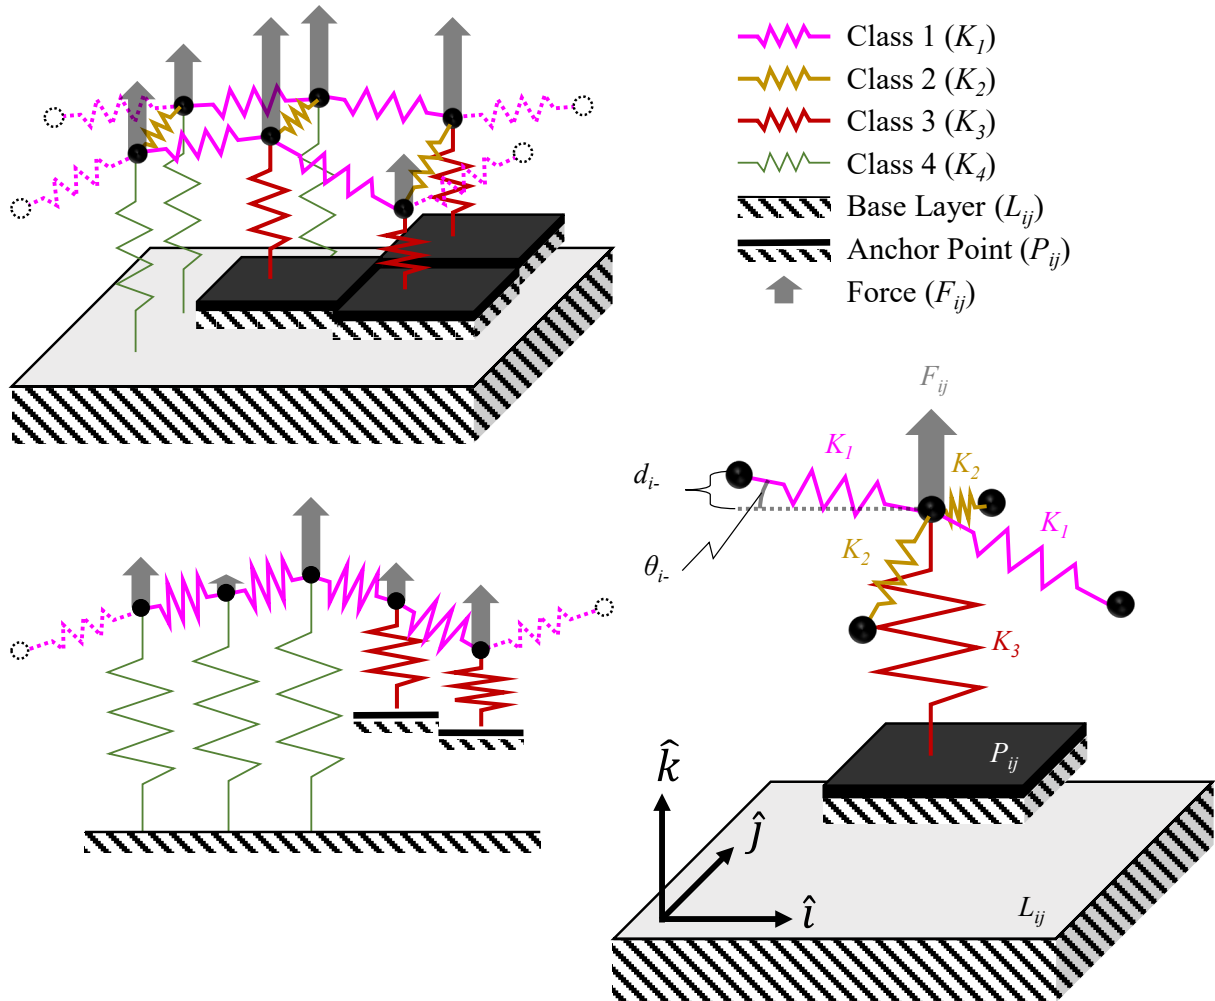


Figure 2.13. Schematic diagram showing the anisotropic linear-elastic mesh system structure. Spring colors indicate class with corresponding stiffness K . Black blocks indicate anchor points (P_{ij}) coinciding with known data points intersecting the column. A light-gray base layer is visible along the bottom (shown here as constant height, $l_{ij}=l$). Semi-transparent gray arrows indicate forces applied to each node (F_{ij}). Adjacent node distance (d) and direction (θ), used in calculating force balance, is shown for a single node.

However, individual assignment increases the complexity of the system, and was not deemed worthwhile in the present case (given the flexibility offered by additional features).

A unique force, F_{ij} , can be applied to each node. The magnitude and direction (positive or negative \hat{k} -direction) of this force may be calculated from original data, spatial position of the column, or any other desired source. This ability to incorporate additional information provides excellent flexibility in application. For example, in the described application to outer border delineation, edges were less likely to be detected with increasing distance (displacement in \hat{k} -direction), so a spring mesh alone would underestimate the surface depth between detected segments. However, by leveraging this *a priori* knowledge, forces can be applied that “inflate” the mesh outwards in regions in which no anchoring surface segments are detected.

A system of linear equations describing this structure can be solved for the position of each node in its equilibrium state:

$$\sum \mathbf{F} = \mathbf{0}. \quad (12)$$

Because each node is, by definition, constrained within its column,

$$\sum F_i \equiv 0, \quad \sum F_j \equiv 0. \quad (13)$$

Only the balance of forces in the \hat{k} -direction must be resolved. Restorative spring force (F_{spring}) for a Hookean linear-elastic spring model, with spring coefficient (K) defining its stiffness, is linearly proportional to its displacement (X) [221]:

$$F_{spring} = -KX. \quad (14)$$

Defining the initial lengths of springs to be 0, forces exerted by the various springs on each node in the \hat{k} -direction are superimposed:

$$\sum F_k^{ij} = K_1 d_{i-} \sin \theta_{i-} + K_1 d_{i+} \sin \theta_{i+} + K_2 d_{j-} \sin \theta_{j-} + K_2 d_{j+} \sin \theta_{j+} + F_{ij}^* + F_{ij} \quad (15)$$

$$\overline{F_{ij}^*} = \begin{cases} K_3 (p_{ij} - S_{ij}) \hat{k} & \text{if } \exists P_{ij} \\ K_4 (l_{ij} - S_{ij}) \hat{k} & \text{if } \nexists P_{ij} \end{cases} \quad (16)$$

where d and θ are, respectively, the distance and angle between adjacent nodes, i and j of the subscript indicate the direction of adjacency (among or between planes, respectively), and +/- signs

of the subscript indicate direction of adjacency (increasing or decreasing node number), as illustrated in Figure 2.13. The sine of angle θ is inversely proportional to spring displacement d ,

$$\sin \theta_{i\pm} = \frac{S_{(i\pm 1)j} - S_{ij}}{d_{i\pm}}, \quad \sin \theta_{j\pm} = \frac{S_{i(j\pm 1)} - S_{ij}}{d_{j\pm}}, \quad (17)$$

So the above relationship (Equation 15) can be simplified:

$$\begin{aligned} \sum F_k^{ij} = K_1(S_{(i+1)j} - S_{ij}) + K_1(S_{(i-1)j} - S_{ij}) + K_2(S_{i(j+1)} - S_{ij}) + K_2(S_{i(j-1)} - S_{ij}) + \\ F_{ij}^* + F_{ij} = 0. \end{aligned} \quad (18)$$

Rearrangement yields:

$$\begin{aligned} 0 = -\left(2 + 2\frac{K_2}{K_1}\right)S_{ij} + S_{(i+1)j} + S_{(i-1)j} + \frac{K_2}{K_1}S_{i(j+1)} + \frac{K_2}{K_1}S_{i(j-1)} + \frac{F_{ij}}{K_1} + \\ \begin{cases} \frac{K_3}{K_1}(p_{ij} - S_{ij}) & \text{if } \exists P_{ij} \\ \frac{K_4}{K_1}(l_{ij} - S_{ij}) & \text{if } \nexists P_{ij} \end{cases}. \end{aligned} \quad (19)$$

All N nodes are numbered sequentially, such that (i,j) corresponds to n and $(i,j+1)$ corresponds to $(n+p)$, and N is equal to $p \times q$, so the full system of equations can be expressed as $\mathbf{Ax} = \mathbf{b}$, where \mathbf{A} is a sparse $N \times N$ square matrix (Figure 2.14):

$$A_{nn} = -2 - 2\frac{K_2}{K_1} - \begin{cases} \frac{K_3}{K_1} & \text{if } \exists P_{ij} \\ \frac{K_4}{K_1} & \text{if } \nexists P_{ij} \end{cases}; \quad (20)$$

$$A_{n(n\pm 1)} = 1; \quad (21)$$

$$A_{n(n\pm p)} = \frac{K_2}{K_1}. \quad (22)$$

Furthermore, \mathbf{x} is a vector of node positions:

$$\mathbf{x} = \begin{bmatrix} S_1 \\ \vdots \\ S_N \end{bmatrix}. \quad (23)$$

Finally, \mathbf{b} is a vector:

$$\mathbf{b}_n = -\frac{F_n}{K_1} - \begin{cases} \frac{K_3}{K_1}p_{ij} & \text{if } \exists P_{ij} \\ \frac{K_4}{K_1}l_{ij} & \text{if } \nexists P_{ij} \end{cases}. \quad (24)$$




Figure 2.14. Visualization of the sparsity pattern for matrix \mathbf{A} ; (Equations 20-22) that expresses the system of linear equations describing the linear-elastic mesh. (Indices indicate node number, n ; nz: number of non-zero elements)

Multiplying both sides of the equation by the inverse of matrix \mathbf{A} , the position of all N nodes can be solved simultaneously with a single operation of computation complexity $O(N^{1.5})$:

$$\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}. \quad (25)$$

This approach may be used in Cartesian coordinates for a non-folded surface or in cylindrical (polar) coordinates for a closed surface, as in the present application. When working in cylindrical coordinates, edge continuity is achieved by connecting nodes at the start and end of each plane, (I,j) to (p,j) , as shown with the dashed-line springs (Figure 2.13). Other boundary conditions can also be readily enforced.

As noted before, this methodological approach for fitting a continuous surface was applied in OCT to (1) smooth a contour in which tentative locations of the entire surface are known, and (2) construct a surface when only limited fragments of the surface have been tentatively located. The spring mesh surface construction method was applied to automatically and fully delineate the borders of a diseased vessel in OCT images.

To smooth the lumen as a continuous 3D surface, an anisotropic linear-elastic mesh was fit to the initial detected points (identified using the method described in Section 2.2.2). Because

each column contained an anchor point ($m = N$), there was no need for a base surface or Class 4 springs. No force function was applied in this scenario, eliminating a third parameter (F_{ij}). The remaining parameters, K_2/K_1 and K_3/K_1 , were determined by minimizing the sum of radial distance magnitudes between the node equilibrium positions and the validation annotations of the lumen:

$$\min_{\substack{K_2, K_3 \geq 0 \\ K_1, K_1}} \sum_{n=1}^N \left\| \mathcal{S} \left(\frac{K_2}{K_1}, \frac{K_3}{K_1}, \mathbf{P}_n \right)_{\mathbf{P}_n^*} - \mathbf{P}_n^* \right\| \quad (26)$$

where \mathbf{P}^* are the validation annotation positions.

Subsequently, it was necessary to fill gaps between detected segments of the outer border (identified in 3D space using the method described in Section 2.2.2) and ensure continuity from frame to frame. Continuity and smoothness are inherently observed in nature, and cross-sectional area does not vary considerably in adjacent frames. It was therefore necessary to fit a smoothed surface with the available information of tentative partial surface locations. This too was accomplished using the novel anisotropic linear-elastic mesh surface fitting method introduced above (Figures 2.13 and 2.15).

In this application, only columns in which border segments were detected include an anchor point ($m < N$). It was observed that base surface position could be uniformly set to 0 ($l_{ij} = 0$) with minimal impact on performance, decreasing complexity and reducing arbitrary offsets. A unique force was applied to each node of the mesh that incorporated information from the original image information:

$$F_{ij} = \alpha \sum_k I(i, j) \quad (27)$$

where α is a scaling coefficient and I is the pixel intensity of the original image. Sum of pixel intensity over a column was utilized because thick, high-intensity regions often indicate the presence of a thick, diseased intima, which typically corresponds to greater outer border depth. This scheme also prescribes near-zero force in regions obscured by the guidewire shadow, a desirable feature resulting in smooth interpolation-like continuity with adjacent regions (Figure 2.15). While multiple force functions were tested, including those proportional to lumen cross-sectional area or simply a constant, the sum of intensities in each column yielded the best result.




Figure 2.15. Diagram illustrating how the surface fitting method was implemented to delineate the outer border within each frame (*top*) and between frames (*bottom*). The detected edges served as anchor points, while the lumen acted as base layer. Each node was constrained to displacement in the radial (\hat{k}) direction, and was acted upon by an abluminal force proportional to summed column intensity. Each node was connected to the nearest nodes in each direction; nodes on either end of a frame were connected to ensure cylindrical continuity. Node density is substantially reduced for illustrative clarity. Schematic representation and color correspondence as in Figure 2.13.

The non-fixed parameters for this system, K_2/K_1 , K_3/K_1 , K_4/K_1 , and α/K_1 , were determined by minimizing the sum of radial distance magnitudes between the node equilibrium positions and the validation annotations of the outer border:

$$\min_{\substack{K_2, K_3, K_4, \alpha \\ K_1, K_1', K_1', K_1'} \geq 0} \sum_{n=1}^N \left\| \mathbf{S} \left(\frac{K_2}{K_1}, \frac{K_3}{K_1}, \frac{K_4}{K_1}, \frac{\alpha}{K_1}, l, \mathbf{P}_n \right)_{\mathbf{P}_n^*} - \mathbf{P}_n^* \right\| \quad (28)$$

where \mathbf{P}^* are the validation annotation positions.

For performance comparison, standard interpolation and surface-fitting to the detected surface segments was also conducted with standard, existing methods.

After the equilibrium node positions were calculated, the resulting surface was reconstituted to the non-flattened state by shifting each column distally by the radial distance of the lumen in the original image (reversing the original flattening process). The reconstituted surface was finally converted from the cylindrical to Cartesian coordinate system—transforming the open surface into a wrapped cylindrical surface—using the standard transformation:

$$x = r \cos \theta = k \cos \left(2\pi \frac{i-1}{p-1} \right); \quad (29)$$

$$y = r \sin \theta = k \sin \left(2\pi \frac{i-1}{p-1} \right); \quad (30)$$

$$z = z = j. \quad (31)$$

As the method was 3D and global, run time was positively, but not proportionally, related to the number of frames in each analyzed pullback. Execution times (Table 2.3) were dominated by initial lumen and visible outer border segment detection, while the lumen smoothing and surface fitting steps were exceptionally fast (2.74 ± 0.28 ms and 40.2 ± 7.5

ms per frame, respectively). Conversion from polar to Cartesian coordinates was also a significant contributor to the run time (contribution not shown). The one-time optimization of the inner and outer border surface fit parameters took 7.8 and 10.1 minutes, respectively, to execute.

Table 2.3. Run times for major steps of the novel 3D surface fitting method

Step	Time per Frame (s) (Mean \pm SD)
Lumen Detection	4.2 \pm 1.5
Lumen Smoothing	0.0027 \pm 0.0003
Image Flattening	0.0064 \pm 0.0013
Edge Detection	5.4 \pm 0.9
Surface Fitting – Set Up	0.0363 \pm 0.0070
Surface Fitting – Solve	0.0039 \pm 0.0013

Executed on machine running Intel Xeon processors with 12 cores (2.80 GHz) and 12 GB RAM

The autonomous border detection method was validated through direct comparison to the manual annotations made by expert interventional cardiologists in the same clinical dataset (Table 2.4). Patients with clinical indication for invasive coronary angiography in whom one or more coronary stenoses with at least intermediate angiographic severity (>30% diameter stenosis by quantitative coronary angiography) were documented could be prospectively included in the study. Exclusion criteria was the presence of tight lesions (>75% diameter stenosis) in which the crossing catheters could impede flow through the lesion and subsequently hamper blood clearance by contrast injection, causing very poor image quality. The study was approved by the Ethics Committee of the institution and all patients gave informed consent.

Table 2.4. Patient characteristics (N = 7)

Characteristic	Mean ± SD or N (%)
Age (years)	61.8±10.8
Male	6 (85.7)
Coronary risk factors	
Diabetes mellitus	3 (43)
Dyslipidemia	4 (57)
Current smoker	2 (29)
<i>Angiographic stenosis characteristics</i>	
Vessels	
LAD	5 (71)
RCA	2 (29)
Diameter stenosis (%)	48.0±7.4
Lesion length (mm)	7.0±2.8

LAD: left anterior descending coronary artery;
RCA: right coronary artery

Angiography was performed following a standard protocol, using radial or femoral access. After the diagnostic angiography, 200 µg of intracoronary nitroglycerin and 5,000 IU of unfractionated heparin were infused. Investigators used the C7-XR FD-OCT optical frequency domain intravascular imaging system and the DragonFly catheter (St. Jude Medical, Lightlab Imaging Inc., Westford, MA, USA). This imaging system offers a maximum frame rate of 100 frames per second, 500 lines per frame, a scan diameter of 10 mm, and axial resolution of 15 µm. The monorail imaging catheter was advanced distal to the target lesion. The automated pullback was performed at 20 mm/s (system nominal pullback speed) while the blood was removed by a short injection of iso-osmolar contrast at 37°C through the guiding catheter. The optimal volume/time intracoronary infusion of contrast was tested to achieve a complete blood clearance in the vessel lumen.

Per protocol, IVUS imaging of the coronary segments imaged with OCT was also conducted. The IVUS catheter, a Revolution 45 MHz rotational IVUS imaging catheter (Volcano Inc., Rancho Cordova, CA, USA) was placed at the same distal position as the OCT catheter, using as reference a fiduciary anatomical point such as a bifurcation or a calcification spot. An

angiographic acquisition was obtained with the imaging catheter in the distal position to confirm matching points for the IVUS and OCT starting runs. An automated motorized pullback at 0.5 mm/s was performed.

Images were digitally stored for off-line analysis, and all imaging data sets were completely anonymized and transferred to the core lab for analysis. For the purposes of this study, OCT datasets in which blood removal was found *post hoc* to be grossly insufficient were not utilized. Delimiting segments of useable pullbacks captured prior to or following injection of contrast were similarly excluded. To ensure the same vessel segments were analyzed with both techniques, matching of the OCT and IVUS pullbacks was performed using landmarks such as side branches and other anatomical features, thereby providing frame-by-frame correspondence. Experts were thereafter blinded to the correspondence of imaging files.

Annotations were compared in 724 frames from 7 unique pullbacks (series of images) acquired from the same number of unique patients (Table 2.4). An average border annotation of two experts was determined by taking the mean radial distance of both annotations from the center of the image around the perimeter in each frame (as in [173]). In affirmation of multiple previous studies showing excellent inter-observer variability in OCT, variability between the two expert annotations in this study was low (Figure 2.20): the average absolute difference between the areas of their annotations was just $0.13 \pm 0.11 \text{ mm}^2$ ($2.5 \pm 2.3\%$) and $0.88 \pm 0.75 \text{ mm}^2$ ($9.7 \pm 8.6\%$) for the inner and outer border tracings, respectively, and coefficients of determination between the two were greater than 0.99 and 0.93, respectively. Additionally, the average absolute radial distance between their inner and outer border annotations was just $0.07 \pm 0.07 \text{ mm}$ and $0.21 \pm 0.23 \text{ mm}$, respectively. The distribution of radial distance between inner border tracings was centered around 0.02 mm with a SD of 0.09 mm, while the distance between outer border tracings was centered nearly symmetrically at 0.10 mm with a SD is 0.30 mm (Figure 2.23), further confirming the observation of strong agreement.

While the experts were able to identify sufficient segments of the outer border to confidently estimate the full border in most frames, the outer wall was often poorly visualized by OCT as a consequence of the patient inclusion criteria that sought high plaque burden vessels. The advanced 3D edge detection process (Section 2.2.2) identified segments of the outer border in only 57% of radial columns (A-lines). (These sparse detected edges were the anchor points, P_n , for the

subsequent surface fit.) This challenge of visibility substantiated the use of IVUS imaging as a secondary reference of vessel geometry in validating the discussed approach; due to the limited penetration depth of OCT, IVUS is the clinical gold standard for measurements involving the outer border, as the border is almost universally visible, regardless of plaque burden [39]. However, evaluation criteria for validating between modalities is limited—only differences between areas enclosed by the algorithm results and expert tracings could be compared; direct overlap and tracing comparisons could only be performed for annotations made within the same modality.

The borders delineated with the presented methodology were directly evaluated against the average OCT border annotations made by the experts. The automated delineations were overlaid on the human annotations (considered ground truth), and the following parameters were calculated:

$$\text{Sensitivity} = \frac{TP}{TP+FN} \times 100 = R_{\text{over}} \times 100; \quad (32)$$

$$\text{Specificity} = \frac{TN}{FP+TN} \times 100; \quad (33)$$

$$\text{Jaccard Index} = \frac{TP}{TP+FP+FN} \times 100; \quad (34)$$

$$\text{Dice Index} = \frac{2TP}{2TP+FP+FN} \times 100; \quad (35)$$

$$R_{\text{nonover}} = \frac{FN+FP}{TP+FN} \quad (36)$$

where TP are true positives, TN are true negatives, FP are false positives, and FN are false negatives. Here, TP are considered those pixels enclosed by both the algorithm results and expert annotations, TN are those pixels of the full image enclosed by neither annotation, FP are those pixels enclosed by only the algorithm result, and FN are those pixels enclosed by only the expert annotation (Figure 2.16). The ratios of overlapping (R_{over}) and non-overlapping (R_{nonover}) areas between the algorithm results and expert annotations (Equations 32 and 36, respectively) were defined similar to Athanasiou *et al.* [173] to represent sensitivity and specificity of automated image segmentation in a way that is not skewed by the image size, particularly when the region of interest is only a subset of the actual image (as in OCT). The distribution of the absolute radial distance between the expert annotations and algorithm results was also determined, the coefficient of determination (R^2) was calculated for the resulting cross-sectional areas, and Bland-Altman analysis of the delineated CSAs was performed.

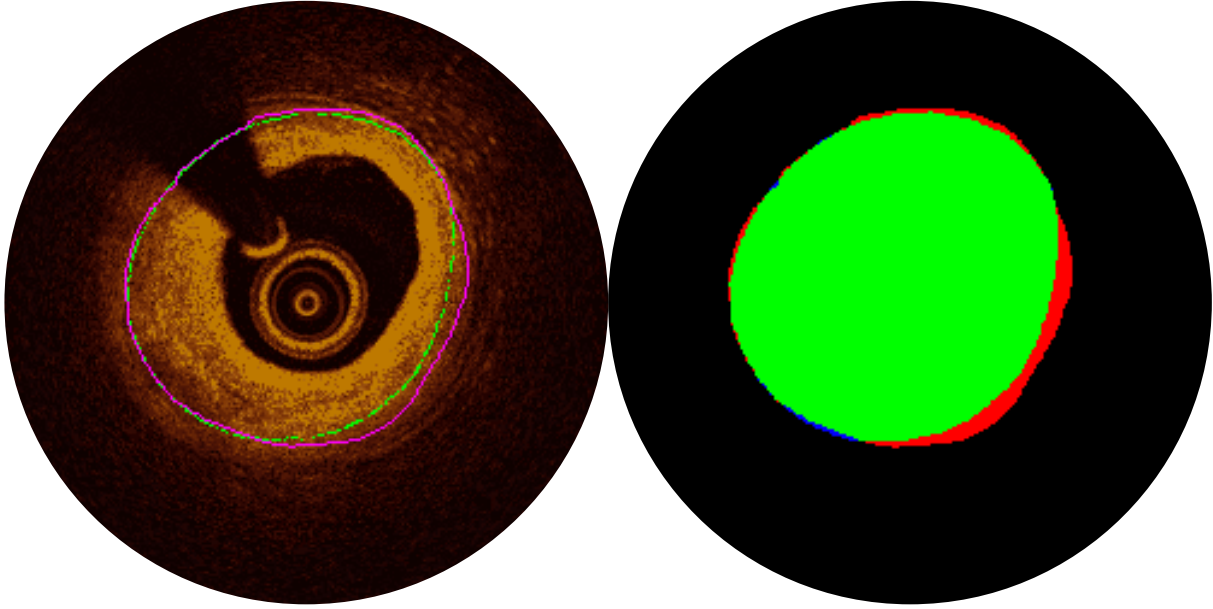


Figure 2.16. Comparison of algorithm result and human expert annotation in a representative OCT frame. (*Left*) Outer border of vessel annotated by both human experts (green) and the presented algorithm (magenta). (*Right*) Corresponding labeled image showing true positive (green), false positive (red), false negative (blue), and true negative (black) pixels used in calculating validation metrics.

The areas delineated by the inner (lumen) border results were first assessed. Here, the initial position of the lumen was determined for each column, and the surface-fitting approach presented above was used just for its smoothing capability. Figure 2.17 shows the error and Bland-Altman plots comparing the algorithm CSA to the average expert annotations. Direct comparison of overlapping regions delineated by the human expert annotations and algorithm results in OCT were made as outlined above. The original lumen trace differed from that of the human expert by 0.12 ± 0.43 mm (0.01 ± 0.44 mm overall, i.e. not absolute distance); the smoothed lumen trace differed by 0.11 ± 0.35 mm (0.00 ± 0.36 mm overall). Sensitivity, specificity, Jaccard index, Dice index, R_{over} , and R_{nonover} parameters are provided for the lumen before and after the smoothing surface-fit operation (Table 2.5). The mean and median percent change in frame-to-frame lumen area of the final smoothed lumen contours (7.5% and 4.4%, respectively) were between that of the human experts (6.0% and 3.8%, respectively) and that of the original lumen contours (9.3% and 5.0%, respectively).

Table 2.5. Performance metrics comparing regions delineated by inner (lumen) border as originally detected and after smoothing surface fit

<i>Metric</i>	<i>Original</i>	<i>Smoothed (Surface Fit)</i>
<i>Sensitivity</i>	98.46	97.54
<i>Specificity</i>	99.47	99.59
<i>Jaccard Index</i>	91.51	92.17
<i>Dice Index</i>	95.57	95.92
<i>R_{over}</i>	0.98946	0.9754
<i>R_{nonover}</i>	0.0914	0.0830




Figure 2.17. Comparison of human expert annotation of the lumen and algorithm detection before and after smoothing. Error plot showing strong correlation between the two areas ($R^2 = 0.94$ before smoothing, 0.96 after smoothing; dashed line shows 1-to-1 ratio); Bland-Altman plot (average: -0.36 mm^2 before smoothing, -0.20 mm^2 after smoothing; standard deviation: 0.82 mm^2 before smoothing, 0.65 mm^2 after smoothing) illustrates no substantial systematic error. Smoothing eliminated serious outliers and improved overall performance.

The full area delineated by the outer border as determined by experts and the algorithm (Figure 2.18) was also assessed. The outer border was determined with the surface-fitting method presented here using only partial detected contour segments (Figure 2.15). The average absolute (net) CSA difference between the algorithm results and human annotations in each frame was 0.93 mm^2 (10.9%; overall mean of 0.07 mm^2) with a SD of 0.84 mm^2 (11.0%; 1.25 mm^2 for overall difference); the algorithm provided a very slight overestimation of the cross sectional area. Figure 2.19 shows the error and Bland-Altman plots comparing the areas determined by the algorithm to those delineated by the average expert annotations. For context, Figure 2.20 shows the distributions comparing the areas delineated by the two average expert annotations, which demonstrates similar deviations. Direct comparison of overlapping regions enclosed by the annotated outer border is quantified in Table 2.6. The mean difference between the radial depths of the algorithm results and expert annotations of the outer border was $0.33 \pm 0.43 \text{ mm}$. The distribution of differences is centered around 0.09 mm , and has a SD of 0.53 mm (Figure 2.23); it largely reflects the distribution of differences between the annotations of the two experts, shown in the same figure.

In addition to direct comparison to expert annotations in OCT, the autonomous border detection method was further validated through indirect comparison to manual annotations made by an expert interventional cardiologist in the matched IVUS images. Agreement between cross-sectional areas delineated in corresponding OCT and IVUS frames show strong but imperfect agreement, as expected. Error and Bland-Altman plots compare areas delineated by an expert in IVUS to those delineated by both experts and the presented algorithm in OCT (Figures 2.21 and 2.22). The average absolute difference between the inner and outer CSAs delineated by the algorithm in OCT and the expert in IVUS was $1.41 \pm 1.17 \text{ mm}^2$ ($25.0 \pm 20.1\%$) and $1.72 \pm 1.43 \text{ mm}^2$ ($19.2 \pm 15.0\%$), respectively (Figure 2.21). For comparison, the average absolute difference between the areas of human expert annotations for the inner and outer border tracings was $1.48 \pm 1.26 \text{ mm}^2$ ($26.9 \pm 21.3\%$) and $1.49 \pm 1.39 \text{ mm}^2$ ($16.4 \pm 13.4\%$), respectively (Figure 2.22).




Figure 2.18. Several examples of fully delineated frames showing expert annotations (green) and algorithm results (magenta). Major swaths of the outer border are not visible in several frames. In frames with underestimated wall thickness, plaque features with media-like appearance were commonly identified falsely as visible IEM.




Figure 2.19. Comparison of algorithm results and human expert annotations of the outer (IEM) border in OCT. Error plot showing strong correlation between the two areas ($R^2 = 0.89$; dashed line shows 1-to-1 ratio); Bland-Altman plot (average: 0.07 mm^2 , standard deviation: 1.25 mm^2) illustrates no clear systematic error. Results demonstrated that the 3D surface fitting approach rivals performance of human experts in this task.




Figure 2.20. Comparison of two expert annotations of the inner (lumen) and outer (IEM) borders in OCT. Error plot showing exceptional correlation between the two areas ($R^2_{inner} > 0.99$, $R^2_{outer} = 0.93$); Bland-Altman plot (average: 0.10 & 0.53 mm², standard deviation: 0.14 & 1.02 mm², respectively) illustrates no substantial systematic error. Human-human variability in outer border appears to largely reflect human-algorithm deviation.




Figure 2.21. Comparison of the inner and outer borders delineated in OCT and IVUS by computer and expert, respectively. Error plot shows correlation between the two delineated areas ($R^2_{inner} = 0.68$, $R^2_{outer} = 0.67$; dashed line shows 1-to-1 ratio). Bland-Altman plot (average: 0.25 & 0.72 mm², standard deviation: 1.82 & 2.12 mm², respectively) illustrates minor systematic offset and greater difference (overestimation) in larger vessels.




Figure 2.22. Comparison of the inner and outer borders delineated in OCT and IVUS by human experts. Error plot shows correlation between the two delineated areas ($R^2_{inner} = 0.66$, $R^2_{outer} = 0.75$). Bland-Altman plot (average: 0.05 & 0.78 mm², standard deviation: 1.94 & 1.88 mm², respectively) illustrates minor systematic overestimation (greater) in larger vessels. Intermodal human variability largely reflects algorithm deviation.

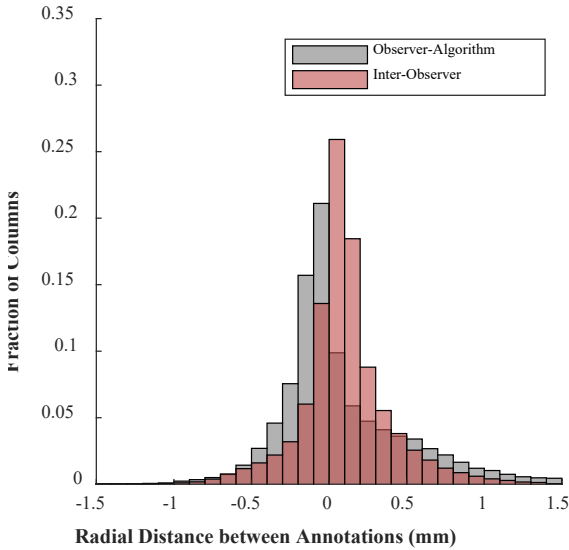


Figure 2.23. Comparison of outer border tracings through measurement of difference in radial depth. Mean difference between the tracing radial depth of the average expert and algorithm was 0.09 ± 0.53 mm (0.33 ± 0.43 mm absolute distance). Mean difference between the tracing radial depth of the two experts was 0.10 ± 0.30 mm (0.21 ± 0.23 mm absolute distance).

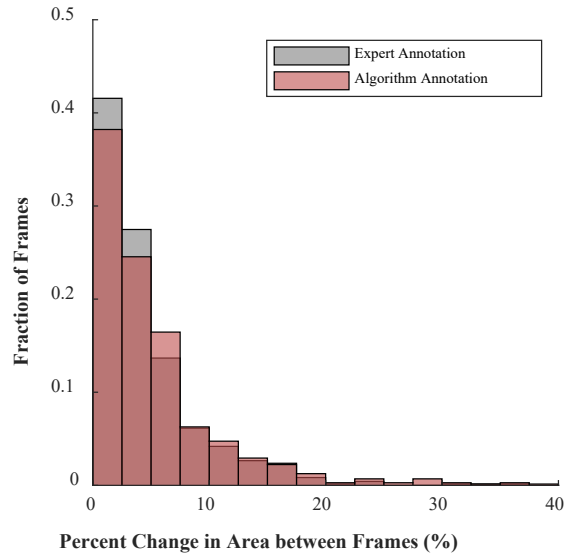


Figure 2.24. Histogram showing frame-to-frame area change for both the expert annotations and algorithm delineations of the vessel outer border. The average area change from frame to frame for experts was 4.5%, while that for the algorithm was 5.5%.

The continuous 3D delineated outer border was relatively smooth. Large frame-to-frame area changes were uncommon, with frequency of occurrence decreasing with magnitude of change for both the expert annotations and algorithm delineations of the vessel outer border (Figure 2.24). The average area change from frame to frame for experts was 4.5%, while that for the algorithm was 5.5%; physiological area change over the corresponding distance should theoretically be around 2% or less [222]–[224]. The median percent change in frame-to-frame vessel area of the automated outer border annotations was also very similar to that of the human experts (3.5% vs. 3.0%).

To demonstrate the value and impressive performance of the new method described here, various alternative interpolation and surface fitting techniques were employed on the same set of detected outer border segments (and with the same automatically-delineated lumen contours). Performance of the existing interpolation and surface fitting methods was clearly inferior to the presented method (Table 2.6 and Figure 2.25), and—as expected—produced annotations which systematically underestimated the vessel area (Figure 2.26).

Table 2.6. Performance metrics comparing vessel wall regions delineated by various interpolation & surface fitting techniques

Metric	Anisotropic Linear- Elastic Mesh	Interpolation Methods			Surface Fitting Methods		
		Linear Interpolation	Nearest Neighbor Interpolation	Cubic Spline Interpolation	Polynomial Surface Model*	Local Linear Regression	Local Quadratic Regression
Sensitivity	90.79	81.89	82.35	81.94	80.74	81.11	80.96
Specificity	99.00	99.19	99.10	99.18	99.42	99.47	99.46
Jaccard Index	83.67	76.61	76.48	76.56	76.90	77.62	77.39
Dice Index	91.11	86.75	86.67	86.73	86.94	87.40	87.25
R_{over}	0.9079	0.8189	0.8235	0.8194	0.8074	0.8111	0.8096
$R_{nonover}$	0.1773	0.2501	0.2533	0.2508	0.2426	0.2339	0.2366

*5th Degree: $z = p00 + p10 \times x + p01 \times y + \dots + p14 \times xy^4 + p05 \times y^5$

Additional studies were undertaken to characterize the method. Optimization was tested for stability with respect to the initial value to ensure robustness and identification of global minimum. For each test, all parameters were identically set to the same initial value prior to object function minimization. Optimization of the surface fit system for the outer border was stable as the initial values were varied over 7 orders of magnitude (Figure 2.27). The most optimal output (resulting in the lowest object function value over the 20 orders of magnitude tested) resulted from an initial value within this stable range.

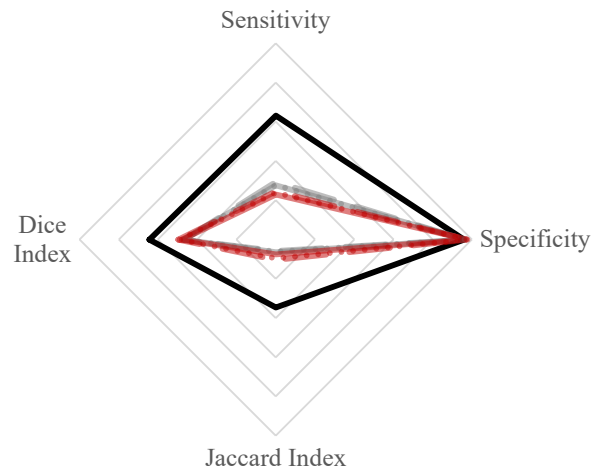


Figure 2.25. Performance metrics achieved by various interpolation and surface fitting techniques demonstrate clear superiority of the novel anisotropic linear-elastic mesh approach (black) over the interpolation (gray) and other surface fitting (red) methods. Axis ranges from 75% to 100%.




Figure 2.26. Comparison of areas resulting from existing interpolation and surface fit methods and human expert annotations of the outer border. Error plot showing correlation between the two areas (R^2 ranges from 0.79 to 0.82; dashed line shows 1-to-1 ratio); Bland-Altman plot illustrates substantial systematic error (error not regularly distributed), with existing interpolation and fit methods systemically underestimating area. Note that these methods less frequently overestimated area in larger vessels (often associated with lower visibility).

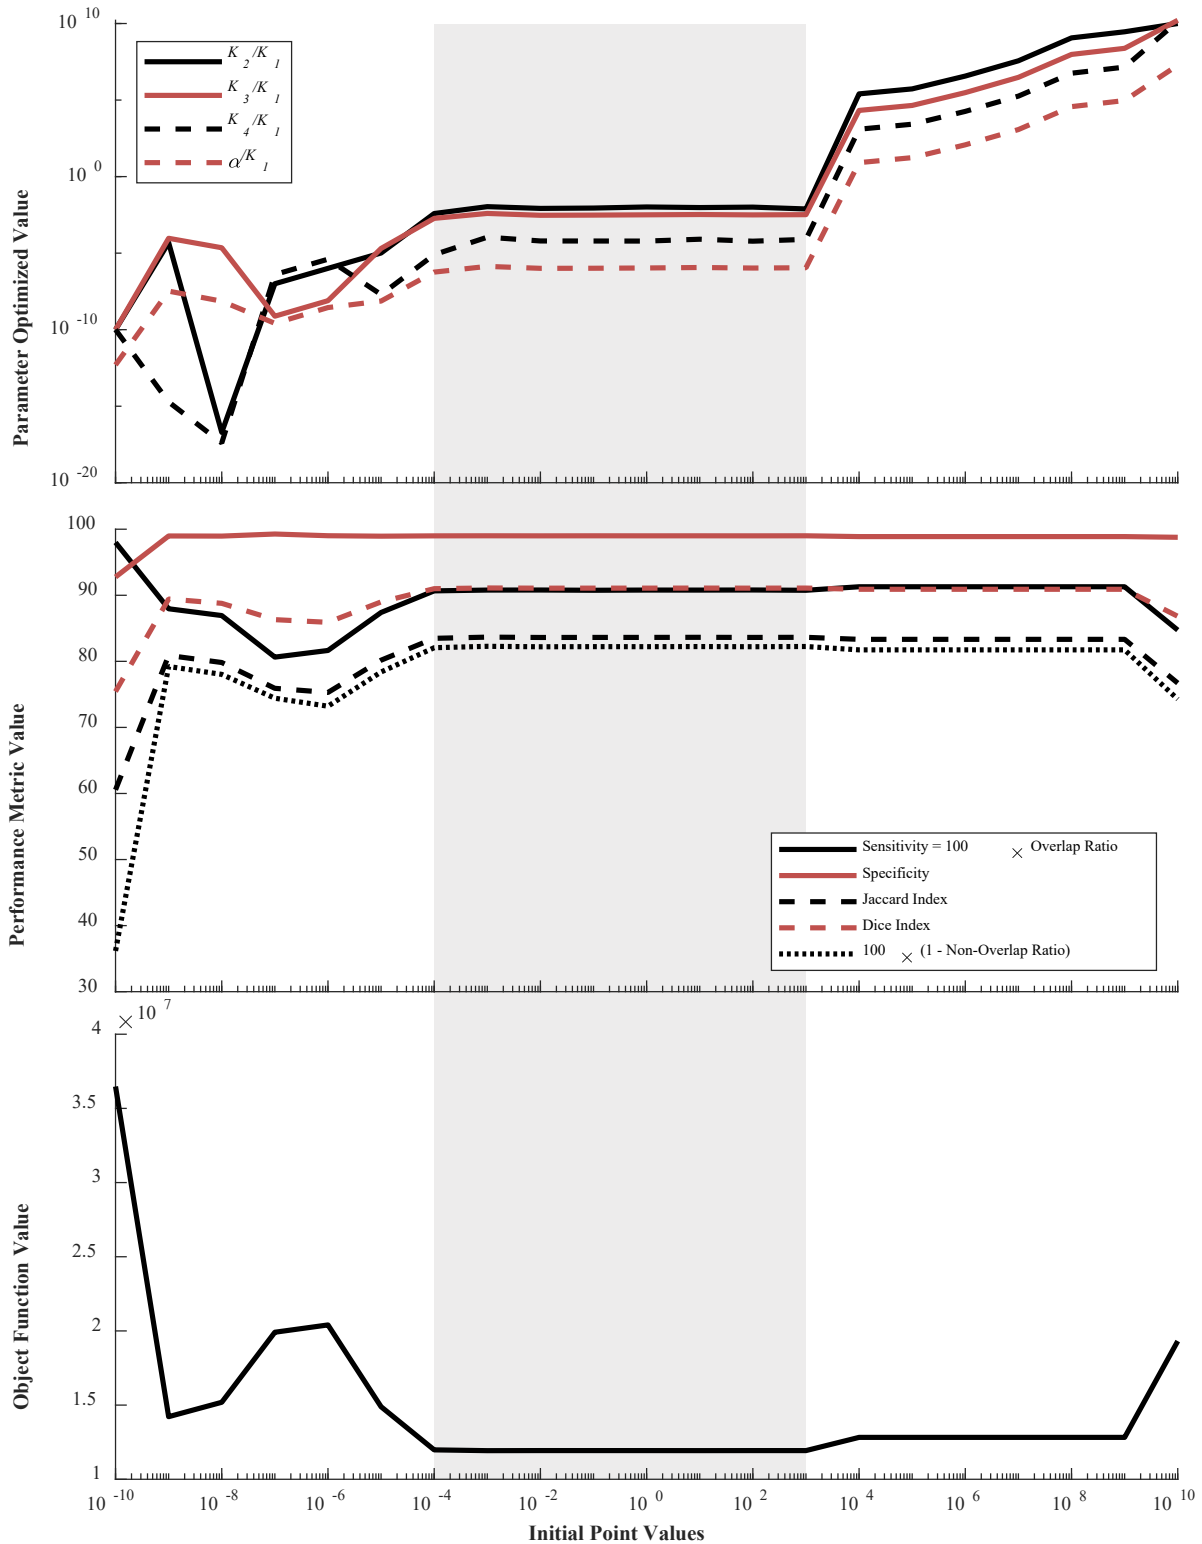


Figure 2.27. Test of (outer border) optimization stability as the initial value of the mechanical element parameters was varied demonstrated robustness over a wide range. Outputs, including parameter output values (*top*) were stable over a range of values from 10^{-4} to 10^3 (gray), coinciding with optimal performance metrics (*middle*) and globally minimal object function values (*bottom*).

To determine the influence of the data used in the optimization on the outcome (optimal parameter values), the optimization process was also tested for convergence with respect to the number of frames used. Ten sets of randomly-selected frames were tested for each number of frames evaluated. As seen in the performance curves (Figure 2.28), all performance metrics improved with increasing numbers of frames included in the optimization, stabilizing around 30 frames; performance improved only slightly when the number of frames used for optimization was increased from 30 (around 4% of available frames) to the full 724. Output parameter values largely converged in the same manner as the number of frames increased (Figure 2.28). However, not only did average values converge as frame numbers increased, but error and variability also reduced profoundly (Figure 2.28). Also noteworthy was the absence of variability when optimization was performed multiple (10) times with all frames; when provided with the same inputs, the optimization algorithm consistently arrived at the same results, demonstrating stability of the underlying optimization framework. Similar testing was performed for the less complex lumen smoothing operation optimization, which demonstrated even greater stability (results not shown).

Having shown method stability and convergence, all frames were used in the optimization with initial values for all parameters set to 10^{-3} , which produced optimal results during the initial condition sensitivity testing. Final optimal parameters for the inner (lumen) border surface fit were: $K_2/K_I = 6.74 \times 10^{-3}$; and $K_3/K_I = 2.02 \times 10^{-2}$. For the outer border surface fit, the final optimal parameters were: $K_2/K_I = 1.09 \times 10^{-2}$; $K_3/K_I = 3.92 \times 10^{-3}$; $K_4/K_I = 1.10 \times 10^{-4}$; and $\alpha/K_I = 1.36 \times 10^{-6}$. These are the values for which the final performance was reported above.

The surface fitting approach inspired by Newtonian physics was highly effective in fitting variably-sparse and questionably-reliable data. Intuitively applying an anisotropic linear elastic mesh with a force function allowed autonomous delineation of the outer border of a diseased vessel in OCT images. However, numerous opportunities exist to enhance and improve the performance of the fitting technique in this application.




Figure 2.28. Test of (outer border) optimization convergence as the number of frames used in the optimization was increased. Parameter output values (*top*), performance metrics (*middle*), and mean radial distance between algorithm and expert annotations (alternative performance metric; *bottom*) converged with improving performance as more frames were used in the optimization process, though minimal change was observed when more than 4% of available frames were used. Note that negative error bars in the top graph are amplified by the logarithmic scale.

Accuracy of this 3D surface fitting approach was primarily limited by two factors: inner and candidate outer border segment detection. Lumen detection was important to the outer border detection process, as it was used to flatten the image prior to detection of visible segments of the outer border—errors in lumen detection thus skewed or distorted the flattened image. However, because the outer surface fitting was performed in the flattened state (to avoid arbitrary penalization of contours arising due to catheter position within the lumen), any error in detecting the lumen was directly introduced to the outer border when the fit surface was reconstituted to the non-flattened state. Therefore, improved techniques to detect the lumen, particularly those robust to artifact presence [179], [225], would have a direct positive impact on the performance of this technique. Detection of visible segments of the outer border was another process step primed for improvement. When outer border segments were falsely identified within the diseased intima (which occasionally occurred in heterogeneous and/or calcified plaques), the existence of the erroneous outer border segments skewed the surface in a way that led to underestimation and misrepresentation of some regions with high plaque burden. However, the fitting approach tempered these negative effects, especially when the false detected segments were small relative to the correctly-identified segments (as they frequently were). Nevertheless, improved approaches would be useful. One possible alternative detection method of note has been reported for segmentation of healthy wall regions [182], and other methods similarly derived from earlier segmentation techniques developed for retinal OCT imaging [191] are sure to follow. A promising feature of this surface fitting method was its compatibility with alternative visible border segment algorithms which may improve upon those developed and utilized here (Section 2.2)—it was agnostic to the source of the points to which it was fit.

An additional opportunity for improving performance was implementation of a more sophisticated force function (F_{ij} ; Equation 27) and/or base layer (L_{ij}). Such functions could leverage additional *a priori* knowledge and insights gained through previous studies of typical plaque and vessel wall morphology (akin to the approach taken by Kok *et al.* to develop a plaque-specific generalized estimation equation model of necrotic core thickness [226]), perhaps incorporating plaque tissue characterization studies [160], [173].

There were certain limits to the applicability of the 3D surface fitting technique, as well as the results of the reported validation studies. The surface fitting technique did not allow for folds

in the surface, which would require multiple nodes in a single column. Similarly, the surface must intersect each column, such that intentional gaps could not be achieved. Modification of the presented approach using the same underlying physical concepts could allow for application in either scenario, however. Another limitation pertained to parameter selection and generalizability. While the dataset included 7 distinct image datasets (pullbacks), convergence testing using information from increasing numbers of randomly-selected image frames (Figure 2.28) showed that optimal parameters—and subsequent overall performance—did vary significantly depending on the number and subset of frames used. Yet, the strong convergent trend, which is generally stabilized when at least 30 (of the 724) frames were used, suggested that the final parameters should be close to optimal for all comparable OCT datasets for this application. Parameters for other applications would surely differ, and would likely require training datasets of their own to determine. Furthermore, optimization for other criteria may have been preferable for this and other applications. For example, rather than minimizing summed radial distance between the surface and the expert annotation of the training set, it could be argued that smoothness or an overlap comparison metric (e.g. sensitivity, specificity, R_{over} , R_{nonover} , etc.) were more important and relevant metrics. Such changes would result in different minimization functions (Equations 26 and 28), and consequently different fit parameters (K_2/K_1 , K_3/K_1 , K_4/K_1 , and α/K_1).

In an exciting advancement, a smooth surface fitting method implementing an anisotropic linear-elastic mesh and force function overcame challenges and shortcomings of existing interpolation and surface fitting techniques [157]. Use of equilibrated mechanical lumped elements helped reconstruct and smooth surfaces in the face of partial, imperfect, or incomplete information. With this efficient, flexible, and intuitive approach, I successfully demonstrated, for the first time, the capability to autonomously delineate the entire outer border of diseased vessels in OCT images. The method is even more intriguing given that extension of this work could assign values and attributes to physical parameters that are extracted from the architecture and morphology of a blood vessel wall; fit parameters derived from the data might well offer deeper meaning and allow inferences about vessel state and function. The approach presented a radically new method for incorporating complex data to achieve smooth surface fitting to incomplete surface segments, like those encountered in imaging diseased coronary arteries. However, it only captured a single surface at a time, even though vessels are multi-layered structures whose geometries are consequently best described by multiple related surfaces.

2.3.4. Simultaneous Three-Dimensional Multiple Surface Fitting

Having applied a mechanical approach to sequentially delineate the luminal and intima-media surfaces (Section 2.3.3) [157], it became clear that such a method could be extended further. Whereas only a single surface was fit at a time, arteries are multi-layered structures. I therefore built upon the previously-introduced anisotropic linear-elastic mesh concept to simultaneously fit multiple surfaces and delineate both the intima-media and media-adventitia boundaries at the same time, demonstrating the flexibility and power of this mechanical approach while extending its utility and capacity [158].

As for the novel 3D method for fitting a single surface (Figure 2.12), the simultaneous fitting of coupled surfaces was employed in concert with advanced 3D edge detection (Section 2.2.2) and utilized the following steps:

- (1) the inner (lumen) border of the vessel was detected in each frame and smoothed as a 3D surface;
- (2) the image was flattened relative to the lumen (Figure 2.2) and visible leading and trailing edges of the media (IEM and EEM, respectively) were identified using image contrast enhancement, 3D filtering, application-specific edge detection kernels, and region-growing (Section 2.2.2; Figure 2.3);
- (3) two coupled, interconnected anisotropic linear-elastic meshes were simultaneously fit to the identified edges with applied forces to generate a continuous, smooth surface along the intima-media and media-adventitia borders; and
- (4) the results were finally reconstituted to the pre-flattened frame of reference and converted to a Cartesian coordinate system.

The system of equations required to solve for the position of the two surfaces was derived similarly to that for the single surface (Section 2.3.3) [157], but with some key complicating differences. Both surfaces were independently “anchored” to their own respective candidate segments, while also being coupled by springs connecting the nodes of each surface in a given column (Figure 2.29).

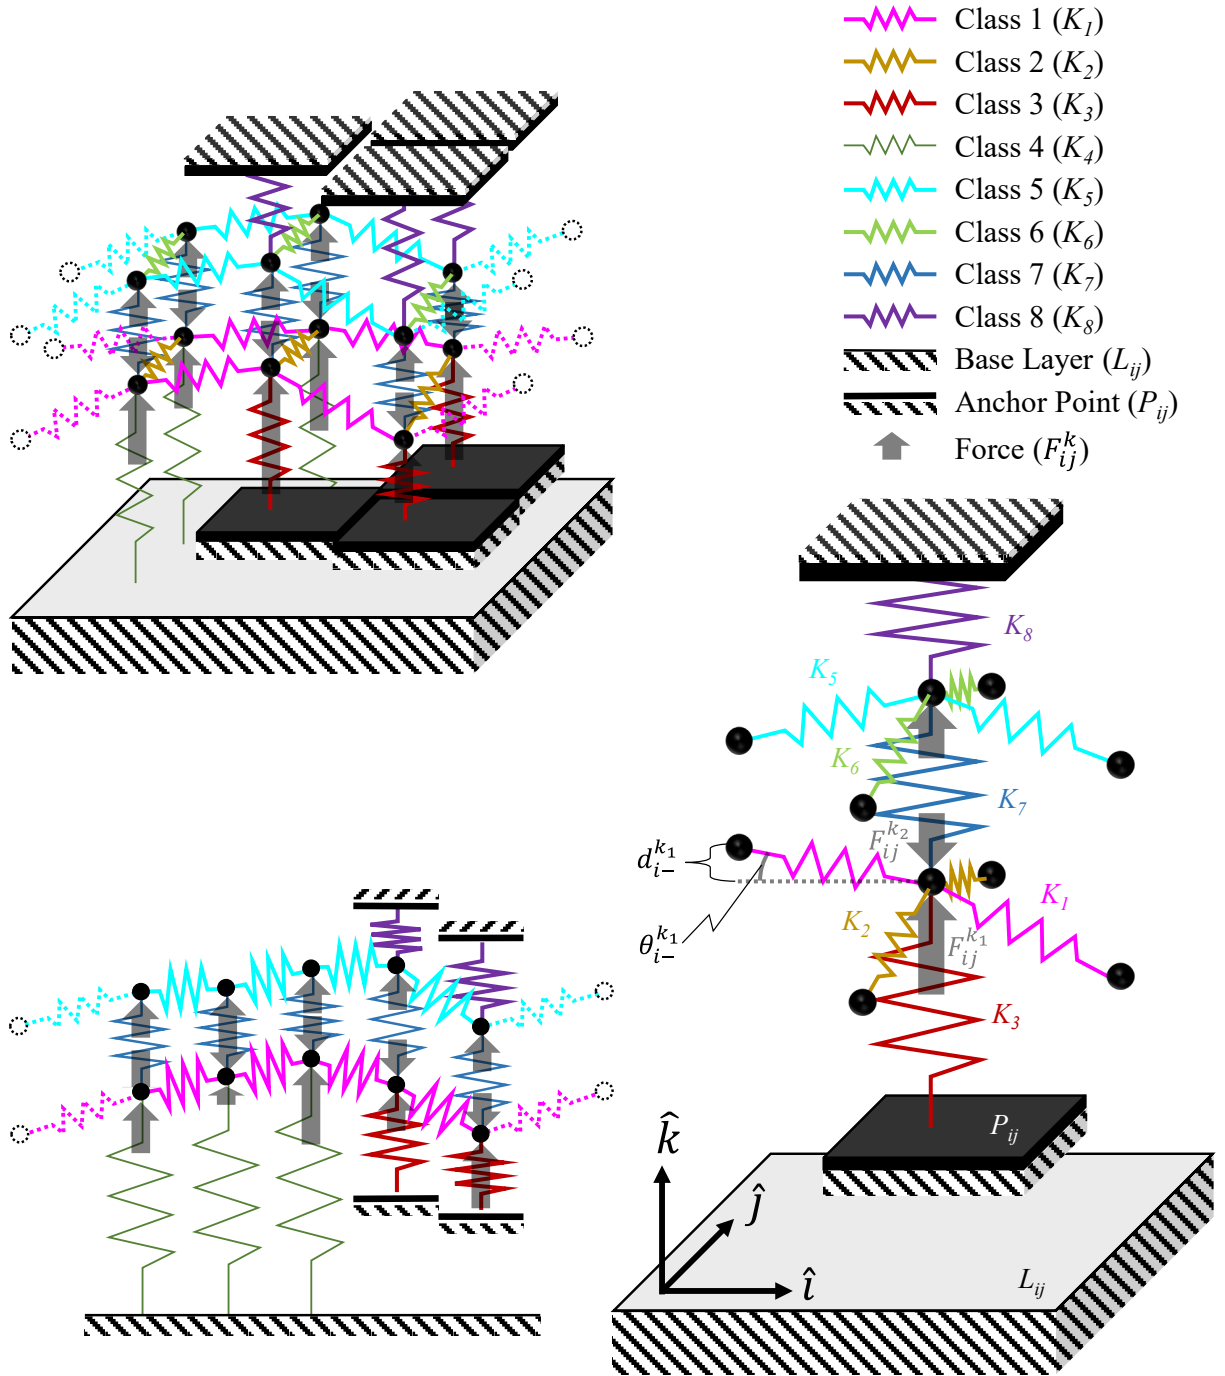


Figure 2.29. Schematic diagram showing the coupled anisotropic linear-elastic mesh system structure. Spring colors indicate class with corresponding stiffness K . Black blocks indicate anchor points (P_{ij}) coinciding with known data points intersecting the column. A light-gray base layer is visible along the bottom (shown here as constant height, $l_{ij}=l$). Semi-transparent gray arrows indicate forces applied to each node (F_{ij}^k). Adjacent node distance (d) and direction (θ), used in calculating force balance, is shown for a single node. Cf. Figure 2.13, which shows the schematic for a mesh system fitting a single surface.

For such a structure, a system of linear equations can be derived that unambiguously solves for the position of each node in the $p \times q$ solution space projection. As for the single surface (Section 2.3.3), this system is solved for its equilibrium state ($\Sigma \mathbf{F} = 0$), and because each node is constrained within its column (i, j), only the balance of forces in the \hat{k} -direction must be computed. Again, defining the springs to be linear-elastic [221] with spring coefficient K and original length 0, forces exerted in the \hat{k} -direction by the various springs acting on nodes of the first surface are combined:

$$\begin{aligned} \Sigma F_{ij}^{k_1} = & K_1 d_{i-}^{k_1} \sin \theta_{i-}^{k_1} + K_1 d_{i+}^{k_1} \sin \theta_{i+}^{k_1} + K_2 d_{j-}^{k_1} \sin \theta_{j-}^{k_1} + K_2 d_{j+}^{k_1} \sin \theta_{j+}^{k_1} + \\ & F_{ij}^{*k_1} + F_{ij}^{k_1} - F_{ij}^{k_2} + K_7 d_k, \end{aligned} \quad (37)$$

$$\overline{F_{ij}^{*k_1}} = \begin{cases} K_3 (p_{ij}^{k_1} - S_{ij}^{k_1}) \hat{k} & \text{if } \exists P_{ij} \\ K_4 (-S_{ij}^{k_1}) \hat{k} & \text{if } \nexists P_{ij} \end{cases}, \quad (38)$$

where d is the distance between adjacent nodes, θ is the angle between adjacent nodes, $i\pm, j\pm$, and k subscripts indicate direction of adjacency, and superscript $k_{[1,2]}$ indicates the surface with which a node, anchor, or force is associated (Figure 2.29). The height of the anchor point (visible, detected surface segment) in column (i, j) is p_{ij}^k , and S_{ij}^k is the position of the surface node in the column. Forces acting on nodes of the second surface are similarly expressed:

$$\begin{aligned} \Sigma F_{ij}^{k_2} = & K_5 d_{i-}^{k_2} \sin \theta_{i-} + K_5 d_{i+}^{k_2} \sin \theta_{i+}^{k_2} + K_6 d_{j-}^{k_2} \sin \theta_{j-}^{k_2} + K_6 d_{j+}^{k_2} \sin \theta_{j+}^{k_2} + \\ & F_{ij}^{*k_2} + F_{ij}^{k_2} - K_7 d_k, \end{aligned} \quad (39)$$

$$\overline{F_{ij}^{*k_2}} = \begin{cases} K_8 (p_{ij}^{k_2} - S_{ij}^{k_2}) \hat{k} & \text{if } \exists P_{ij} \\ 0 & \text{if } \nexists P_{ij} \end{cases}. \quad (40)$$

The sine of angle θ is inversely proportional to spring displacement. As such, substitution and application of the equilibrium conditions allows for simplification and elimination of the θ and d variables. Further rearrangement and grouping of the balance of forces yields:

$$\begin{aligned} 0 = & - \left(2 \frac{K_1}{K_7} + 2 \frac{K_2}{K_7} + 1 \right) S_{ijk_1} + \frac{K_1}{K_7} S_{(i+1)j}^{k_1} + \frac{K_1}{K_7} S_{(i-1)j}^{k_1} + \frac{K_2}{K_7} S_{i(j+1)}^{k_1} + \frac{K_2}{K_7} S_{i(j-1)}^{k_1} + \\ & S_{ij}^{k_2} + \frac{F_{ij}^{k_1}}{K_7} - \frac{F_{ij}^{k_2}}{K_7} + \begin{cases} \frac{K_3}{K_7} (p_{ij}^{k_1} - S_{ij}^{k_1}) & \text{if } \exists P_{ij} \\ \frac{K_4}{K_7} (-S_{ij}^{k_1}) & \text{if } \nexists P_{ij} \end{cases}; \text{ and} \end{aligned} \quad (41)$$

$$0 = -\left(2\frac{K_5}{K_7} + 2\frac{K_6}{K_7} + 1\right)S_{ij}^{k_2} + \frac{K_5}{K_7}S_{(i+1)j}^{k_2} + \frac{K_5}{K_7}S_{(i-1)j}^{k_2} + \frac{K_6}{K_7}S_{i(j+1)}^{k_2} + \frac{K_6}{K_7}S_{i(j-1)}^{k_2} + S_{ij}^{k_1} + \frac{F_{ij}^{k_2}}{K_7} + \begin{cases} \frac{K_8}{K_7}(p_{ij}^{k_2} - S_{ij}^{k_2}) & \text{if } \exists P_{ij} \\ 0 & \text{if } \nexists P_{ij} \end{cases}. \quad (42)$$

Note that there are $N = p \times q$ nodes in each of the two surfaces. If all $2N$ nodes are numbered sequentially, such that (i,j,k_1) corresponds to n , $(i,j+1,k_1)$ corresponds to $(n+p)$, and (i,j,k_2) corresponds to $(n+N)$, the full system of equations can be expressed as $\mathbf{Ax} = \mathbf{b}$, where \mathbf{A} is a sparse $2N \times 2N$ square matrix (Figure 2.30):

$$A_{nn} = \begin{cases} -2\frac{K_1}{K_7} - 2\frac{K_2}{K_7} - 1 - \begin{cases} \frac{K_3}{K_7} & \text{if } \exists P_{ij} \\ \frac{K_4}{K_7} & \text{if } \nexists P_{ij} \end{cases} & \text{for } n \leq N \\ -2\frac{K_5}{K_7} - 2\frac{K_6}{K_7} - 1 - \begin{cases} \frac{K_8}{K_7} & \text{if } \exists P_{ij} \\ 0 & \text{if } \nexists P_{ij} \end{cases} & \text{for } n > N \end{cases}; \quad (43)$$

$$A_{n(n\pm 1)} = \begin{cases} \frac{K_1}{K_7} & \text{for } n \leq N \\ \frac{K_5}{K_7} & \text{for } n > N \end{cases}; \quad (44)$$

$$A_{n(n\pm p)} = \begin{cases} \frac{K_2}{K_7} & \text{for } n \leq N \\ \frac{K_6}{K_7} & \text{for } n > N \end{cases}; \quad (45)$$

$$A_{n(n+N)} = 1, n \in [1, N]; \text{ and} \quad (46)$$

$$A_{n(n-N)} = 1, n \in [N+1, 2N]. \quad (47)$$

Furthermore, \mathbf{x} is a vertical vector of node positions S_1 through S_{2N} , and \mathbf{b} is a vector:

$$b_n = \begin{cases} -\frac{F_n^1}{K_7} + \frac{F_n^2}{K_7} - \begin{cases} \frac{K_3}{K_7}p_{ij}^{k_1} & \text{if } \exists P_{ij} \\ 0 & \text{if } \nexists P_{ij} \end{cases} & \text{for } n \leq N \\ -\frac{F_n^2}{K_7} - \begin{cases} \frac{K_8}{K_7}p_{ij}^{k_2} & \text{if } \exists P_{ij} \\ 0 & \text{if } \nexists P_{ij} \end{cases} & \text{for } n > N \end{cases}. \quad (48)$$




Figure 2.30. Visualization of the sparsity pattern for matrix (A ; Equations 43-47) that expresses the system of linear equations describing the coupled (two-surface) linear-elastic mesh. (Indices indicate node number, n ; nz : number of non-zero elements; cf. Figure 2.14)

A simple constant force function ($F_n = c$) was used for each surface. Note that, in the representation presented in Section 2.3.3 [157], the base layer (L_{ij}) was set to 0. Multiplying both sides of the equation by the inverse of matrix A , the positions of all $2N$ nodes were solved simultaneously with a single operation: $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$. Results were then reconstituted to the pre-flattened state and converted from polar to Cartesian coordinates. As for the single surface (Section 2.3.3), optimization of spring coefficients K and force c was achieved by minimizing summed radial distance between expert annotations and algorithm results.

Validation was performed by comparing algorithm-delineated surfaces to contours of the target features annotated by human experts in the same OCT images. The same clinical dataset used to validate the simple 2D shape fitting (Section 2.3.1), which constituted a subset of that used in the more extensive validation of the other methods described above, was utilized again here. Indeed, the same lumen and intima-media border annotations manually traced by two interventional cardiologists experienced in reading OCT images was reused. However, to validate the simultaneous surface fitting, additional “ground truth” data were needed; a third experienced reader of OCT images annotated the media-adventitia border. While the best available reference

for the labeled media, this reference of the fine feature was imperfect; in the available dataset, 5.1% of expert IEM annotations were abluminal to the EEM annotations—a physiological impossibility. Nevertheless, results were analyzed through comparison of distance between algorithm- and human-annotated contours and corresponding layer thicknesses, with outcomes reported as absolute differences (mean \pm SD).

One-time optimization of the 9 parameters took 25.5 minutes, but setting up and solving the system of equations took just 0.8 seconds. Computational speed and preliminary image processing performance reflected that reported for the single-surface procedure (Section 2.3.3) [157]. Results for a sample frame are shown in Figure 2.31 alongside the human expert annotations.

The distributions of deviations of the two surfaces from those annotated by human experts are shown in Figure 2.32. The overall (combined) deviation for both surfaces was 0.10 ± 0.18 mm. The average difference in intima thickness (i.e. radial distance of the IEM surface), as well the average difference in the combined thickness of the intima and media layers (i.e. radial distance of the EEM surface), were both 0.10 ± 0.18 mm, while the average difference in media thickness (i.e. distance between IEM and EEM surfaces) was 0.038 ± 0.029 mm.




Figure 2.31. Representative image superimposing the human annotation (*left*) and algorithm results (*right*) on the same OCT cross-sectional image demonstrating moderate atherosclerotic plaque burden. *Green*: Human expert annotation (ground truth); *magenta*: IEM (intima-media border); *cyan*: EEM (media-adventitia border).




Figure 2.32. Histograms displaying deviation of algorithm results from human expert annotations for the intima, media, and combined thicknesses (reflecting the absolute and relative positions of the two surfaces).

The primary limitation of this work developing and implementing simultaneous 3D multiple surface fitting was the extent of validation. As noted before, due to its limited penetration depth, OCT is not the gold standard reference for measurements of far-field features—it is difficult even for experts to identify such features when plaque is present (despite impressively ability to approximate feature location [62]). Unfortunately, IVUS was not a feasible validation reference for this method as it was for the single outer border delineation; the IEM is not visible in IVUS images, so the media cannot be delineated using this modality [19].

Representing the culmination of geometry extraction of vessel walls using contextual information and flexible 3D approaches approaching performance of humans (Figure 2.4), this new method for simultaneous surface fitting allowed delineation of the intima and media layers in atherosclerotic arteries imaged by high-resolution OCT for the first time [158].

2.4. Extracting or Tracking Geometry from Disparate Datasets

Having achieved unprecedented ability to automatically extract vessel geometry from OCT images, it became highly desirable to be able to automatically align multiple vascular image datasets, thereby allowing direct comparison or combination of the geometric information contained by each. As discussed in Section 2.1.4, geometric information of interest is often distributed between multiple datasets. These datasets may be comprised of acquisitions from several complementary modalities, thereby providing a more complete description of vascular geometry, may be spatially sequential and partly overlapping acquisitions (effectively extending the modality’s range), or may be acquired at different time points to capture longitudinal changes. Due to inconsistencies in imaging modality through longitudinal care, evolution in algorithm availability and use, and variability in human annotations (when performed), properly aligning the disparate sources of geometric data remains a challenge; there is an unaddressed need for a simple method that utilizes universally identifiable features which can be automatically extracted from the image datasets and robustly applied in the task of co-registration.

Dataset alignment is typically performed on an individual basis using visible fiducial markers—clear landmarks visible in both datasets, such as side-branches, distinct calcium deposits, indwelling devices, or similar unique and uncommon features [18], [19]. However, across modalities, the most clearly and reliably identified feature is universally the lumen itself. Indeed, consistency between expert observers and modalities was found to be excellent, as was performance of automated methods for its extraction from OCT (Figures 2.10, 2.17, 2.20, 2.21, and 2.22). Yet, while previous (semi-)automated methods for inter- and intra-modality co-registration of intravascular image acquisitions used the lumen among sets of other features [206]–[211], it was unclear when—or whether—the lumen itself offered sufficiently distinct and reliable information to accurately align such acquisitions. Along with co-authors Karmakar *et al.* [159], I undertook work to develop a method to perform this task and assess the utility and limits of lumen-based co-registration.

The technical problem statement is as follows: given two tomographic image acquisitions, A and B, which capture at least part of the same vessel segment, the objective of co-registration is to identify corresponding frame pairs in the overlapping region of the datasets (i.e. those frames appearing in both sets which correspond to the same physical location in the vessel). To co-register the sets, the sequences must have at least partial overlap; one entire dataset may lie within the span of the other (which must be of equal or greater length), or the first M frames of one dataset may overlap with the M trailing frames of the other (Figure 2.33). To accomplish the co-registration task, sequence B, of length M_B , was registered to (aligned according to) sequence A, of length M_A . To represent the overlapping region between the two datasets, defining the first point of correspondence (FPC) is sufficient. Therefore, registration of dataset B to dataset A was represented as an ordered pair; the FPC between two datasets was either of the form $(a,1)$ or $(1,b)$, where a and b are frame numbers in datasets A and B, respectively. (For example, the ordered pair $(1,51)$ indicates that frame number 51 of dataset B corresponds with frame number 1 of dataset A.) With an assumption of constant frame spacing—a condition intended by modality developers and assumed clinically—the M overlapping frame pairs of datasets A and B corresponding to (FPC_A, FPC_B) constitute a set:

$$M = \min(M_A - FPC_A, M_B - FPC_B); \quad (49)$$

$$\text{Overlap}(FPC_A, FPC_B) = \{(FPC_A + l, FPC_B + l)\}, l \in [0, M - 1]. \quad (50)$$

(Note that, for simplicity, these relationships are expressed for sequences with equal longitudinal resolution or frame spacing. Equivalent general relationships can be expressed with proper scaling to accommodate sequences with differing, but internally consistent, longitudinal resolutions.)




Figure 2.33. Co-registration of tomographic datasets requires identification of common frames between two image volumes. Image sequences to be co-registered depict some common segment of the imaged vessel (shown in magenta). One acquisition (A, blue) is shifted relative to another (B, red) until an optimal match in common segments is found (magenta).

Within this framework, the task of the registration algorithm is to identify the ordered pair (FPC_A, FPC_B) defining the overlapping set of images, and subsequently to identify the angle θ by which the overlapping subset of B is rotated relative to that of A.

To identify the overlapping region between two datasets, several features and metrics of the vessel lumen depicted by each were compared between the sets to identify the optimal matches between them. These included the lumen area, lumen shape-signature, and shape difference metric, calculated for the lumen tracings of each dataset. As described by Karmakar *et al.* [159], longitudinal registration, refinement, and subsequent rotational alignment was achieved, in essence, by minimizing the difference in lumen areas and shape profiles between the matched (co-registered) frame pairs.

To assess the method and interrogate the applicability and limitations of co-registration utilizing the lumen alone, the algorithm was applied in several clinical datasets. These sets included the co-registered OCT and IVUS datasets with expert annotated lumen described and utilized in previous work [155]–[158], as well as sets of temporally longitudinal OCT acquisitions of heavily calcified vessels imaged before and after an atherectomy procedure, with lumens annotated automatically as in Section 2.3.3 [157].

When presented with the full useable and annotated OCT and IVUS datasets, the lumen-based automatic co-registration was excellent (Figure 2.34). All seven patient datasets, ranging

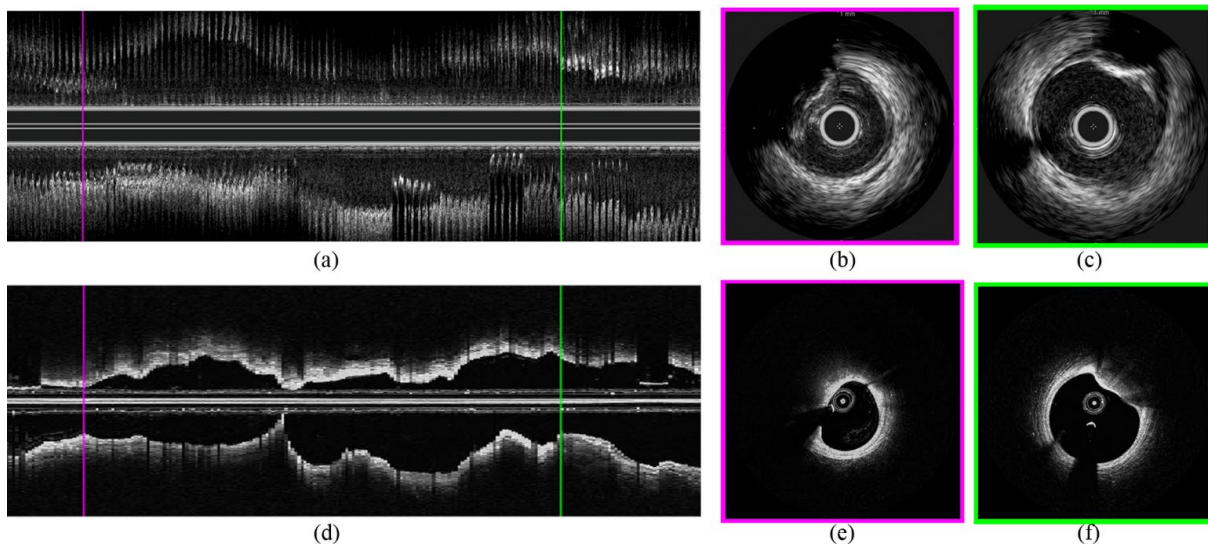


Figure 2.34. Automatically co-registered IVUS (a-c) and OCT (d-f) acquisitions acquired in the same diseased coronary artery. Cross-sectional images from the same position in each pullback show strong agreement in both longitudinal and rotational alignment, with lumen shape and (calcified) plaque location coinciding.

from 81 to 110 frames each, were aligned within 3 frames of the expert manual co-registration benchmark. Following an automated refinement process, offset from the expert alignment dropped to 0 for nearly all sets. Of greater scientific interest, the limitations of such lumen-based automatic co-registration was interrogated. In particular, we assessed the accuracy of the approach with respect to acquisition overlap length and lumen profile variability within the overlapping region; our hypothesis was that co-registration using lumen features and metrics alone would fail as the overlapping region between image sets became shorter or less variable (i.e. included fewer distinct features). Subsets of acquisitions—of varying length and depicting vessel segments of varying variability—were extracted and aligned with the acquisition of the complementary modality, and we indeed observed the trend which we hypothesized (Figure 2.35). We concluded that, at least in cases with minimal lumen change between acquisitions, the lumen is sufficient as a standalone feature to accurately co-register acquisitions longitudinally with at least 55 frames of overlap or lumen area monotonicity of less than approximately 0.8 in the overlapping region, though good performance is often (though less reliably) achieved with far less. Rotational alignment, however, was far less reliable in all permutations of overlap length and area profile variability. One particularly interesting observation can be seen among long overlapping segments of high

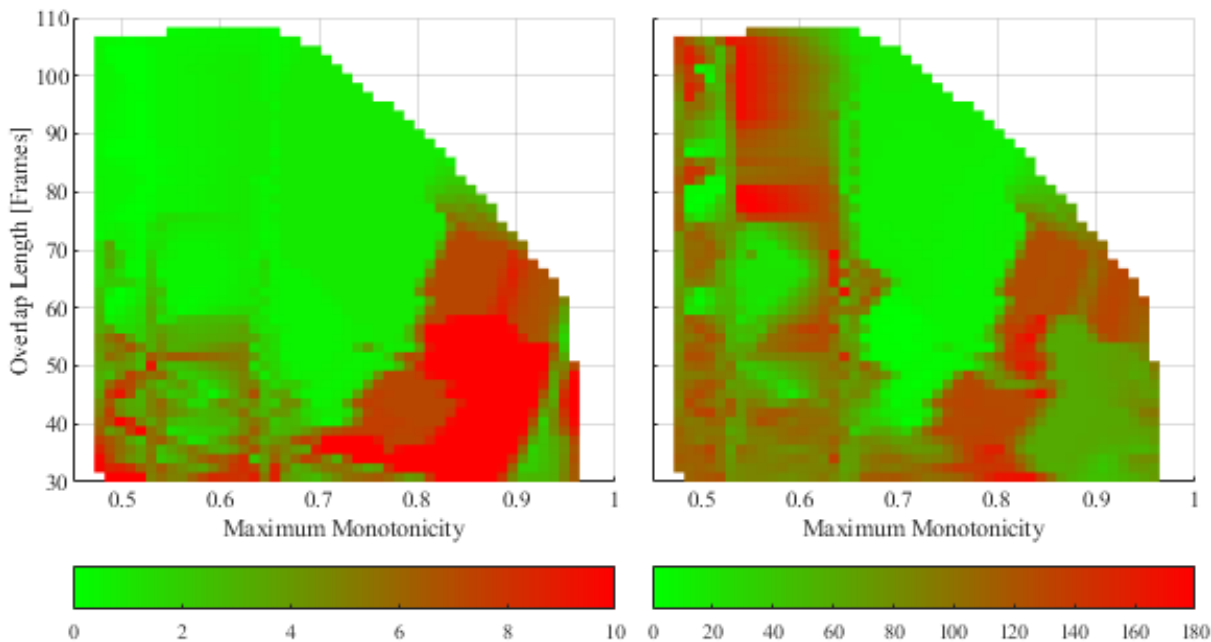


Figure 2.35. Analysis of the limitations of lumen-based automated co-registration examining the impact of overlap length and lumen area monotonicity (a metric of profile regularity in the overlapping region) on match accuracy. As the lumen area profile variability (negative correlate of monotonicity) and amount of overlap increases, the probability of acceptable longitudinal alignment results also increases (*left*). Performance of rotational alignment (*right*) is not as readily predicted, with overall performance being worse than longitudinal alignment. Colors indicate frame offset (in mm) and rotational offset (in degrees), respectively, from human expert co-registration.

variability (low monotonicity), which should yield the best results. Among these instances, there is a largely bimodal distribution of performance clustered around 0° and 180° deviation from the expert-determined alignment, indicating that even when eccentricity is present, lumen profiles alone are not sufficient to distinguish between opposite rotational alignments (Figure 2.35).

When presented with a more challenging dataset, the performance was nearly as accurate, and all the more impressive. Here, five sets of temporally longitudinal OCT pullbacks acquired before and after an atherectomy procedure, and again after stent deployment, were annotated automatically as in Section 2.3.3 [9]—adding additional error and uncertainty to the lumen tracings relative to those manually annotated by human experts in the multimodal datasets. Given these sets of exceptional clinical relevance, the algorithm achieved excellent qualitative matching between the datasets, allowing for direct qualitative and quantitative comparison of geometry to assess the impact of the atherectomy and stent deployment thereon (Figure 2.36).

The development of this flexible new automatic co-registration approach, and knowledge of its valid application, offers exciting opportunities where previous inter- and intra-modality co-registration of intravascular image acquisitions has been found lacking [206]–[211]. In particular, it offers intriguing possibilities to extract or track geometry from disparate datasets—quantitatively and thoroughly evaluating at scale the impact of interventions, such as atherectomy, stenting, and pharmacotherapy, or merging and combining geometric information extracted from complementary modalities.

For a more complete description of the methods, results, and discussion for the lumen-based method of co-registration and its applicability, see Karmakar *et al.* [159].

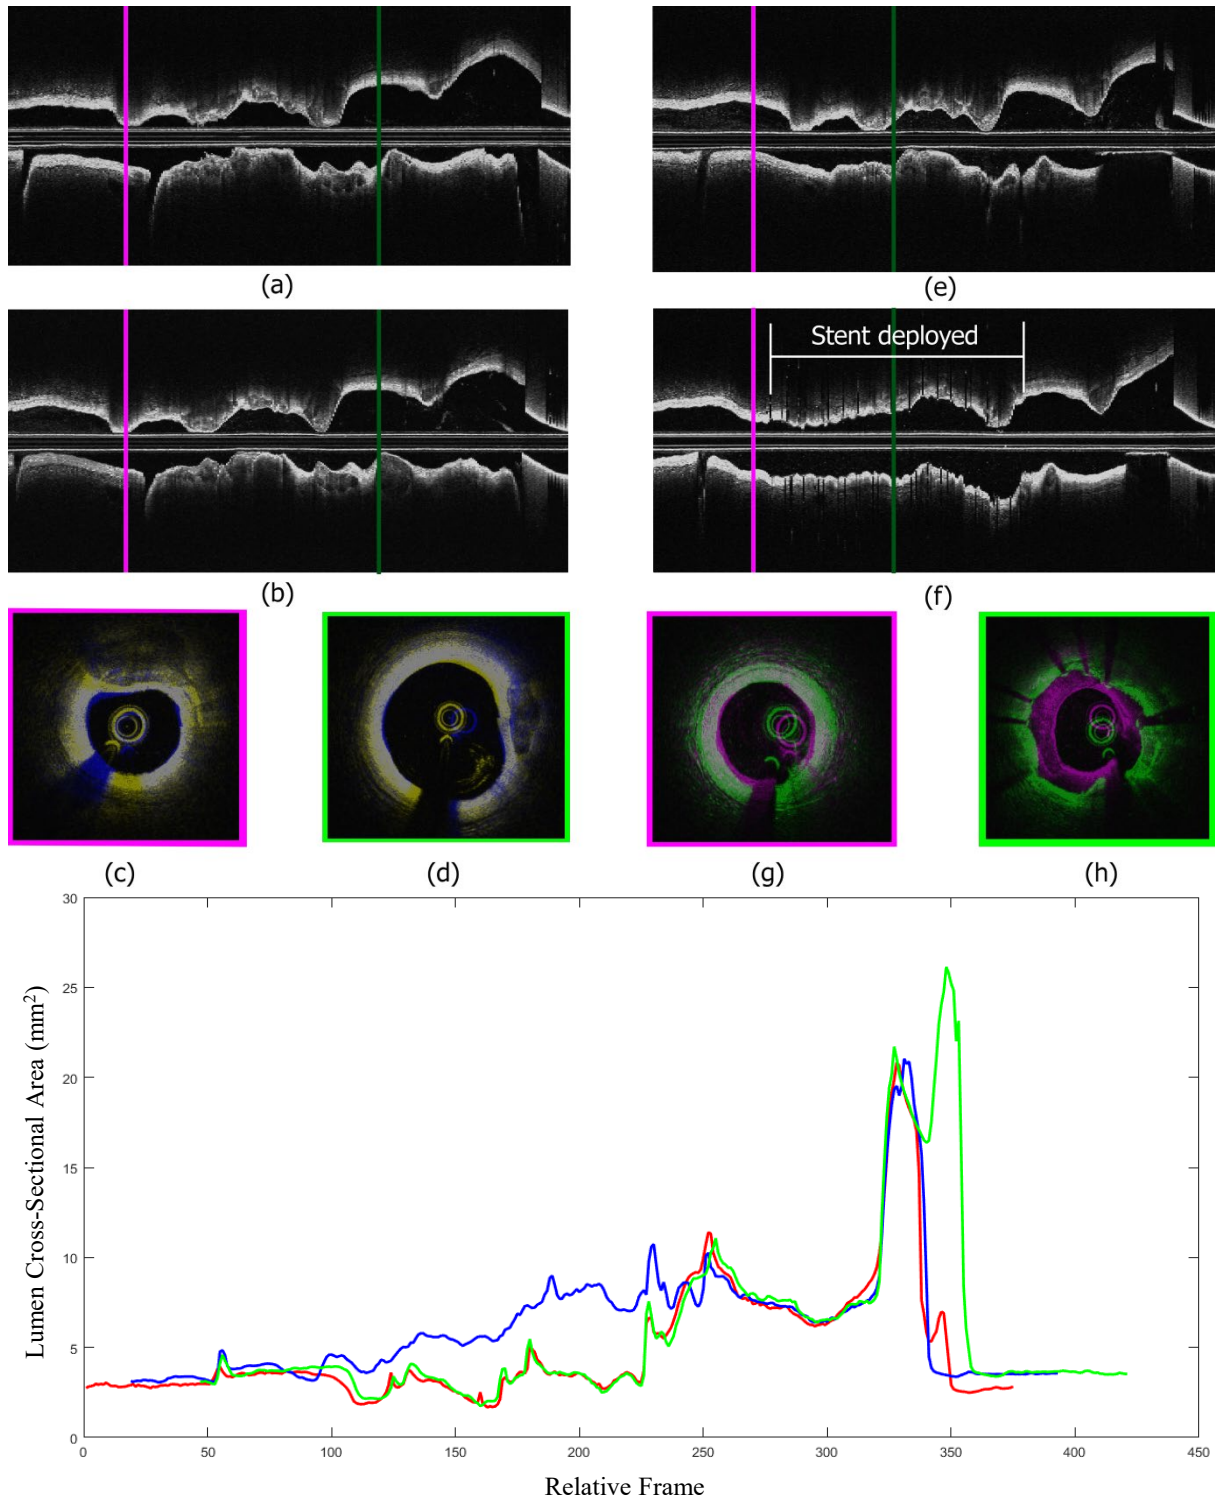


Figure 2.36. Co-registration of OCT acquisitions acquired in a heavily diseased vessel (a) before intervention (green in plot), (b, e) following atherectomy (red), and (f) following stenting (blue) showed surprising robustness despite lumen alterations. Unexpectedly, little change in lumen area (vertical axis) is observed following the orbital atherectomy procedure, but stenting has a clear impact (increasing lumen area). In this patient, a bifurcation (seen here as a sudden and drastic increase in lumen area) is observed just before the catheter sheath (a small lumen of constant size into which the imaging catheter is retracted at the end of its pullback). Note how the three pullbacks begin and end at different points, and also how they enter the sheath at different points in their acquisition.

2.5. Implications & Future Directions

My work and findings in extracting vascular geometry from intravascular imaging has technical, as well as scientific and clinical, implications.

Prior to this work, automated extraction of comprehensive vessel geometry in clinical intravascular imaging was limited to IVUS. IVUS does not allow for discrimination between the vessel's intima and media layers, so only lumen and EEM contours can be extracted therefrom. While OCT consensus standards indicate that measurements made using the IEM better reflect pathological definitions of atherosclerotic diseases, and dictates its use in most measurements [18], IVUS more clearly shows the nearby media-adventitia border associated with the EEM. Annotations in this modality therefore measure EEM, not IEM, CSA [19]. This area generally exceeds IEM CSA by 0.3 to 1 mm², depending on vessel size and media thickness, making measurements of plaque less reliable. Furthermore, IVUS has far lower in-plane resolution than other modalities, such as OCT, is slower and therefore subject to more motion artifacts, and has been shown to provide less accurate geometric measurements than other modalities. In studies comparing IVUS and OCT in measuring luminal areas either in phantom models or in coronary arteries of patients, a systematic overestimation of lumen areas has been observed with IVUS (10-15%), with OCT being more accurate [40]. With the development of the methods described herein to automatically extract vessel geometry—lumen, IEM, and EEM—from OCT, as well as relatively recent studies showing that experts can derive accurate measurements of vessel area in arteries with high plaque burden [62] (especially when images are computationally enhanced [65]), the combination of IVUS shortcomings challenges the prevailing exclusive use of IVUS in many centers worldwide. The results of my studies highlighting the presence of inconsistencies between OCT and IVUS also reinforce the need for clinicians to consider possible systematic IVUS measurement inaccuracies in decision making, as well as the benefit of utilizing normalized measurements whenever possible.

The novel concept for simultaneous delineation of vessel layers in diseased arterial segments imaged with OCT, allowing for the first time isolation of the media of diseased vessels, may improve understanding of disease progression and improve clinical utility of OCT. Because media thickness correlates with both disease status [164] and progression [183], future clinical application to automate layer thickness measurement warrants exploration. Furthermore, the

ability to delineate the media could enhance the diagnostic accuracy of OCT plaque characterization systems [160], which, along with VH-IVUS, currently rely upon contrived approximation of the media location and thickness independent of any patient-specific factors. The most promising applications in this field, therefore, may improve characterization of vessel wall plaque by subsequent computational routines, offer quantitative assessment and prognosis of patient health, and enable scalable longitudinal monitoring and study of disease progression.

The capability of extracting detailed vessel geometry also has exciting implications because it allows for the construction of more accurate 3D models of patient-specific vessel geometry from high-resolution OCT data for research and future clinical applications. Such models have been used to quantify and/or visualize 3D vessel properties, simulate interventions and their consequences through finite element and computational fluid dynamics analyses, and identify some cases of vulnerable plaques [68], [70], [103], [172]. However, as described in Section 2.1.1, previous coronary artery anatomic reconstructions have been subject to limitations in application and/or fidelity arising from methodological approaches. The new methods described in this chapter allow for the construction of models from *in vivo* geometry, without the limitations of idealization or unrealistic geometric assumptions, and with the benefit of the high resolution of OCT. These methods can also be used in conjunction with existing reconstruction techniques that presuppose knowledge of the outer border location [226].

The surface-fitting method ultimately developed to determine the full vessel wall on the basis of dispersed and potentially unreliably identified visible segments holds great potential for wide use—in its current application and well beyond. The method described in this chapter is flexible, and may be applied to various other applications within and beyond image processing. In situations where gaps exist in data spaces, where no or limited information is available, data points are noisy, or data are unevenly-distributed, the surface fitting technique can be used where others may struggle, providing estimates for the missing information and smoothed assessments of the known data. Moreover, as shown in the smoothing of the lumen, which was nominally detected with a previously validated method [173], the presented method can also be used to smooth complete sets of noisy, yet otherwise reliable, data. Furthermore, while optimization using curated datasets improves fit performance, parameters may be selected intuitively or heuristically to achieve reasonable fits. This feature sets the approach apart from others, such as polynomial fits,

which provide drastically-varying results with slight modifications of input parameters. Finally, the intriguing ability to specify a force function provides a unique opportunity to incorporate complex features of all available data to improve the fit and resulting evaluation of underlying information. This is particularly useful in image analysis, where features may not be readily detected but hints about the likely position of such features may be extracted from the rich set of information embedded within the image. Thereby, insights from other studies, including machine learning, pattern recognition, and other analyses, may be incorporated.

This method, using equilibrated mechanical lumped elements, is even more intriguing given that extension of this work could assign values and attributes to physical parameters that are extracted from the architecture and morphology of a blood vessel wall. Fit parameters derived from the data might well offer deeper meaning and allow inferences about vessel state and function. Either through representative indices of geometry coupling disease progression and vessel mechanics, or direct mechanical representation by the lumped parameters akin to simplified inverse finite element analysis (explored in Chapter 4), the method may offer direct metrics by which to cross between geometric and constitutive properties of the vessel. While preliminary work in the relatively small dataset of seven patients lacked the power to reveal clear underlying trends, it remains feasible that mechanical parameters, reflecting geometric relationships, may well vary with disease state or phenotype and offer new insight into the disease-driven changes undergone in the vessel. Indeed, the outer surface method, utilizing spring parameters optimized for the original published cohort of CAD patients, has been subsequently found to perform sub-optimally in a new cohort of patients with very high calcified plaque burden undergoing atherectomy. While visible border segment detection contributed to the diminished performance (as sheets of calcium displaying visual similarity with the media were sometimes misidentified as the latter), the geometric and structural regime of the heavily calcified phenotype may be better served, and reflected, by a different set of mechanical parameter values.

The dramatic improvement in outer border delineation achieved through the progression from 2D to 3D also reveals insights into the human experience of determining this information from visually incomplete datasets. The profound differences suggest an imperative role for context—an observation obvious at face value. However, availability of information from immediately adjacent segments of the dataset is not sufficient; rather, it seems likely that humans

take a more global approach, identifying overall trends in the dataset and visualizing a solution in 3D space that is projected to a 2D plane. This insight should inform both the presentation of datasets to humans for evaluation and the assessment of human reliability of such evaluations within the context of the data made available to them (and the format thereof).

Insights and opportunities from the automated lumen-based image co-registration method are also noteworthy. The methods for extracting geometry described in Sections 2.2 and 2.3 were designed for application to OCT image datasets (though the mechanical approach to smooth surface fitting is notably applicable to a wide range of uses). Various approaches to automatically segment the inner and outer border of vessels imaged with IVUS have been validated [175], [180]. The most accurately, reliably, and conveniently extracted feature of all modalities, however, is the lumen. The ability to align automatically extracted geometries on the basis of this feature alone therefore offers intriguing possibilities to track or combine geometry from disparate datasets. Tracking geometry allows for comparison between state or time points, enabling scalable quantitative and thorough evaluation of intervention impact. Already, application of the method to OCT datasets acquired before and after atherectomy have led to surprising observations about the minimal change in lumen area resulting in some cases. This approach, along with other tools and methods developed in this thesis, may help to better ascertain the impact and driving factors which lead to divergent interventional outcomes. Merging and combining geometric information extracted from multiple datasets may be used to extend the effective spatial range (i.e. length) of pullbacks when necessary, and could also improve geometric measurements or models when complementary modalities are utilized. However, of immediate application are the insights regarding the applicability and limitations of lumen-based co-registration. Major features involving the lumen, such as side-branches, remain important, but surprisingly are not strictly necessary when sufficient alternative sources of lumen profile variability exist in the imaged vessel. While it is unclear whether humans can utilize such nuanced and subtle information alone, we have shown that computational methods can do exactly that. Moreover, while it may be expected that differences in physiological state (e.g. point in the cardiac cycle) or differences caused by interventions (e.g. stenting) may prevent alignment of non-concurrent acquisitions, it was found that underlying and adjacent anatomic variability dominates such deviations when a sufficiently long or undulating segment is considered. In other cases, more complex approaches

utilizing supplementary features, such as plaque morphology or indwelling devices, may be needed.

All methods described in this chapter should be further improved and validated prior to deployment in a clinical setting. Refinement will require, in part, more imaging datasets from additional patients and more reliable expert annotations thereof. An improved annotation user interface that allows finer control of contour tracing should reduce annotation errors. Algorithmic refinement should improve lumen detection robustness in the face of significant imaging artifacts (e.g. residual blood) and other challenges (e.g. catheter contact with the vessel wall). Improved visible feature (e.g. media) detection, and optimization of specificity relative to sensitivity of such methods, is also desired. To further the validation of these approaches, one unavailable but potentially useful tool in this endeavor is a phantom model displaying pathophysiologically relevant imaging properties. If a phantom model could be developed with physiological optical properties, it could serve as an ideal platform for validation, as ground truth measurements would be known with relative certainty. Currently, errors arising from physiological state disparities between non-concurrent acquisitions and imperfect co-registration amplify and compound inaccuracies in the evaluated methods. To avoid these magnifying errors, other work in the field has leveraged phantom models of known, directly measureable dimensions for validation purposes [40], [170]. While exceedingly useful for validating imaging technology and evaluating measurements of inner surfaces and their contents (e.g. stents), measurement of the outer border does not offer a useful validation mechanism due to the lack of physiologically-relevant optical properties and feature patterns [18], [227].

2.6. Conclusion

In this chapter, I have discussed a breakthrough method of extracting vessel geometry which overcomes current limitations of intravascular imaging. By employing *a priori* knowledge of arterial geometry and image acquisition procedure, and progressing from 2D feature detection to 3D integration of adjacent and contextual information, existing barriers arising from limited penetration depth is largely surmounted to determine the geometry of diseased vessel wall geometry in high-resolution OCT. Insights on the application and limitation of lumen-based image co-registration were also presented.

Penetration depth has been a significant limitation of OCT since its introduction. The methods detailed herein enable key functions of OCT that were previously available only from older imaging modalities. Mitigating this limitation may promote greater adoption of this technology and free researchers and clinicians from choosing between modality tradeoffs. The capability to autonomously delineate the entire outer border of diseased vessels in OCT images was successfully demonstrated, for the first time, and the results were successfully co-registered automatically on the basis of lumen profile alone. The approaches present radically new methods for incorporating and merging complex data to achieve smooth surfaces from incomplete surface segments, like those encountered in imaging diseased coronary arteries, offering exciting clinical and technical applications and revealing important insights into human interpretation of spatial data.

With vessel geometry successfully extracted from intravascular imaging, next steps for these methods include their integration with machine-learning methods used to characterize plaque. This approach, addressed in the next chapter, allows for the construction of heterogeneous models from *in vivo* geometry, without the limitations of idealization or unrealistic geometric assumptions.

Chapter 3

Determining & Characterizing Vessel Morphology

While vessel geometry drives functional capacity and yields critical metrics of vascular health, it is an incomplete descriptor of vascular disease state. The consequences of atherosclerotic plaque deposition and accumulation are variable, and evolve with the plaque morphology, distribution, milieu, and complex interactions between these factors. Morphology, including composition, structure, and form, can be highly heterogeneous, with lesions comprised of calcifications, fatty lipid deposits, fibrotic tissue, and variants thereof. Variations in morphology consequently drive differences in lesion risk profile (e.g. rupture risk), optimal treatment strategy, and interventional success. Morphology can be determined visually from intravascular ultrasound (IVUS) and optical coherence tomography (OCT) images, and such assessment is frequently performed qualitatively and informally in clinical settings [17]–[19], [163], [228], but the process is inherently limited. In this chapter, the development and critical assessment of a unified framework to automatically determine and characterize vessel wall morphology (and indwelling devices) in OCT [160] and IVUS [161] is described, with studies highlighting the importance of enriching such methods with domain knowledge of vascular geometry and imaging [161] and carefully considering the curation of training data [162].

ASM	Angular Second Moment
BVS	Bioresorbable Vascular Scaffold
CNN	Convolutional Neural Network
DC	Dense Calcium
FCNN ..	Fully Connected Neural Network
FFT	Fibrofatty Tissue
FT	Fibrous Tissue
GLCM	Gray-Level Co-Occurrence Matrix
GPU	Graphics Processing Unit
IVUS ...	Intravascular Ultrasound
LT	Lipid Tissue
MT	Mixed Tissue
MLP	Multilayer Perceptron
NC	Necrotic Core
NPV	Negative Predictive Value
OCT	Optical Coherence Tomography
PPV	Positive Predictive Value
ReLU ..	Rectified Linear Unit
RF	Radio Frequency
ROI	Region of Interest
SGD	Stochastic Gradient Descent
VH	Virtual Histology
2D	Two-Dimensional

3.1. Diseased Vessel Morphology: Needs, Progress, & Challenges

The morphology of a vessel, like its geometry, can be distinctly altered by disease and significantly impact its function. Plaque composition and structure can also drastically alter the optimal course of treatment. Vessel morphology is therefore a critical consideration in clinical diagnosis, prognosis, and treatment planning, as well as research. This need, and the tediousness of manual annotation, has motivated efforts to automate the characterization of diseased vessels, though success and clinical translation to date has been limited.

3.1.1. The Need for Accurate Morphology

As previously described, atherosclerosis is a complex process involving inflammation, thrombosis, and cellular infiltration and proliferation culminating in rupture and/or sclerosis to ultimately obstruct arterial flow [229], [230]. The hallmark of atherosclerosis are mural plaques composed of lipids, inflammatory cells, and calcium that progressively impinge on the lumen, reducing distal perfusion. Tissue insufficiency that follows causes diseases that are the leading cause of morbidity and mortality globally [1].

A primary step in diagnosing and treating atherosclerosis is imaging the arterial lumen and vessel wall. While dye angiography can easily outline the former, observing the latter is more challenging. Several techniques can visualize the lumen border and roughly ascertain the constitution of the arterial wall, but intravascular imaging is the current method of choice at divining deeper tissue content [175], [231], [232]. IVUS and OCT are invasive techniques which provide two-dimensional (2D) tomographic views of the arterial lumen and vessel wall, allowing comprehensive visualization of present plaque. Generated images can provide reliable geometric measurements and estimates of plaque composition [18], [19]. A well-trained expert can manually determine the dimensions of the lumen and media-adventitia border. Together these metrics delineate the limits of the arterial wall and primary region of interest (ROI), which can include various plaque constituent types.

The plaque constituent types identifiable in each modality differ somewhat in their classification, which results from variation in imaging mechanism (mechanical or optical interaction), tissue penetration, and levels of available evidence (e.g. definitive published correlative studies), and are reflected in the respective consensus standards for IVUS [19] and OCT [18]. Plaques identified in IVUS, particularly VH-IVUS, are typically classified as dense

calcium (DC) or calcific, necrotic core (NC), fibrous tissue (FT), and fibrofatty tissue (FFT), sometimes alternatively called mixed tissue (MT; though such a classification also includes fibrocalcific and other multi-subtype combinations) [19], [233]–[235]. DC is composed of compact calcium crystals, while NC consists of high levels of lipids with many necrotic cells. While both FT and FFT include collagen fibers; the former is mainly bundles of fibers [236], and the latter loosely packed fibers with lipid accumulations [235]. In OCT images, FT can similarly be identified, as can MT, though the latter has a greater emphasis on fibrocalcific (rather than fibrofatty) plaque. Larger calcifications (e.g. nodules; not differentiated from “dense calcium”) are also identifiable, as is lipid tissue (LT)—“lipid pool” is a broader histopathologic category including both NCs and regions containing extracellular lipid, and its confident identification in OCT is better supported by correlative studies than the more narrow NC classification [18], [237]. Due to their varying composition, each plaque type has unique echoreflectivity and optical characteristics, and consequently differentiable appearance within an IVUS or OCT image. For example, calcium regions uniquely appear dark with well-defined boundaries in OCT images [18], [237], whereas in IVUS it appears as very bright echoes but casts an abluminal acoustic shadow [19].

Knowledge of the plaque morphology, visualized through intravascular imaging, has important tangible benefits in both clinical treatment of atherosclerosis and translational research. The composition and morphology of plaque is an important determinant of adverse events and interventional success. Certain morphological features assessed by intravascular imaging, for example, are associated with plaque vulnerability⁵, including thin fibrous cap overlaying a lipid pool or necrotic core [238]. The presence and distribution of lipid is also a predictor of future nonculprit lesion-related major adverse cardiac events [239], and is related to intervention outcome, as stents should be placed so as to avoid edges over lipid formations, which may rupture and embolize [240]. The impact of morphology on intervention is particularly profound as it relates to calcium presence, prevalence, and distribution, which have demonstrable prognostic and treatment implications [27], [241]. For example, lesion calcification can significantly impact stent expansion [30], [242], [243], and thick, eccentric calcified plaque can especially oppose optimal

⁵ Though in widespread use, the concept of “vulnerable plaque” is a disputed one, now largely considered to be a useful but incomplete framework. It should be considered not in isolation, but within the context of multifactorial patient vulnerability [352], [353], as well as the evolution of disease phenotype prevalence in modern medicine and society [340].

expansion [244]. Calcified plaque has also been shown to limit intravascular drug delivery which would otherwise limit intimal hyperplasia and restenosis of treated vessels [245]. Therefore, pre-intervention evaluation of plaque type and morphology often dictates decisions on lesion preparation and plaque modification, such as cutting or scoring balloons, excimer laser angioplasty, rotational or orbital atherectomy, or lithotripsy, prior to stenting to improve acute outcomes [240]. Further decisions may also be swayed by the identification of calcium, such as a preference for metallic, rather than more brittle polymeric, scaffolds, as well as the use of higher balloon inflation pressures [241].

Beyond current clinical use, plaque morphology plays a key role in ongoing translational and fundamental research into atherosclerosis pathology and treatment, with additional future applications in the clinic. Intravascular imaging studies have illuminated clinical implications of coronary calcification and shed light on the role of calcium in atherosclerotic pathology and treatment [246]. Other avenues of investigation have included as key elements biomechanical computational modeling. For example, plaque structural stress, evaluated by finite element models of patient-specific plaque morphologically, has been found to predict major adverse cardiovascular events [247]. Such studies rely on accurate plaque-specific morphology to make predictions and compute stress, which is substantially altered by intima heterogeneity [76] and material stiffness [248].

Due to the actionable insights derived from plaque morphology and its research importance, this information is commonly extracted from intravascular imaging, which captures plaque composition and structure in far greater detail than alternative imaging modalities. Indeed, manual ROI and tissue detection has been used since the introduction of intravascular imaging. However, a single acquisition sequence can contain several thousand individual frames (images) in IVUS [19] and several hundred frames in OCT, so manual processing of a full acquisition is prohibitively time-consuming and laborious. It is also subject to high inter- and intra-observer variability [249]. Morphological assessment is therefore often performed qualitatively in just a few frames. Moreover, discrimination of FT from FFT plaque constituents is limited, since the two plaques share similar characteristics. These limitations, and the promise of greater insight derived from plaque morphology, has driven a demand for automated, objective methods to extract vessel morphology from intravascular images.

3.1.2. Progress in Automatic Morphology Extraction

Image processing and segmentation is increasingly being automated to overcome the subjectivity, slowness, and training requirements of human annotation. This is particularly true of medical imaging, though the field continues to struggle with the natural variability of native anatomy and pathologic feature presentation and prevalence. Such is the case of classifying patient arteries visualized through intravascular imaging and identifying various plaque components, which vary in appearance and prevalence.

A pre-requisite for extracting morphology from vascular imaging is the segmentation of vascular geometry to delineate the vessel wall, which constitutes the ROI for characterization routines. Various automated ROI detection algorithms have been developed for IVUS [175], [176], [180], [250]–[254]. Previous limited efforts in OCT [182], [183] were introduced in the previous chapter alongside the novel methods developed in this thesis [155]–[158], though some superficial characterization methods took a more simplistic approach to ROI delineation by distally expanding the lumen border some fixed distance [173], [175].

Since manual plaque characterization is time consuming and relies on well-trained readers, several studies have attempted to automatically detect the various plaque components in vessel wall ROI using intravascular images [255]. Several narrowly-focused approaches have been proposed to detect calcium automatically in OCT images. Methods have included grayscale k-means clustering [256], fully connected neural network (FCNN) binary segmentation [257], and convolutional neural network (CNN) multi-class segmentation [258]. However, reported performance metrics vary considerably, with sensitivities ranging from 26 to 83% and positive predictive value (PPV) varying between 72 and 88%. Other approaches in OCT have attempted to characterize a broader set of tissue types. Xu *et al.* correlated the backscattering and attenuation coefficients with CT, LT, and FT [259], while, in a similar attempt, van Soest *et al.* correlated the attenuation coefficients with healthy vessel wall, intimal thickening, lipid pool, and macrophage infiltration [260]. However, neither could fully define any clear threshold values between the different tissue types. Going one step further and using machine learning, Athanasiou *et al.* presented a fully-automated OCT plaque characterization method which classified plaque as CA, LT, FT, or MT, with 83% accuracy [173]. More recently, deep learning approaches using CNNs to characterize OCT images [257], [258], [261]–[263] have achieved an overall accuracy of up to

95.8% [263]. However, the best performing method achieving the highest overall accuracies have binned tissue with highly simplified classes of lipid, fibrous tissue, and “background” [263], distinguishing between fewer levels than those identified by experts and defined in modality standards [18], [19]. This common lumping strategy yields less precise and detailed information on plaque morphology, but at the same time typically attains higher accuracies because the set of evaluation metrics—also defined by segmented bins rather than precise classes—intrinsically allows for greater leniency by neglecting class distinctions.

Numerous plaque characterization methods using IVUS images have also been reported in the literature. The majority of these methods are based on machine learning approaches. Zhang *et al.* were of the first to automatically extract image texture features and classify pixels using a learned piecewise linear discrimination function [264]. Since then, many have followed, using different feature sets and classification algorithms [265], [266]. Such methods follow the same general pattern: grayscale images are used as input and pixels are classified by a machine learning algorithm according to pixel intensity and imaging characteristics (e.g. acoustic shadows) or a supplementary set of extracted texture and geometric features. The gold standard for those methods was human expert manual annotations, which limited available data and suffered from inter- and intra-observer variability; subsequent implementation of the methods in clinical practice was hindered in part because validation and training relied upon such manual annotations. Therefore, Taki *et al.* [267]–[269]—followed by others [270]–[272]—proposed similar machine learning approaches trained and validated using the results of a commercially available software: virtual histology (VH) IVUS [235].

VH-IVUS was introduced to surmount the limitations of manual labeling of diseased vessels [235]. By processing the frequency spectrum of backscattered radiofrequency (RF) signal [273], rather than just the reflected signal amplitude, a more detailed assessment of the plaque can be generated with high accuracy confirmed through histology validation [233], [235], [274]–[276]. VH-IVUS can classify plaque into the four subtypes defined for IVUS [235], and treats the non-pathological tissue and media—the concentric layer separating the disease-prone intima from the outer adventitia layer—as a separate combined class (M). VH-IVUS processed images are presented as color-coded plaque characterization maps, often overlaid on the corresponding grayscale images (Figure 3.1). The technology is the current standard for *in vivo* and *in situ* examination of atherosclerotic arterial lesions [233], [235]. Although VH-IVUS provides




Figure 3.1. Sample VH-IVUS frame illustrating a highly heterogeneous plaque lesion. A grayscale IVUS image (*left*) is overlaid with plaque types characterized by VH (*right*) as dense calcium (white), necrotic core (red), fibrotic tissue (green), fibrofatty tissue (light green), and media or non-pathological tissue (gray).

relatively accurate plaque characterization, its main disadvantage beyond the lacking in-plane resolution of IVUS itself is the fact that it requires acquisition of RF signal and proprietary software to process this signal. As a consequence, the plaque composition of grayscale IVUS frames acquired without the full RF signal (or without the proprietary software) cannot be characterized by this technique. Moreover, the RF signal is available only in the ECG-gated R-peak IVUS frames [277]—~1 of every 30 frames—resulting in significant information loss and large segments of uncharacterized vessel. Thus, alternative methods able to characterize the plaque in a similar manner as VH-IVUS using grayscale images alone remain attractive and highly relevant.

3.1.3. The Endeavor to Extract Vascular Morphology from Intravascular Images

Several gaps in capabilities and knowledge persist in the extraction of vascular morphology from intravascular images. In both OCT and IVUS, existing methods either lack sufficient accuracy or wide applicability (e.g. VH-IVUS). Previous OCT methods utilized rudimentary ROIs which did not accurately reflect vessel geometry, and also did not appropriately address image artifacts, particularly the guidewire shadow. Furthermore, tissue characterization lagged behind in the application of new technological advances in machine learning, and lacked a unifying framework for the development and consideration of new techniques and applications. There was also limited understanding of tradeoffs between alternative machine learning methods, the

importance of enriching learning methods with domain knowledge, and the full impact of training data curation.

In characterizing tissue in OCT images, although CNN-based methods outperformed older machine learning methods, they all failed to characterize the full arterial wall, resulting in methods which could not compete with widely-used VH-IVUS [235], limiting the imaging detail superiority benefits of OCT when compared to IVUS, and weren't appropriate for the development of image-based computational models. The primary drawbacks hampering automated OCT plaque characterization were the lack of large amounts of annotated images and the non-realistic tissue (ROI) segmentation caused by the difficulty of automatically detecting the outer border. To generate geometrically and morphologically correct representations, this thesis describes an approach which leverages a recently developed lumen detection algorithm [225] and the outer border detection method described in the previous chapter [157] to automatically characterize the whole arterial wall. Innovative aspects of the proposed tissue characterization method included its use of CNNs with a large amount of annotated data to detect atherosclerosis, detection of normal tissue and shadowed areas within the OCT images, and classification of the full vessel. The novel and accurate method detected and characterized, for the first time, the whole arterial tissue in a way similar to VH-IVUS. Building upon this work, additional characterizations methods for OCT were explored, illustrating key tradeoffs to inform future development. Furthermore, the flexibility of such methods was extended to also classify indwelling devices in the vessel, specifically bioresorbable vascular scaffold (BVS) struts.

In characterizing tissue in IVUS images, poor performance of grayscale image analysis and limited applicability of RF-based VH were overcome. Recent developments in deep learning and CNNs made possible characterization tools in different imaging modalities which outperform methods deploying traditional machine learning or image processing. Indeed, none of the existing IVUS plaque characterization methods, which required explicit feature set design, selection, and extraction through pre-processing, had achieved overall label assignment accuracy >90% [175] (Table 3.1). However, deep learning had been applied to IVUS only for delineating inner and outer boundaries of the arterial wall (i.e. ROI) [180], [254] and to select frames containing calcification [278]; no method had applied CNNs to grayscale IVUS imaging data to improve plaque characterization and generate information akin to VH-IVUS. This thesis describes work that did just that, achieving significant accuracy. Moreover, fundamental insights were also gained through

Table 3.1. Comparison of current methods & previous methods reported in literature for the characterization of atherosclerotic plaque in IVUS

Method	Tissue Classes*	Gold Standard	Validation Dataset	Reported Accuracy
Zhang <i>et al.</i> (1998)	DC, NC, & FT	Expert annotations	12 IVUS images	Overall: 89.9%
Brunenberg <i>et al.</i> (2006)	DC, NC, & FT	Expert annotations	30 IVUS images	Overall: 75.82%
Athanasίου <i>et al.</i> (2013)	DC, NC, & FT	Experts annotations	40 IVUS images	Overall: 86.17%
Taki <i>et al.</i> (2010)	DC, NC, & FT-FFT	VH-IVUS	500 VH-IVUS frames	DC: 79%, NC: 52%, & FT-FFT: 81%
Taki <i>et al.</i> (2013)	DC, NC, & FT-FFT	VH-IVUS	500 VH-IVUS frames	DC: 84.87%, NC: 80.57%, FT: 77.4%, & FFT: 63.47%
Hwang <i>et al.</i> (2018)	DC, NC, FT, & FFT	VH-IVUS	252 VH-IVUS frames	Overall: 82.8%
Kim <i>et al.</i> (2018)	DC, NC, FT, & FFT	VH-IVUS	252 VH-IVUS frames	Overall: 85.1%
Athanasίου <i>et al.</i> (2012)	DC, NC, FT, FFT, & M	VH-IVUS	300 VH-IVUS frames	DC: 84.87%, NC: 80.57%, FT: 77.4%, & FFT: 63.47%
Domain Enriched Deep Learning Approach Olender <i>et al.</i> (2020) Section 3.3	DC, NC, FT, FFT, & M	VH-IVUS	200 VH-IVUS frames	DC: 98.5%, NC: 88.6%, FT: 91.1%, FFT: 90.0%, & M: 99.4%
Naïve Deep Learning Approach Olender <i>et al.</i> (2020) Section 3.7	DC, NC, FT, FFT, & M	VH-IVUS	200 VH-IVUS frames	DC: 98.7%, NC: 87.3%, FT: 87.5%, FFT: 89.5%, & M: 75.9%
VH-IVUS Validation by Nair <i>et al.</i> (2002)	DC, NC, FT, FFT, & M	Histology	889 histological regions	DC: 96.7%, NC: 95.8%, FT: 93.5%, & FFT: 94.1% (M not reported)

* DC: Dense Calcium; NC: Necrotic Core; FT: Fibrous Tissue; FFT: Fibrofatty Tissue; M: Media/Non-Pathological

analyses examining the impact of enriching such deep learning methods with domain knowledge of vessel geometry and IVUS imaging technology.

To analyze the suspected impact of training data on characterization performance variability, a novel method of automatic calcified plaque detection was devised that allowed precise control of class distribution in the dataset. By varying the relative representation of pixels belonging to either class, the impact of the training data on overall model performance was quantified. Segmentation results of the proposed method were validated against expert annotations. While calcium in OCT images could be automatically detected with meaningful precision using

local texture measures and traditional neural networks in a limited dataset, training data balance was found to be remarkably consequential.

Together, this chapter details technological advancements and insights that enable more widespread use of plaque morphology, derived from intravascular imaging, to inform prognosis and treatment of atherosclerosis and advance fundamental knowledge of pathology and intervention mechanisms through emerging research platforms.

3.2. A Flexible Approach for Intravascular Image Characterization: Fundamental Methodological Framework

Leveraging the powerful, modern tools of deep learning tools, and integrating knowledge of vascular anatomy, pathology, and imaging, a flexible and well-performing approach was devised and applied to characterize OCT and IVUS images. In this way, plaque morphology, and with slight modification also state and distribution of indwelling devices, was determined from *in vivo* clinical data.

The framework developed for characterizing intravascular images was built around a CNN-based domain enriched method that classified the imaged arterial tissue. The method detected the ROI and then subdivided the ROI into pathological and non-pathological tissue based upon basic spatial and geometric constraints informed by physiology. Pathological areas of the ROI were partitioned into patches and fed through a CNN architecture. As such, this automated plaque characterization method consisted of three primary steps (Figure 3.2, left and center):

- 1) The ROI (i.e. vessel wall area) was first detected using lumen and outer border detection methods.
- 2) Pathological tissue was partitioned from the rest of the vessel wall (M) based upon domain knowledge of spatial constraints enforced by arterial physiology and pathology, imposing physically-relevant limits on the location and dimensions of this tissue class while also reducing the number of classes to be subsequently segmented by the CNN.
- 3) Pixels of the ROI in the pathological area were classified into one of the plaque types using a CNN algorithm.

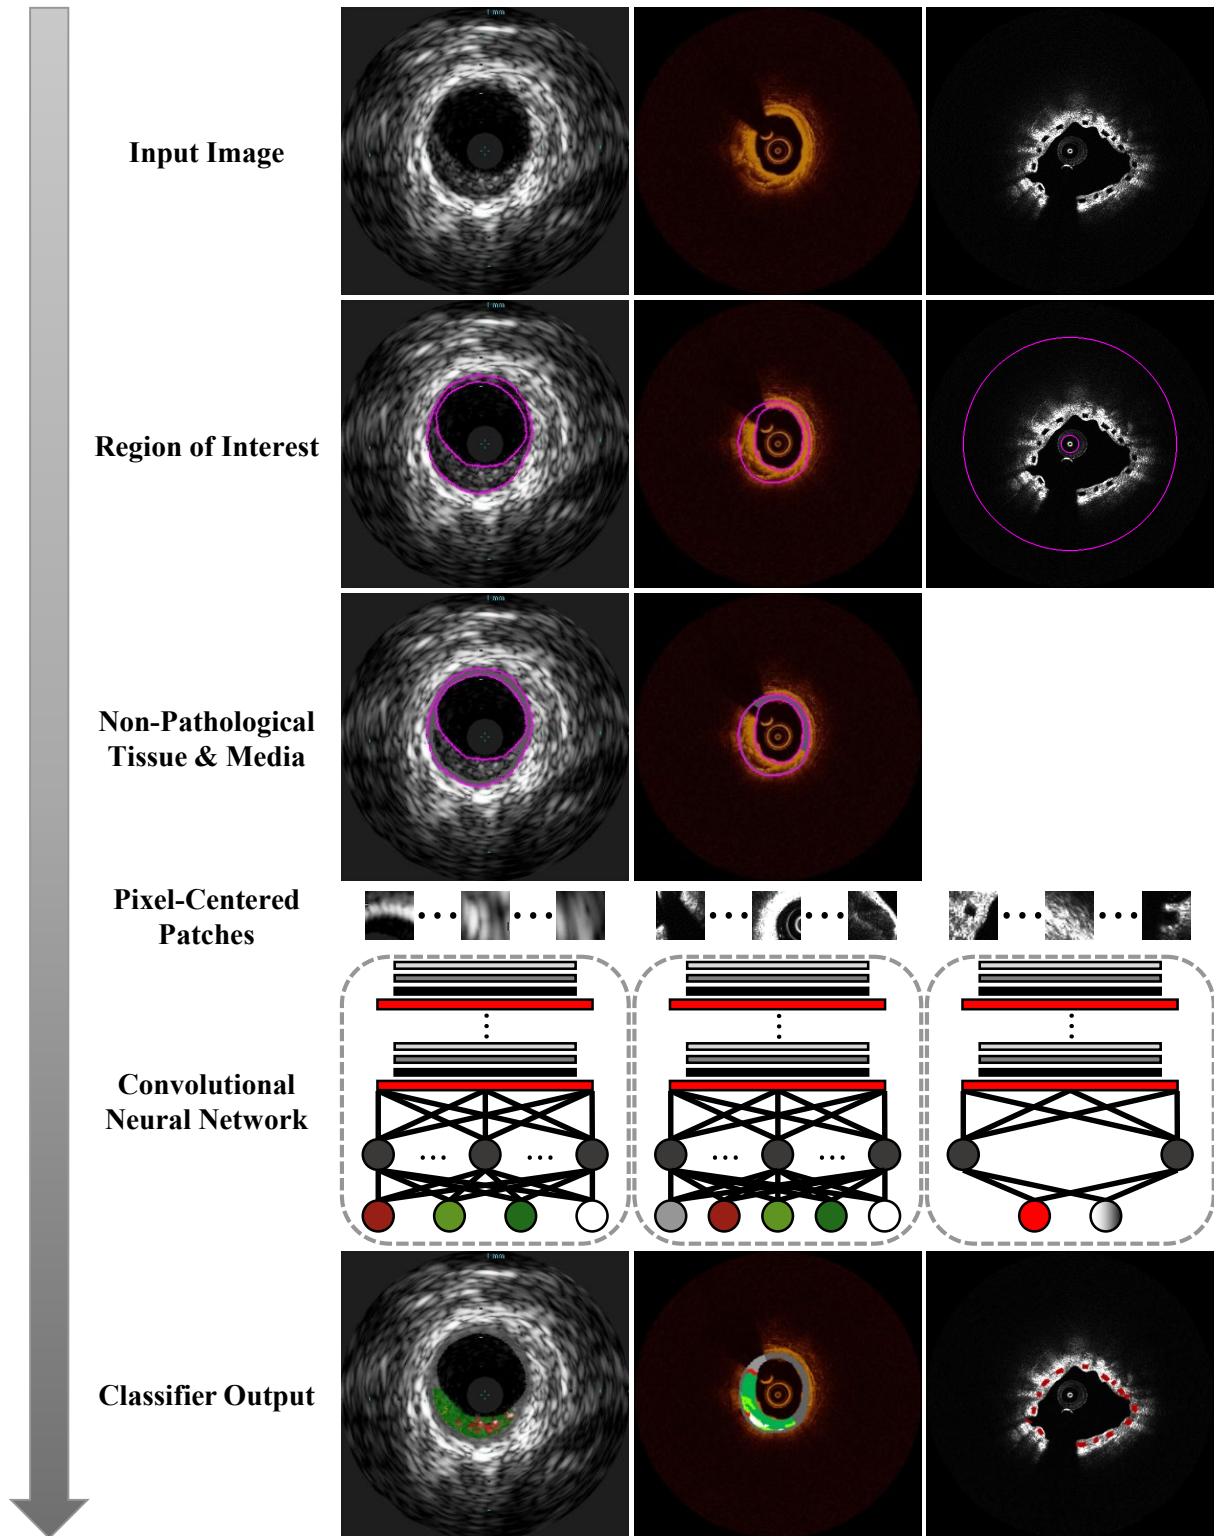


Figure 3.2. Flowchart of the characterization framework enriched by domain knowledge. The method determines inner and outer borders of the ROI (magenta) for a given input image. For tissue characterization, the method then segments the non-pathological tissue and media (dark gray) based upon vascular physiology and pathology constraints. The remaining ROI is input to the CNN, which segments the pathological vessel wall (intima) into various tissue classes (white, green, and red) and shadow artifact (light gray; OCT only) or segments scaffolding.

Segmenting indwelling devices, such as a bioresorbable vascular scaffold, was achieved with only minor adaptation of the framework (Figure 3.2, right). Rather than specifying inner and outer borders of the vessel wall to delineate the ROI, a large standardized area of the image was instead selected to include scaffold struts residing superficially within the lumen and those deep within the vessel wall. Furthermore, pathological tissue was not partitioned and treated differently from the other tissue within the ROI. Finally, rather than classifying a pixel into one of several tissue classes, the pixel was assigned to one of two binary classes indicating presence or absence of a strut.

The framework was applied and refined for both OCT and IVUS. The general principles of the fundamental methodological framework are first presented, then the details of the three unique use cases are described. In each application, the method was implemented in MATLAB (MathWorks, Natick, MA) using the Deep Learning Toolbox running on a NVIDIA TITAN Xp graphics processing unit (GPU; PG611) with 12 GB RAM.

3.2.1. Region of Interest Delineation

ROI segmentation is a prerequisite for subsequent methodological steps (Figure 3.2), though succeeding procedures are agnostic to the means of segmentation (of which there are a large and growing number). For plaque characterization, the ROI was defined as the vessel wall area between the lumen and the media-adventitia transition (i.e. between the inner and outer border of the vessel wall; Figure 3.3)—the region in which atherosclerotic plaques develop. Therefore, within each image $I(i, j)$, the inner (lumen) border b_i and outer (media-adventitia) border b_o fully delineate the ROI (intima and media region) $r_{im}(i_{r_{im}}, j_{r_{im}})$. For the task of indwelling device (i.e. BVS strut) segmentation, the ROI was defined as a fixed area of the image in which the objects may reasonably be expected to be found. Because such features may be present either within the lumen or embedded in the vessel wall, and because the classifier output was a binary present/absent designation independent of tissue presence, this conservative approach to delimiting the ROI for analysis was a reasonable choice.

3.2.2. Non-Pathological Tissue Detection

The presented methods focused on the evaluation of vessel wall morphology and the characterization of its phenotype, distinguishing not only plaque subtype but normal from pathological tissue, as well as non-tissue (device). In extracting plaque morphology, once the

vessel wall is delineated and ROI established, the non-pathological tissue and media layer are determined. This concept has similarly been implemented in VH-IVUS, where each tissue type is highlighted as a specific color and the media portrayed in gray along the rim of the vessel wall (Figure 3.1). In the presented framework, dimensional limits were imposed, leveraging expert recommendations for interpreting intravascular images; intima was deemed normal of its thickness was $<360 \mu\text{m}$ [18], [235], and the media was assumed to have nominal thickness of 250-350 μm [8], [164], [277]. Thus, the location and thickness of non-diseased and media tissue was defined such that wall regions thinner than threshold were not to be considered diseased or analyzed as such, and the media layer approximated by a band of constant thickness around the outer edge of the ROI.

To determine the normal wall and the media layer locations and dimensions, two geometrical parameters were computed for each pixel in the ROI (r_{im}):

$$D_{thick}(i_{r_{im}}, j_{r_{im}}) = D_1(i_{r_{im}}, j_{r_{im}}) + D_2(i_{r_{im}}, j_{r_{im}}), \text{ and} \quad (51)$$

$$D_{outer}(i_{r_{im}}, j_{r_{im}}) = D_1(i_{r_{im}}, j_{r_{im}}), \quad (52)$$

where D_1 and D_2 are the Euclidian distances of the pixel $(i_{r_{im}}, j_{r_{im}}) \in r_{im}$ from the outer (media-adventitia) border b_o and the inner (lumen) border b_i , respectively (Figure 3.3). Threshold values for D_{thick} and D_{outer} were calculated to determine whether a pixel was in a section of sufficient thickness (Th_{path}) to be considered pathological or sufficiently close (Th_{media}) to the media-adventitia border to lie within the media. Pixels of the ROI were classified as pathological tissue (ROI_{path}) if $D_{outer} \geq Th_{media}$ and $D_{thick} \geq Th_{path}$ (Figure 3.3). The remaining pixels of the ROI were assigned to the non-pathological tissue or media (M) class.

This pathological tissue detection procedure was the primary mechanism by which domain knowledge enriched learning to address the image classification problem. Following this step, classification for tissue characterization was only required for the four remaining tissue types (and guidewire shadow, for OCT).




Figure 3.3. Schematic presentation of the pathological tissue segmentation, shown for a representative IVUS frame. Given the inner (b_i) and outer (b_o) borders (*top left*), Euclidean distances from a pixel ($r_{im} \in \text{ROI}$) to the lumen border (D_2 ; black arrow) and the media-adventitia border (D_1 ; black arrow) were calculated (*bottom*). Pixels within the ROI for which $D_{outer} < Th_{media}$ and $D_{thick} < Th_{path}$ correspond to media and non-pathological tissue, respectively (*right, inset; white*). Other pixels within the ROI correspond to pathological tissue (ROI_{path} ; *top left, highlighted*). Color in distance maps indicates relative magnitude of values (blue: small, red: large).

3.2.3. CNN-Based Classification

Pixel-centered patches were created for those pixels to be classified—all pixels within the full ROI for strut segmentation (i.e. $r_{im} \in \text{ROI}$), and for tissue characterization those pixels of the

ROI not belonging to the non-pathological tissue or media (i.e. $r_{im} \in \text{ROI}_{path}$). These patches were then automatically classified using a CNN.

CNNs are a class of deep neural networks [279] commonly applied to image classification because they can leverage spatial locality and translational invariance to dramatically reduce the number of weighted network connections requiring optimization (cf. FCNNs). Their architecture can be described by multiple layers, which can be categorized as input, output, or hidden: the input layer received the 2D (grayscale) image patch; the hidden layers were formed by multiple functional layers combined in series in which the compound image features were calculated and strategically pooled; and the output layer produced the classification result. To classify the pixels, the functional layers executed a sequence of convolutions, batch normalizations, rectified linear unit (ReLU) activations, and pooling operations (Figure 3.4). To achieve the best classification results, different patch sizes, numbers of convolution sequences, filters, and filter sizes were tested.

The transformation of data executed by a such a CNN (Figure 3.4) can be represented by a non-linear function, $P(I; \Theta) = p$, which maps an image $I \in \mathbb{R}^{H \times H}$ of $H \times H$ size to a vector $p = (p_1, p_2, \dots, p_c)^T$. The probability of I belonging to one of the c target classes $i = \{1, \dots, c\}$ is represented by $p_i \in [0,1]$, and $\Theta = \{\Theta_1, \Theta_2, \dots, \Theta_K\}$ are the K parameters (weights and biases) used to map I to p . CNN training is an optimization problem for a non-linear function with K degrees of freedom:

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \mathcal{L}_{\{I^{(1)}, I^{(2)}, \dots, I^{(N_{train})}\}}(\Theta), \quad (53)$$

where $\mathcal{L}(\Theta) \in [0,1]$ is a loss function and N_{train} is the number of training images.

In this framework, (multiclass) cross-entropy loss (also known as negative log likelihood), the most popular choice for probabilistic classification problems, was used:

$$\mathcal{L}_{\{I^{(1)}, \dots, I^{(N_{train})}\}}(\Theta) = - \sum_{n=1}^{N_{train}} \sum_{i=1}^c y_i^{(n)} p_i^{(n)}. \quad (54)$$

This loss function measures the performance of the classifier P relative to the binary class label vector y_i for each image $I^{(n)}$ with the given set of parameters Θ .




Figure 3.4. Progressive data processing performed by the 26-layer CNN to classify pixels within the pathological region of interest in an IVUS image. The square $H \times H \times 1$ image patch is iteratively convolved, normalized, operated on by an activation function, and pooled, resulting in the intermediate data structures of decreasing height and width but increasing depth. Each layer in depth (channel) is the result of a convolution of the previous array with a single kernel, or filter. Therefore, with each successive sequence, more—and more complex—features across the input image can be identified. Weights of the filter kernels are learned through optimization during training as discriminable features of the various output classes. On the basis of these features, a final array of probabilities, indicative of a probability mass function, is generated to reflect the likelihood that the input belongs to each of the possible classes. These operations constitute the CNN module of the characterization framework (Figure 3.2).

To reduce the training time for the CNN, the stochastic gradient descent (SGD) iterative method was used. This method approximates the dataset with a subset of samples randomly-drawn from the full training dataset, called a mini-batch, and uses the gradient calculated for the mini-batch to update the model in each iteration. SGD is known to sometimes oscillate along the path of steepest descent (maximum gradient) towards the optimum, rather than directly along the path toward the optimum, since the gradient always points towards the opposite side of this optimum from the current position. A solution to this problem is the addition of a momentum term to the parameter update to reduce oscillations and aberrant sensitivity to noise and intermittent fluctuations [280]:

$$\Theta_{l+1} = \Theta_l - \alpha \nabla \mathcal{L}(\Theta_l) + \gamma(\Theta_l - \Theta_{l-1}), \quad (55)$$

where l is the iteration number, $\alpha > 0$ is the learning rate, and the momentum term γ determines the contribution of the previous gradient step to the current iteration. Thus, the SGD algorithm selects a subset of the training set \mathcal{D}^{train} , evaluates the mean gradient of the loss function \mathcal{L} for this mini-batch, then updates the network parameters Θ . Each evaluation is an iteration, and at each iteration the loss function is minimized further. The full pass of the training process over the whole training set, in mini-batch increments, forms an epoch.

In training the network described herein, a SGD with momentum optimizer was implemented with standard generic parameters—a constant learning rate (α) of 0.03 and momentum value (γ) of 0.9. A mini-batch size of 3,000 patches was utilized; data were shuffled after each epoch. Weight decay (L₂ regularization) by a factor of 10^{-4} was used to reduce overfitting. Weights were initialized with a Glorot initializer, which independently samples from a uniform distribution centered around zero; biases were initialized to zero.

3.3. Characterizing Tissue in Grayscale IVUS

The framework presented above was applied to establish a novel CNN-based domain enriched method that classified arterial tissue imaged through IVUS [161]. The method detected the ROI using recently developed software [281], and then subdivided the ROI into pathological and non-pathological tissue based upon basic spatial and geometric constraints informed by physiology. Pathological areas of the ROI were partitioned into patches and fed through a CNN architecture. Corresponding VH-IVUS images served as the comparative control. The method

offered several meaningful benefits stemming from its independence from the RF signal data, which increased the clinical utility and research applicability of the method. In particular, the method could be applied to grayscale IVUS data, including previously-acquired images that had not been characterized by the VH technique due to a lack of RF signal or proprietary software, or to intermediate frames of VH-IVUS acquisitions between ECG-gated frames, thereby increasing the effective rate at which meaningful information on plaque morphology could be attained and reducing procedure time.

3.3.1. Materials and Methods

The methods implemented to characterize tissue in grayscale IVUS images followed the fundamental methodological framework described in the previous section (Figure 3.2, left).

The region between the lumen border and the media-adventitia border where atherosclerotic plaques develop was denoted as the ROI. To detect the ROI in each frame, a previously validated method [250] recently incorporated into a user-friendly software suite [281] was utilized. In brief, initial contours for the lumen and media-adventitia borders were estimated using basic image processing: the image was binarized using Otsu's automatic thresholding algorithm [282], and the tentative borders were found by scanning radial projections for binary state transitions. The method subsequently refined the borders using active contour models [283].

As noted before, the method focused on the evaluation of vessel wall morphology and the characterization of its phenotype, distinguishing not only plaque subtype but normal from pathological tissue (similar to VH-IVUS, Figure 3.1). Physical and dimensional limits leveraging expert recommendations were imposed; intima was deemed normal if its thickness was $<360\ \mu\text{m}$, and the media was assumed to have nominal thickness of $250\text{-}350\ \mu\text{m}$ [8], [164], [277]. Thus, the location and thickness of non-diseased and media tissue were defined such that wall regions thinner than threshold were not to be considered diseased or analyzed as such, and the media layer approximated by a band of constant thickness around the outer edge of the ROI. Though media thickness does vary somewhat, its range is largely negligible relative to that of the inner intima layer, and is furthermore at the horizon of VH-IVUS imaging resolution ($100\text{-}200\ \mu\text{m}$) [19], [234], [277].

Non-pathological wall and the media layer locations and dimensions, and consequently pathological and media thickness thresholds, were established on the basis of observed tissue

distributions in the training dataset. Threshold values were calculated as the optimal cutoffs to discriminate whether a given pixel was within an area deemed pathological. To determine these optimal cutoff values, D_{thick} (Equation 51) and D_{outer} (Equation 52) were calculated for each pixel in the ROIs of all N_{train} VH-IVUS images from the dataset. All ROI pixels labeled by VH as belonging to the media or non-pathological class (r_{im}^M) were considered. The pathological thickness threshold (Th_{path}) was calculated as the maximum r_{im}^M section thickness (D_{thick}) immediately adjacent to the inner border (b_i), which is also equivalent to the maximum distance of a pixel belonging to this class from the outer border (D_{outer}):

$$Th_{path} = \max_{N_{train}} \left(D_{thick}^{r_{im}^M \in b_i} \right) = \max_{N_{train}} \left(D_{outer}^{r_{im}^M} \right). \quad (56)$$

The maximum media thickness threshold was calculated as the mean maximum thickness of r_{im}^M sections in which pathological tissue is present (i.e. $D_{thick} \geq Th_{path}$):

$$Th_{media} = \frac{1}{N_{train}} \sum_1^{N_{train}} \max \left(D_{outer}^{r_{im}^M (D_{thick} \geq Th_{path})} \right). \quad (57)$$

Th_{path} was found to be 30 pixels, and Th_{media} was 11 pixels. Therefore, pixels of the ROI were subsequently classified by the method as pathological tissue (ROI_{path}) if D_{thick} (local wall thickness) exceeded 30 pixels and D_{outer} (distance from the outer limit of the wall) exceeded 11 pixels (Figure 3.3).

Pixel-centered patches were created for those pixels of the ROI which remained after segmenting the non-pathological or media class (i.e. $r_{im} \in ROI_{path}$). These patches were automatically classified into one of the four VH-IVUS plaque types using a CNN. To achieve the best classification results, different patch sizes, numbers of patch convolution sequences, filters, and filter sizes were tested. A patch size of 41×41 was determined to perform best through parameter sensitivity analysis (Figure 3.10). The network found to perform best, and utilized in this work, is shown in Figures 3.4 and 3.5.

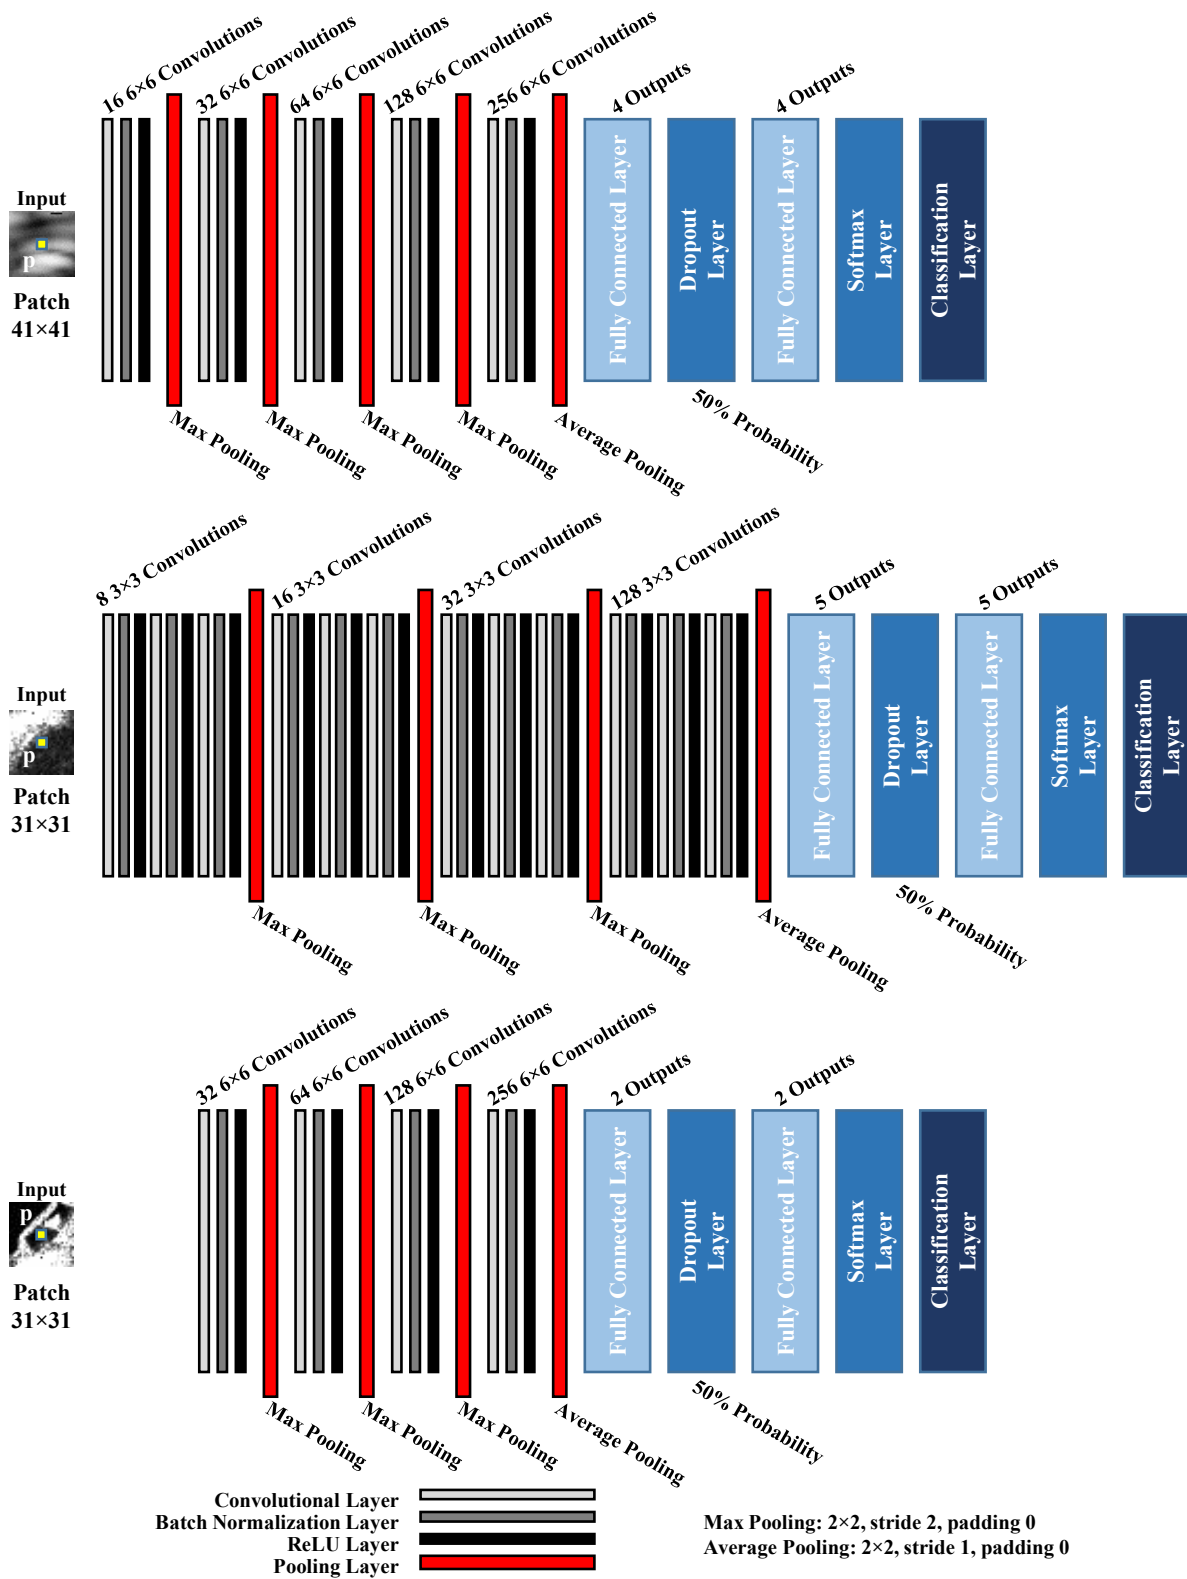


Figure 3.5. Detailed architecture of the multi-layer CNNs used to classify pixels within the region of interest for tissue characterization in IVUS (*top*) and OCT (*middle*), and BVS strut segmentation in OCT (*bottom*). See Figure 3.4 for representations of intermediary data structures and activations of the hidden layers for the IVUS classifier.

3.3.2. Dataset

To train and test the IVUS plaque characterization algorithm, 553 VH-IVUS frames and the corresponding grayscale IVUS frames were acquired from eight patients. The data were acquired at 20 MHz using a 3.5 F electronic probe with synthetic aperture (Eagle Eye Gold Catheter, Philips Healthcare, Andover, MA), in accordance with clinical standards [19], [277]. As intravascular images were collected from a clinical population in the course of treatment, the data were inherently inhomogeneous and imbalanced. Figure 3.6 shows the average quantity and relative distribution of tissue types present in the imaged vessel walls of each patient, as well as the number of VH-IVUS frames available from each patient’s acquisition); the patients represented a spectrum of lesion morphologies and sizes. Though sometimes present in small numbers, all five tissue types were detected by VH-IVUS in each patient (though notably not in each frame).




Figure 3.6. Plaque distribution of patients comprising IVUS dataset. *Top:* Prominence of plaque and composition for each patient, reported in average number of pixels per frame corresponding to each characterized tissue type as determined by VH-IVUS. *Bottom:* Number of VH-IVUS frames available for each patient.

From the dataset, 200 frames were withheld exclusively for testing while the remaining frames were sampled for training and validation. From the 353 frames used for training and validation, 3.4×10^5 41×41 pixel patches of each class were extracted and augmented 7-fold (through reflection and rotation in 90° increments) for use in training the full networks. For the pixel sensitivity study, a smaller sub-set of 10^4 patches of each class was selected to accelerate training; ten-fold cross-validation was performed. From the 200 frames withheld for testing, 5×10^4 patches of each class were randomly selected from bulk regions of tissue and used for final testing and evaluation. For the ablation study, a smaller sub-set of 10^4 patches of each class was randomly selected to accelerate classification and analysis.

While patches used in training were sampled from all labeled training data to avoid potential loss of signal, in sampling the labeled testing data, regions immediately adjacent to boundaries between tissue types and edges of tissue regions were avoided. Doing so limited uncertainties arising from VH resolution and image data degradation due to file compression; enhanced label certainty for inclusion in testing was also achieved by requiring agreement between two automated color-based methods used for determining categorical labels from colored VH images (in which color-coded labels have been integrated with the underlying grayscale image). Furthermore, avoiding transition zones and borders ensured higher certainty by targeting to a greater extent the “bulk” of a tissue region—VH-IVUS validation has been conducted by histology and atherectomy [233], [235], [274]–[276], which allows comparison of bulk tissue type identification rather than pixel-by-pixel comparison. Therefore, an advantage of this more restrictive approach to sampling data for testing against VH-IVUS is that it more reliably assesses accuracy of the novel characterization method with respect to underlying tissue (by transitive considerations). A sample frame from the testing subset, alongside the corresponding “bulk” regions and actual randomly-selected pixels used in quantifying test performance, is shown in Figures 3.7.




Figure 3.7. Selection of testing data from a typical frame withheld from training. (a) Sample VH-IVUS frame from the testing subset. (b) Masked region showing regions of bulk tissue (avoiding boundaries between tissue types and edges of tissue regions). (c) Specific pixels included in the balanced testing dataset (selected from the bulk tissue region shown in (b)). Note how, due to the relative scarcity of dense calcium and necrotic core, a greater portion of each calcified and necrotic region is sampled in constructing the balanced set used to quantify the test performance of the networks.

3.3.3. Results

Overall Performance

Image segmentation accurately replicating VH-IVUS classification was successfully achieved using only grayscale IVUS images. Table 3.2 provides the error (or confusion) matrix for the presented enriched method, showing that it achieved an overall accuracy of 93.5%. It can be appreciated in a representative example of a classified image (Figure 3.8) and additional detailed regions (Figure 3.27) that the method accurately captured major tissue morphology and features within the pathological region. It captured fine features and provided sharp distinctions between various plaque types; it generated images that were very similar to gold standard VH-IVUS. Furthermore, due to the spatial constraints imposed prior to CNN classification, the method addressed non-pathological and media tissue very accurately, and was not disposed to violating physiological constraints.

Table 3.2. IVUS characterization method error matrix

		Target Class					<i>Precision</i>
		DC	NC	FT	FFT	M	
<i>Output Class</i>	DC	49247	4664	0	0	3	91.3%
	NC	664	44310	1208	0	30	95.9%
	FT	0	996	45564	4761	159	88.5%
	FFT	0	0	3107	45015	84	93.4%
	M	89	30	121	224	49724	99.1%
<i>Recall</i>		98.5%	88.6%	91.1%	90.0%	99.4%	93.5%




Figure 3.8. Representative classified IVUS image segmented by VH-IVUS (ground truth) and the CNN-based method. Both methods identify the same major pathological tissue morphology features. The latter provides sharp distinctions between various plaque types and consequently captures fine features, and is similar to VH-IVUS.

The overall performance metrics for the domain enriched CNN-based method (Table 3.2) depend on (four-class) classifier performance and reliability of pathological tissue detection, which together share responsibility for the full segmentation procedure. The CNN classifier achieved generally high precision (i.e. PPV) and recall (i.e. sensitivity). Tables 3.3 and 3.4 show the error matrix and mean predicted class scores for the enriched method’s four-class CNN classifier when presented with the testing set. The model, trained on just pathological tissue classified by VH-IVUS, achieved an accuracy of 92.3%. Progression of training for this network, which took several weeks, is shown in Figure 3.9. Training was halted after 50 epochs, at which point accuracy and

Table 3.3. IVUS pathological tissue classifier error matrix

Output Class	Target Class					Precision	M
	DC	NC	FT	FFT			
DC	49336	4664	0	0	0	91.4%	858
NC	664	44340	1208	0	0	95.9%	12320
FT	0	996	45657	4771	0	88.8%	24893
FFT	0	0	3135	45229	0	93.5%	11929
Recall	98.7%	88.7%	91.3%	90.5%		92.3%	-

Table 3.4. IVUS pathological tissue classifier mean predicted class score

Output Class	Target Class					Mean
	DC	NC	FT	FFT		
DC	0.9855	0.1779	0.0023	0.0000	0	0.2914
NC	0.0145	0.7969	0.0271	0.0000	0	0.2096
FT	0.0000	0.0204	0.7158	0.1012	0	0.2093
FFT	0.0000	0.0048	0.2549	0.8988	0	0.2896
Mean	0.2500	0.2500	0.2500	0.2500	0	0.2500

loss had plateaued (Figure 3.9); with further training, validation metrics began deteriorating, indicating overfitting of the model to training data.




Figure 3.9. CNN validation accuracy (*top*) and loss (*bottom*) through 50 epochs of training for the enriched model. Accuracy and loss plateaued by around the 50th epoch; deviation between training and validation accuracies—indicating overfitting—became apparent soon afterwards when training was allowed to proceed beyond 50 epochs.

The error matrix of the classifier illustrated some general trends. The four-class network supporting the domain enriched method struggled to differentiate FFT from FT and, unexpectedly, DC from NC. Furthermore, the enriched model was subject to compounding uncertainties arising from pathological tissue delineation. While delineation of pathological tissue, as defined by VH-IVUS, was very accurate, the CNN of the enriched method was incapable of classifying M tissue it encountered (and typically identified it as FT; Tables 3.2 and 3.3).

Performance of the enriched deep learning method described here is compared with that reported for existing methods in Table 3.1, which also includes validation metrics reported for VH-IVUS. Included methods include those addressed and described briefly in the introduction: Zhang *et al.* [264], Brunenberg *et al.* [265], Athanasiou *et al.* [266], Taki *et al.* (2010) [268], Taki *et al.* (2013) [269], Athanasiou *et al.* [270], Hwang *et al.* [271], and Kim *et al.* [272]. Performance of VH-IVUS is reported for results from an *ex vivo* validation study carried out by Nair *et al.* [275].

Speed and Efficiency

Execution time of the characterization method was dominated by the pixel-wise network classification of the ROI. Each patch took 7.4 ± 0.4 milliseconds (mean \pm standard deviation) to classify, though this value was found to be sensitive to the machine on which classification was performed. Each ROI contained 37801 ± 22455 pixels, of which the enriched method determined that 26776 ± 20805 pixels were pathological and subsequently classified by the network (at the above rate of approximately 7.4 milliseconds per pixel). Calculation of D_1 and D_2 , and subsequent designation of the media and non-pathological tissue in a frame, took just 25.5 ± 0.9 milliseconds per frame. Because the ROI delineation method is considered interchangeable for this method, execution time of this step was not determined, but several methods report execution times significantly less than 1 second per frame [180], [250], [251], [254]. Consequently, characterization of full frames took 200 ± 150 seconds with the method. The range of execution times corresponds to the drastic variability in plaque content between frames; while segments with high plaque burden took several minutes to characterize, frames depicting cross-sections without diseased tissue (just media and/or non-pathological tissue) took just a fraction of a second to characterize. However, per-frame characterization time is reported here for a scenario in which every individual pixel of the ROI is characterized, rather than a strategically selected subset, and furthermore neither software nor hardware were optimized for execution time. As such, these times should be interpreted as an upper bound.

Patch Size Sensitivity Test

The CNN takes as input not a full image, but rather a patch, or neighborhood of pixels surrounding the pixel to be classified. The size of the patch can impact the subsequent performance of the trained network—if a patch is too small, it may not capture sufficient context or information for a robust classification to be made, but if a patch is too large, additional parameters (θ) must be learned, increasing solution space complexity, training time, and data and network storage size, among other potential pitfalls. Therefore, a preliminary study of patch size was performed to assess performance achieved by networks utilizing square neighborhoods of various sizes ranging from 33 to 71 pixels in dimension (1,089 to 5,041 total pixels).

Ten-fold cross-validation was performed for each patch size using a relatively small dataset (4×10^4 patches equally distributed among the four tissue types). That is, the dataset was randomly split into ten equally-sized balanced subsets, and the training procedure was repeated ten times; each time, a different one of those subsets was withheld from training and used to evaluate the final performance of the network trained on the other nine. The same network architecture was used for all 12 patch sizes with the exception of the input layer, which was modified to accept the desired patch size; the number of inputs to the first fully-connected layer differed accordingly. Training time was recorded for a single training run executed on the same hardware in consistent conditions for all 12 patch sizes. (Due to the amount of training required, tasks were otherwise distributed across different hardware and conditions, not allowing for direct comparison. A range of times is therefore not reported.)

Results of the sensitivity study indicated that the overall domain enriched deep learning approach and method to classify atherosclerosis using IVUS wasn't largely dependent on the patch size used (Figure 3.10). However, a patch size of 41×41 was shown to offer desirable performance. In particular, this patch size had among the highest average test accuracies, lowest standard deviation in accuracy, and lowest training times of those patch sizes assessed. Though not the best performing by any of those metrics alone, taken together this patch size demonstrated the most desirable holistic performance profile. Furthermore, the smallest patch size yielding acceptable performance is generally preferred, as increasing patch size, and consequently the number of trained network weights, increases not only training and execution time, but also susceptibility to overfitting. Here, decreasing patch size further diminished performance (lowered accuracy and increased performance variation in cross-validation). However, while helping to select a well-




Figure 3.10. Results of the sensitivity study demonstrate the effect of patch size on a small dataset. The averages and standard deviations of final test accuracies generated by ten-fold cross-validation are plotted as a solid line (left vertical axis). Time to execute the training process once is plotted as a dashed line (right vertical axis) trending upwards with patch size. Patch size as listed corresponds to a single dimension (i.e. the square root of the total number of pixels contained in a given patch of that size). The diamonds call out the performance of the network utilizing 41×41 patches like that used in the work presented herein. It can be seen that this patch size was among the best performing with regards to average final testing accuracy (high), standard deviation in final testing accuracy (low), and training time length (low).

performing patch size, the study also provided confidence that the overall framework was robust, and that the quality of the data and fidelity with regards to the underlying population from which the set is drawn is more important than the specific details of patch size and similar network design features.

Ablation Study

To interrogate the organization and relationships within the network structure, and assess the distribution and similarity of class-associated features, an ablation study was carried out on the 4-class CNN classifier of the domain enriched method. In turn, each of the 504 nodes or channels was “ablated” by negating a filter’s weights and associated bias term (in a convolutional layer) or all input connection weights and bias term (in a fully-connected layer), thereby depriving the node or channel of input and inhibiting its function by mandating an output of zero. Systematically “ablating” small segments of the network in this way probed where and how “knowledge” was represented in the network, how robust and redundant or centralized the network was, and how important certain parts of the network were to overall performance.

Two major analyses were performed to investigate characteristics of the network: the impact of individual channels and nodes on final class sensitivity (i.e. recall) and the relationship between changes in class sensitivity caused by the various ablations. The results of the first analysis are shown in Figure 3.11, which illustrates the absolute drop in class sensitivity resulting from ablation of channels and nodes, arranged by layer. On the left side of this figure, each horizontal line represents a channel or node of the corresponding layer which was ablated, with colors indicating the extent and direction of change in sensitivity for each class resulting from its ablation. Note that each convolutional layer has increasing numbers of channels (16, 32, 64, 128, and 256) while the fully connected layers have 4 nodes apiece (Figure 3.5)—each corresponding to a horizontal line in Figure 3.11. Changes are also compiled for entire layers to examine importance of network layer in the outcome for each class, as shown on the right side of the figure; note that the vertical axes (scales) vary for the boxplots of each layer.

The ablation of most functional units had a very small effect on end output, suggesting the network was fairly robust in that classification was generally distributed. A few notable exceptions existed, however. First, a few units played an outsized role in the successful classification of FFT pixels, particularly in the first convolutional layer; when these units were ablated, sensitivity for FFT dropped by up to 61%p; the output of these early channels were important to the downstream processing performed in later layers to positively identify pixels showing FFT tissue. The other units with a major role were the nodes of the second fully-connected layer, the output of which entered the softmax and subsequently classification layers. As the layer whose outputs were directly compared to establish class assignment, it was expected that each node would have profound importance to the classification of a single corresponding tissue type, as observed. An interesting observation is the impact on the three other classes; ablation of the node corresponding to DC resulted in an improvement in NC sensitivity of just over 10%p. Ablation of the nodes corresponding to FT and FFT similarly resulted in more modest improvements in sensitivity for the other class. One interpretation is that these units not only contributed to classification of one class, but effectively suppressed the others. While most clear in the final fully-connected layer, examples of competing performance driving confusion between DC and NC and FT and FFT were visible throughout the network. For both pairs, an ablation causing a decrease in sensitivity for one class in the pair was often accompanied by an increase in the sensitivity of its counterpart. These relationships and trends emerged, and were explored more explicitly, in Figure 3.12.

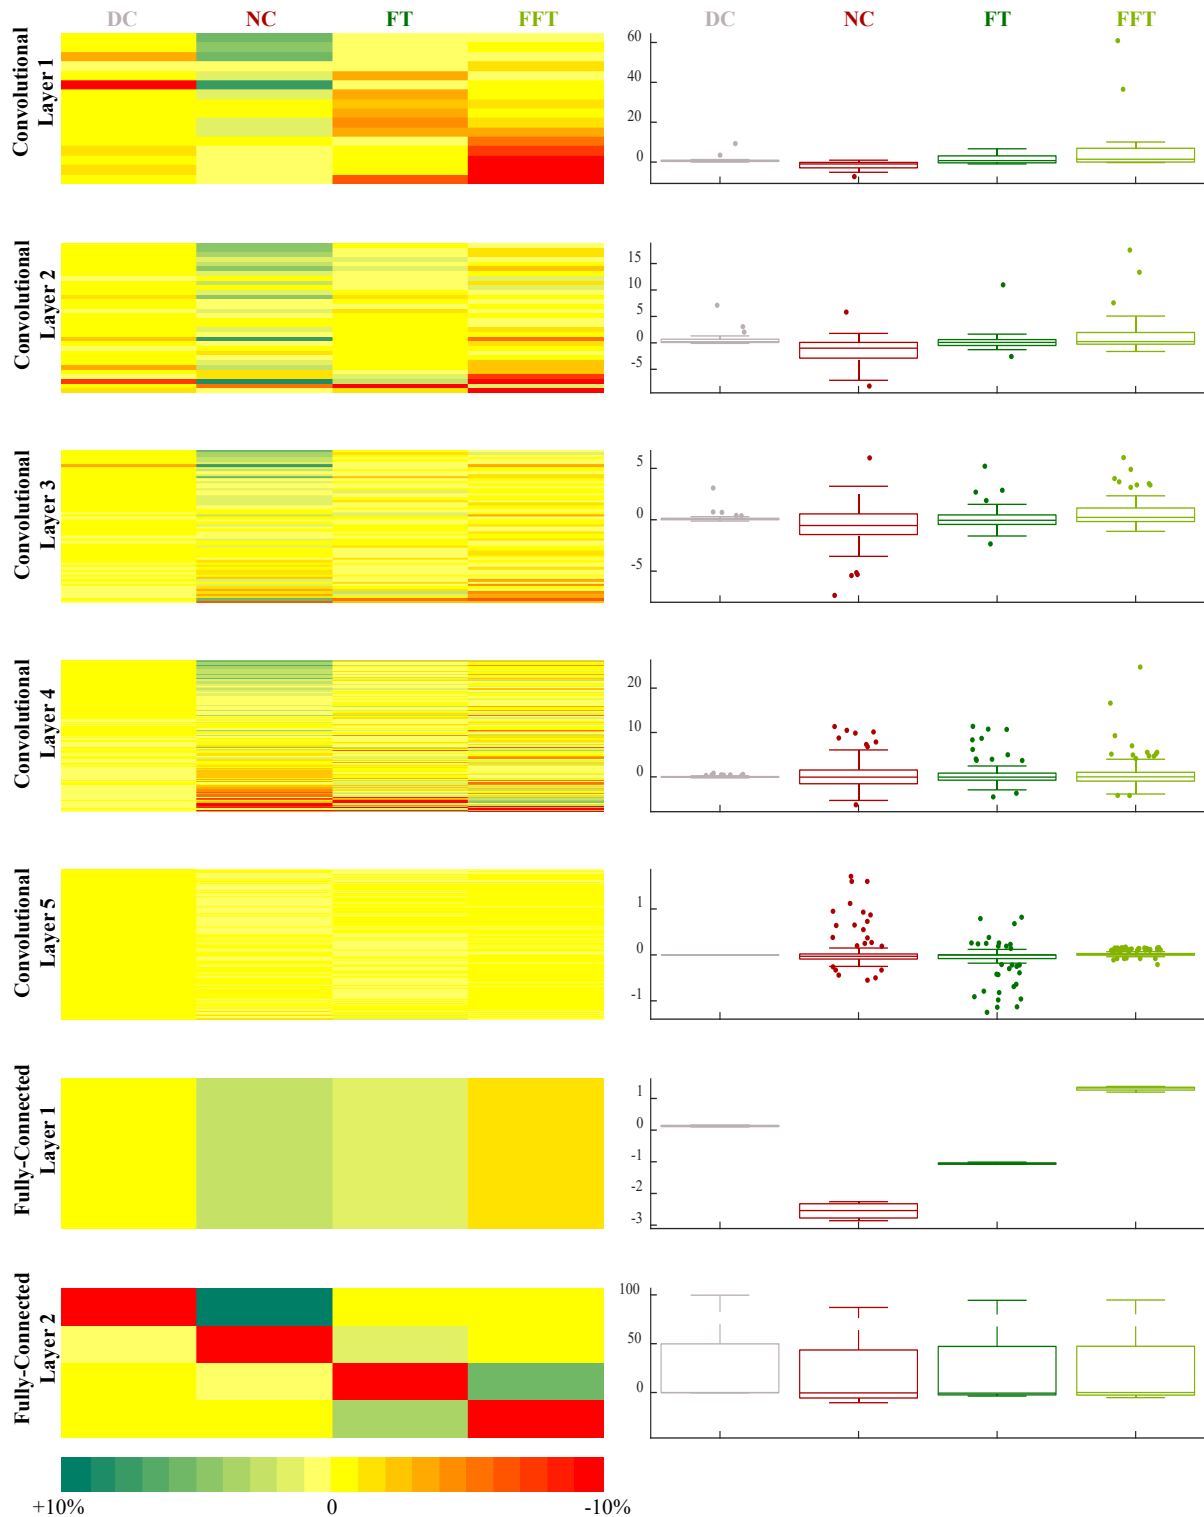


Figure 3.11. Results of the neural network ablation study illustrate distributed responsibility for classification. The absolute drop in class sensitivity resulting from ablation of channels and nodes showed that most individual ablations do not drastically impact overall performance. *Left:* Each horizontal line represents an ablated channel or node; colors indicate change in sensitivity for each class resulting from its ablation. (The color scale does not span the entire observed range.) *Right:* Changes are compiled for entire layers, showing importance of network layer in the outcome for each class. Boxplots illustrate drop (in percentage points) in class sensitivity; axes differ.

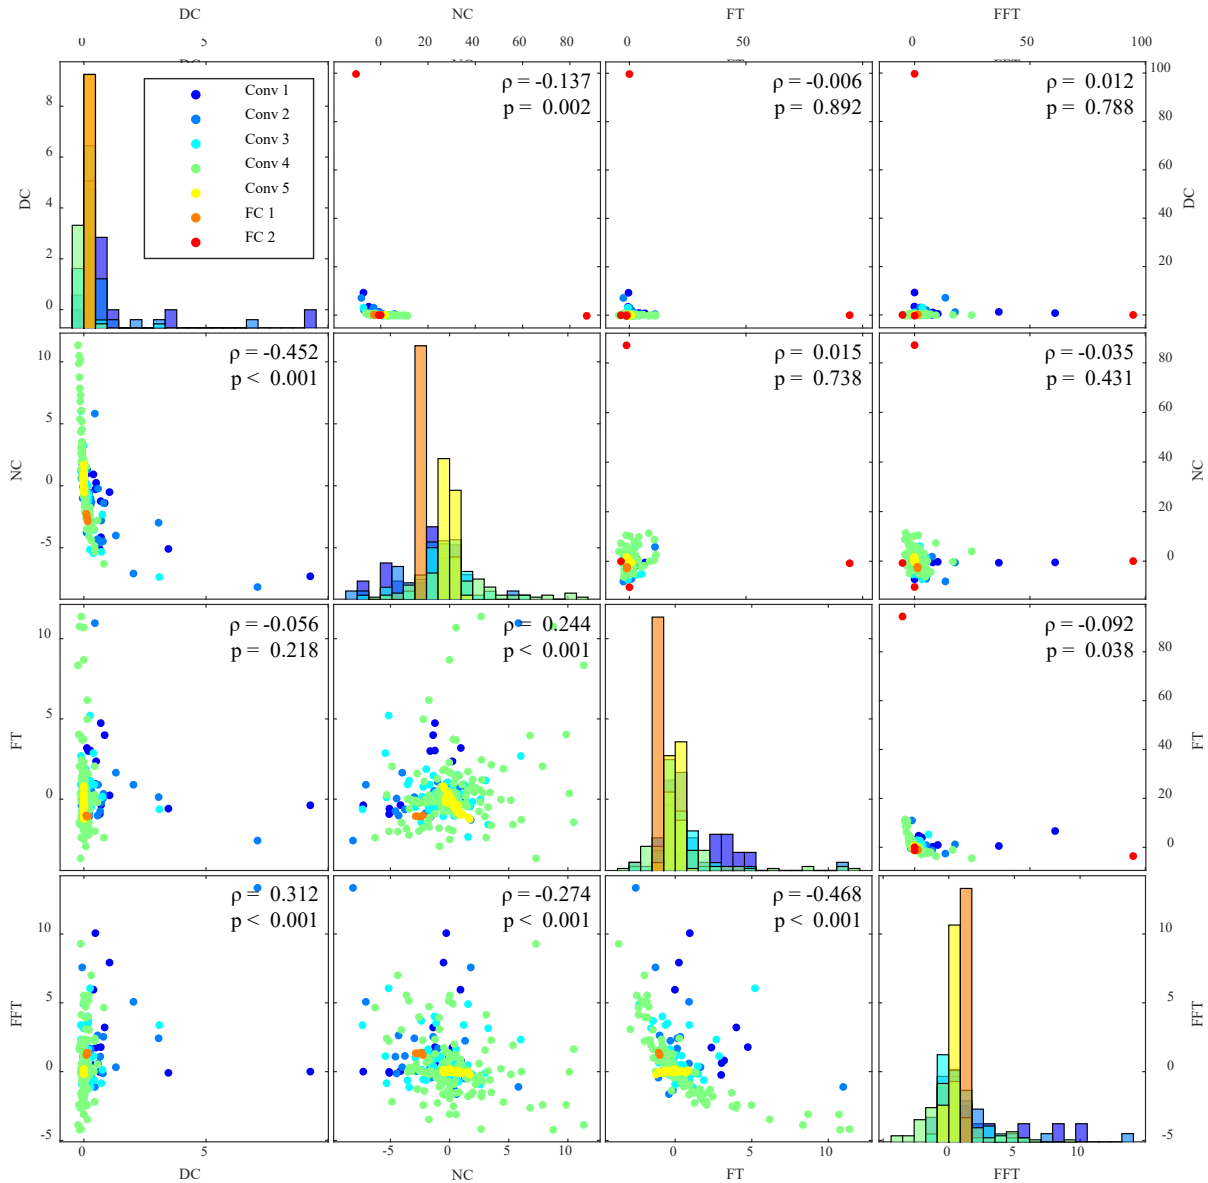


Figure 3.12. Results of neural network ablation study suggested that network representations and classification pathways of DC and NC, as well as FT and FFT, were closely related. Each point represents a single unit ablation, with position indicating drop in class sensitivity (in percentage points); as such, each point corresponds to a single horizontal line in Figure 3.11. Color indicates the layer to which the ablated unit belonged. Plots mirrored across the diagonal show the same information in transpose but with different axis ranges; plots on the upper right include 9 outlier points (maximum change resulting from ablation greater than 3 standard deviations from the mean) that are excluded from those on the bottom left (as well as the histograms). Note that the scales differ for the plots above and below the diagonal. Pearson's linear correlation coefficient (ρ) and p-value (p) are displayed for each set of observed changes. (Values below the diagonal are calculated excluding the 9 outliers while those above the diagonal include all 504 points.) The p-values less than the displayed precision are as follows: $p_{DC-NC} = 2.67 \times 10^{-26}$; $p_{DC-FFT} = 1.32 \times 10^{-12}$; $p_{NC-FT} = 3.85 \times 10^{-8}$; $p_{NC-FFT} = 5.43 \times 10^{-10}$; $p_{FT-FFT} = 2.68 \times 10^{-28}$. The strongest correlations are negative ones between DC and NC and FT and FFT, supporting observations in Figure 3.11 and described in the text. Conv: Convolutional Layer; FC: Fully-Connected Layer.

The results of the second analysis, shown in Figure 3.12, illustrated relationships between the representation of different tissue classes. Positive correlations may suggest that the tissue classes share many features encoded by the network such that ablation of the units strongly activated by these features similarly impacted both classes. However, negative correlations may also suggest shared features; if two classes share many features, ablation of units strongly activated by distinguishing features could cause pixels of the feature-sharing class to be selected with greater frequency, thereby driving changes in sensitivity of the two classes in opposite directions. In such a case, the unit activated by the distinguishing feature could be considered to be a suppressor or inhibitor of the related class which was inactivated by the ablation. Given this understanding, the results indicate that network representations of DC and NC were closely related, as were representations and classification pathways of FT and FFT. Both pairs demonstrated negative correlations. However, the complexity of these relationships were neither fully characterized nor done justice by so simple a summary. The results of the ablation study illustrated the inextricable, distributed, and interdependent nature of the interrogated neural network.

3.3.4. Discussion

While previous methods have classified tissue in grayscale IVUS images, the method presented here surpassed performance of the current state-of-the-art (Table 3.1). Previous work trained and validated on the same dataset implemented several varieties of classification algorithms, including support vector machines, neural networks, and random forests, with the latter achieving greatest performance. This method achieved an overall accuracy of 85.65%; sensitivity for the five classes ranged from 63.47% to 97.31%, while specificity ranged from 93.34% to 99.29% [270]. Because neural network training data can dramatically impact intravascular image segmentation performance metrics [162] (as shown later in this chapter), direct comparison with other work is tenuous, though performance meets or exceeds all comparable methods reported in literature (Table 3.1). Standardized datasets and methods to benchmark, analyze, and thereby fairly compare methods of intravascular tissue characterization are still needed, as has been previously established for evaluating lumen and media segmentation in IVUS by Balocco *et al.* [180]. To enable independent evaluation, and in anticipation of a future community standard for performance assessment, full confusion matrices have been reported here in order to allow computation of evaluation measures that are likely to be determined for such purposes.

Comparison of computational cost is similarly tenuous due to variation in both data and execution environment (hardware and software). Furthermore, execution time is sparsely [160], [264]–[266], [269]–[271], incompletely [268], or ambiguously [272] reported. Taki *et al.* reported feature extraction times for a typical frame between 7 and 300 seconds for different methods, but reported no overall process times [268], while Kim *et al.* reported test times between 189 and 673 seconds for different feature selection methods [272]. These wide ranges appear to bound the execution time of 200 ± 150 seconds per frame measured here.

Results showed that FFT and FT were confused by the model at much higher rates than other pairs of classes. This can be appreciated and anticipated through knowledge of the class tissue constitution. As noted before, fibrofatty and fibrotic tissue both contain collagen fibers, but configured differently. The former contains collagen bundled in fibers [236] and collagen in the latter are loosely packed fibers embedded in lipid accumulations [235]. It is expected, then, that the similarities in composition would result in similar echoreflective properties that would consequently make them difficult to distinguish from each other. Indeed, several previous methods have reported similar difficulties in distinguishing FFT or mixed tissue from FT, and some have forgone the distinction altogether and lumped several classes into larger, more easily differentiated groups [175].

Another pair of tissue classes confused with moderate frequency was NC and DC, though not in equal portions. While just over 9% of NC pixels were misclassified as DC, only around 1% of DC pixels were misclassified as NC. Further insight is offered by the ablation study performed on the CNN, which suggested that DC and NC shared features in network representation (Figures 3.11 and 3.12). When DC class output was inhibited, NC sensitivity increased, though the conjugate was not true. This observation prompted an investigation of activation strength for each class, which revealed that the predicted class score for calcium was, on average, 19% higher than that for necrotic core (Table 3.4). Due to the strong network response invoked by calcium, mild deviation in necrotic core appearance could be enough for the response to be eclipsed. Calcified and necrotic tissue often appear in tandem, and calcified structures are associated with acoustic shadowing [19], [277]; the imbalanced misclassification phenomenon could potentially be explained by such shadowing confounding the CNN as it identifies features of necrotic core that vary in appearance depending on its spatial position relative to the calcium. Accommodating such variation may result in the overall weaker activation for individual observations of NC tissue and

consequent non-reciprocated misclassification as DC. This uncertainty is remarkably similar to human experience, wherein shadowing and resulting variation in tissue appearance diminish certainty in an observation (e.g. identifying the adluminal edge of calcified tissue). Even when the CNN was trained without human intervention, it arrived at similar performance strengths and challenges; it is another example, amplifying lessons from Chapter 2, of how the human experience can inform likely scenarios in which trouble can be anticipated in automated processing methods. Future methods should similarly look to human experience for inspiration in mitigating these challenges by integrating contextual knowledge and an appreciation for the 3D continuity of systems.

As noted previously, segmentation of the vessel's inner and outer border, which together circumscribe the ROI, was a critical prerequisite to extract the geometric information necessary for the enrichment of the deep learning approach, and limits the accuracy of its results. This was a limitation shared with VH-IVUS; just as VH-IVUS relies upon—indeed assumes—an accurate inner and outer border to determine plaque composition within the vessel wall [235], so too did the presented method. Media and non-pathological tissue characterization was consequently particularly sensitive to ROI delineation. Any diminished performance in the ROI delineation degrades overall vessel characterization performance and compounds the final classification error, and as such contributions of this step were included in the reported errors. Indeed, a former study of cumulative error propagation in plaque image characterization found that image formation and border detection errors contribute to and increase plaque characterization error (i.e. decrease accuracy), but that these contributions are in acceptable limits and would not affect clinical decision [284]. Furthermore, accurate automated border detection algorithms are available, and because this segmentation is an interchangeable module on which the method builds, new or specialized methods may be utilized at will in concert with the presented domain enriched approach.

Taken together, these observations demonstrate and explain why lumping increases accuracy. As noted in the introduction, lumping multiple sub-classes together increases method accuracy, but at the expense of granularity. By neglecting differences between tissues with shared compositional structure and appearance (e.g. FT, FFT, and M), or between tissues which typically appear in proximity, the difficulty of the segmentation problem is mitigated substantially. These differences, however, may hold relevant signal which are lost by these dimension-reducing

simplifications. The minimal lumping of classes in the presented approach, and consequent granularity of the method output, further highlights the impressive performance both in terms of comparative performance metrics relative to previous methods and potential utility to drive detailed investigation and insights.

Work is warranted to extend validation of this method to ground truth histology. In the present work, the method has been both trained and validated against VH-IVUS. While VH-IVUS has itself been validation through *in vitro* histopathology [274], [276], it remains a step removed from the ultimate aim of classifying the tissue underlying the image. Furthermore, expert recommendations on intravascular RF data analysis maintain that media thickness cannot, in fact, be measured using either grayscale IVUS or VH-IVUS; media labels in the VH-IVUS images are themselves based on histological studies [277]. In a way, the domain enriched method emulates this approach; use of VH-IVUS for validation may therefore somewhat exaggerate the true benefit of the approach in considering the goal of tissue characterization. For example, because media thickness actually varies [8], [164], [277], a more sophisticated method of approximating media thickness (rather than assuming a fixed threshold thickness) may better reflect the underlying imaged tissue. However, in achieving the goal of replicating the utility of VH-IVUS without its associated restrictions and burdens, VH-IVUS itself presents a desirable, useful, and well-validated reference. Still, vigilance and transparency is prudent to avoid reinforcing potentially unfounded or weak assumptions that have guided development of VH-IVUS and the medical field more broadly.

Further work should also address the execution speed of the method. As currently implemented, the method cannot be applied in real time, limiting its usefulness. Immediate and drastic improvements could be achieved by exploring strategies to tactically select subsets and/or ordered progressions of pixels to be classified, rather than classifying every single pixel in the ROI sequentially by index. Updates to software, possibly including programming language, may also be accompanied by optimization of hardware.

Finally, as with any classification system, appropriateness of the model must be considered for any specific application. In particular, work discussed later in this chapter demonstrated that neural network training data profoundly impacts intravascular image segmentation [162]. Here, equal representation across all classes was enforced in the training dataset, and the CNN model

was consequently optimized for balanced accuracy across all classes, rather than weighted by prevalence in the dataset or overall population. Therefore, other models may prove more appropriate for the detection of specific plaque types or in patient populations with plaque phenotype profiles which deviate significantly from a balanced distribution. Furthermore, IVUS images can vary significantly in texture and appearance depending on the specific imaging system (hardware and software), system settings (e.g. transducer frequency), and acquisition protocol; performance of analysis algorithms can vary commensurately [180]. Generalizability of the specific network and quantitative performance reported should not be assumed for other datasets, though general trends are expected to hold.

By leveraging domain knowledge and recent technological advances, a domain enriched method of classifying plaque morphology using only grayscale IVUS images achieved higher accuracy, as well as granularity and precision, than that of others previously reported. By first imposing geometric constraints based upon pathological studies and normal vessel morphology, segmented images were produced that replicate VH-IVUS characterization with exceptional fidelity—without use of RF signal data. The method could therefore be applied to any grayscale IVUS data, including previously-acquired images that have not been characterized by the RF-based VH technique and images in VH-IVUS acquisitions occurring between characterized ECG-gated frames, thereby increasing the effective information acquisition speed. While care must be taken to consider and convey assumptions which may be reinforced or perpetuated through the application of domain knowledge to learning methods for medical imaging, this method offers practical, translational opportunities for immediate application-specific deployment.

3.4. Characterizing Tissue in Intravascular OCT

The framework presented in Section 3.2 was also applied to generate a novel CNN-based method to classify arterial tissue imaged through OCT (Figure 3.2, center), as described by Athanasiou *et al.* [160]. Unlike application in IVUS, where RF-based VH offered voluminous fully-classified segments for reference, penetration depth is rarely limiting, and well-established methods offered reliable vessel wall geometries, OCT presented a unique challenge due to the availability of only piecemeal manual annotations, guidewire shadow artifacts, limited penetration depth, and (consequently) less dependable geometric information. In its application to OCT, the framework detected the ROI using recently developed methods, including the outer border

detection method described in the previous chapter, then subdivided the ROI into pathological and non-pathological tissue based upon basic spatial and geometric constraints informed by physiology. Pathological areas of the ROI were partitioned into patches and processed by a CNN architecture. Human expert annotations served as the comparative control. The method offers several meaningful benefits stemming from its ability to characterize the full arterial wall. Whereas human experts may partly annotate a handful of frames from a full pullback, and other methods may identify a single tissue type (e.g. calcified) throughout a pullback, both approaches provide scarce and incomplete information, excluding information which can be important in many cases to the comprehensive quantification and computational modeling of diseased arteries.

3.4.1 Materials and Methods

As noted in the general description, ROI was defined as the area between the lumen and outer border (i.e. the media-adventitia transition; Figure 3.3). The lumen detection method accepted as input 2D cross-sectional OCT images, produced sagittal cuts (i.e. longitudinal cross-sections; which more accurately represent the sequential area of the OCT pullback), detected the lumen contour by applying bilateral filtering and k-means algorithms, and projected the detected contours back to the 2D OCT images [225]. The outer border detection method, described in the previous chapter, detected visible segments of the outer vessel border within the OCT pullback, then completed the full contour (including non-visible segments) using the unique three-dimensional surface-fitting method [157].

The ROI was subsequently divided into pathological and non-pathological or media regions. OCT differs from IVUS in its pixel size (i.e. scaling), so threshold values for pathological tissue had to be determined independently. Th_{path} was found to be 16 pixels, and Th_{media} was 5 pixels. Therefore, pixels of the ROI were classified as pathological tissue (ROI_{path}) if $D_{thick} \geq 16$ pixels and $D_{outer} \geq 5$ pixels (Figure 3.3).

Pixel-centered patches were created for those pixels of the ROI which remained after segmenting the non-pathological or media class (i.e. $r_{im} \in ROI_{path}$). These patches were automatically classified using a CNN as one of the four plaque types specified for OCT or as shadowed area (i.e. guidewire shadow artifact; assigned a class label of “N”). To achieve the best classification results, different patch sizes, numbers of input patch convolution sequences, filters,

and filter sizes were tested. A patch size of 31×31 was determined to perform best. The network found to perform best, and utilized in this work, had 45 layers and is shown in Figure 3.5.

3.4.2 Dataset

Twenty-eight patients who underwent OCT examinations gave informed consent to participate in the study, which was approved by the Ethics Committee of the institution. Medical experts used the FD-OCT C7XR optical frequency domain imaging system and the DragonFly catheter (St. Jude Medical, St. Paul, MN)—which offers a maximum frame rate of 100 frames per second, 500 lines per frame, a scan diameter of 10 mm, and axial resolution of $15 \mu\text{m}$ —to image 28 coronary vessels. All images were digitally stored in raw format for off-line analysis, and all imaging data sets were anonymized and transferred for further analysis. Data from 22 patients were used for training, while data from the other 6 patients were withheld for testing.

Two medical expert manually reviewed frames, selected images corresponding to diseased segments, annotated the lumen and intima-media borders, and independently labeled regions depicting clearly identifiable bulk regions of the various tissue types. Image areas were marked as calcium (assigned a class label of “C”), LT, FT, MT (a combination of calcified and lipidic tissue), or N (non-visible, i.e. catheter shadow). In total, 300 plaques were annotated in the training dataset and another 50 plaques were annotated in the test set. Image patches generated for pixels in these labeled regions were augmented through rotation (90° and 180°). Within the training dataset, 450,000 image patches were used for training, while 30,000 image patches were used for concurrent validation during the training process to assess training progression and guard against overfitting. Equal numbers of patches drawn from each of the five classes were used in the training process, resulting in a balanced training dataset. The test dataset was not balanced, instead reflecting the actual underlying tissue distribution visible in this patient cohort. Of the 843,778 patches used in testing, 1.1% were of C, 20.7% of LT, 25.6% of FT, 4.2% of MT, and 48.4% N.

3.4.3 Results

Image segmentation accurately identified plaque types labeled by medical experts, and successfully emulated clinically utilized classifications that categorize the entire section on the basis of (grayscale) OCT images. Three representative examples of classified images resulting from the method—demonstrating its ability to produce integrated plaque characterization maps categorizing the entire vessel wall cross-section using OCT—are shown in Figure 3.13. The




Figure 3.13. Representative classified OCT images segmented by the presented plaque characterization method. The method identifies major pathological tissue morphology features visible in the initial images (*top*) within the corresponding color-coded images (*bottom*). C: white, LT: red, FT: green, MT: light green, N: light gray, and M: dark gray.

method accurately captured major tissue morphology and features within the pathological region (Figure 3.13). Due to the spatial constraints imposed prior to CNN classification, the domain enriched method addressed non-pathological and media tissue accurately, and was not disposed to violating physiological constraints.

The CNN classifier achieved generally high precision (i.e. PPV) and recall (i.e. sensitivity). Table 3.5 shows the error (or confusion) matrix for the method’s five-class CNN classifier when presented with the (imbalanced) testing set. The model, trained on regions of pathological tissue manually labeled by experts, achieved an overall (combined) classification accuracy of 96.1% for the full testing dataset.

Progression of training for this network, which was conducted on a balanced dataset, is shown in Figure 3.14.

Table 3.5. OCT pathological tissue classifier error matrix

		Target Class					Precision
		C	LT	FT	MT	N	
Output Class	C	6831	8835	1451	1738	1259	34.0%
	LT	959	162214	2425	1782	2629	95.4%
	FT	163	1378	208421	2072	872	97.9%
	MT	1497	1803	4038	29709	168	79.8%
	N	0	218	1	0	403315	99.9%
Recall		72.3%	93.0%	96.3%	84.2%	98.8%	96.1%

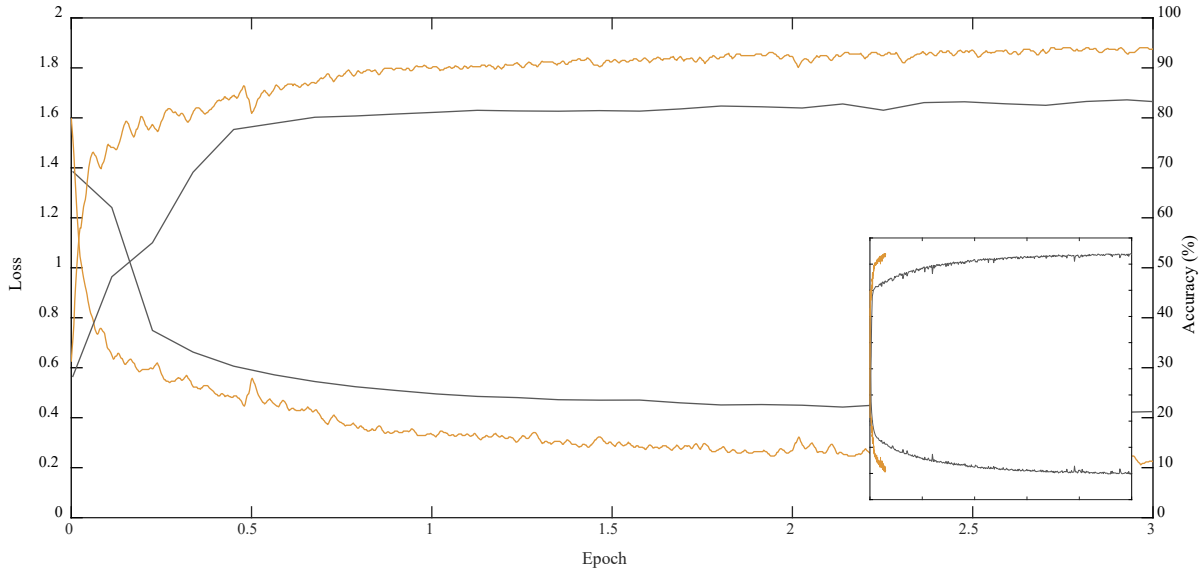


Figure 3.14. CNN validation accuracy (*top*) and loss (*bottom*) through 3 epochs (10,545 iterations) of training for the model. Progress of the OCT characterization model described in Section 3.4 is shown in gold; progress of the IVUS characterization described in Section 3.3 is shown in gray. The inset shows the 3 epochs in the context of the full 50 epochs employed in training the IVUS characterization network (Figure 3.9). Accuracy and loss plateaued for the OCT network by around the 3rd epoch, and the relatively small datasets motivated cautioned against overfitting through excess training.

Training was halted after 3 epochs, at which point accuracy and loss had plateaued, having reached a validation accuracy of 94.0% (Figure 3.14). Due to the relatively small dataset arising from the difficulty of acquiring annotated data (cf. VH-IVUS), a conservative approach to training extent was imposed to avoid overfitting (non-generalizable model optimization) to the data through extended training. Relative to training for IVUS, accuracy initially improved more quickly in training for OCT (Figure 3.14). One hypothesis for this observation is that the higher resolution and less noise of the OCT data relative to IVUS made the appearance of classes more consistent within the dataset, facilitating more rapid identification of characteristic features (i.e. signal) while decreasing errant optimization steps during SGD.

The error matrix of the classifier illustrates some general trends. The five-class network supporting the domain enriched method struggled to differentiate calcium (C) from LT and MT, resulting in conspicuously low precision in identifying C. The network also had difficulty, though less extensive, in differentiating MT from FT. The CNN was highly capable of classifying pixels in which no visible tissue was present (N), which represented a sizeable portion of the testing dataset drawn from all image labels in the withheld patients. Due to the nature of the ground truth data, compounding uncertainties arising from pathological tissue delineation were not assessed as they were for the IVUS characterization method.

3.4.4 Discussion

Utilizing the flexible framework established above, an automated method was developed for characterizing intravascular OCT images that automatically segments the non-pathological tissue or media (M) and classifies the imaged atherosclerotic plaque into to 5 different classes (C, LT, FT, MT, or N). The method combined advanced ROI detection (geometry extraction) algorithms and a CNN, and was validated using the manual annotations of expert observers as gold standard in a clinically-relevant dataset. Results demonstrate fairly reliable tissue detection and characterization, aside from calcium detection precision, even in images having artifacts.

The method described herein is more accurate and realistic than methods previously presented in the literature. Over the past decade, several methods have been presented for the detection and characterization of atherosclerotic plaque using OCT images [173], [257], [258], [260]. The methods have been primarily based on machine learning algorithms [173], [259], [260], and more recently on deep learning approaches using CNNs [257], [258]. These methods can detect atherosclerotic tissue within the arterial wall with reasonable accuracy. However, while CNN-based methods outperformed the machine learning predecessors, they could not characterize the full arterial wall as does VH-IVUS [235]. The framework described here allows this functionality, making it notably qualified for use in the clinical and research arenas.

The method can be improved further, and warrants efforts to address the lower accuracy when differentiating certain tissue classes (Table 3.5). Indeed, the error matrix of the classifier highlights the most pressing needs. Results showed that C was confused by the model with LT and MT at much higher rates than other combinations of classes, followed by MT-FT confusion. This can be appreciated and anticipated through knowledge of the class tissue constitution and CNN structure. This limitation is due in part to the nature of MT, which includes calcium and lipid and shares image characteristics of both C and LT; as noted before, MT consists of variable combinations of calcified, lipidic, and fibrous tissue, so it may be expected, then, that the similarities in composition would result in similar optical properties that would consequently make them difficult to distinguish from each other. Variable appearance of the inhomogeneous MT further complicates its successful identification. Furthermore, while calcium and lipid pools display different attenuation properties, both appear as locally dark (low intensity) regions. Without additional global or spatial context—of which CNNs are notoriously poorly suited to

integrate—it is understandable that these two classes may be difficult to distinguish. Expanding the clinical dataset and incorporating histological findings in the training and testing phase of the proposed method may help to address these limitations. Furthermore, beyond improving plaque classification, integrating new methods for directly segmenting the media [158] (described in the previous chapter; not possible in IVUS), rather than assuming a constant thickness based upon histology studies (as implemented in VH-IVUS), could offer more accurate segmentation and richer information on vascular morphology.

By leveraging domain knowledge and recent technological advances, a domain enriched method of classifying plaque morphology using only (grayscale) intravascular OCT images achieved higher accuracy and greater functionality than that of others previously reported. By first applying new ROI-delineating algorithms and imposing geometric constraints based upon pathological studies and normal vessel morphology, segmented images were produced that emulate VH-IVUS (i.e. full section) characterization with unprecedented detail. This method offers practical, translational opportunities for immediate application-specific deployment, and could be used in the clinical and research arenas to advance the field of computational cardiology [67].

3.5. Segmenting Indwelling Devices in Intravascular OCT

Segmenting indwelling devices, such as a BVS, was achieved with only minor adaptation of the framework presented in Section 3.2 (Figure 3.2, right). Rather than specifying inner and outer borders of the vessel wall to delineate the ROI, a large standardized area of the image was instead selected. Furthermore, pathological tissue was not partitioned and treated differently from the other tissue within the ROI. Blending pathologies and structures in this way leveraged some of the aforementioned benefits of lumping—while decried earlier for the resulting loss of granularity and precision, here the aim was only the segmentation of the device and not the differentiation between tissues. Therefore, including a spectrum of “non-strut” presentations in a single binary class reduced complexity and makes the discrimination task easier. Finally, rather than classifying a pixel into one of several tissue classes, the pixel-centered patch provided to a CNN was assigned to one of two binary classes indicating presence or absence of a strut. Several network architecture variations and input patch sizes were assessed to identify the combination which worked best for the given application, and data sensitivity studies assessed the

generalizability and robustness of the method. Human expert annotations served as the comparative control. By utilizing the presented framework with deep learning, the method offers benefits stemming from its independence from rigid, hand-crafted rule sets describing strut size, shape, and distribution, which may evolve with device resorption or be violated by outlier cases, image artifacts, or anomalies. Furthermore, whereas human experts may annotate or make measurements in only a handful of frames from a full pullback, the current approach enables comprehensive quantification to track resorption and vascular response, as well as reconstruction for computational modeling purposes.

3.5.1 Materials and Methods

The methods implemented to segment indwelling BVS struts in intravascular OCT images followed the fundamental methodological framework described earlier (Section 3.2; Figure 3.2, right).

A large standardized area of the image was designated as the ROI. Because scaffold struts can be located in various positions—ranging from malapposed struts within the lumen space to deeply embedded struts within the vessel wall—an assessment of the available dataset was undertaken to determine the typical radial range in which the search for struts should be focused. Struts were observed between 18.6 and 180.1 pixel widths from the center of the image in the available labeled dataset. Conservatively extending the range observed to accommodate uncertainty and expected variability, the ROI was defined as the region 15 to 185 pixels from the center of the image; this ROI (106,840 pixels) included 42.7% of the entire 500×500 pixel image (250,000 pixels) and 54.4% of the circular image data (i.e. excluding the bounding frame; 196,364 pixels). As noted before, pathological tissue was not partitioned and treated differently from the other tissue within the ROI, forgoing this step in the general framework.

Pixel-centered square patches were created for all pixels of the ROI (i.e. $r_{im} \in \text{ROI}$). These patches were automatically classified into one of two binary classes—indicating presence or absence of a strut—using a CNN. Different patch sizes, ranging from 11 to 61 pixels across in increments of 10, as well as varying CNN architectures, ranging from 3 to 5 sequences of convolution, batch normalization, and ReLU activation, were tested in all possible permuted combinations to achieve the best classification results for the given task. The networks utilized in this work are shown in Figures 3.5 and 3.15. Further testing of the generalizability and robustness

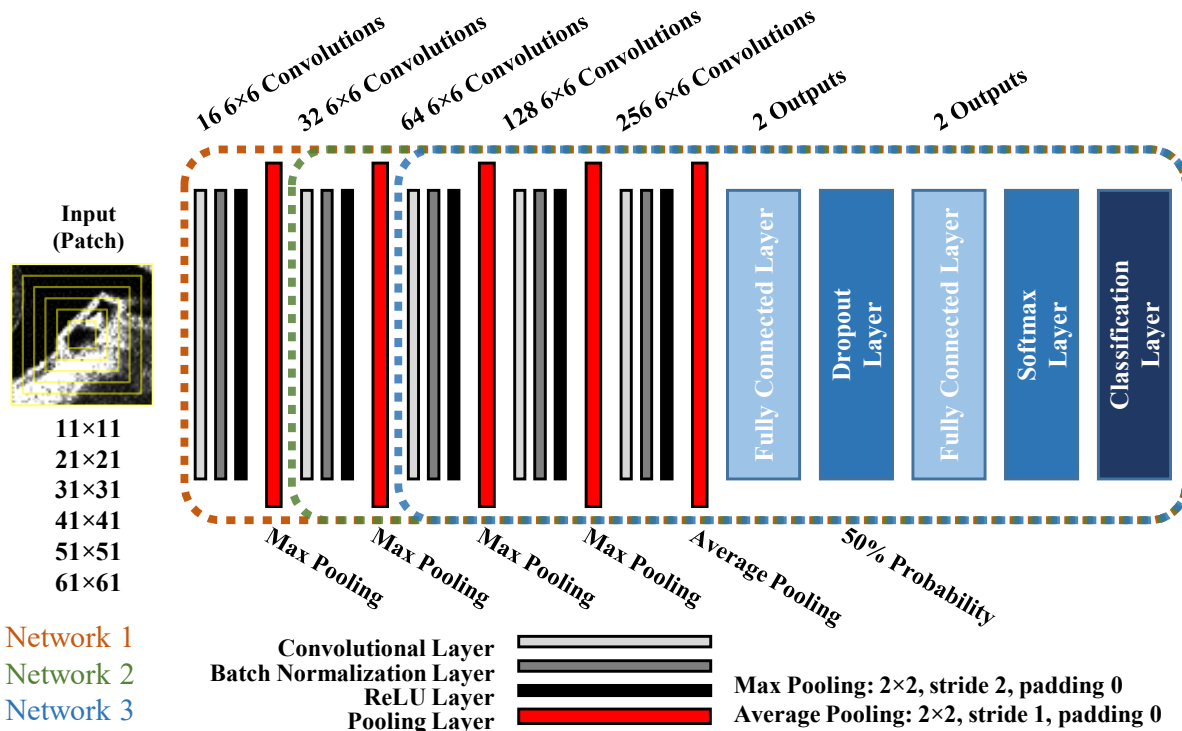


Figure 3.15. Detailed architecture of the multi-layer CNNs used to classify pixels within the ROI for BVS strut segmentation in OCT. Various patch sizes, ranging from 11 to 61 pixels, were tested, as were various numbers of functional sequences, ranging from 3 (network architecture 3; blue dashed box) to 5 (network architecture 1; orange dashed box). See Figure 3.5 for the final network architecture (2) and a comparison to the CNN architectures applied to the tissue characterization task.

of the method using the preferred patch size and CNN architecture was achieved by performing extensive cross-validation in which each patient was iteratively withheld from training and used exclusively in testing. All networks were trained over 50 epochs.

3.5.2 Dataset

Clinical OCT images showing vessels implanted with BVS (i.e. endovascular polymeric implants) were used to train and validate the classification network. The dataset is described in detail by Amrute *et al.*, who used the same data to develop a polymeric endovascular strut and lumen detection algorithm using basic image processing and k-means clustering [179]. Pullback sequences from 14 patients were acquired in accordance with clinical standards [18] following BVS implantation, ranging from 1 day to 9 months post-procedure. Consequently, struts varied in their presentation and included malapposed, well-apposed, and embedded struts; due to variable resident times, struts also varied from fully intact to mildly degraded. Medical experts identified and marked each visible strut. Only frames acquired within the scaffold (i.e. depicting at least one strut) were utilized. In total, there were 1156 frames containing 6.10×10^5 pixels labeled as showing

a strut, with each frame including 528 ± 182 such pixels— $0.5 \pm 0.2\%$ of the ROI (average \pm standard deviation; Figure 3.16).

Training and testing was performed in balanced datasets with equal representation of strut and non-strut data points. As the intravascular images were collected from a clinical population in the course of treatment, and endovascular scaffolds are—by design—very sparse, the data were inherently inhomogeneous and imbalanced. Figure 3.16 shows the average quantity and relative distribution of scaffold struts present in the imaged vessels of each patient, as well as the number of OCT frames available from each patient’s acquisition. Though sometimes present in small numbers, at least one strut was identified by the experts in each frame. However, the dataset displayed extensive variability in strut prevalence, image quality, and available frames.




Figure 3.16. Strut distribution of patients comprising dataset. *Top:* Prominence of struts for each patient, reported in average number of pixels per frame corresponding to BVS struts as determined by manual expert annotation. Error bars indicate standard deviation. *Bottom:* Number of strut-containing OCT frames available for each patient. Overall, from the 14 patients there were 1156 frames containing 6.10×10^5 pixels labeled as showing a strut, with each frame including 528 ± 182 such pixels— $0.5 \pm 0.2\%$ of the ROI (average \pm standard deviation).

To increase speed and efficiency, datasets of increasing sizes were used in progressive stages of method refinement and evaluation. For preliminary studies, data from 9 patients were used for training, from which 10^5 patches of each class were randomly selected; 3×10^4 patches of each class were randomly selected from the remaining 5 patients for validation. Subsequently, for the comprehensive assessments for which results are reported here, patches were generated for every pixel labeled as depicting a strut (6.10×10^5 total), and equal numbers of non-strut pixels were randomly selected from the ROI for each patient, such that the data extracted from each patient was balanced, though the quantity of data contributed by each patient differed somewhat ($4.36 \times 10^4 \pm 1.06 \times 10^4$, ranging from 2.55×10^4 to 5.66×10^4 , of each class; Figure 3.16). Data from 11 patients, which included 875 frames and 4.52×10^5 pixels labeled as strut, were used exclusively for training; the remaining 3 patients, which included 281 frames and 1.58×10^5 pixels labeled as strut, were withheld and used exclusively for testing. In testing for generalizability and data dependency, 14-fold cross-correlation was performed in which all but one patient was included in the training dataset, and the withheld patient was the sole patient used in validation. To maximize generalizability for future application in new datasets, a final network was trained with all patients pooled and every 14th frame withheld for validation. In all cases, 7-fold data augmentation was performed through reflection and rotation in 90° increments.

Various CNN architectures and patch sizes were systematically tested to optimize the performance of the method. Patch sizes of $(10n+1) \times (10n+1)$ pixels for $n = \{1, \dots, 6\}$ were tested. The network architecture with the greatest number of convolutional sequences (network 1) was incompatible with the smaller patch sizes ($n = \{1, 2, 3\}$); the smallest patch size (11×11 ; $n = \{1\}$) was also too small to allow the number of convolutional sequences of network 2. All other permutations were assessed.

3.5.3 Results

All compatible combinations of the three network architectures and six patch sizes were successfully trained and validated with the limited dataset. Similar outcomes for many of the permutations demonstrated the robustness of the framework, though some important trends emerged (Figure 3.17). The smallest patch size (11×11 pixels) did not perform well, with the lumen border often being falsely identified as a strut, indicating that this neighborhood scope did not provide sufficient context for the network to successfully identify struts. Conversely, large patch

sizes did not perform well with the smaller CNNs. Due to the large input size and relatively few convolutional sequences, a very large number of parameters had to be trained in the first fully-connected layer (Figure 3.15). This feature elevated the risk of over-fitting while inflating the computational cost and memory requirements of training. Without an increased amount of data required to achieve accurate performance, these permutations tended to have poor performance and particularly high false positive rates in classifying the large information- and context-rich patches. Due to their non-competitive performance and high computational burden, these permutations were abandoned following the preliminary studies.

The full dataset helped to distinguish somewhat the performance of the various permutations of patch size and network architecture. Metrics of accuracy, sensitivity, specificity, PPV, and negative predictive value (NPV), evaluated relative to the expert manual annotations, were all strong, but varied slightly (Table 3.6; Figure 3.18a). A patch size of 31×31 pixels provided all-around strong performance, including a desirable balance between sensitivity and specificity. Network architecture 2, implementing four convolutional sequences (Figures 3.5 and 3.15), with an input patch size of 31×31 was ultimately selected.

Table 3.6. Device segmentation performance for various permutations of patch size and network architecture

Patch Size	11	21	31	41	51	61
<i>Accuracy</i>						
Network 1				98.12%	98.35%	97.41%
Network 2		97.96%	98.22%	98.27%	97.32%	
Network 3	96.94%	97.65%	98.46%	97.83%		
<i>Sensitivity</i>						
Network 1				97.05%	97.74%	95.72%
Network 2		96.69%	98.22%	97.48%	96.64%	
Network 3	95.64%	96.28%	97.78%	96.65%		
<i>Specificity</i>						
Network 1				99.18%	98.96%	99.09%
Network 2		99.23%	98.95%	99.07%	98.00%	
Network 3	98.23%	99.02%	99.14%	99.01%		
<i>PPV</i>						
Network 1				99.17%	98.95%	99.06%
Network 2		99.21%	98.95%	99.05%	97.98%	
Network 3	98.18%	98.99%	99.13%	98.99%		
<i>NPV</i>						
Network 1				97.12%	97.77%	95.86%
Network 2		96.77%	98.23%	97.52%	96.69%	
Network 3	95.75%	96.38%	97.81%	96.73%		

Color ranges from worst performance (red) to best performance (green) for each metric




Figure 3.17. Strut segmentation performance of the CNN classifier in a representative frame with various combinations of network architecture and patch size. Pixels are color-coded by classifier outcome relative to manual annotation: true positive (green), false positive (red), false negative (blue), or true negative (gray, i.e. initial image). Results are arranged in a grid indicating the network architecture (row) and patch size (column) that produced the displayed results with a training set comprised of 10^5 patches of both classes. The grayscale image provided to the workflow is shown in the top left corner (in place of the incompatible Network 1- 11×11 patch size combination). The three network architectures are shown in Figure 3.15.




Figure 3.18. Variation in validation performance metrics with changes in (a) network architecture and patch size (Table 3.6) and (b) withheld patient (for the combination of network architecture 2 with 31×31 patch size). Each box (blue) spans the 25th and 75th percentiles, with the central mark (red) indicating the median. Whiskers (black) extend to the most extreme non-outlier values; outliers are individually marked with plus (+) symbols (red). The range of all 5 performance metrics is greater for CNN variation (a) than data variation (b).

Having selected a final network, the impact of data and robustness of the method was assessed. Results of the cross-validation, in which each patient was, in turn, withheld from the training set and used exclusively in testing, showed reassuring consistency (Figure 3.18b). All 5 performance metrics assessed demonstrated smaller ranges among the 14 dataset partitions than among the 11 network-patch permutations. Average performance by all 5 metrics also exceeded that achieved by the same network with the somewhat smaller dataset (in which 3 patients were withheld for testing), as hoped.

Reassured by the generalizability and data independence, a final network was trained with all patients pooled to maximize generalizability for future application in new datasets. Every 14th frame was withheld for validation to achieve approximately the same train-test data ratio as withholding one of the 14 patients. The final network, trained in this manner, exceeded the average performance of the cross-validation, as anticipated, and achieved a validation accuracy of 99.37%, sensitivity of 99.48%, specificity of 99.27%, PPV of 99.27%, and NPV of 99.48%. As a proof of

concept, the method was applied to an entirely new dataset, acquired under different conditions and at a different medical center (but with the same commercial system). Though performance was somewhat diminished, initial results proved promising across a range of presentations (Figure 3.19).




Figure 3.19. Strut segmentation performance of the CNN classifier in two images from a new dataset acquired independent of the training and validation set. Initial grayscale images (*left*) and their annotated counterparts (*right*) reveal a stent well apposed shortly after implantation (*top*) and following vessel recovery with neointimal grown covering partly degraded struts (*bottom*).

3.5.4 Discussion

Utilizing the flexible framework established in Section 3.2, an automated method was developed for annotating an indwelling device in intravascular OCT images. The method combined simplistic ROI delineation and a CNN to automatically segment the imaged vessel into one of 2 classes (strut or non-strut). It was validated using the manual annotations of expert observers as gold standard in a clinically-relevant dataset, and results demonstrated fairly reliable strut segmentation, with some improvement possible in reducing the false positive rate.

The method can be improved further, and warrants efforts to address both the efficiency and the false positive rate. For example, the crude and conservative ROI results in a very large area to be classified. Except in (rare) cases of vessel dissection, struts should not be found beyond the intima. Therefore, leveraging geometric information on the vessel wall should allow for a narrower, refined ROI to be determined for each frame, improving both the efficiency and reducing opportunities for false detections. Additionally, the inner limit of the ROI could be based upon the catheter artifact in the given acquisition, rather than a fixed value which sometimes extended into the artifact itself, unnecessarily extending the ROI into a region with some strut-like visual patterns in which no strut could exist. Additional screening methods, possibly considering grayscale intensity, to mask and narrow the ROI may also prove fruitful.

Beyond refinement of the ROI, challenges introduced by the underlying data should also be considered further. First, due to the relative scarcity of struts within the images, the “population” dataset is a highly imbalanced one; pixels belonging to the “strut” class represent only a small fraction of all pixels, even within the ROI. As discussed in a later section of this chapter (Section 3.8), such imbalance warrants careful consideration of the training dataset curation. Here, a balanced training dataset was used, but depending on the specific application of this method, it may be preferable to utilize a different distribution—or a modified loss function weighting inaccuracies differently by class—to drive performance towards a desired profile (e.g. greater sensitivity or specificity; greater PPV or NPV). Second, struts appear with a range of presentations depending on where they are located within the vessel, the extent of their degradation/resorption, and potentially also their specific model and material formulation. Despite the strong performance of the CNN in the context of variable data, it should be considered whether this challenge is best addressed by the collection and training on additional data, consideration of several differing classes

(rather than monolithic “strut” and “non-strut” classes), or training of multiple distinct CNNs (each dedicated to a narrower application scenario), depending again on the specific application and intended use cases. Enriching the method with additional domain knowledge, 3D context, or spatial information may also help to address this challenge of variable appearances and improve performance.

By leveraging recent technological advances and a proven framework [160], [161], an optimized method of segmenting indwelling devices using only (grayscale) intravascular OCT images achieved high accuracy and substantial functionality with only minor modification of the framework developed for the characterization of tissue. This method offers practical, translational opportunities for application-specific use, and could be used in the research and clinical arenas, particularly in the field of cardiovascular interventions.

3.6. Variations on the Framework

One of the many benefits of the framework presented herein is its modularity, which makes it incredibly flexible and allows the most relevant techniques to be leveraged and interchanged with the advent of new advancements. This includes changes in imaging modality (input image), algorithms to delineate the ROI, approaches to distinguishing pathological tissue from non-pathological tissue and media, and the neural network-based classification. To explore a possible avenue for improved performance and assess tradeoffs between methods, a substantially different alternative neural network structure—one which employed a form of semantic segmentation—was implemented for the classification of the vessel ROI.

As an alternative to the pixel-wise CNN-based classification, a U-net was implemented and evaluated. U-nets, so named for the progression and shape of the intermediary (hidden) layers, consecutively downsample and upsample the full input image to generate a full map of class labels (rather than a single class label corresponding to the full image) as an output [285]. U-nets and similar architectures have been shown to be efficient and well-performing in biomedical image analysis [285]. Here, a grayscale intravascular image and a corresponding mask of the ROI was provided as input to the tissue classification network, which output a label map for the entire image, encoded as a one-hot structure (Figure 3.20). Of note, the media and non-pathological tissue were not first separated from the pathological tissue in the ROI. Rather, providing the ROI mask as an input to the network provided the necessary context for the network to learn itself the relative




Figure 3.20. U-net architecture used to characterize intravascular images. Blue boxes indicate multi-channel feature maps resulting from convolution, normalization, and activation operations; map sizes of hidden layers are symmetric. Skip connections (black arrows) concatenate channels of equal size.

distribution and geometric relationships governing this tissue class and serve the functions of both modules. An additional network was developed for the BVS strut segmentation task; a grayscale OCT image (but no ROI) was provided as input to the network, which output a binary label map for the entire image.

The network built upon a model by Isola *et al.* [286], with substantial modification and adaptation for the application at hand, and was implemented in Google Colaboratory (Alphabet, Mountain View, CA) using the TensorFlow library running on varying virtual machines with GPU or tensor processing unit resources. The network (Figure 3.20) iteratively downsampled in sequences of 2D convolutions (4×4 kernel size with stride of 2), batch normalizations, and leaky ReLU activations, then upsampled in sequences of 2D transposed convolutions (4×4 kernel size with stride of 2), batch normalizations, and ReLU activations (with dropout). Skip connections concatenated all channels of a downsampling layer to the channels of the equally-sized upsampling layer. The output layer performed a final 2D transposed convolution (4×4 kernel size with stride of 2). For tissue characterization, the final convolution produced 5 or 6 channels (for IVUS and OCT, respectively) which underwent softmax activation, resulting in a final one-hot structure with each channel representing one of the same tissue classes implemented in the standard pixel-wise CNN-based classification scheme described above. For strut segmentation, the final convolution produced just a single channel which underwent sigmoid activation, resulting in a final 2D structure of values between 0 and 1, with indices exceeding 0.5 being classified as strut and others as non-strut pixels.

Two forms of loss were considered in the optimization of this classifier. One form considered just the categorical multiclass cross-entropy loss (Equation 54) between the known tissue class labels and the predicted labels generated by the network. Another form implemented an adversarial model [287], whereby a discriminator network (Figure 3.21) was trained concurrently with the classifier (a form of conditional generative network). The discriminator, a CNN, is tasked with distinguishing between the “ground truth” tissue label map and the classifier-generated label map for a given image. Provided with a tissue map and a corresponding image, this CNN (Figure 3.21) iteratively downsampled in sequences of 2D convolutions (4×4 kernel size with stride of 2), batch normalizations, and leaky ReLU activations. The output layer performed a 2D transposed convolution (4×4 kernel size with stride of 2) with a single channel, resulting in a scaled confidence map assigning a continuous value to discrete regions of the image indicating

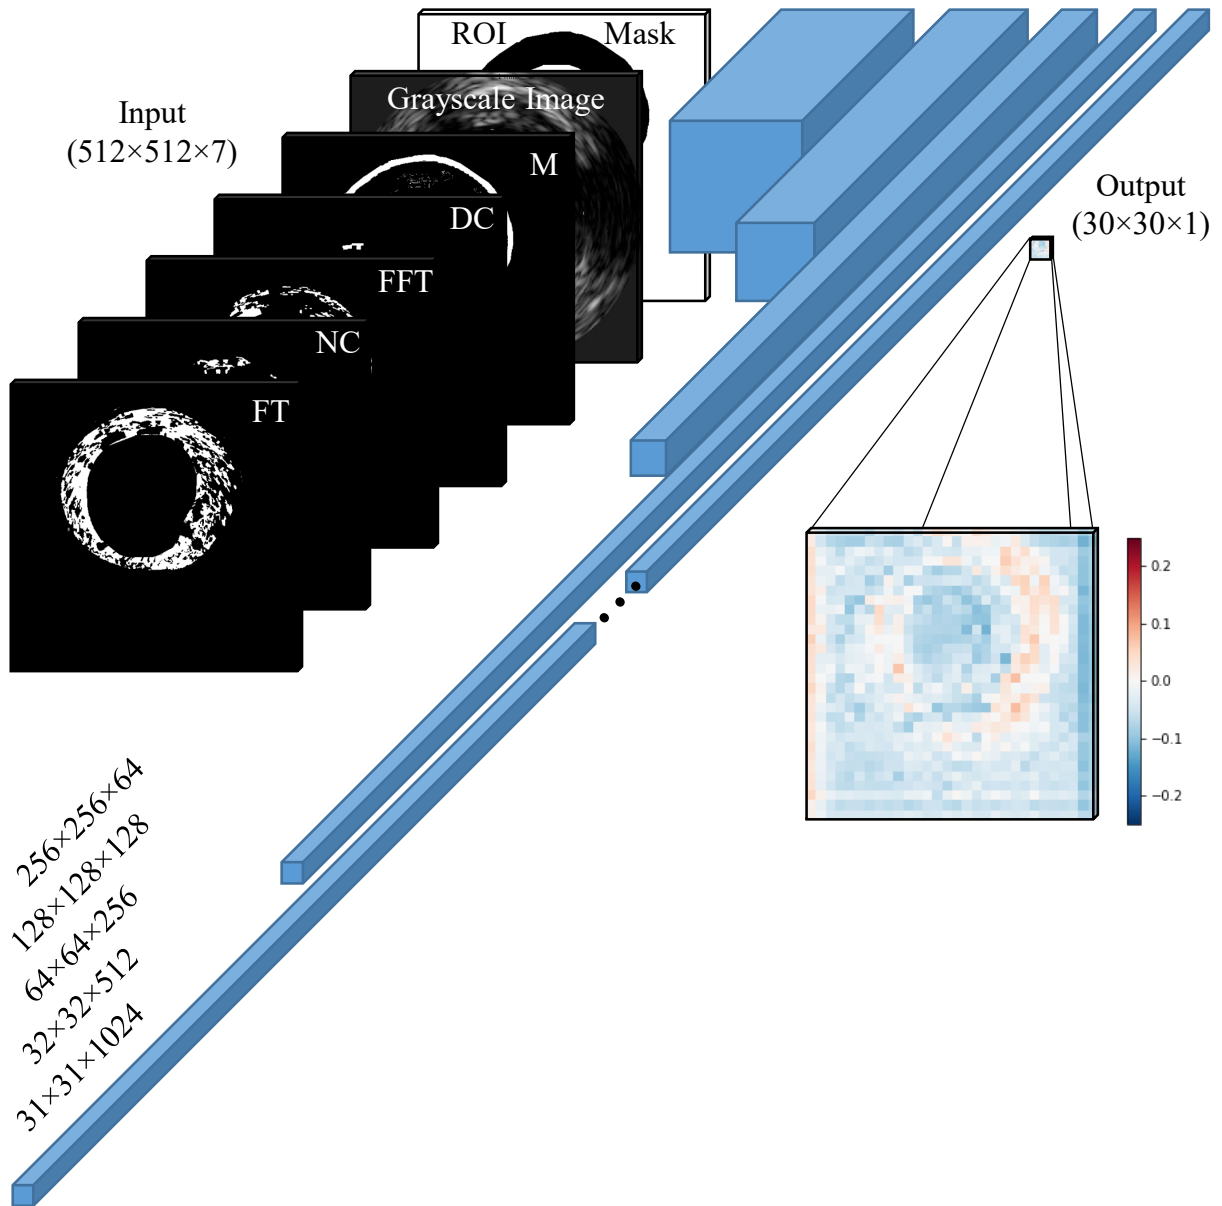


Figure 3.21. CNN architecture of the discriminator used in adversarial training of the U-net to characterize intravascular images. Blue boxes indicate multi-channel feature maps resulting from convolution, normalization, and activation operations. A tissue map, image, and ROI mask are provided as input; a scaled confidence map is output. During training (concurrent to a generator), the discriminator is penalized by binary cross-entropy loss for incorrectly designating a classifier-generated map as authentic or a ground truth label map as inauthentic.

confidence in the authenticity of the corresponding tissue map provided. The ability of the discriminator to successfully identify classifier-generated maps as such penalized the classifier (“generator”) through binary cross-entropy loss. (See Chapter 5 for more detailed and technical descriptions of adversarial and generative networks.) Therefore, the first, more straight-forward form of loss considered just cumulative pixel-wise error (as the pixel-by-pixel classifiers did), whereas the second form considered context and more holistic appearance of the maps comprising

the pixel-wise labels. Three models were trained for each modality: one which considered only loss driven by pixel-wise label differences (“non-adversarial”), one which considered only loss driven by discriminator success (“fully adversarial”), and one which considered a weighted combination of the two (“adversarial”).

The classifier was applied to tissue characterization in both IVUS and OCT images, as well as BVS strut segmentation in OCT. For application to IVUS, the same dataset was used as that for the approach described in Section 3.3; tissue label maps were extracted from VH-IVUS as the target distribution, and characterization was performed on the corresponding grayscale IVUS frames. For application to OCT tissue characterization, the same imaging dataset was used as that for the approach described in Section 3.4. However, as this segmentation approach required full label maps for training (rather than the piecemeal region labels available from the manual annotation of OCT images), the *output* of the CNN-based method—the fully characterized images—were used as the target distribution (i.e. “ground truth”) for training the network described here. For application to strut segmentation in OCT, the same imaging dataset was used as that for the approach described in Section 3.5, for which full label maps were available.

Because full images were utilized, and tissue types and struts were present in vastly unequal quantities (Figures 3.6 and 3.16; Table 3.5), all three datasets were significantly imbalanced. However, unlike pixel-wise classification schemes, training datasets could not be balanced to include equal numbers of each class. As highlighted and quantitatively explored in Section 3.8, networks trained with such imbalanced datasets are strongly biased towards classes prevalent in the training sets, to the detriment of less-prevalent classes. Therefore, each network was trained with and without compensatory loss weighting, whereby the loss incurred for incorrect classification of a pixel was inversely proportional to the prevalence of the pixel’s class—as labeled in the target (“ground truth”) distribution—in the training dataset. As such, the cumulative potential error contributed by each class was balanced.

Results revealed some intriguing and promising attributes for this variation on the morphology-extraction framework, but also remaining shortcomings to be considered and addressed. The fully adversarial scheme did not perform well in either modality, as insufficient feedback was available to the classifier (generator) from the discriminator alone to meaningfully improve during training. Figure 3.22 and 3.23 show illustrative examples for the characterization

of IVUS and OCT frames, respectively, by the other two U-net tissue classifiers. Following random initialization, the network quickly gained proficiency and learned to generate reasonable tissue distribution maps using semantic segmentation. However, several details never emerged. In IVUS, boundaries between fibrotic and fibrofatty tissue and between necrotic core and dense calcium were rather amorphous and lacked the sharpness of VH and pixel-wise CNN segmentation. In OCT, the identification of calcium, a relatively scarce constituent, never emerged meaningfully in the models trained with unweighted loss. Furthermore, when trained with unweighted loss, mixed tissue was segmented in non-negligible quantities only by the network trained in adversarial manner against a discriminator, and even then with poor performance. For both IVUS and OCT characterization, compensatory loss weighting facilitated the emergence of classification of less prevalent tissue, particularly (dense) calcium and necrotic core/mixed tissue, to the extent where these tissues were over-expressed in the classifier outputs. These qualitative and anecdotal observations are reinforced and quantified in the error (i.e. confusion) matrices for each set of methods (Tables 3.7 and 3.8).




Figure 3.22. U-net output for the semantic segmentation of tissue imaged by grayscale IVUS. The input to the network was the grayscale input image (*top left*) and a mask of the ROI (not shown; Figure 3.20). The target (“ground truth”) label map (*top right*) was extracted from the VH-IVUS frame corresponding to the grayscale image. The network was trained for 50 epochs, but advanced quickly, as seen in the output of the network (predicted maps) for the input image after 1, 25, and 50 epochs (*bottom*). However, it took several epochs to refine the media and non-pathological tissue (M; dark gray) classification and for finer details to emerge. Compensatory loss increased identification of less prevalent plaques (e.g. calcium; white) at the expense of precision.

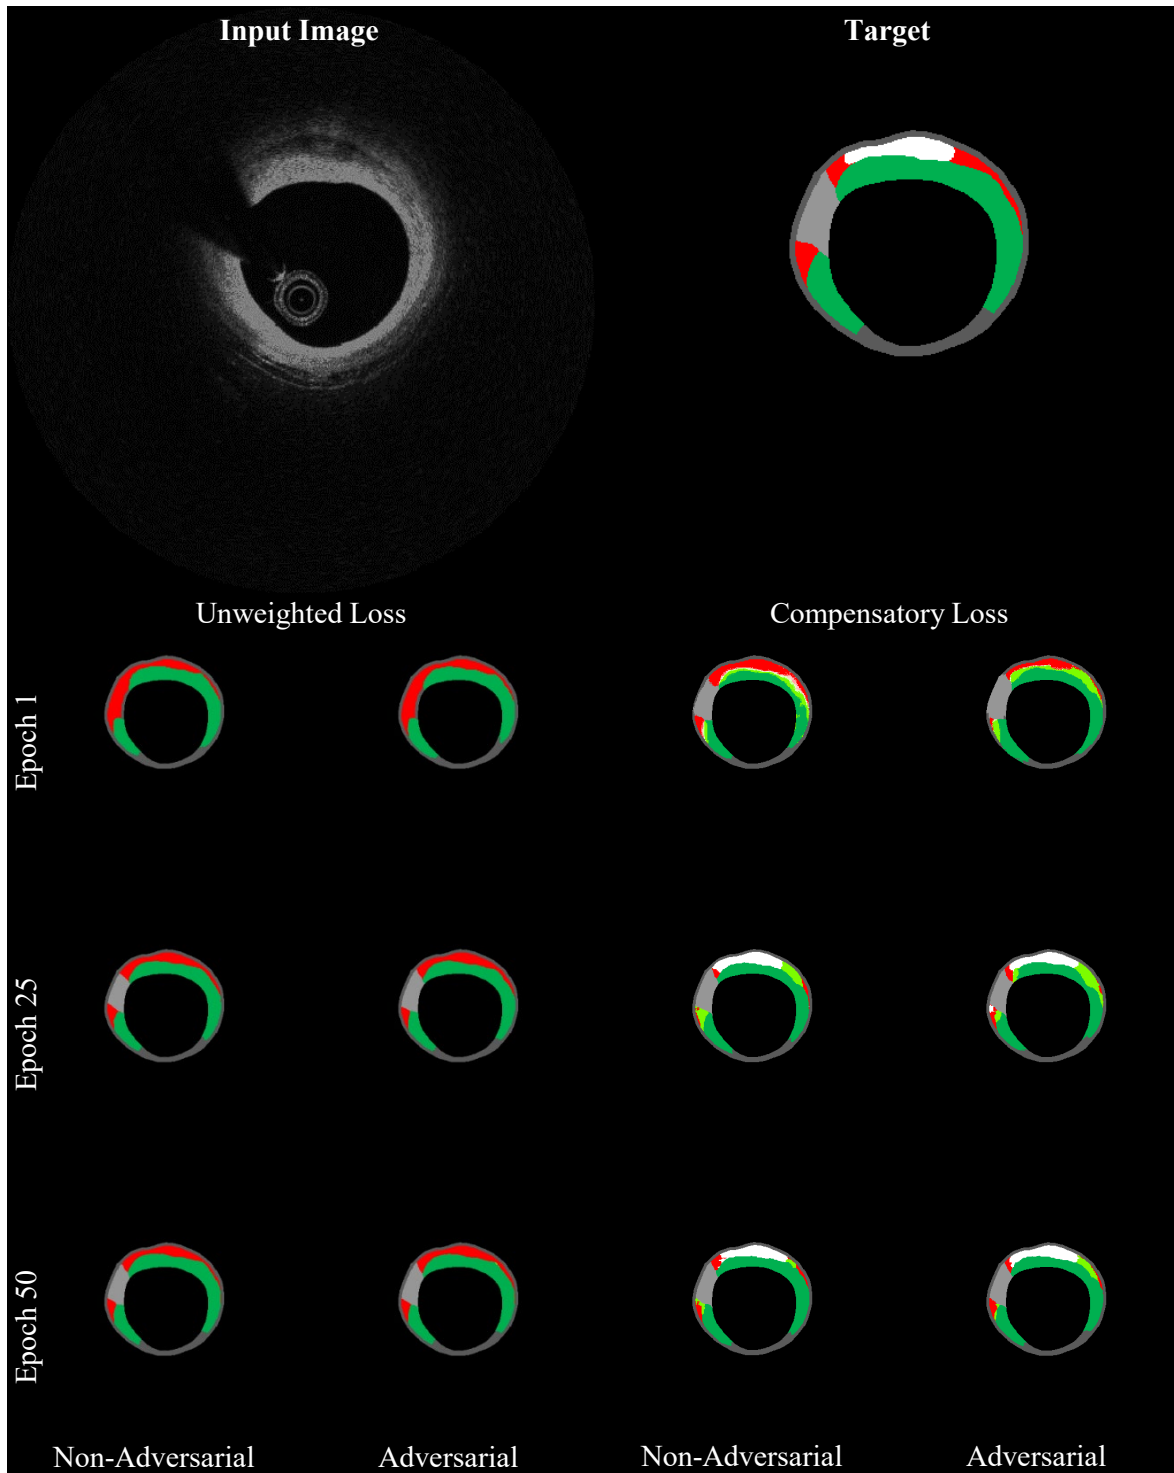


Figure 3.23. U-net output for the semantic segmentation of tissue imaged by OCT. The input to the network was the grayscale input image (*top left*) and a mask of the ROI (not shown). The target (“ground truth”) label map (*top right*) was determined by the fundamental framework described in Section 3.4 using CNN-based classification of the grayscale image. The network was trained for 50 epochs, but advanced quickly, as seen in the output of the network (predicted maps) for the input image after 1, 25, and 50 epochs (*bottom*). However, without compensatory loss, it took several epochs to refine the media and non-pathological tissue (M; dark gray) classification, for shadowed region (N; light gray) to be identified, and for calcium (C; white) and mixed tissue (MT) to emerge.

Table 3.7. IVUS semantic characterization method error for tissue classification following training with unweighted loss and compensatory loss weighting

<i>Unweighted Loss</i>								<i>Compensatory Loss Weighting</i>							
<i>Non-Adversarial</i>															
		Target Class					Precision			Target Class					Precision
Output Class	DC	NC	FT	FFT	M	DC		NC	FT	FFT	M				
DC	20900	5514	652	0	3151	69.2%	DC	26762	10006	1583	0	16669	48.6%		
NC	5066	61889	25135	6	7038	62.4%	NC	2549	69282	49939	47	10412	52.4%		
FT	381	20211	383093	74533	7540	78.9%	FT	96	9309	289309	21474	4947	89.0%		
FFT	103	1	54576	154389	2642	72.9%	FFT	80	67	123683	208284	4514	61.9%		
M	3497	2289	10346	3884	373056	94.9%	M	460	1240	9288	3007	356885	96.2%		
Recall 69.8% 68.8% 80.9% 66.3% 94.8%						81.4%	Recall 89.4% 77.1% 61.1% 89.5% 90.7%						77.9%		
<i>Adversarial</i>															
		Target Class					Precision			Target Class					Precision
Output Class	DC	NC	FT	FFT	M	DC		NC	FT	FFT	M				
DC	18182	2781	247	0	1854	78.8%	DC	26664	9468	1125	0	16548	49.6%		
NC	7319	56647	15598	8	6694	65.7%	NC	2765	70507	53562	60	11325	51.0%		
FT	647	27896	389674	72104	9178	78.0%	FT	115	8956	317963	40596	5055	85.3%		
FFT	66	3	58020	156534	1982	72.3%	FFT	72	12	91417	188551	3815	66.4%		
M	3733	2577	10263	4166	373719	94.7%	M	331	961	9735	3605	356684	96.1%		
Recall 60.7% 63.0% 82.2% 67.2% 95.0%						81.5%	Recall 89.0% 78.4% 67.1% 81.0% 90.7%						78.7%		

Table 3.8. OCT semantic characterization method error for tissue classification following training with unweighted loss and compensatory loss weighting

<i>Unweighted Loss</i>								<i>Compensatory Loss Weighting</i>							
<i>Non-Adversarial</i>															
		Target Class					Pre.			Target Class					Pre.
Output Class	C	LT	FT	MT	N	M		C	LT	FT	MT	N	M		
C	0	0	0	0	0	2	0.0%	C	10264	15401	5479	4624	381	975	27.6%
LT	18658	242235	5914	21950	2141	1054	83.0%	LT	8569	216859	6066	5509	879	4168	89.6%
FT	5137	16485	356211	29422	1976	2949	86.4%	FT	1857	10871	341924	7455	872	15102	90.4%
MT	0	0	2	1	0	1	25.0%	MT	3184	12897	9353	34304	875	2135	54.7%
N	113	5802	820	549	75406	450	90.7%	N	64	8477	1564	218	76762	3448	84.8%
M	84	1103	3309	388	504	513103	99.0%	M	54	1120	1870	200	258	491731	99.3%
Rec. 0.0% 91.2% 97.3% 0.0% 94.2% 99.1%						90.9%	Rec. 42.8% 81.6% 93.4% 65.6% 95.9% 95.0%						89.7%		
<i>Adversarial</i>															
		Target Class					Pre.			Target Class					Pre.
Output Class	C	LT	FT	MT	N	M		C	LT	FT	MT	N	M		
C	0	0	3	0	0	3	0.0%	C	10559	27639	5449	3983	193	686	21.8%
LT	19731	245969	11359	23893	1488	1300	81.0%	LT	5880	195372	4390	3149	678	3422	91.8%
FT	3524	9940	348550	22529	1383	2173	89.8%	FT	1582	8637	327224	3607	363	5757	94.3%
MT	514	1711	753	4841	171	26	60.4%	MT	5795	23316	20968	40885	769	1953	43.6%
N	157	6969	1097	574	76679	800	88.9%	N	104	8778	2747	230	77665	3262	83.7%
M	66	1036	4494	473	306	513257	98.8%	M	72	1883	5478	456	359	502479	98.4%
Rec. 0.0% 92.6% 95.2% 9.3% 95.8% 99.2%						91.1%	Rec. 44.0% 73.6% 89.3% 78.2% 97.0% 97.1%						88.4%		

Rec.: Recall; Pre.: Precision

Results also revealed some intriguing and promising attributes for this variation of the device segmentation framework, alongside remaining challenges. Of particular interest, when complementary loss weighting was not used in training the classifier, the network learned to identify *all* pixels as non-strut. With just over 0.2% of all pixels in an image belonging to this class, such a strategy actually yielded the best overall accuracy, despite offering no value (Table 3.9). With the introduction of compensatory loss weighting, behavior shifted dramatically. Figure 3.24 show illustrative examples for the resulting segmentation of BVS struts in an OCT frame by the two U-net classifiers trained through adversarial and non-adversarial means. Similar to behavior observed with respect to less prevalent tissue, with substantially higher penalty accrued for inaccurately classifying a strut than a non-strut pixel, a conservative strategy emerged that overestimated strut location. Consequently, sensitivity of strut achieved very high rates even as specificity remained tenuous (Table 3.9). Still, the results were promising, and suggested that there is likely a more effective loss weighting strategy which could achieve excellent results.

Table 3.9. OCT semantic segmentation method error for BVS strut labeling following training with unweighted loss and compensatory loss weighting

<i>Unweighted Loss</i>				
		Target Class		
		Strut	Non-Strut	Precision
Output Class	Strut	0	0	-
	Non-Strut	238078	73424386	99.7%
Recall		0.0%	100.0%	99.7%
<i>Non-Adversarial with Compensatory Loss Weighting</i>				
		Target Class		
		Strut	Non-Strut	Precision
Output Class	Strut	233832	295543	44.2%
	Non-Strut	4246	73128843	100.0%
Recall		98.2%	99.6%	99.6%
<i>Adversarial with Compensatory Loss Weighting</i>				
		Target Class		
		Strut	Non-Strut	Precision
Output Class	Strut	236326	603896	28.1%
	Non-Strut	1752	72820490	100.0%
Recall		99.3%	99.2%	99.2%

Overall, semantic segmentation performance was worse in the identification of all classes (Table 3.10). Many of the challenges encountered by CNN-based classification, such as the differentiation of DC and NC in IVUS, FT and FT tissue in IVUS, C and LT in OCT, or MT, FT, and LT in OCT, and the tradeoffs between scarce strut segmentation sensitivity and specificity, were similarly encountered by the U-net-based classification (Tables 3.7, 3.8, and 3.9, cf. Tables 3.2, 3.5, and 3.6). However, there were some classes for which the alternative classifier achieved strong recall and precision, and for the tissue class distribution encountered in the testing dataset, overall accuracy was fairly strong—around 81% for IVUS and 91% for OCT. The use of adversarial or non-adversarial training made little systematic difference overall, with the notable exception of MT emerging as a non-negligible class in the OCT classifier network that underwent

adversarial training. Compensatory weighting of loss contributions, however, had a profound impact on overall performance, particularly in the classification of less prominent classes.




Figure 3.24. U-net output for the semantic segmentation of BVS struts imaged by OCT. The input to the network was the grayscale input image (*top left*). Unlike for the tissue segmentation task, no ROI mask was provided. The target (“ground truth”) label map (*top right*) was determined by manual annotation, as described in Section 3.5. The output of the network (predicted maps) are shown for the input image after 50 epochs of training under adversarial and non-adversarial regimes (*bottom*). Compensatory loss was implemented in which loss for strut-labeled pixels was magnified by a weight equal to the inverse of strut pixel prevalence. When unweighted loss was used, the network learns to never positively identify pixels and consistently produces empty maps.

Table 3.10. Summary of tissue classifier and strut segmentation performance for various networks and training regimes implemented in the characterization framework

Network	Recall					Precision					W c	
	CNN	U _{NA}	U _A	U _{NA-C}	U _{A-C}	CNN*	U _{NA}	U _A	U _{NA-C}	U _{A-C}		
IVUS	DC	98.5%	69.8%	60.7%	89.4%	89.0%	91.3%	69.2%	78.8%	48.6%	49.6%	36.4
	NC	88.6%	68.8%	63.0%	77.1%	78.4%	95.9%	62.4%	65.7%	52.4%	51.0%	13.1
	FT	91.1%	80.9%	82.2%	61.1%	67.1%	88.5%	78.9%	78.0%	89.0%	85.3%	2.7
	FFT	90.0%	66.3%	67.2%	89.5%	81.0%	93.4%	72.9%	72.3%	61.9%	66.4%	5.2
	M	99.4%	94.8%	95.0%	90.7%	90.7%	99.1%	94.9%	94.7%	96.2%	96.1%	3.1
OCT	C	72.3%	0.0%	0.0%	42.8%	44.0%	34.0%	0.0%	0.0%	27.6%	21.8%	48.4
	LT	93.0%	91.2%	92.6%	81.6%	73.6%	95.4%	83.0%	81.0%	89.6%	91.8%	5.1
	FT	96.3%	97.3%	95.2%	93.4%	89.3%	97.9%	86.4%	89.8%	90.4%	94.3%	3.3
	MT	84.2%	0.0%	9.3%	65.6%	78.2%	79.8%	25.0%	60.4%	54.7%	43.6%	20.9
	N	98.8%	94.2%	95.8%	95.9%	97.0%	99.9%	90.7%	88.9%	84.8%	83.7%	14.9
	M	-	99.1%	99.2%	95.0%	97.1%	-	99.0%	98.8%	99.3%	98.4%	2.7
BVS	S	99.5%	0.0%	0.0%	98.2%	99.3%	99.3%	-	-	44.2%	28.1%	372.4
	NS	99.3%	100.0%	100.0%	99.6%	99.2%	99.5%	99.7%	99.7%	100.0%	100.0%	1.0

*Note that the pixel-based CNN was evaluated in a class-balanced testing dataset, so precision is skewed relative to true performance in an imbalanced population; U-nets were evaluated in the full population testing dataset.

CNN: Pixel-wise CNN classification (Sections 3.3 & 3.4); U: U-net semantic segmentation (Section 3.6);

NA: Non-adversarial training; A: Adversarial training; C: Compensatory (weighted) loss; W: Compensatory loss class weight; S: Strut; NS: Non-strut.

Colors span from red (0%) through white (median value; 89.0%) to green (100%).

Hyphen (-) indicates value is not defined.

A major tradeoff of the degraded performance of the semantic segmentation technique was a massive increase in both training and execution speed. While the CNN network took weeks to train (over 50 epochs), the tissue characterization U-nets took just 171.4 ± 0.8 seconds per epoch (less than 2.5 hours for the full 50 epochs). Once trained, full-frame segmentation took only 34.2 ± 5.4 milliseconds (cf. 200 ± 150 seconds for full-frame characterization with the pixel-based methods), representing a 99.98% speed increase (decrease in execution time). The strut segmentation U-nets were even faster due to their simplified input and output structures. Training took just 73.9 ± 0.4 seconds per epoch (1.0 hours for the full 50 epochs). Once trained, full-frame segmentation took only 26.5 ± 1.5 milliseconds, a remarkable benefit with respect to the pixel-based method.

In addition to gaps in classification performance, several challenges remain to be addressed. Training required fully-annotated images. This approach was therefore very amenable to training with VH, which provides such label maps, but will require modification to accommodate piecemeal manual annotations like those typically available for OCT. Here, the target classification map utilized to train on OCT was derived itself from the pixel-wise classification. This remains a clear limitation, though provided a useful opportunity to compare attainable performance.

Additionally, the impact of adversarial training remains unclear. Networks were trained just a single time on a single set of data. Repeated training and cross-validation, and a parametric analysis of the weighting factor between categorical- and discriminator-driven losses, would help to illuminate the effect of adversarial training and, more broadly, the stability and generalizability of this technique within the context of the characterization framework. Refined methods to (compensatory) weigh loss contributions of various plaque types should also be explored to improve accuracies of tissue classes less prevalent in the dataset while mitigating negative impacts on specificity. Finally, with a constant eye towards application and its context, an assessment of sufficient performance (e.g. details and changes which would effect a change in diagnosis or treatment strategy) would frame the applicability and focus future efforts to address current shortcomings of the method.

The development and implementation of a U-net to classify the designated region of interest within the characterization framework demonstrated the flexible modularity of the approach while also highlighting an exciting area for further development. While additional research is needed to substantiate and address several shortcomings of the neural network-based classification scheme presented here, there are immediate opportunities for application. As noted earlier, the slow pixel-wise CNN-based classification process was not optimized, and it was suggested that a more strategic approach to pixel selection could dramatically improve execution speed without substantial loss of information. Rather than selecting between classification strategies, the semantic segmentation process could provide an initial estimate of the plaque morphology which is subsequently selectively refined by the more accurate, slower network, thereby achieving potentially orders-of-magnitude improvement in execution speed.

With the effective methodological framework for intravascular image characterization applied to several applications and its flexibility firmly established, studies addressing the importance of domain enrichment and training data balance offered further insight to inform future development and application, with broader implications.

3.7. Importance of Domain Enrichment

Many algorithmic learning frameworks are entirely data-driven, which makes them particularly susceptible to degraded quality and limited quantity of training data. These challenges are particularly prevalent in medical imaging. Growing recognition of these obstacles has led to the increased development and adoption of approaches which enrich learning frameworks with domain knowledge and signal structure [288]. In the framework presented in the previous sections (Section 3.2), the primary mechanism by which domain knowledge enriched learning to address the image classification problem was the pathological tissue detection procedure. Pathological tissue was partitioned from the rest of the vessel wall based upon domain knowledge of spatial constraints imposed by arterial physiology and pathology. Physically-relevant limits were thereby enforced for the location and dimensions of this tissue class while also reducing the number of classes to be subsequently segmented by the CNN; following this step, classification was only required for the remaining (pathological) tissue types (and shadow, for OCT). To investigate the utility of leveraging domain enrichment, an equivalent “naïve” method was implemented for IVUS where non-pathological tissue was not first segmented from the pathological tissue prior to CNN segmentation, but was instead segmented as an additional class (Figure 3.25) [161].




Figure 3.25. To investigate the utility of leveraging domain enrichment, an equivalent “naïve” method was implemented where non-pathological tissue was not first segmented from the pathological tissue prior to CNN segmentation, but was instead segmented as an additional class.

For the naïve method developed to assess the importance of the contribution, division of the ROI into pathological and non-pathological tissue was not completed. Instead, subsequent classification routines were taught to detect non-pathological and media tissue directly from the image patch data. That is, for the domain enriched method, pixel-centered patches were created for pixels of the ROI remaining after segmenting the M class ($r_{im} \in ROI_{path}$), then automatically classified into one of the four plaque types using a CNN. However, for the naïve method, patches were created for all pixels of the ROI ($r_{im} \in ROI$) and sorted into one of the five tissue types by the classifier.

As anticipated, the domain enriched method provided better results than the naïve one in image segmentation of grayscale IVUS images emulating VH-IVUS classification. Tables 3.2 and 3.11 provide the error (or confusion) matrices for the enriched and naïve methods, respectively, showing that the former achieved an overall accuracy of 93.5% and the latter just 87.8%. Direct comparison of the performance, evaluated in the exact same dataset, is shown in Table 3.12. Performance metrics by tissue class are summarized and compared in Figure 3.26.

Table 3.11. Naïve IVUS tissue classifier error matrix for full-ROI segmentation

		Target Class					Precision
		DC	NC	FT	FFT	M	
Output Class	DC	49371	1840	1	0	123	96.2%
	NC	506	43652	978	0	2040	92.5%
	FT	0	477	43747	2168	4222	86.4%
	FFT	0	0	1408	44738	5686	86.3%
	M	123	4031	3866	3094	37929	77.3%
Recall		98.7%	87.3%	87.5%	89.5%	75.9%	87.8%

Table 3.12. Relative performance of naïve IVUS classifier compared to domain-enriched characterization method

		Target Class					Precision
		DC	NC	FT	FFT	M	
Output Class	DC	124	-2824	1	0	120	+4.90%
	NC	-158	-658	-230	0	2010	-3.40%
	FT	0	-519	-1817	-2593	4063	-2.10%
	FFT	0	0	-1699	-277	5602	-7.10%
	M	34	4001	3745	2870	-11795	-21.80%
Recall		+0.20%	-1.30%	-3.60%	-0.50%	-23.50%	-5.70%

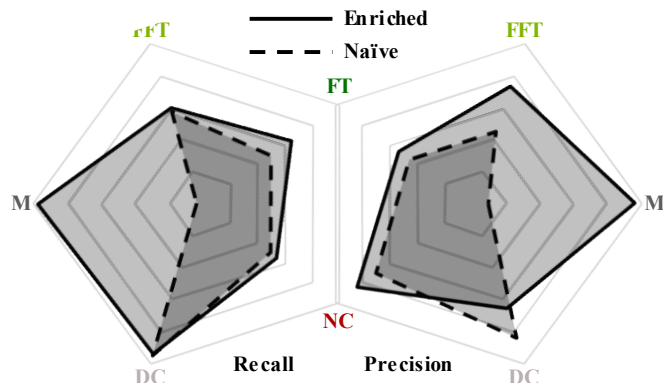


Figure 3.26. Comparison of recall (i.e. sensitivity) and precision (i.e. positive predictive value) achieved by the enriched and naïve methods (shown with solid and dashed borders, respectively). The enriched method demonstrates clear superiority, particularly, but not exclusively, in categorizing M class tissue. Axes range from 75% to 100% (linear scale from center to perimeter).

Performance of both the naïve and enriched deep learning methods described here were compared with that reported for existing methods in Table 3.1, which also includes validation metrics reported for VH-IVUS.

Representative examples of classified images resulting from each method are shown in Figure 3.27. Both methods accurately captured major tissue morphology and features within the pathological region. However, the naïve method struggled to identify non-pathological and media tissue, and occasionally generated physiologically implausible configurations (Figure 3.27). Due to the spatial constraints imposed prior to CNN classification, the domain enriched method addressed non-pathological and media tissue very accurately, and was not disposed to violating physiological constraints. It captured fine features and provided sharp distinctions between various plaque types, generating images that were more similar to gold standard VH-IVUS.

While the naïve method performance metrics (Table 3.11) reflect only the five-class CNN classifier, as the classifier itself performs all segmentation operations, the overall domain enriched method metrics (Table 3.2) depend both on (four-class) classifier performance and reliability of pathological tissue detection, which together share responsibility for the full segmentation procedure. Both CNN classifiers, trained only on pixels classified by VH-IVUS, achieved generally high precision (i.e. PPV) and recall (i.e. sensitivity). Table 3.3 shows the error matrices for the enriched method's four-class CNN classifier—the model achieved an accuracy of 92.3%. As noted above, results of the CNN for the 5-class model are the same as those for its corresponding complete segmentation method (Table 3.11)—the model achieved an accuracy of just 87.8% overall (90.8% for just the same four tissue classes identified by the enriched method's CNN classifier). CNN training took several weeks (roughly 3 days per epoch for the 5-class model and somewhat less for the 4-class model). Training was halted once accuracy and loss plateaued; with further training, validation metrics deteriorated, indicating overfitting of the model to training data.




Figure 3.27. Sample classified regions of IVUS images segmented by VH-IVUS (ground truth) and both the domain enriched and naïve methods. Both presented methods identified major pathological tissue morphology features quite well, but the enriched method demonstrated clear superiority. In these examples, the naïve method misclassified much of the non-pathological and media tissue and proposed several variations of physiologically non-feasible morphologies. These physiological impossibilities included islands of non-pathological tissue embedded within a diseased region (A–E), exaggerated, thick segments of healthy (normally-thin) intima or media tissue (C), and calcified and lipid deposits within exceptionally thin wall segments (A, B). Light blue hash marks within each image demarcate 1 mm increments.

Error matrices of the classifiers illustrate some general and model-specific trends. Both classifiers—the five-class network supporting the naïve method and the four-class network supporting the domain enriched method—struggle to differentiate FFT from FT and DC from NC. Notably, while class confusion trends were universally observed for both models, performance was worse in all cases for the 5-class CNN except in the task of identifying calcium (DC). Furthermore, classification of the media by this model was only mediocre—pixels belonging to the M class were often misclassified as FT, FFT, or NC, and these tissues were conversely misclassified as M with moderate frequency (Table 3.11). These findings show that imposing

spatial constraints to determine non-pathological and media tissue prior to CNN classification, and excluding this class from classification, not only improved segmentation of this non-diseased tissue type, but that of the classified plaque as well. However, the enriched model was still subject to compounding uncertainties arising from pathological tissue delineation. Notably, while delineation of pathological tissue was very accurate, the CNN of the enriched method was incapable of classifying M tissue it encountered (unlike its naïve counterpart).

Performance with respect to execution speed and efficiency was also impacted by domain enrichment. As noted for the enriched method (Section 3.2), execution time of the characterization method was dominated by the pixel-wise network classification of the ROI. The enriched method determined that only a subset of pixels within the ROI were pathological and subsequently classified by the network, whereas the naïve method classified all pixels within the entire ROI. Efficient and rapid vessel wall border detection methods exist, and designation of the media and non-pathological tissue in a frame was exceptionally fast (25.5 ± 0.9 milliseconds per frame), resulting in a net reduction in execution time with domain enrichment; characterization of full frames took 200 ± 150 seconds and 280 ± 170 seconds with the enriched and naïve methods, respectively. Due to the additional class of data needed for training and the additional parameters of the larger CNN, training of the naïve method's CNN was also slower than that of the domain enriched method's CNN. Therefore, both preparation and subsequent performance was made faster and more efficient by domain enrichment.

In many ways, the benefits of applying the domain knowledge to segment the non-pathological and media tissue were foreseeable and expected. Clinical expert consensus reported by the American College of Cardiology and developed in collaboration with the European Society of Cardiology maintains that, while the trailing edge of the media (media-adventitia border) is generally well delineated in IVUS images, the leading edge is not [19]. Automated edge detection therefore only extracts lumen (lumen-intima) and media-adventitia borders, and the resulting wall area analyzed is consequently the plaque plus media area [19]. It is not surprising, then, that a CNN would have difficulty distinguishing the media from surrounding tissue within this region of a grayscale image, since the echoreflectivity profile is not conducive to distinctive transitions and the region is typically not distinguishable even by trained experts. Furthermore, the spatial invariance intrinsically assumed by CNNs—generally one of their great assets in image processing—here is a liability, as the media is spatially constrained between the intima (where

plaque develops) and the adventitia layers of a blood vessel. Therefore, utilizing *a priori* knowledge, derived previously from studies using alternative visualization modalities and mechanisms (e.g. histology [8], [164]), provided strong benefit. Furthermore, imposing geometric constraints based in physical reality made the method more robust to poor image quality and artifacts by preventing impossible class configurations. And finally, reducing the number of classes improved classification accuracy, precision, and specificity by the CNN for all but one of the remaining classes while also reducing the number of pixels to be classified, thereby decreasing execution time.

Generalizability of the specific benefits and quantitative gains reported should not be assumed for other datasets or modalities, though general trends regarding the impact of domain enrichment are expected to hold. For example, while the media is not distinguishable in IVUS, this is not the case in OCT, where the media is visible and evident in regions without extensive plaque burden (which can attenuate the signal such that the media is located beyond the local penetration depth). Therefore, a trained CNN may be more adept in classifying the media in OCT images than one trained for IVUS images. However, distinguishing non-pathological from pathological intimal tissue without imposing clinical knowledge of pathology would remain an unresolved challenge, as would the avoidance of non-feasible morphologies. Furthermore, the inclusion of an additional class would be expected to reduce performance in segmenting the original classes, even if the new class could be identified with reasonable accuracy. And finally, benefits in training time, data requirements, execution speed, and efficiency are all expected to hold across applications.

The impact of leveraging domain knowledge to distinguish pathological from non-pathological tissue prior to CNN classification was assessed, and was found to offer substantial benefit. In particular, enforcing physiologically-imposed spatial constraints to assign the non-pathological and media tissue class not only improved classification performance for this class, but also benefited classification of the remaining pathological tissue types and decreased execution time. Application of this domain knowledge further prevented various forms of unrealistic morphologies that arose in the unconstrained naïve model. Although reasonably strong performance was achieved without domain enriched learning, leveraging domain knowledge in learned algorithmic tasks offers several benefits and should be strongly considered even when not strictly necessary, particularly when designing approaches for analyzing medical imaging.

3.8. Importance of Training Data

Tremendous progress has been made in machine learning for health and biology, particularly as applied to biomarkers, radiomics, and other imaging. However, work is often presented without adequate reflection on the models and performance metrics. The impact of training data on neural network classification, particularly applied to medical image segmentation, is a phenomenon that is often grappled with, but infrequently quantified. For example, among the previous methods for characterizing tissue in intravascular OCT images, there was a tremendous range of performance reported (Figure 3.32), yet the cause of the differences—the relative contributions and importance of different approaches and models or the underlying data—has not been examined and is difficult to discern. An entirely new model was developed to address this uncertainty. While the results were largely intuitive, the stark contrast of the findings, and the magnitude of the observed effect, offered a challenging lens through which to interpret work in the field.

To analyze the suspected impact of training data on characterization performance variability, a novel yet simple method of automatic calcified plaque detection was devised that allowed precise control of class distribution in the dataset, as described by Gowrishankar *et al.* [162]. Texture measures were used as features to train a neural network to automatically identify calcified regions of an arterial wall; a multilayer perceptron (MLP), a type of FCNN, was created to map the desired output to the input in the training set. In an MLP, the value at each node in a layer is based on the input from nodes in the layer above it, and the single node in the final layer yields an output based on the input parameters. During training, an error value based on the difference between the desired output and the output of the network is computed and back-propagated to adjust the weights and biases of the MLP to minimize the error. This process enables the network to converge to a model that predicts pixel class for any input set of local texture features. However, the convergence point is highly dependent upon the training data; the error that is minimized in the course of training depends entirely on this data, such that the model is optimized for these. Because the local texture of each pixel was considered individually, pixels of known class could be assigned to the training dataset in any desired quantity. By varying the relative representation of pixels belonging to either class, the impact of the training data on overall model performance was quantified—segmentation results of the proposed method were validated

against expert annotations. While calcium in OCT images can be automatically detected with meaningful precision using local texture measures and traditional neural networks in a limited dataset, training data balance was remarkably consequential.

3.8.1 Materials and Methods

The process of creating the classification model to detect calcified plaques is outlined in Figure 3.28. Pixel texture features from cross-sectional OCT images of calcified coronary artery segments were used as input to a neural network. Outputs were compared with manually assigned pixel designations, and were counted as correct if they matched and incorrect otherwise.

Region of Interest Delineation

The inner boundary of the artery wall and a region of calcium were annotated in each image by a medical expert (Figure 3.29, center). The artery wall was approximated in each image as the region between the annotated inner boundary of the artery wall and its 1.5 mm radial projection with sufficient grayscale intensity (Figure 3.29, right) [173].

Local Texture Features

Various image feature measures were tested for use in classifier training. Edge detection methods, gradients, and standard deviation were all assessed, but texture features from the gray-level co-occurrence matrix (GLCM) ultimately yielded the best results. Therefore, a set of six local texture features [289] were used to distinguish calcified regions from non-calcified regions in the OCT image [290], [291]:




Figure 3.28. Flow and architecture of the texture-based calcium detection method. The annotated lumen was extended radially outwards to delineate the ROI within the OCT image, from which a GLCM was computed for each pixel. Six texture features of the GLCM served as inputs to the MLP, which output a binary classification.




Figure 3.29. Determining and labeling the region of analysis for calcium detection in OCT images. The original image (*left*) was annotated with lumen and calcium (*center*). A mask of the artery wall/ROI (*right*) was based upon a radial projection of the lumen and grayscale intensity.

$$\text{Angular Second Moment (ASM)} = \sum_{i,j=0}^{N-1} (P_{i,j})^2, \quad (58)$$

$$\text{Contrast} = \sum_{i,j=0}^{N-1} P_{i,j} (i - j)^2, \quad (59)$$

$$\text{Correlation} = \sum_{i,j=0}^{N-1} P_{i,j} \frac{(i - \mu_i)(j - \mu_j)}{\sqrt{\sigma_i} \sqrt{\sigma_j}}, \quad (60)$$

$$\text{Dissimilarity} = \sum_{i,j=0}^{N-1} P_{i,j} |i - j|, \quad (61)$$

$$\text{Energy} = \sqrt{\text{ASM}}, \text{ and} \quad (62)$$

$$\text{Homogeneity} = \sum_{i,j=0}^{N-1} \frac{P_{i,j}}{1 + (i - j)^2}, \quad (63)$$

where N is the number of gray levels and $P_{i,j}$ is defined as:

$$P_{i,j} = \frac{C_{i,j}}{\sum_{i=0}^{N-1} \sum_{j=0}^{N-1} C_{i,j}}. \quad (64)$$

Here, $C_{i,j}$ represents the number of occurrences of gray levels i and j within the window and μ_i , μ_j , σ_i , and σ_j are defined as:

$$\mu_i = \sum_{k=0}^{N-1} k \sum_{l=0}^{N-1} P_{k,l}, \quad (65)$$

$$\mu_j = \sum_{k=0}^{N-1} l \sum_{l=0}^{N-1} P_{k,l}, \quad (66)$$

$$\sigma_i = \sum_{k=0}^{N-1} (k - \mu_i)^2 \sum_{l=0}^{N-1} P_{k,l}, \text{ and} \quad (67)$$

$$\sigma_j = \sum_{k=0}^{N-1} (l - \mu_j)^2 \sum_{l=0}^{N-1} P_{k,l}. \quad (68)$$

A 7×7 pixel neighborhood around each pixel in the ROI was extracted from the OCT image. For each neighborhood, a GLCM was computed for distance 1 and angle 0. From the GLCM, the texture measures defined above in Equations 58 through 68 were calculate for each pixel.

Neural Network Architecture and Training

Different classification algorithms were initially tested, including a k-nearest neighbor classifier, a limited-memory BFGS neural network, and a neural network from the Keras API. The Keras sequential neural network model yielded the highest accuracy, so this model was selected for further development.

An MLP classifier neural network was constructed with three densely connected layers (Figure 3.28). The MLP used the six texture measures as input and output a value between 0 (non-calcium) and 1 (calcium). The first layer contained 12 neurons and the second contained 8 neurons, each of ReLU type (giving a positive output equal to the maximum of the input or zero). The output of the hidden layer was fed to a final layer with a sigmoidal activation function. This gave an output in the range of 0 to 1; if the output was greater than 0.5 the pixel was classified as calcium, and non-calcium otherwise.

The model used binary cross-entropy loss to fit the training data. Cross-entropy (Equation 54) measures how much the output of the neural network differs from the labeled output: a low value of cross-entropy loss means that the output of the neural network is close to the true value. The Adam optimization algorithm—a modification of the SGD procedure—sped the iterative updating of weights. Training data were fit with a batch size of 50 over 5 epochs.

3.8.2 Dataset

Imaging data were obtained from 8 patients undergoing examination of coronary stenosis of at least intermediate severity. After obtaining patient consent, invasive angiography and frequency domain OCT were performed in accordance with standard protocol [18], in some cases following implantation of a metallic stent. The OCT system acquired 500-line frames with $15 \mu\text{m}$ resolution at 100 frames/second. Only frames depicting at least one region of calcium were utilized, resulting in a dataset of 30 images. Ground truth was provided by a medical expert; the region of calcified plaque, as well as the inner wall of the artery, was annotated in each image by a cardiologist.

Pixels in the ROI of all images were used. Based on the expert annotations, each pixel was assigned a value of either 1, designating a pixel that resided within a region of calcified plaque, or 0, designating a pixel depicting non-calcified tissue. Within the 30 OCT images, only 3.4×10^4 pixels fell within regions classified as calcium; the other 5.1×10^5 pixels did not, resulting in a calcium to non-calcium ratio of 15:1. Pixels of each type were then pooled and randomly divided into testing and training sets. Half of the calcium pixels (1.7×10^4) was used as a training set and the other half as a testing set; non-calcium pixels were similarly divided evenly between training and testing sets. The model was subsequently trained with different numbers of pixels from non-calcified regions; the number of pixels from calcified regions was constant throughout and the set size changed accordingly.

3.8.3 Results

Performance of the neural network model for automatic classification of calcium regions was measured in terms of accuracy and loss. Figure 3.30 shows three examples of the classification of calcified regions in the image by the neural network. Images in the center column show the regions annotated by human experts; they were not used in generating the images in the right column. The pixels that were assigned 1 (probability output >0.5) in the output layer of the neural network are shown in white in the classified image (right). Most of the calcified regions in the annotated image were correctly identified by the neural network model, but the output did contain some regions of false positive and false negative (misclassified) pixels. To quantitatively validate the local texture-based neural network method, PPV, sensitivity, and F_1 score were calculated for both classified and non-classified pixels, and their averages (weighted by prevalence in the testing dataset) were found to be close to 95% when all available training data was used (Table 3.13). PPV value denotes precision of the method, while F_1 score is the harmonic (or subcontrary) mean of precision and sensitivity.

Table 3.13. Sensitivity, PPV, and F_1 score of the calcium detection method obtained with the complete training dataset

Pixel Class	PPV	Sensitivity	F_1 Score
1 (classified)	0.69	0.35	0.47
0 (non-classified)	0.95	0.99	0.97
Average (overall)	0.94	0.95	0.94




Figure 3.30. Three sets of original image (*left*), manually annotated image (*center*), and neural network texture-based classifier output (*right*) showing regions of calcium in white. Classification results (*right*) are shown for the full training dataset (15:1 ratio of non-calcium pixels to calcium pixels, shown diagrammatically in the top right).

The calcified regions were much smaller in area than non-calcified regions in the artery walls: the total number of pixels in calcified regions was 15 times smaller than the number of pixels in non-calcified regions in the full dataset. This made the data in the complete training set, as well as the testing set, severely imbalanced. To quantitatively assess how training data balance affects precision and sensitivity, performance metrics resulting from different ratios of classified pixels were compared (Figures 3.32 and 3.33). Figure 3.31 illustrates the associated qualitative performance in a demonstrative classified image. As fewer non-calcium pixels were included in




Figure 3.31. Classification results with different ratios of non-calcium pixels to calcium pixels in the training dataset: 1:1 (*left*), 2:1 (*center*) and 4:1 (*right*). The subset of the total available training data used is shown diagrammatically in the top right of each result. (Compare with Figure 3.30.)

the training set, the algorithm identified more pixels as calcium; as a result, both true and false positive rates increased. That is, as the number of non-calcium pixels is increased, the model classified fewer pixels as calcium overall, decreasing its sensitivity for identifying calcium (Figures 3.31 and 3.33).




Figure 3.32. Performance comparison of the presented method (local texture classification) trained with extreme distributions of class representation ratios (15:1 and 1:1) and other calcium detection methods previously reported in the literature. PPV was not reported by He *et al.*; sensitivity is that reported for calcium.




Figure 3.33. As training sets were balanced (representation ratio approached 1), sensitivity and PPV for calcium improved while overall accuracy worsened.

3.8.4 Discussion

Coronary calcified plaque is an important indicator of atherosclerosis and thrombotic syndromes [292]. At present, however, these lesions are manually identified in intravascular images, which is time consuming and highly dependent on well-trained readers. As for many other biomarkers embedded in dense digital data, automated detection has attracted significant interest, though universally satisfactory methods remain elusive. Galvanized by the disparate reported performance of previous approaches [160], [256]–[258], a pixel-based neural network method was developed for automatic detection of calcified plaques in OCT images that allowed modulation of well-controlled training datasets to achieve a range of pixel type (class) representation ratios. Even using small datasets, calcification in arterial walls could be automatically detected with high precision and sensitivity using neural network approaches with the proper data distribution, but at a high cost to specificity. Results were strongly subject to training data class distribution.

The range of sensitivity and PPV achieved by varying training datasets nearly spanned the entire range reported for previous calcium classification methods [160], [256]–[258] (Figure 3.32). Of note, performance achieved with a balanced training dataset more closely resembles that of

methods tuned to calcium [256], [257]—training dataset representation was not reported for these methods—while that achieved with the imbalanced dataset resembled that of methods utilizing multiple classes [160], [258], which similarly report using imbalanced training datasets (which more closely reflect the underlying imbalanced distribution of the tissue itself). However, each method utilized different training datasets with varying class representation, prohibiting true direct comparison.

There are some minor confounding factors to consider in evaluating the performance of the presented method. For example, the expert annotated only a single large region of calcium in each image. However, there could have been smaller regions of calcification in the artery wall which may not have been included in the annotation. It is possible that the neural network was able, in some cases, to identify these regions as well, but was penalized for doing so. Additionally, the region of analysis was defined artificially using a projection of the lumen by a fixed distance. This could lead to undesired inclusion of perivascular tissue in the dataset and exclusion of plaque (including the characteristic distinct outer border of calcified plaque). However, future work to improve the region of analysis could utilize a method for delineating the full wall of diseased arteries like that described in the previous chapter [157]. Finally, in this demonstration, only 30 images were utilized, but performance would likely improve for all class distributions with a larger dataset.

Nevertheless, the domineering impact of training data balance should not be overlooked. Indeed, performance achieved with datasets of varying distributions nearly spanned the full range reported for the various previously-reported methods, suggesting that the training dataset might even surpass segmentation approach or model architecture as the dominant factor in determining performance. It is clear, as expected, that having many data points of one type makes it more likely that a classifier will sort uncertain data into that overrepresented class. This is particularly challenging in medicine, where physiological variability produces precious few clear-cut distinctions. It is therefore necessary to determine the optimal performance tradeoff for any given application, acknowledging that there likely will not be a single universal solution for all scenarios (or even for all datasets within a single scenario).

The global population, and physiological findings therein, are not equally distributed among medically relevant classes. While this will often suggest use of an imbalanced dataset

(reflecting the global class distribution) to achieve overall optimal performance, it may not always be so, depending, for example, on the consequence of missed detection or false discovery on patient treatment decisions. It may therefore be preferable to utilize a method that does not achieve overall optimal performance, or even a range of methods (or a range of weights for a given model) from which the preferred option is selected after a first pass of the data to be characterized and its context. Given anticipated tradeoffs, there will typically be no single “best” method or model—one that will work in all scenarios or meet all needs. The optimal solution always depend on what metrics and aims matter, and competing interests may imply very different approaches to training data selection.

In order to allow well-informed decisions on these matters, reported methods should be careful to convey not only complete performance metrics, but also the underlying distribution of, and sensitivity to, data used to train (or optimize) and test the computational method. The results shown here reinforce that it makes very little sense to talk about the results or performance of a method without context; the single method presented here spanned the full range of performance of others reported in literature, simply by tweaking the training data distribution. Unfortunately, very often there is insufficient information reported to fully assess and understand the performance of a method. Publically available, standardized datasets for training and evaluation will also be a critical tool in this effort. Such resources are the only reliable way to directly compare and evaluate different methods, which has become increasingly important as new methods and architectures are developed.

It is known *a priori* that, without deliberate countermeasures, having many data points of one type during neural network optimization makes it more likely that a classifier will sort uncertain data into that highly-represented class. This is particularly challenging in medicine—natural physiological variability produces very few unambiguous cases of anything, and on top of that, classes are not equally distributed among the dataset or population. Therefore, those working to develop machine learning methods, particularly for application in healthcare, must carefully, deliberately, and transparently consider the implications of training data curation decisions.

3.9. Implications & Future Directions

Building upon my progress in extracting vascular geometry from intravascular imaging (Chapter 2), my work and findings in extracting vascular morphology from intravascular imaging extends, and is intrinsically imbued with, technical, scientific, and clinical implications.

The confluence of domain knowledge in vascular pathology and physiology and intravascular imaging, and advancements in machine learning, has enabled enhanced deep learning approaches to classify atherosclerosis and indwelling devices using grayscale intravascular image data. With a clear modular framework providing structure to the approach, neural networks were leveraged in applications to related tasks of tissue characterization and device segmentation in IVUS and OCT. Each application introduced unique challenges and obstacles, but through flexible and considerate application of framework, strong results were achieved for all three. By appropriately constraining the problem with geometric information and selectively partitioning the solution space, successful automated methods to interpret rich imaging data were cultivated which, even in their shortcomings, strongly reflect—and inform—human experience in interpreting the images.

These approaches exceeded the performance of previously-reported methods for plaque segmentation in both IVUS and OCT without the use of spectral signals [175], and produced maps of tissue morphology that closely resembled clinical gold standard VH-IVUS. Of great importance, the methods offer some attributes that even exceeded those of VH-IVUS. Because no RF (spectral) data were required, the method's applicability was not limited to ECG-gated frames or RF-enabled systems, but could be used to extract plaque morphologies in any grayscale intravascular image acquired with the system and parameters on which the network was trained. To acquire the same lateral resolution of plaque morphology using VH-IVUS would require extensive procedural time; the methods were also not subject to the loss of temporal resolution that limits VH-IVUS [277]. Furthermore, VH-IVUS offers lower axial spatial resolution than even its relatively low resolution grayscale IVUS counterpart [19], [234], [277], suggesting that a classification method based upon grayscale information alone—particularly from OCT—could offer superior detail and information on fine features. All of these benefits were achieved without the need for specialized hardware or proprietary software.

Elsewhere, automated extraction of vessel morphology in clinical intravascular imaging is sometimes approached in the absence of a unifying framework, which can result in disjointed contributions failing to reach full potential. Furthermore, the endeavors sometimes lack key features needed for granular analysis (especially for OCT, where little work has attempted full wall characterization), and frequently pursue globally optimal performance at the expense of other performance metrics without critical consideration of subsequent application and impact on utility. These shortcomings have led to often repetitive and duplicative work, unrealized utility, and low rates of adoption. The results of my studies demonstrate that these obstacles can each be overcome. In the context of the presented framework, similarities in challenges and solutions across applications can be readily appreciated, as can the ability to serve comparable and compatible outputs that enable utilization of, and comparison between, any pertinent available datasets.

For example, learnings from work in tissue characterization were applied to the segmentation of indwelling vascular scaffold devices. The framework described here allows for complete, pixel-wise annotation of BVS struts, with potential opportunities to apply the framework to other indwelling devices. This capability offers exciting prospects to not only reconstruct the device for use in 3D computational models, but to comprehensively and quantitatively assess intervention success (e.g. extent of apposition and stent area), track the longitudinal progression of vascular response (e.g. neointimal growth or restenosis) and device status (e.g. degradation), and to study the relationship of these processes with the disease and lesion phenotype. These latter studies could be undertaken with the integration and joint use of the strut segmentation application with the tissue characterization methods developed in this thesis (Sections 3.3, 3.4, and 3.5). The range of possible studies may offer benefit to clinical practice, fundamental insight on disease and treatments, and improved treatment guidelines, intervention planning, and device design.

Furthermore, the impact of leveraging domain knowledge was assessed, and was found to offer substantial benefit. The findings provide compelling motivation to leverage domain knowledge in other learned algorithmic tasks—especially in medical imaging—to reap benefits in the form of improved results and efficiency, even when such enrichment is not strictly necessary.

The impact of training data was also assessed, and was found to be remarkably consequential. Unfortunately, in machine learning applications for cardiovascular imaging data, wherein often data is variable, cases ambiguous, and findings sparse—that is, when such impact

is likely to be greatest—adequate reporting of data distribution and comprehensive performance metrics to inform selection for specific applications was found to be lacking. Findings provided a strong argument for more careful and deliberate considerations of the implications of data curation and more transparent and comprehensive reporting of model performance. The findings also dispelled the notion of universally applicable models. Due to tradeoffs inherent to the characterization task, it is unlikely that a single method will be widely applicable across a wide range of applications, use scenarios, or populations. Rather, it may be preferable to develop a range of methods or models—from which the preferred option is selected after an initial assessment of the data—to perform well by selective and narrow criteria. These more realistic expectations may help to guide future development of technologies for commercial and clinical translation.

Taken together, the work described in this chapter offers the means to conduct meaningful new research, and potentially better inform clinicians. Given the critical role of plaque morphology in lesion risk profile and optimal treatment strategy, the enabled availability of more comprehensive information from a wider range of sources could elevate consideration of this data to a more central position in the clinical consciousness and workflow. The role of plaque morphology in device and drug performance may also refine treatment strategies and guidelines, moving the field closer to the promise of highly-personalized medicine and computational cardiology [67].

Beyond extended validation with histology, improved efficiency and execution speed, refined ROI delineation, and establishing performance sufficiency (all discussed previously), there are several additional exciting avenues for future work. As noted earlier (Section 3.6), promising opportunities exist in the further development and assessment of semantic segmentation methods to classify full intravascular images, rather than pixel-wise classifiers. While largely effective, the pixel classifiers described in the central framework presented in this chapter—which apply a convolutional neural network to a pixel-centered patch to identify the tissue class depicted by the given pixel—were slow and failed to incorporate potentially useful contextual information. Initial work in applying a U-net to perform semantic segmentation yielded exciting preliminary results, but also generated many open questions. Advances in developing semantic segmentation models to classify all tissue depicted in an intravascular image could improve both accuracy and clinical applicability of tissue characterization methods, commensurately advancing the potential benefits thereof. Application to segmentation of indwelling devices (e.g. scaffold struts) may prove

particularly effect and also offers sizeable gains as well, given the large ROI in which the feature may reside and the current inefficiency of the presented strut segmentation method. Systematic assessment of hyperparameters, techniques to compensate for low class representation, and methods to leverage available manual annotation data for training would all be invaluable in these pursuits.

Another immediate area which invites investigation, and builds upon the application of domain knowledge in the characterization tasks, is the impact of attenuation compensation on the characterization (tissue classification) of intravascular images. As noted previously, the light signal of OCT is attenuated as it traverses through the tissue, such that images appear darker at greater radial depths. Methods for compensating for this attenuation have been developed, such as the popular method described by Girard *et al.* [218], and its use has been shown to be beneficial in visual inspection and interpretation by human experts (clinicians), and shown by Gerbaud *et al.* [65] and Teo *et al.* [293]. As signal intensity is associated with tissue class, performing attenuation compensation prior to classification may have a similarly beneficial effect on classifier performance. Indeed, anecdotal observations indicate that classifiers discussed in this chapter tended to falsely identify highly-attenuating tissue with higher incidence at greater radial depth. Future directions should investigate this phenomenon. Consideration extends to shadows cast by calcium in IVUS, which was suspected to inhibit NC classification accuracy. Work may apply attenuation compensation to images prior to classification—requiring the retraining of existing networks (or development of new networks)—and quantify the impact of such a pre-processing step.

An additional direction for future work may address the application and assessment of the use of transfer learning across imaging modalities in the characterization of intravascular images. There is not an abundance of labeled data available with which to train networks in the classification of tissue in intravascular images. Furthermore, data availability varies across among modalities, including OCT and IVUS. In particular, VH-IVUS provides a relative abundance in pixel-wise tissue labels and is fairly well validated. However, significant differences between the modalities, and even within each modality (as hardware systems and acquisition parameters, e.g. ultrasound frequency, change) limit the utility of this data. Utilizing geometry-based co-registration to construct multimodal corresponding datasets—an imaging Rosetta Stone of sorts—may further facilitate the utilization of disparate training data. Future work could seek to leverage

all available data to the greatest extent possible by implementing transfer learning, allowing training in one modality or system to improve final performance in another.

The application and assessment of the use of unsupervised training in the characterization of intravascular images should also be further explored. Two of the primary shortcomings of intravascular image characterization are the dearth of reliable labeled data and the necessity to impose prior assumptions regarding tissue classes, which themselves are defined inconsistently between imaging modalities. Endeavors to use unsupervised training may help to overcome both of these shortcomings, allowing networks to identify and segment tissue classes without reliance on label annotations and class specifications, instead relying upon inherent, underlying features and properties of the imaged tissue.

Drawing upon lessons from findings described in the previous chapter on extracting geometry from intravascular imaging, future work is also warranted to implement and assess the use of 3D image data in the characterization of tomographic intravascular images. Both OCT and IVUS are tomographic in nature, meaning that the data is pseudo-volumetric. Therefore, while noise and artifacts may hinder classification of some images using in-plane information alone, the incorporation of information from adjacent planes (i.e. frames or slices) may improve classifier performance. Work to incorporate multi-dimensional imaging data from the tomographic intravascular sequences and quantify the impact of such inclusion of contextual information holds promise.

Though several additional directions exist in this exciting, rapidly growing and evolving field, the final future direction noted here relates to the subject matter of the first chapter, geometry. Just as geometric information improved characterization of vessel morphology, morphology may improve or refine determination of vessel geometry; an iterative approach may yield better results for both. Morphological features, often utilized as fiduciary features by experts in the performance of manual co-registration and alignment tasks, likely would also prove useful in refining automated co-registration, enhancing the lumen-based method presented in the previous chapter. Rather than a unidirectional pipeline, bidirectional flows of information based on deeper relationships between geometry and morphology may prove mutually beneficial.

3.10. Conclusion

In this chapter, I have described the development and critical assessment of methods to determine and characterize vessel wall morphology. Aided by constraints imposed by vessel geometry and effective artificial intelligence tools, plaque components of diseased vessels are labeled in OCT and IVUS images [160], [161]. The approach is extended as well to labeling intravascular devices (BVS), and modular substitution of the neural network classifier further illustrated its flexibility and adaptability while highlighting an innovative and promising approach to tissue classification. The importance of enriching the method with domain knowledge and information on vessel geometry was quantitatively assessed [161], as was the impact of training data on overall performance [162].

Plaque morphology has long been assessed visually in the clinic, but qualitative, sparse sampling and informal guidelines for use of this information has limited the utility of this trove of data. The methods detailed herein enable more comprehensive assessment and interpretation of lesion morphology, with the potential to inform clinical guidelines, device development, and fundamental research on disease pathology. The capability to autonomously characterize diseased vessels at scale was successfully and comprehensively demonstrated with unprecedented performance and flexibility. The approaches present intriguing methods for integrating domain knowledge—an understanding of anatomy, pathological presentation, and imaging mechanisms—to achieve robust results, offering exciting clinical and technical applications and revealing important insights into the increasingly popular classification of medical images with neural networks.

With vessel morphology and structure successfully extracted from intravascular imaging, next steps for these methods include their integration with computational methods to analyze and simulate plaque in patient-specific models. This approach, addressed in the next chapter, allows for the biomechanical assessment of heterogeneous models from *in vivo* geometry and morphology, without the limitations of idealization, unrealistic assumptions, or *ex vivo* mechanical testing.

Chapter 4

Computational Modeling & Simulation of Patient-Specific Vessels

Vessel geometry and plaque morphology yield critical insights and useful metrics for consideration in patient care and treatment of atherosclerosis. However, the consequences of these geometric and morphologic features are not intrinsic, but rather are contingent upon their effect on states with direct impact on physiological function or treatment efficacy—flow, stress, inflammation, and similar facets which dictate system viability and to which the system responds. These facets are, to an extent, intuitively understood, but difficult to quantify or assess with certainty. Much interest has therefore been generated by the prospect of analyzing atherosclerotic plaques and diseased arteries from a biomechanical perspective using computational approaches [72]. Indeed, *in silico* models offer quantification of such states, as well as intriguing insights that cannot be obtained by direct measurement, particularly *in vivo*. However, study results can be highly sensitive to model fidelity. To obtain three-dimensional (3D) plaque geometry and heterogeneous morphology from patients, imaging is required [103]. However, image-based models are often highly simplified and fail to capture important structural details. Intravascular optical coherence tomography (OCT) offers an ideal data source to overcome these challenges, as it provides the highest resolution *in vivo* images currently available in medical practice [39]. Leveraging the methods developed in the previous two chapters, geometry and morphology was determined for lesions imaged with OCT, and this 3D information established the parameters for 3D computational models. In this chapter, use of clinical OCT data to construct highly detailed finite element (FE) models of an *in situ* diseased artery is described, opening the way for exciting new research and potential future clinical applications. In particular, a novel approach to estimate patient-specific biomechanical tissue characteristics is demonstrated *in silico* and validated *in vitro*, highlighting potential for future *in vivo* application.

CFD	Computational Fluid Dynamics
FE	Finite Element
FEA	Finite Element Analysis
FSI	Fluid-Structure Interaction
IVUS ...	Intravascular Ultrasound
OCT	Optical Coherence Tomography
VH	Virtual Histology
2D	Two-Dimensional
3D	Three-Dimensional

4.1. Diseased Vessel Modeling: Needs, Progress, & Challenges

Geometry, structure, and material constitutive properties together describe fully the mechanical behavior of a physical object. All three of these parameters are altered for arteries in disease. However, if known, this information can be implemented in computational models to perform measurements and simulations which provide new and unique insights into the state of the vessel.

A more comprehensive general discussion and review of the applications and approaches employed in modeling diseased vessels can be found in Section 1.3 of the Introduction, “Arterial Modeling & Simulation: Promise & Challenges.”

4.1.1 The Need for Computational Modeling and Simulation in Cardiovascular Medicine

As previously described, the course of atherosclerotic disease is in many ways driven by mechanical forces [9]. Symptoms and physiological effects are mediated by pressure-driven blood flow, changes to which can be precipitated by acute catastrophic events caused by the stress-driven rupture of mural tissue. Biology-mediated alterations to physiology in the course of disease progression (or regression) similarly respond to mechanical forces, including phenotype modulation by flow-induced shear stress along the endothelium and macroscopic stress-induced remodeling of the vascular wall.

Mechanical state, however, is notoriously difficult to measure directly *in situ*. Pressure, for example, can be measured directly with an invasive catheter-mounted probe. However, such tools allow only for measurement at the single point in the vessel at which the sensor (or fluid channel outlet) is located—and even then with little control and low spatial precision. Furthermore, the presence of the measurement catheter itself induces errors in the vessel pressure measurements [294]. Catheter access also limits the vessel segments in which measurements can be made (as the segment must be sufficiently large and navigable), and precludes simultaneous presence of other tools and devices. Therefore, pressure cannot be measured distal to narrow stenosis, during some interventions, or with differential precision through a vessel cross-section. Additionally, the injection of contrast, common for the placement of such catheters and the guidance of concurrent activities, impacts the intraluminal pressure [44], [48], [49]. In measuring flow and hydraulic resistance, direct measurement is made by injecting contrast, obtaining coronary cineangiography, and observing the extent and speed of antegrade contrast perfusion. Initial approaches to

assessment utilized rigid yet qualitative criteria to sort cases into one of four flow grades ranging simply from “no perfusion” to “complete perfusion” [295]. Subsequent approaches intended to increase precision and quantification adopted a proxy for flow; the number of cineframes (obtained at consistent, discrete time increments) over which contrast dye reached a standardized landmark after injection was counted [296]. While more objective and less categorical in nature, and shown to be a useful index for trial end point evaluation [296], this metric of flow and perfusion offers limited insight. Principally, the metric is not vessel- or lesion-specific, but instead yields only a single value for the entire coronary tree. This makes interrogation of local flow-mediated effects impossible to investigate in this manner. However, given appropriate boundary conditions, computational methods simulating pressure and flow distribution can precisely estimate these features at any arbitrary point within the modeled blood vessel.

Beyond flow and pressure, stress—within the vessel wall and experienced by devices—is an important driving mechanical factor. Unlike the former, however, there is no means of direct observation or measurement of stress state and distribution. Coupled to stress state is the matter of material constitutive properties, which describe the strength of materials and their response to loading. While these properties can be inferred—to an extent—by their general composition, material properties vary dramatically, even among the same class of tissue or plaque [109], [297]. The directly observable state which provides insight into both stress and mechanical properties is displacement, or strain. Specifically, given a measured strain, either loading condition (stress) or material properties can be determined if the other is known. In the field of solid mechanics, this concept is exemplified by the tension test commonly used to determine fundamental constitutive mechanical properties of a material. A known load is applied to a material of unknown properties—shaped into a known geometry—and the displacement experienced by the material is measured; from these results, mechanical properties are determined. However, such controlled testing is, naturally, not feasible for living tissue *in situ*. However, computational methods offer some alternative approaches to determining either stress distribution within a structure or constitutive properties of the composition.

Without a sound understanding of the mechanics of a system, which often cannot be directly or reliably measured with high resolution, the function, risk profile, and mechanisms of diseased vessels, devices, and interventions cannot be fully appreciated. Computational modeling and simulation offer both qualitative and quantitative windows into these states, enabling their use

in clinical prognostication, treatment planning, device design, and establishment of intervention guidelines.

4.1.2 Progress in Computational Modeling of Diseased Vessels

The promise of applying computational modeling to diseased vessels was recognized not long after the establishment of FE methods for aeronautical applications [71]. Understanding that the approach could be used to investigate stress states in vessels, early approaches first used idealized two-dimensional (2D) models of vessel cross-sections to investigate the impact of various basic morphological features on stress [75], [298]. Subsequent approaches advanced to model lesion-specific morphology [72], [108] by manually annotating and extracting 2D geometry and tissue structure from histology of excised diseased vessels [299], [300], and later from intravascular ultrasound (IVUS) [301] and OCT [302]. These studies identified some general trends and associations, such as the inverse relationship between fibrous cap thickness and peak circumferential stress, and the colocalization of high stress and intimal tears, plaque fracture, and rupture. However, the simplistic models lacked specificity and substantial details needed for robust assessment of the complex systems, and more recent work has demonstrated that study results can be highly sensitive to model fidelity. For example, 2D models are inferior to 3D models and fail to capture important out-of-plane effects and interaction [303]. Similarly, in the blood flow modeling realm, initial work in idealized models gave way to patient-specific rigid-wall computational fluid dynamics (CFD) models [69]. However, models have moved to adopt increasing complexity, with implementation of fluid-structure interaction (FSI) models assuming a more prominent position in the field [69], [103]. Here, too, work has demonstrated that the increased complexity is important to model fidelity, with studies verifying the importance of wall compliance and displacement—implemented in FSI but not CFD modeling—in accurately computing wall (or endothelial) shear stress [81]. Despite much attention and continued use of simplified models for select applications, reliable methods for constructing detailed 3D computational models of diseased arteries, incorporating both geometry and morphology of the wall, have remained a challenge for the field.

In addition to the construction of detailed 3D models, another important area of intense investigation, but limited achievement, has been in determining mechanical properties of tissue *in situ*. A robust body of work has demonstrated the fundamental impact of material property on

stress values and distribution in modeled arteries [109], [248], [304]–[307], which has motivated efforts to identify appropriate properties to use in such computational models. Mechanical testing, such as the tensile testing described briefly above, has been performed on excised tissue through painstaking separation of various plaque components and tissue classes [109], [297]. However, while such tests—which cannot be performed non-destructively for tissue *in situ*—provide some quantitative ranges of feasible values to implement in generic models, the high degree of variability observed among those samples highlighted the need to identify lesion-specific constitutive properties in order to yield reliable results from lesion-specific models.

Alternative approaches to non-destructive, *in situ* material characterization have been explored for application to diseased vessels relying upon the same fundamental principles as the *in vitro* testing. Namely, by observing the displacement response of tissue to a distorting force, governing constitutive relationships can be estimated. This general approach has taken many forms.

Relying upon microscopic behavior, laser speckle imaging has leveraged variation in Brownian motion with material viscosity and viscoelasticity. Because light scatterers are subject to such motion, decorrelation of laser light reflected within tissue can provide an index of viscoelasticity without imposition of external forces [111], [308], [309]. While robustness to small deformation conditions was demonstrated *ex vivo* and the feasibility for use *in vivo* is described [111], necessary hardware remains to be developed, and additional shortcomings inherent to the method exist. First, measurements are taken for a single 75 μm diameter spot at a time, with each spot requiring up to 200 milliseconds of signal acquisition for best performance [111]. Therefore, characterization of a full vessel segment seems prohibitively slow (though strategic, selective sampling may offer value even without a complete distribution). However, the primary limitation for use in computational modeling is that the image-based method reports only an index of viscoelasticity, but not a constitutive material property which could be implemented in a mechanical model.

Shear wave elastography has been demonstrated in phantoms and *ex vivo* setups to characterize atherosclerotic plaques [310]–[313]. By inducing wall displacement and monitoring the speed of displacement wave progression along the length of the wall, the material's shear modulus, which determines the speed of propagation, can be estimated. However, while able to

differentiate between different plaque groups, there was limited non-relative (i.e. absolute or independently-defined) quantification of properties which could be subsequently implemented in a mechanical model. The method also did not offer an ability to interrogate differences within the cross-section of a wall, instead reflecting some amalgamation of properties in the directions transverse to the direction of propagation [310]–[312]. Furthermore, the method relies upon the use of ultrasound (linear array) transducers and acoustic radiation force impulse imaging [310]–[312] which is not currently amenable to intravascular use, so while feasible for superficial vessels like the carotid arteries, significant advances in hardware would be required for use in coronary and other deep arteries.

The most common approach to *in situ* mechanical characterization has been displacement (i.e. non-shear wave) elastography. Given a known or controllable applied force or pressure, displacement (i.e. strain) can be observed and related to the governing constitutive relationship. This has typically been achieved through strain field estimation. Several methods have successfully utilized IVUS to track in-plane (2D) tissue displacement over time in response to changes in loading state in phantoms and excised tissue [314]–[317], with a small number of implementations even advancing to *in vivo* validation in animal models and clinical evaluation (though not widespread use) [313]. OCT elastography, or optical coherence elastography, has also been used for strain field estimation [110], [318]. While some methods have only been achieved in simulated data (reflecting a range of practical feasibility) [112], [319]–[321], the earliest approaches utilized basic 2D cross-correlation to map displacement in phantom models and tissue, even demonstrating *in vivo* capability (for easily accessible and manipulable skin) [322]. Others have utilized advanced acquisition and signal processing techniques to accurately acquire strain rate and strain mapping in a phantom while slowly imposing compression [323], [324]. All of these methods, however, advanced no further than deformation field estimation through image processing, yielding a strain map qualitatively implying information on relative stiffness of adjacent tissues within a given vessel cross-section.

To yield values of material stiffness, approaches had to be extended beyond deformation tracking to also incorporate stress. The simplest approach to do so was to apply a known, well controlled force to the surface of samples and utilize simplifying assumptions. OCT elastography, for example, was performed on slabs of phantom materials and excised tissue to which weights were applied, compressing the samples [325], [326]. The elastic modulus could be subsequently

determined by making the assumptions that the material was uniform, isotropic, incompressible, and elastic, and that the load applied to the sample resulted in a uniform uniaxial stress. Under these conditions, moduli could be computed at each point in the deformation field as the quotient of stress (applied force per unit area) and strain (local relative displacement) [325], [326]. Advancing the framework somewhat to consider a more realistic intravascular application, IVUS palpation was performed by securing phantom vessels and excised artery specimens in a water tank and applying static intraluminal pressures during IVUS imaging [327]. Similar to the other *in vitro* tests, assumptions were made that the material was homogeneous, isotropic, incompressible, and linear elastic, resulting in a uniform stress; local incremental stress-strain modulus of elasticity was computed as the quotient of half the change in intraluminal pressure and radial strain [327]. These attempts at vascular elastography, while reasonably valid in some instances and still yielding qualitative insight in others, rely on assumptions which break down in most clinical applications. In particular, assumption of uniform material and constant stress—which yields stiffnesses inversely proportional to displacement/strain, thereby making derivation straightforward—are invalid in heterogeneous mechanical environments such as plaques, where much stiffer components (such as calcium) bear greater loads than more compliant ones (such as lipid). One alternative approach which does not require such an assumption is inverse finite element analysis (FEA).

The inverse FEA approach utilizes a computational FE model of an imaged object's geometry and morphological structure to iteratively simulate the imposed loading, updating material properties with each iteration until the simulation yields results matching the observed outcome. Early work on such methods assessed very simple models; optical coherence elastography displacement was simulated and coupled with an inverse FE model of a soft tissue block with a single heterogeneous inclusion (of varying size) [320]. Subsequent work advanced from simple, generic geometries to vessel-replicating models. Displacement was still simulated, but 2D OCT-derived vessel cross-section models and idealized 3D models were implemented, each including three distinct materials; inverse FEA was performed to estimate the mechanical properties of each [321]. Another method simulated noninvasive ultrasound images from high resolution multicontrast magnetic resonance imaging of *in vivo* carotid arteries; FE-simulated displacements and strain elastograms were utilized to generate the target for subsequent inverse FEA [328]. Moving beyond simulated displacement, tissue phantoms and excised tissue samples

were loaded under OCT imaging, with inverse FEA model material parameters optimized to recapitulate the displacement fields [329], [330]. However, while one method used simple models and reported only relative moduli [329], the other required a calibration layer of known stiffness to constrain the solution set [330]. Both were performed only in 2D [329], [330]. Relatively recent approaches have abandoned OCT in favor of noninvasive ultrasound, which, as noted before, is not currently amenable to intravascular use but feasible for superficial vessels (e.g. carotid arteries). One method used a linear array probe to estimate 2D strain in phantoms—homogeneous or with heterogeneous inclusions—and estimate maps of Young’s modulus for each through inverse FEA [328]. Another such framework derived displacement data of a cannulated iliac artery from an external ultrasound transducer and vessel structure from subsequent histology of the excised atherosclerotic vessel, and characterized 4 materials through inverse FEA [331].

All of the aforementioned approaches to characterizing mechanical properties of tissue relied on accurate acquisition of displacement fields. However, such data is difficult to obtain reliably, and even so is often very noisy, which is why many groups have resorted to simulating this input (e.g. as output of a forward FEA simulation) rather than actually measuring it. Specialized noise reduction has been the focus of several of the previous works [329], and cross-correlation kernel size has been shown to be exceptionally important in the performance of OCT elastography of vascular tissue [332]. Furthermore, catheter position has also been shown to substantially influence estimated strain in intravascular elastography, though specialized routines have been proposed to partially attenuate these effects [333]. Beyond issues of noise and bias due to arbitrary acquisition state, elastography is typically performed in plane in order to record the reaction of tissue over time (in response to loading). Consequently, such methods rely upon 2D observations, which only allow for consideration of material stiffness in a single direction, even though the mechanical properties of tissue components in human atherosclerotic plaques, like most tissue, can be highly anisotropic [297].

A recent publication has described work in characterizing atherosclerotic tissue on the basis of virtual histology (VH) IVUS imaging independent of elastography [334]. Noble *et al.* performed repeated image acquisition sequences in *ex vivo* cadaveric femoral arteries pressurized at 6 different pressure levels. A basic FE model of each vessel was constructed with elements assigned to material classes on the basis of VH designation. Inverse FEA was performed to optimize the properties of each class in order to minimize the difference between the maximum and minimum

observed and simulated vessel diameters (rather than minimizing the overall difference in displacement field, as in the aforementioned approaches) [334]. This process allowed for patient-specific characterization of artery and plaque material properties in the diseased peripheral arteries.

While previous work has utilized common imaging modalities to perform arterial elastography, often coupled with inverse FEA, this work was conducted *in vitro* or simulated *in silico*. As such, loading conditions were imposed at will and well defined, and sequential series of images of the same tissue under varying loading conditions allowed for continuous tracking of the same tissue. Even the method which did not rely on elastography was not well formed and utilized full acquisitions at a large number of known, well-controlled, constant (i.e. not time varying) pressures. These luxuries are not available in the clinic. Consequently, while several approaches to *in vivo* mechanical characterization have been proposed, none have been adopted for clinical use, and a reliable method for identifying appropriate constitutive parameters remains an outstanding challenge. Coupled with the dearth of reliable methods to construct detailed 3D models of diseased arteries, prospects of providing real-time patient-specific modeling and clinical decision-making guidance, as well as fundamental research insights, remain tenuous.

4.1.3 The Endeavor to Model and Simulate Patient-Specific Vessels

This chapter discusses thoughtful new approaches to modeling diseased, atherosclerotic vessels on the basis of imaging data, with a focus on determining lesion-specific mechanical constitutive properties. Building upon the new methods and approaches described in the previous chapters, geometry (Chapter 2) and plaque morphology (Chapter 3) were integrated into a continuous, 3D computational representation which allowed biomechanical assessment and simulation, with potential to extend to chemical and biological realms.

Whereas most previous computational models of diseased arteries utilized simplified, 2D, or histology-derived structure, the work presented herein develops such models using intricate, 3D structure derived from *in vivo* clinical images. Considering heterogeneous and multidimensional plaque morphology enabled us to explore complex interactions and gain more confidence in observational results due to the model's elevated fidelity. Furthermore, developing such models from intravascular imaging data rather than histology (or other method requiring tissue excision) promotes its consideration for future use with patients and in longitudinal studies examining changes of the tissue *in situ* over time.

Next, departing from previous work, an approach to material mechanical characterization was developed and validated which relied neither on elastography measurement nor on several sequences of highly-controlled image acquisitions. The hypothesis driving the approach was that, by blending morphological assessment with displacement observations, two tomographic image acquisitions with pressure data should provide sufficient information to estimate 3D vessel stiffness so long as the sequences are acquired at different pressures. Rather than requiring a full displacement field, only the initial and final states were compared. Here, the key assumption made was that all tissue of a single class, as determined through image analysis, displayed the same material properties within the imaged lesion, and that plaque features are consistently identifiable in sequential acquisitions. These assumptions alleviate the need for full displacement maps—displacement of mural structures (plaque features) can be tracked instead of individual points, and rather than considering each individual point as an independent agent requiring individual characterization, a constrained number of constitutive parameters define the loading response of the entire vessel. Furthermore, by tracking mural structure, rather than discarding all displacement data other than inner and outer diameter, as others have done, a richer set of data can be extracted from each acquisition, thereby reducing the number of sequences needed to determine the desired information. While these strategic assumptions are simplifying, the approach may mitigate effects of noise, allow for characterization of tissue volumes (rather than discrete planes—elastography is performed in one or two dimensions), and enable assessment of anisotropic properties.

Major contributions to this chapter were made by Narayanan [335], whose supervised master's thesis, "Image-based Mechanical Characterisation of Atherosclerotic Plaques – An Inverse Finite Element Approach," furthered this body of work substantially and includes several detailed, technical descriptions of model construction and *in vitro* validation testing not repeated here. Subsequent, ongoing work in honing the modeling approach, which is also briefly introduced in this chapter, benefits from major contributions by Kadry, whose supervised master's thesis builds upon and further this body of work through refined model construction.

With improved accuracy, fidelity, and robustness, the computational models, mechanical characterization methods, and other approaches described in this chapter may offer unprecedented insights into not only risk assessment and treatment planning for the modeled patient, but through high-volume analysis also general characteristics of plaque risk, disease progression, and intervention efficacy.

4.2. Constructing Detailed Image-Based Volumetric Models

Having developed algorithms which could delineate and classify vessel tissue imaged by OCT or IVUS with high detail and accuracy, the next stage of the work focused on volumetric model construction. While the image-derived geometric and morphologic information is itself valuable, translating this series of 2D data (characterized intravascular images) into a 3D volumetric model—in an effective way—remained an unresolved challenge. Therefore, leveraging the detailed information derived from intravascular images, and with forethought towards future application in mechanical property estimation, techniques were developed to construct such FE models from the classified tomographic images.

Geometry and morphology of a diseased coronary artery were first extracted from *in vivo* imaging data. A series of cross-sectional images acquired during an OCT pullback were processed using the automated method described in Chapter 3 [160]. Pre-processing delineated the vessel wall area, then designated non-diseased and media tissue using physiological geometric parameters. A convolutional neural network subsequently characterized the plaque within the vessel wall and assigned to each image pixel a label: calcium, lipid tissue, fibrous tissue, mixed plaque, or no visible tissue (i.e. guidewire shadow) [160]. Plaque label maps were then post-processed to make the data more amenable to FE representation (Figure 4.1); lumen centroids were aligned to mitigate catheter position variation and reduce motion artifacts (4.1A), obscured regions were assigned tissue labels through region growing (4.1B), and 3D mode filtering was performed




Figure 4.1. Classified label map pre-processing facilitated model generation from characterized images. After lumen centroids were aligned among the initial label maps (A), regions obscured by the guidewire catheter shadow artifact (light gray) were first assigned to tissue classes by estimating the obscured content (B). The resulting classification distribution was then smoothed to improve consistency by removing very small standalone regions of a distinct class and performing 3D mode filtering (C). Finally, non-diseased/media class labels were inserted to separate any plaque or lumen labels from immediately adjacent perivascular labels (D), which represented a non-physiological arrangement which could occur in matrix format due to frame-to-frame variation in the context of limited longitudinal resolution. White arrows indicate instances of features impacted by the subsequent operation. The color scheme for the tissue classes is the same as that in Chapter 3.

to improve spatial consistency (4.1C). Image characterization and label pre-processing was performed in MATLAB (MathWorks, Natick, MA).

The series of processed plaque maps was then transformed into a 3D FE model (Figure 4.2). From the maps (4.2A), point clouds of the inner and outer surfaces were generated using spatial characteristics of the images (4.2B). Additionally, point clouds of the full vessel wall were generated, with each point tagged with its corresponding plaque class. Next, a closed surface mesh of the vessel wall was constructed from the bounding surface points (4.2C). This surface was then meshed with tetrahedral elements to produce a volumetric FE mesh (4.2D). Subsequently, the tagged point nearest the centroid of each element was identified, and its class tag was assigned to the element (4.2E). The entire mesh could then be transformed to follow a 3D contour, and material-specific properties were assigned to the various element class sets. Point cloud processing and surface meshing were performed in MeshLab (Consiglio Nazionale delle Ricerche, Rome, Italy) [336] and Rhinoceros (Robert McNeel & Associates, Seattle, WA). Volumetric meshing and related processing was performed in Abaqus (Dassault Systèmes, Vélizy-Villacoublay, France).

To quantify resolution and fidelity retained through the meshing procedure, slices of the model corresponding to plaque map locations were extracted for comparison to the input (4.2F). Boundary conditions and loads were applied to perform proof-of-concept simulations, as well.




Figure 4.2. Finite element models were generated from characterized intravascular images. (A) OCT images were characterized with the method described in Chapter 3, and the classification maps were post-processed (Figure 4.1). (B) Spatial point clouds of the vessel's inner and outer surfaces were extracted from the tomographic maps. (C) Surface meshes of the vessel wall were constructed from the point cloud. (D) Volumetric (FE) meshes of the vessel wall (volume delineated by the surface meshes) were constructed. (E) Elements of the mesh were assigned to a material class on the basis of the tomographic label maps. (F) Cross-sections of the volumetric mesh were extracted for comparison to the tomographic label maps used to generate the model.




Figure 4.3. Comparison of an OCT image (*left*), resultant tissue label map (model generation input; *center*), and tissue label slice extracted from the corresponding location of the generated FE model (output; *right*). Relative to the basis of the model (*center*), the resulting material distribution (*right*) has more sharp boundaries between material regions (due to the shape of the discrete tetrahedral elements), but retains many of the fine features, relative distributions, and spatial relationships of the input.

Comparison of the plaque label map inputs and the extracted outputs showed tremendous agreement (Figure 4.3). Overlap of wall presence exceeded 98.0%, as evaluated by binary pixel-wise comparison. In the coincident wall area, class label agreement reached 94.4%. Initial simulations carried out in Abaqus demonstrated corresponding FE model viability and functionality (Figure 4.4).

The presented method retained exquisite detail of vessel wall morphology while enabling detailed assignment of heterogeneous tissue properties as well as 3D spatial consistency and continuity. By first generating a mesh for the entire wall and subsequently assigning mesh elements to material classes, interface and posterization issues of most other model generation approaches were avoided while demonstrating detail retention and flexible material segmentation on the spatial scale of individual volumetric elements ($\sim 10^5 \mu\text{m}^3$). Using OCT enabled exceptional resolution, but the method could be applied to other data sources, such as IVUS (VH or characterized grayscale). Through this flexible new approach, a highly detailed *in vivo* diseased artery model was successfully generated, offering unprecedented fidelity for research and future clinical application.

Though simple and powerful, this method of model generation did have some drawbacks. The jaggedness and lack of refinement at interfaces between material regions (Figures 4.3 and 4.4) was non-physiological and led to stress concentrations and large stress gradients in these regions which could not be adequately resolved. Consequently, while the model yielded smooth and continuous displacement fields (Figure 4.4, *bottom*)—critical for later mechanical property




Figure 4.4. Initial simulations demonstrated coarse FE model viability and functionality. As shown in the cutaway views, the heterogeneous FE model (*top*) yielded heterogeneous distributions of von Mises stress (*middle*) and displacement (*bottom*). The five tissue classes were each modeled as incompressible neo-Hookean materials. A static pressure was applied to the lumen surface, and the boundary conditions were set such that the vessel was free to expand while the ends of the model were constrained in the longitudinal direction. Mesh colors (*top*) indicate material (class) assignment, with colors adhering to the scheme of Chapter 3.

characterization—due to the inherent spatial continuity and constraints, the corresponding stress fields were rather coarse and disjointed (Figure 4.4, *middle*). Therefore, while stresses in bulk regions likely provided reasonable estimates of stress, precise values, particularly at interface regions, and extreme (maximum and minimum) values were not reliable. To address this shortcoming, more precise methods were needed to construct heterogeneous models that provided accurate measures of mural stress.

While not suitable or necessary for all uses, alternative approaches to generate refined volumetric computational modeling which built upon the initial strategy were developed. Rather than constructing a monolithic volumetric mesh for the full vessel wall—irrespective of plaque features—and subsequently assigning the resulting elements to material classes based upon the space the element happened to occupy, meshing was carried out to accommodate the plaque feature shapes. Geometries of significant plaque and wall features were determined and meshed independently, then merged and aligned to generate a full mesh of the vessel wall which conformed to the defined plaque component geometries. The various steps of the process were executed in MATLAB, Gmsh (Geuzaine & Remacle, Louvain-la-Neuve, Belgium) [337], and SpaceClaim (ANSYS, Canonsburg, PA). First, 3D geometries were generated by extracting plaque components from tomographic images, transforming 2D image indices to 3D spatial dimensions, and connecting discrete planar objects corresponding to the same component into continuous volumes. Volume meshes of these 3D plaque components were then meshed and merged. Unlike the coarse mesh (described earlier in this section), the resulting full vessel wall mesh was refined around interfaces between regions such that elements aligned along the interfaces between geometric features rather than straddling them. Element regions were then assigned to the material classes to which they belonged. Simulations were then conducted in Abaqus on models with increasing levels of refinement to assess mesh dependence and sensitivity of the results, which confirmed the validity of the simulations. Initial results incorporating a subset of identified plaques (Figure 4.5) highlighted the impact of various plaque components on stress distribution within the wall.

While refined, this method of model generation also had drawbacks. The smoothness which was achieved came at the expense of manual processing time and effort, and simplification to define continuous geometries for plaque components (Figure 4.6). However, stress profiles—in addition to displacement fields—were more continuous and behaved as generally expected. Work is ongoing to reduce the simplifications imposed to generate the preliminary models created in this




Figure 4.5. Initial simulations demonstrated refined FE model viability and functionality, and illustrate trends in stress distribution and plaque composition. As shown in the cutaway views, the heterogeneous FE model (*top*) yielded heterogeneous distributions of von Mises stress (*middle*) and displacement (*bottom*). The three tissue classes were each modeled as hyperelastic materials. A static pressure was applied to the lumen surface, and the boundary conditions were set such that the ends of the model were completely constrained in the all directions.

way, though it remains to be determined what level of detail is needed and how tradeoffs between detail and refinement impact simulation results.



Figure 4.6. The two modeling approaches exhibited tradeoffs in refinement and detail. Whereas generating a monolithic mesh and subsequently assigning material properties (*left*) allowed for substantially more detail to be retained, features were coarse. To the contrary, generating separate geometries for each plaque feature before generating an accommodating cohesive mesh (*right*) was far more laborious and enabled inclusion of fewer features, but included features were much more refined. Note that the two models shown here do *not* correspond to the same patient, and the latter had not been constructed to include distinct fibrous or mixed tissue—as preliminary (and laborious) work, only a subset of the plaque has been implemented in the model. The resulting differences in plaque representation are reflected in the stress and displacement distributions (Figures 4.4 and 4.5).

4.3. Determining Patient-Specific Constitutive Parameters

A reliable approach to estimate mechanical properties of vessels from imaging data remained an outstanding challenge in the field. In particular, reliance on fickle elastography data or infeasible numbers of repeated acquisitions limited the feasibility and applicability of previous approaches (as discussed in Section 4.1.2). In addressing these challenges, inverse FEA presents an ideal option for its flexibility in allowing optimization to match any number of target metrics. For example, by aiming to match overall displacement of internal mural structures (alignment of plaque features) to determine material properties, the approach does not necessarily require a full displacement field to be captured, but rather images at just two states. Such an approach can be employed using the same imaging data and imposed constraints as the modeling approaches described in the previous section (Section 4.2), which can serve as the basis for mechanical property estimation. Work was therefore undertaken to develop methods of discerning mechanical parameters from imaging data by estimating mechanical properties for a given volumetric model of a heterogeneous diseased artery.

Building upon previous work in extracting geometry and morphology of diseased vessels from intravascular images (Chapters 2 and 3), detailed 3D patient-specific FE models were constructed of a coronary artery. The first method described in Section 4.2 was utilized, whereby the vessel wall was meshed and elements subsequently assigned to materials based upon the spatial plaque class distribution determined from the tomographic image sequence (Figure 4.2). As described in detail by Narayanan [335], an inverse FEA routine was developed in which constitutive parameters were optimized for each of five different tissue classes to minimize the difference in deviation between surfaces in a simulated deformation and an observed (target) deformation. Unlike other approaches, which sought to match only inner and outer vessel surface dimensions—which substantially distills the rich displacement information, offering limited data on deformation and therefore requiring many image acquisitions—the interfaces between all tissue domains (as well as inner and outer vessel surface) were compared. Both linear elastic and neo-Hookean material models were implemented. Tests were also carried out in which zero-pressure geometry was first estimated to account for initial stresses [85].

Mechanical property estimation was validated both through *in silico* simulated tests and *in vitro* phantom testing. A patient-specific model was forward-simulated with feasible material

properties specified to assume a deformed shape at an elevated pressure, which served as the target state. Inverse FEA was then executed for the initial model with the prescribed pressure increase and unknown material properties, which were optimized to achieve a final state which minimized deviation from the target (forward simulated) state. The final material properties differed from the specified properties (used in the forward simulation) between 3 and 21% among the five tissue types.

Testing of the method *in vitro* was performed in phantoms comprised of materials which were directly tested by established means. Two phantoms, one constructed from homogeneous manufactured flexible polymer tubing, and another developed especially for this purpose made from a medical grade liquid plastic, were utilized (Figure 4.7). The former was comprised of Tygon tubing (Saint-Gobain, Courbevoie, France) with an inner diameter of approximately 3 mm and a wall thickness of around 0.8 mm. The latter was comprised of Steralloy elastomer E-2036 (Hapco, Hanover, MA) and formed in a custom 3D-printed mold shaped with a nominal 3 mm inner diameter, 2 mm wall thickness, and minor focal stenosis in which a heterogeneous inclusion of UV-cured glue resided. Both phantoms were imaged by OCT while subjected to various known internal pressures (Figure 4.7). Inner and outer surfaces of the phantom were automatically detected in each acquisition using approaches similar to those for extracting vessel geometry described in Chapter 2. When present, the heterogeneous inclusion was automatically segmented as well. With image-based geometry and structure fully defined, FE models were constructed in the manner described for atherosclerotic vessels in Section 4.2. The inverse FEA routine was then executed to determine the material properties which yielded the minimum deviation between a vessel imaged at an elevated pressure and a vessel imaged at a lower pressure and subjected to computational simulation of pressure inflation to the elevated pressure.





Figure 4.7. OCT acquisitions of two styles of phantoms were acquired. The first was constructed from a homogeneous manufactured flexible polymer tubing (*left*), and the other made from a medical grade elastomer in a specially-designed mold which allowed for inclusion of a secondary material in a stenotic region (*right*).

To establish a reference for evaluating method accuracy, mechanical characterization was carried out for the materials comprising the two phantoms. Testing was performed following, to the extent possible, the ASTM standard test method for tensile properties of plastics (D638-14) [338]. Specimens were prepared for the materials of both phantoms. For the phantom prepared in house, a dumbbell-shaped test specimen was prepared to match the standard's designated dimensions for a Type IV specimen (for thin nonrigid plastics). A mold was designed in SolidWorks computer-aided design software (Dassault Systèmes, Vélizy-Villacoublay, France), 3D printed with a stereolithography Form printer (Formlabs, Somerville, MA), and post-processed (Figure 4.8). Specimens were prepared alongside the phantoms, with material mixed and cured under shared conditions (Figure 4.9). A small segment of the tube (approximately $30.5 \times 2.75 \times 0.7$ mm) was cut from the same source as the phantom (Figure 4.9). Testing was performed on a load cell (Instron, Norwood, MA; 2580-2KN), with elongation prescribed at a speed of 50 mm/min. Unfortunately, no extension indicator was available, so strain was measured from the reported grip separation. Specimens were fixed within aligned grips and seated to remove slack in the system. Both specimens were tested three times (Figure 4.10): twice to an extension of 20%, and a final




Figure 4.8. A mold was designed for creation of the mechanical testing specimen to characterize the elastomer mixture used in the custom phantom. The central image is a rendering of the model created in computer-assisted design software with dimensions conforming to published standards. The inset image in the top right shows the actual mold, which was 3D printed from the model design.




Figure 4.9. Test specimens were characterized in mechanical testing. Both specimens were subject to tensile testing along their longest dimension. Note that the scale of the two images differ—specimen (a) is approximately the length of the narrow region of specimen (b) demarcated by the inner set of markings. The outer set of markings on (b) indicates the grip positions.

time to a much higher strain (75-100%). Data fitting and material stiffness estimation was performed in MATLAB according to standard definition; the slope of the stress-strain curve was determined for low stress/strain ($\leq 10\%$ strain) to estimate the modulus of elasticity (Figures 4.10; i.e. Young's modulus). Poisson's ratio could not be determined from the available data (as no transverse strain was measured).

The *in vitro* phantom testing yielded promising results for the material characterization approach. The inverse method determined that the flexible tubing had a Young's modulus of 3.37 ± 0.18 MPa, whereas the tensile testing reported a modulus of 2.98 ± 0.03 MPa (13% error). However, as noted above, because the tensile testing specimen was retrieved from the small tube itself, it did not conform to shape standards and therefore may not be entirely reliable. Furthermore, the modulus reported by the supplier for the material was a much lower 1.9 MPa, though the specimen was quite old and had likely stiffened substantially since its manufacture. The phantom elastomer was determined through the computational approach to have a modulus of 788 ± 69 kPa. Tensile testing yielded a modulus of 863 ± 11 kPa (8.7% error). The supplier reported a range of moduli from approximately 800 kPa. Unfortunately, due to delamination between the hard inclusions and the primary elastomer, the inclusion material could not be characterized by the inverse method. However, the combined success of the *in silico* validation of heterogeneous material characterization and *in vitro* validation of homogeneous material characterization, both yielding errors far less than the variation observed in mechanical testing of tissues, suggests that the approach could hold promise for future *in vivo* applications when at least two intravascular imaging acquisitions at different pressures are available.

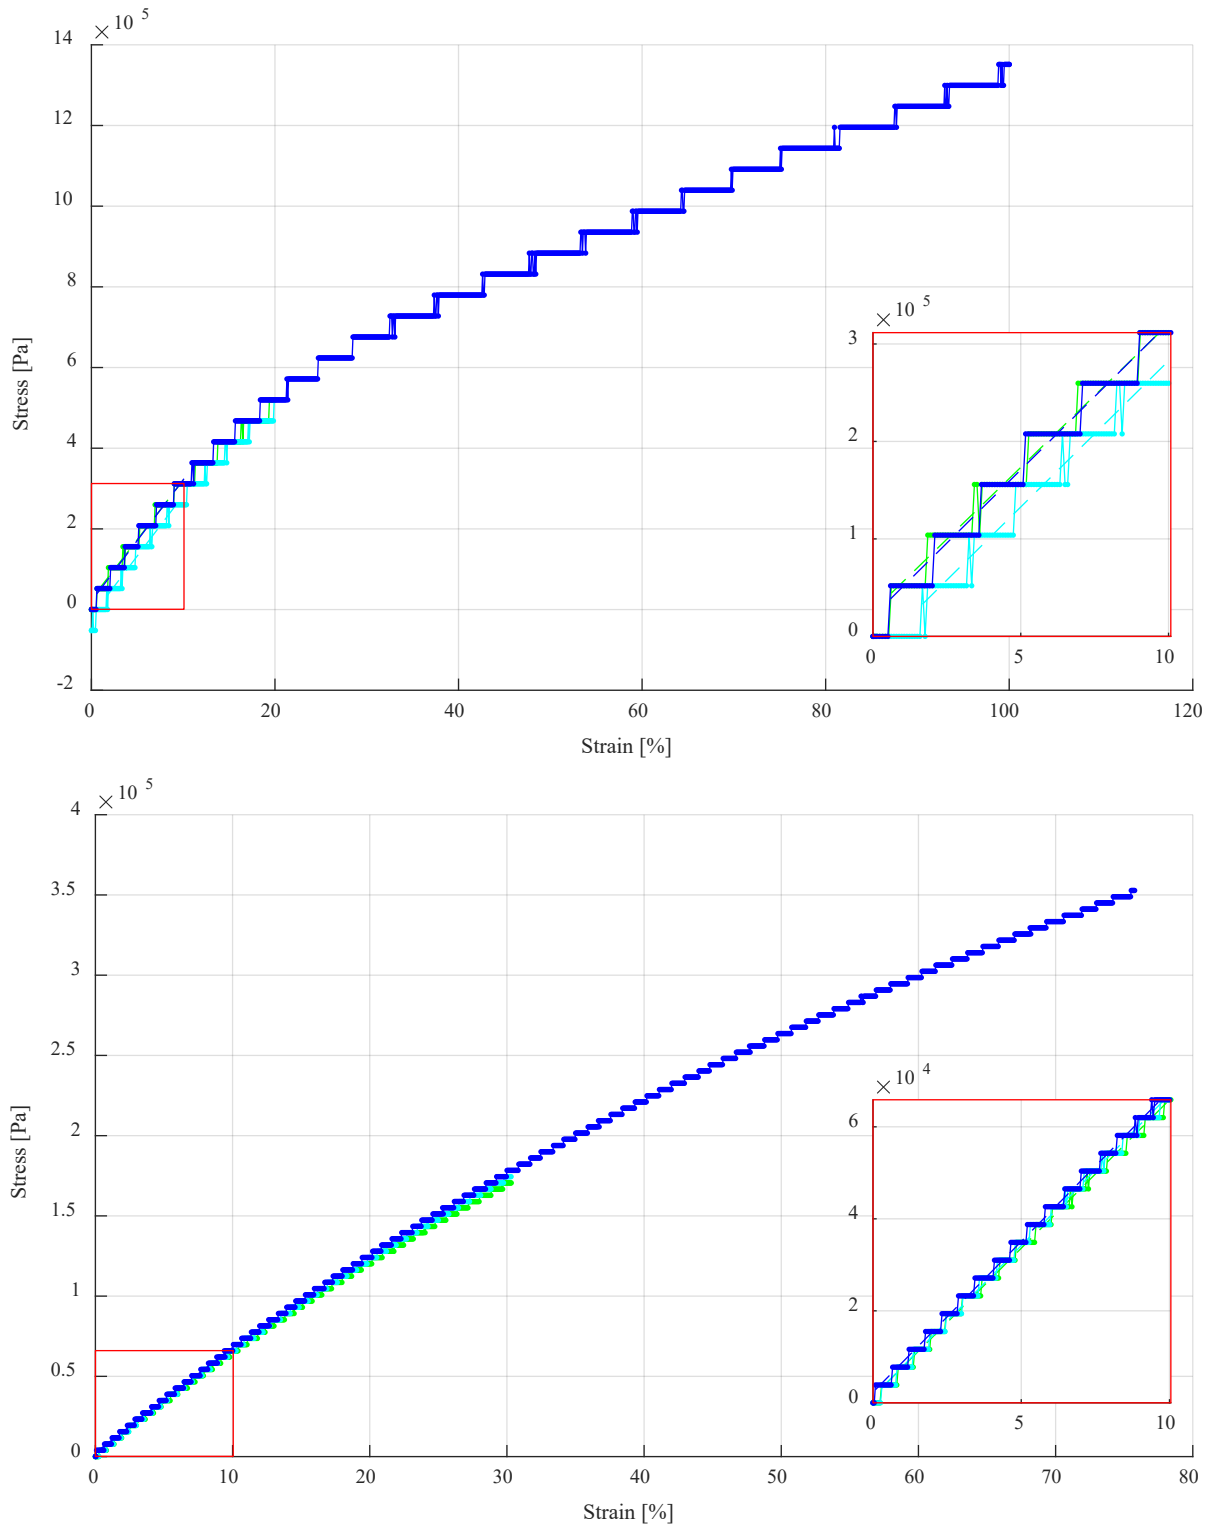


Figure 4.10. Tensile test results for the homogeneous flexible polymer tubing material (*top*) and the custom phantom's elastomer material (*bottom*). The red boxes and corresponding inset plots show the first 10% of strain and the region of elastic modulus fits (slope of dashed line). The best fit lines yielded Young's moduli of 2.98 ± 0.03 MPa and 863 ± 11 kPa for the flexible tubing (*top*) and liquid elastomer (*bottom*) materials, respectively.

4.4. Implications & Future Directions

Building upon my progress in extracting vascular geometry (Chapter 2) and morphology (Chapter 3) from intravascular imaging, my work and findings in image-based computational modeling and simulation extends and offers technical and clinical potential. Various avenues for future work are also implied.

Initial studies demonstrated the feasibility of constructing highly detailed patient-specific computational models from characterized intravascular images. Coarse meshes, readily constructed directly from geometry and pixel-wise class assignment information extracted from images, can successfully simulate displacement, though stress distributions resulting from such models are more tenuous. However, for application such as material mechanical characterization, which depends on displacement, these straightforward models are viable and have been shown to yield relatively accurate stiffnesses. Refined meshes formed by the assimilation of individually-meshed plaque structures are more laborious to construct and reliant on morphological simplifications, but may be better suited for other applications, particularly those for which stress distributions and extremes are of central importance. However, the impact of the simplifying assumptions remains to be explored.

One critical and ongoing direction of work is analyzing the impact of modeling assumptions. While previous work has shown that 3D modeling is superior to 2D models, and intima heterogeneity is critical to accurate stress calculations, the sensitivity and impact of more detailed techniques—including simplifying assumptions and choices made regarding retained resolution—must be quantified with respect to mechanical (simulation) outcomes. Comparisons of detailed models constructed from OCT data would also ideally be compared to models constructed from older, lower-resolution but well-validated imaging modalities (such as IVUS). Furthermore, the approaches presented here apply some degree of coarse posterizing of plaque morphology, and assume that all tissue of the same class share the same mechanical properties. These simplifying assumptions, which differ from elastography, warrant exploration. Additionally, while the ability to transform the coarse model to follow a 3D trajectory was developed in this work, it was not implemented in the demonstrative simulations, as accompanying angiographic data indicating imaged segments were not readily available. The impact of vessel trajectory is critical to fluid flow (and consequently CFD and FSI simulations), but it is unclear

whether, or under what conditions, the effects are as important in modeling the solid mechanics of the system. Finally, the impact of motion artifacts, non-uniform rotation, and non-tangential catheter alignment on 3D geometry reconstruction, modeling, and simulation should be examined.

Beyond assessing the validity of assumptions and quantifying their impact on mechanical simulation results, a more fundamental question regarding model complexity must be considered. A compelling case has been made in favor of reducing patient-specific model complexity, rather than striving for ever-increasing detail, in order to achieve clinical translation and its concomitant benefits [339]. Indeed, that may well be true, but these different approaches need not be pursued or used exclusively—no single model will be applicable in all scenarios. Hybrid approaches, such as a lumped-parameter mesh model (like the non-isotropic linear-elastic spring mesh developed for smooth surface fitting in Chapter 2 [157]) could replace the far more complex FE model, with optimization of component parameters (e.g. spring constants) replacing the very computationally demanding inverse FEA process. Use scenarios demanding patient-specific model complexity, amenable to reduced model complexity, or possibly requiring a hybrid approach, should be identified to guide future work and improve likelihood of clinical impact.

Additional information, which was assumed to be known in this work, is needed to perform mechanical characterization. In particular, to realize clinical and research utility, intraluminal pressure must be determined. As noted earlier, some form of load and displacement information is necessary in order to determine mechanical properties of a material. Further work must therefore undertake the task of measuring or estimating the pressure in an artery during its imaging. Given the frame rate and standard length of an OCT pullback, this data acquisition process lasts approximately 2.5 seconds. Therefore, it is to be expected that the patient's heart will beat 2-4 times during the imaging sequence, generating corresponding pressure waveforms. Given assumptions about the expected tapered shape of the coronary artery, as well as assumptions regarding standard pressure waveforms of the coronary artery, it may be possible to reconstruct a relative pressure trace corresponding to the luminal pressure experienced within the artery in each frame. Additional patient data (e.g. systemic blood pressure and other vital measurements), possibly in conjunction with lumped parameter models, may be necessary. Complications may be introduced by stenoses, which lead to diminished distal pressures, loading introduced by bulk motion and squeezing of the heart wall in which the arteries are embedded, and superimposed pressures induced by non-occlusive flush of blood-clearing media. Therefore, a strong preference

would be for direct pressure measurement to be acquired within the vessel concurrent to imaging, preferably with integrated hardware.

Several improvements to the models are also worthy of exploration. First, in the theme of pressure variability, methods for obtaining zero-pressure geometry and comparing displacement throughout a dynamic loading process will likely be critical to future clinical use, where well-controlled and constant intraluminal pressures would be difficult, if not impossible, to achieve. Extension of current approaches to achieve these capabilities should be possible, but will require more complex technical and logistical implementations. Second, refined models should incorporate more tissue classes (beyond calcium, lipid, and a third general class). Work exploring the utility of retaining mixed tissue as a separate class, or incorporating it with the fibrous class, should be considered. Finally, looking to applications beyond mechanical simulations, tissue-specific properties of mass transfer (e.g. drug delivery) and biological response (e.g. inflammation, lipoprotein accumulation, calcification, and remodeling) can be incorporated into the model as a coupled, multi-physics platform.

Even with current capabilities, the lesion-specific computational models developed in this chapter offer several exciting applications. One such application is the modeling of patient-specific lesions to determine clinically relevant characteristics. While designation of a plaque or lesion as “vulnerable” is generally an oversimplification (as discussed briefly in Chapter 3), certain characteristics have been shown to make plaques more or less susceptible to rupture. However, much less is known about the features which may be important to the gradual and catastrophic progression of plaque erosion, which is becoming an increasingly important and prevalent mode of acute coronary syndrome [340]. Furthermore, the relationship between plaque morphology and residual vasomotor capacity is not well understood. Through FE and CFD or FSI simulations and analysis of associated patient outcome data, the various morphological and mechanical features of plaques which make diseased vessels more or less stable and/or vasoactive can be explored and quantified.

Another enticing application of the methods described herein is the assessment of devices and interventions in patient-specific and pathologically-relevant models. Despite tremendous progress made in the treatment of atherosclerosis and the devices used therein [2], complications and failures continue to dog the field. Using patient-specific arterial models and clinical

information on patient treatment and progression, it may now be possible to identify presently-unknown factors which may impact success of treatment and mechanisms of device or treatment failure. Instances of devices of particular interest include bioresorbable vascular scaffolds, drug coated balloons, and atherectomy devices, the structure and deployment of which may be modeled and simulated within the context of the vessel model. For example, complications such as stent malapposition and strut fracture can be examined. The sensitivity and impact of various morphological and mechanical features of plaques may be quantified with respect to simulation and clinical outcomes.

4.5. Conclusion

In this chapter, the computational modeling and simulation of patient-specific vessels was described. Enabled by greater geometric and constitutive information extracted from clinically available intravascular OCT imaging, improved patient-specific models of individual artery segments were developed, offering the promise of more comprehensive assessment to guide clinical decision-making. An inverse approach utilizing the new models was also demonstrated for determining lesion-specific material properties, with promising results of *in silico* and *in vitro* validation shared.

Models were developed with intricate, 3D heterogeneous structure on the basis of information which can be acquired non-destructively from patients. Offering greater model fidelity and wider applicability than previous methods, the approach was particularly amenable to use in *in situ* evaluations of mechanical constitutive properties through inverse FEA. By melding previous work on geometry and morphology segmentation with the simulation workflow, 3D material properties could be estimated for each tissue class on the basis of two image sequences acquired at different pressures. By segmenting mural structure in this way, constraints were imposed that negated the need for full displacement maps, which are susceptible to noise, typically one-dimensional or 2D, require homogeneity and uniform stress assumptions for direct stiffness estimates, and pose poorly constrained optimization problems. However, by leveraging this mural structure in defining the optimization target, richness in the dataset was preserved, thereby limiting the number of sequences required.

The work laid out in this chapter implements and operationalizes the results of work reported in previous chapters in extracting greater information from data currently collected in the

clinical setting. With additional work, including the assessment of modeling assumptions and acquisition of intraluminal pressure data, the proposed work may well open new windows into the biomechanical state of atherosclerotic coronary arteries and provide expanded awareness of the mechanical context of interventions. A wide range of future directions for this work offers exciting and enticing opportunities for clinical, translational, and fundamental research applications.

Chapter 5

Bridging the Chasm between Virtual & Physical Medicine with Synthetic Imaging

Vessel geometry, morphology, and computational models hold great promise in interrogating fundamental questions in cardiovascular disease pathology. Integration of computation with medicine also offers great potential in refining and personalizing the diagnosis, prognosis, and treatment of patients. However, despite significant strides in modeling physiological systems, resulting from both advances in computing and understanding of governing mechanisms, translation to clinical use and patient care has been scarce. One impediment to crossing the chasm from computer to clinic is the discordance in data representation and visualization—the output of computational models doesn't fit within the workflows and decision-making processes of clinical medicine. There are also impediments to crossing the chasm in the other direction, from clinic to computer. This represents a mismatch in priorities—while the clinic often cannot justify the complication risks and costs of testing and data acquisition not directly advancing patient treatment, computational models and researchers have an insatiable thirst for ever more data. In this chapter, a paradigm is presented to surmount the schism between physical and virtual medicine through the generation of synthetic medical images. Techniques developed to construct such images are described and compared, and potential applications and opportunities highlighted.

AI	Artificial Intelligence
cGAN ..	Conditional Generative Adversarial Network
CNN	Convolutional Neural Network
CT	Computed Tomography
GAN	Generative Adversarial Network
IVUS ...	Intravascular Ultrasound
ML	Machine Learning
MR	Magnetic Resonance
OCT ...	Optical Coherence Tomography
ReLU ..	Rectified Linear Unit
RMSE ..	Root Mean Squared Error
SGD	Stochastic Gradient Descent
VH	Virtual Histology
2D	Two-Dimensional
3D	Three-Dimensional

5.1. Synthetic Medical Imaging: Needs, Progress, & Challenges

Even as increasingly capable and complex technologies and techniques are developed with the aim of advancing clinical practice, there remains a divide between engineers, scientists, and the medical doctors who are the target of growing sources of data and information. Concurrently, photo manipulation and generation facilitated through the application of artificial intelligence (AI) has prompted a reckoning in modern culture; convincing synthetic photos and videos force observers to question the reliability of their own discerning faculties while also providing opportunities to explore alternative realities that never came—or have not yet come—to pass. At the intersection of these trends lies unexplored opportunities for synthetic medical imaging. The implications of such technology in clinical medicine are both underexplored and profound; this potential solution to the translational gap offers both exciting prospects and foreboding admonitions.

5.1.1 The Need for Synthetic Imaging in Medicine

Integration of computation with medicine offers great promise in improving the diagnosis, prognosis, and treatment of patients. Significant strides have been achieved in modeling physiological systems, driven by both advances in computational technologies and mechanistic biological understanding. However adoption for clinical use and patient care has been disappointingly scarce [113]. One impediment to translation is incompatibility in data representation and visualization. The output of computational models, which can be highly technical and require specialized training to fully appreciate, isn't readily integrated into the decision-making processes and workflows of medical practice. Furthermore, the disconnect between modeling and clinical practice acts in the opposite direction, inhibiting the medical knowledge and data from informing technological development. While scientists and engineers desire vast amounts of labeled image data with which to develop their models, clinicians often cannot justify the acquisition of additional data not directly advancing patient treatment, nor spare the time needed to laboriously label the data that is generated. Even so, the quality of the images that are acquired varies, and may compromise the utility to either group. These challenges are particularly prevalent in cardiology.

Cardiovascular medicine is a data-rich specialty, capturing and communicating vast amounts of information numerically and visually as images. These modes of communication

establish a common language and facilitate holistic consideration—humans are particularly adept at interpreting information-dense images and rendering holistic assessment and judgment—and clinicians rely on visualization in diagnostic, prognostic, and intervention-planning workflows. However, their clinical utility is susceptible to variability in image quality, interpretability, and completeness. Specialization, training, and expertise also vary by imaging modality among clinicians and institutions. In some cases, honed clinical sense and intuition may enable experts to overcome these challenges, but this ability is case- and clinician-specific. Exacerbating these potential issues, the addition of new sources of information, even if complementary, may induce an information overload, whereby the additional information may provide limited benefit to, or even detract from, informed deliberation and consideration. Consequently, there is an unaddressed need to assimilate and fuse disparate complementary data, provide more coherent perspective from parochial images, and integrate information intuitively across or beyond imaging modalities.

Furthermore, many technological advances, including computational methods, do not speak this “common language” central to clinical care, and so are not readily interpreted or integrated by those in the cardiovascular field. For example, patient-specific virtual phantoms—*in silico* representations of physical systems—promise to offer greater insight and foresight by computing parameters that are prohibitive to measure directly and simulating interventions or disease progression to inform treatment decisions, as discussed in Chapter 4. The outputs of such models, though, are often expressed in time-varying numerical solutions to partial differential equations, tensors of stress or strain distribution, or similarly abstruse structures, at best consolidated into arrays of metrics or somewhat more intuitive heat maps. Even such elaborate representations are not readily integrated into the image-based workflows and decision-making processes of personalized medicine [67], and so fail to reach their potential. Synthetic clinical imaging offers a potential means to provide the necessary “translation” to bridge this gap.

In addition to facilitating clinical care and integrating technical advancements into medical practice, the lack of key clinical data obstructs the advancement of the medical technology itself. This is particularly true in the training and validation of image processing and image analysis tools to aid medical care providers in digesting and interpreting image data. Several challenges exist in this realm. First, training and validation of such methods typically require vast amounts of labeled data from which to learn and against which to be compared. This type of data is difficult to acquire in large quantities, as data must typically be labeled manually and painstakingly by experts with

substantial training (and even then is subject to human error and inter- and intra-observer variability in labeling). Beyond the issue of establishing reliable standards is the challenge of data imbalance, discussed in Chapter 3. Some features, phenotypes, or presentations are relatively scarce, which makes the training of methods to identify them difficult. Standard augmentation methods, such as mirroring or rotating existing images is inherently limited, yet increasing the data pool for these classes can take significant time (to encounter naturally in due course). Even when classes are sufficiently represented in the population and labeling can be reliably performed at scale, there is still the matter of collecting the data, which invokes potential cost, risk, ethical, and privacy considerations (particularly if not part of routine patient care). The ability to synthetically generate new data, particularly of scarce presentations, with inherent labels obviates each of these impediments.

Technology development can be hindered by its unproven nature, cost, safety, and ethical considerations, and the acquisition of new or additional imaging data needed for its advancement or demonstration may not be justified. For example, computational methods (such as that used for constitutive property recovery discussed in Chapter 4) often require and consider multiple image acquisitions. While successful demonstration of such methods may unlock benefits for clinicians and patients, current (proven) benefit, clinical guidelines, and funding limitations do not substantiate the acquisition of additional imaging sequences. Such pragmatic considerations restrict exploration of new methods which could be employed if such data were available, and which have the potential to provide value justifying its acquisition in the future. Furthermore, methods (or devices) developed must be robustly validated in any use scenario. This important safeguard can significantly increase costs and delay deployment, as the method must be applied redundantly for extended periods and in a sufficiently diverse population of phenotypes. The desire to challenge and fully assess methodologies in a vast array of diverse presentations, while reducing the time and cost of such studies, has birthed an emerging paradigm of computer-simulated clinical trials for regulatory evaluation [341], which utilizes simulated patients and images but remains in its infancy.

In various ways, synthetic imaging presents rich opportunities to address each of the above needs and challenges, as explored later in this chapter and by predecessors in the fledgling field.

5.1.2 Progress in Medical Synthetic Imaging

AI and machine learning (ML) have gained tentative but growing popularity in medicine, and continue to offer great promise for their ability to tirelessly integrate vast amounts of data [122]. Applications to medical imaging are particularly attractive, especially in the field of cardiology where information-dense imaging has assumed a keystone role. Moving forward, AI to augment human intelligence and offer decision support offers perhaps the most challenging and beneficial future applications [133], [134]. Yet, while automated diagnosis, risk-stratification, decision-making, and intervention guidance are vaunted goals, more immediate contributions of AI will be to ancillary tasks and care delivery workflows [117]. Garnering particular interest are applications to collaboratively support clinicians in the amelioration of routine tasks [123], [125] and processing of incomprehensibly vast troves of medical data [125]–[129]. Residing at the intersection of routine task automation, big data processing, and augmented decision support, deep learning for medical image analysis is a booming, promising specialty [144]–[147]. A niche corollary in medical image generation and synthesis—which inverts prevailing paradigms of AI as an integrative tool of distillation—has attracted some interest in the research community as well, but is yet to find a foothold in any clinical applications [153]. However, recent exciting advances in ML addressing image processing and computer vision offer new possibilities for this specialty field.

Interest has grown in medical image generation alongside progress in image segmentation within the medical imaging field. Motivation has been driven by difficulties in comparing data acquired with differing image modalities, as well as ethical considerations surrounding patient health and unnecessary risk exposure limiting availability of data for various research applications. Physics-based simulations can be conducted given a sufficiently simple system, sound understanding of the governing physics, and known, reliable physical behavior of the imaged tissues. For example, image generation has been accomplished with multiple radiation-based imaging modalities using ray tracing projections and Monte Carlo simulations, but these simulations required sophisticated anatomical models—digital or computerized phantoms—which are generated through constrained manipulation of intricate generic precursors [342], [343]. Even so, these existing anatomic models provide resolution on the organ, rather than tissue (or even cellular) scale, limiting their applicability to large scale imaging uses [343]. Furthermore, such physics-based models are not amenable to imaging systems governed by more complex physical

phenomena, which demand alternative approaches. The use of AI, which does not rely on a deep understanding of the physical principles underlying the imaging modality, are very attractive in such cases, and have motivated the application of ML, particularly generative adversarial networks (GANs), to the generation of medical images [153].

In GANs, two models—a generator and a discriminator—are trained concomitantly and in adversarial fashion, wherein the generator learns to construct more and more realistic images while the discriminator learns to distinguish original images from synthesized ones with greater and greater precision; the output of the generator serves as an input to the discriminator, and the output of the discriminator contributes to the loss of the generator [287]. In conditional GANs (cGANs), the two models are also provided with a structured input, which the network learns to map to the desired output image. Real images serve as the target for the generator and labeled input data to the discriminator. (See Section 5.2.1 for a more comprehensive description.)

GANs have found several applications in medical image generation. Researchers have used computational techniques to generate synthetic images independent of any reliance on physical imaging principles [153]: magnetic resonance (MR) images have been directly translated into corresponding x-ray [344] and computed tomography (CT) images [345], [346] without a virtual phantom intermediary or physics-based simulator. However, these methods require an existing image dataset to transform; they would not facilitate generation of image data from a manipulated model—that is, one for which there may exist no physical basis in reality and consequently no directly corresponding imaging data. Intriguing work by Shin *et al.* generated MR images using a cGAN from a brain atlas manipulated to introduce various tumor labels [347]. Though able to introduce variation by imprinting new features (i.e. tumor label) onto the atlas, the map was otherwise static and did not facilitate further interaction or computational modeling. Still, highlighting both the protection of personal health information and the ability to augment training data for the detection of relatively scarce features, the images were subsequently used to augment a dataset in training a tumor segmentation network, improving its performance [347]. However, investigation of comparable applications to the intravascular imaging sub-field have been underexplored.

Intravascular optical coherence tomography (OCT) and intravascular ultrasound (IVUS) provide cross-sectional images of a blood vessel wall which provide rich information to

interventional cardiologists, as discussed extensively throughout this thesis. OCT images have been previously simulated for research purposes through physical first principles by modeling the propagation and coherence of optical signals [112]. However, they have not been of complex physical systems such as arterial walls, but rather exceptionally simple physical systems (specifically gelatinous cubes with a hard inclusion) [112]. Simulating more complex systems has previously not been achieved, nor applied in a larger system. Indeed, interactions between highly-variable heterogeneous materials, internal reflections resulting from complicated geometry and morphology, and the unintuitive nature of coherence measurement combine to make precise physics-based simulation of OCT images nearly intractable. IVUS images have been generated by a physics-based ultrasound image simulator and subsequently refined through several iterative applications of GANs to improve image texture and overall authenticity of appearance [348]. While producing convincing images, the method required both physics-based modeling to obtain general patterns and ML methods to improve texture and realism of the image. The more severe limitation, however, was that it only considered a homogeneous wall comprised of a single material [348], and thus was not applicable in clinically relevant diseased states. No solutions have been proposed for simulating intravascular imaging of diseased vessels relevant to clinical treatment.

5.1.3 The Endeavor to Bridge the Chasm between Virtual and Physical Medicine

This chapter details technological advancements and insights that inform an optimistic and enthusiastic perspective on how computational modeling and assessment of diseased vessels, developed from intravascular imaging, can be integrated into the clinical workflow. Integrating and building upon the achievements, lessons, and insight of work presented in the previous chapters, the work of this chapter seeks to bridge the chasm between virtual and physical medicine through synthetic medical imaging. Such a bridge enables emerging research platforms, including those presented in this thesis, to inform prognosis and treatment of atherosclerosis and advance fundamental knowledge of pathology and intervention mechanisms.

A series of methods was explored with the purpose of generalizing the concept of image processing to not just deduce, but also communicate, information of interest. By learning the unspoken nomenclature of images in conveying geometric and morphological information, an unconventional approach to image processing was developed to express such information in

familiar visual vernacular from any digitized source of choice. In particular, synthetic medical images were generated of tissue or systems which may or may not have a basis in a physical, existing tissue or system. The work primarily applied ML methods to generate realistic OCT and IVUS images of diseased arteries given only cross-sectional tissue class label maps. The generation approaches were largely agnostic to imaging physics, and relied only on the basic assumption—arising *a priori* and thoroughly confirmed through the previous work on tissue characterization and segmentation (Chapter 3)—that different tissues appear differently (that is, interact differentially with the imaging system). Indeed, the goal of this work to generate realistic intravascular images of diseased arteries given only cross-sectional tissue class label maps can be essentially framed as the reversal of the tissue classification procedure.

In this chapter, an AI tool is developed to synthesize enhanced yet familiar clinical images to augment the cardiologist’s visual clinical workflow. A departure from other AI approaches in cardiology [115], [150], it is adept at handling very real clinical scenarios by accommodating preconceived notions of appearance to quickly convey relevant information not available from individual raw images. This platform can fuse disparate complementary data to provide more coherent perspective from parochial images; anatomic inputs can be reconstructed as composites of several sources of data and information, resulting in a single comprehensible and intuitive visual representation. Images acquired with one modality can also be rendered as those of another. The tool demonstrates the potential to assimilate otherwise overwhelming sets of data with high speed and volume capacity. As such, the AI tool could potentially advance—rather than simply automate—honed clinical sense by integrating augmented information into enhanced yet familiar clinical images. The methodological framework is presented here, drawing in intravascular imaging to illustrate the potential of this approach.

Furthermore, maps used by the image generation models can be extracted from an *in silico* model, like those discussed in the previous chapter. Such models may be built and/or manipulated independent of existing source image data; computational models—virtual and *in silico* representations of the physical system—not only enable generation of different projections and views, but, through simulation and modeling, visualization of the system at different time points and conditions. The capability to project in a familiar way information from tissue class spatial maps extracted therefrom may be key to closing the loop between clinicians and virtual models by transforming technical information contained within the computational models into clinically

relevant visual representations for interpretation by medical experts. This capability also offers unique, fully-controlled and inherently labeled datasets to aid in technology advancement.

The following sections describe work in generating synthetic intravascular images and potential applications for such technology in the cardiology field at large. Initial foundational and informative work applied a simplified and intuitive physics-based model, following the basic principles of ray tracing, to simulate OCT images; results served as a useful baseline for comparison and an opportunity to appreciate the limitations of such approaches. Recent advancements in applying deep learning, a form of AI and ML, were subsequently used to generate synthetic medical images by learning the underlying relationship between system morphology and appearance in varying image modalities. Alongside the technological platform, the following section focuses on select proposed applications and clinical utility, including integration of augmented morphological information for clinical interpretation, merging of hybrid or multimodal image datasets for unified visualization, as well as translation between modalities. Consideration of future applications suggest that synthetic imaging could become a mainstay clinical tool for diagnostic, prognostic, and intervention-planning workflows in which cardiologists regularly render holistic assessment and judgement of visual data. The technological framework furthered herein offers promise in enhancing ancillary support of diagnosis and prognosis, improving training and education of clinicians, and advancing further technology development and testing.

5.2. Synthetic Image Generation

Integrating knowledge of vascular anatomy, pathology, and imaging, a well-performing approach was devised through various iterations and applied to synthesize realistic images replicating intravascular OCT and IVUS on the basis of vessel geometry and morphology. A preliminary approach optimized a mechanistic (physics-based) model to generate OCT images on the basis tissue distribution. Subsequent approaches leveraged the powerful, modern tools of deep learning; in brief, a neural network was trained on characterized images of coronary artery cross-sections (posterized representations of the various classified tissue types comprising the wall; Chapter 3) and the associated IVUS or intravascular OCT images. The generative networks were then able to construct realistic synthetic IVUS or OCT images given just a classification map (which may be entirely contrived, determined through characterization of alternative imaging

modalities, defined by a computational model, or estimated from other data sources). The training and performance of these methods were critically assessed.

5.2.1 Problem Definition and Technical Structure

Translation of a class label map into an image is not a new problem in ML. However, details of the implementation are not universally applicable—though they are largely independent of the underlying imaging physics, factors such as data structures, loss functions, architectures, and hyperparameters should conform to the specific problem being addressed, as well as the specific aims and goals of the image generation. For example, the image generation goal can be cast as any of the following: (1) minimizing the intensity difference between real and generated images resulting from the same underlying physical composition; (2) minimizing the error of subsequent reclassification routines performed on the generated image (i.e. deviation from the input classification map); or (3) minimizing the ability of a human or machine to determine whether or not the generated image is real.

In this work, such map-to-image translation was performed for OCT using a simplified physics-modeling approach and two different ML methods: basic pixel-by-pixel regression with hand-picked features, and cGAN image generation. All three leveraged knowledge of the dataset, underlying physiology and pathology, and basic understanding of the imaging modality to output estimates of grayscale (intensity) values for each pixel of the synthetic image. The actual grayscale value of corresponding pixels in the original images—discrete positive integers between 0 and 255—represented the target values during the model training process. Note that, while OCT images are usually presented as gold-colored images, they are, in fact, just grayscale intensity images to which a known color map is applied to aid shade differentiation by the human eye. Therefore, to reduce the complexity of the problem, target images were converted to their corresponding grayscale representations and the models similarly generated grayscale images (to which the known color map could subsequently be applied). This aspect of the data structure across methods, which leveraged contextual knowledge of the problem and dataset, allowed the transition from three-dimensional (3D) to two-dimensional (2D) outputs through exclusion of color information.

MATLAB (MathWorks, Natick, MA), running on a local workstation (2 Intel Xeon processors, 12 GB RAM), was utilized to implement the physics-based approach, which relied

upon functions in the Optimization Toolbox. Colaboratory (Google, Mountain View, CA) was utilized as the primary computational resource for the ML methods, with model construction, training, and testing executed on a virtual machine running Python 3 code with a graphics processing unit hardware accelerator. NumPy and TensorFlow libraries were used in implementing the regression and cGAN models, respectively. MATLAB was used for data pre- and post-processing.

Physics-Based Model Formulation

In the basic physics-based approach, pixel-wise grayscale intensity values were calculated by explicitly modeling physical processes which give rise to images using optimized tissue optical properties. In brief, a method akin to ray tracing explicitly modeled light propagation, attenuation, and reflection within the vessel wall, and found optimal tissue parameters governing these processes for each class. Grayscale values were predicted for each individual pixel. In this case, the only aforementioned goal that could be readily expressed in the formulation of this model was (1), minimizing the intensity difference between corresponding pixels of the real and generated images.

As for ray-tracing methods employed to simulate other modalities (e.g. CT), this approach was broadly guided by the underlying physics governing the imaging. OCT and IVUS imaging are performed in polar (cylindrical) coordinates, with the images being constructed in polar columns (A-lines) radiating from the catheter; the signal is generated and subsequently detected from within the catheter—located at the center of the image. The signal penetrates, and is attenuated, scattered, and reflected within, the vessel wall, whose constituent plaque tissues display distinct optical and echoreflective properties (Figure 5.1). Considering this physical process, tissue distribution was considered in polar coordinates, with each of 7 tissue types (calcium, fibrous, lipid, mixed, healthy or media, perivascular, and obscured/shadowed) being assigned a unique set of material optical properties, μ —one lumped attenuation and scattering term (as both mechanisms prevent signal from returning to the detector) and one reflection term. From these values, probability maps were constructed for each given plaque distribution, with each indexed pixel corresponding to the properties of the tissue class label at that point; p_{att} indicated the probability that a signal would be attenuated (or scattered) while passing through the region of tissue represented by the pixel, while p_{ref} indicated the probability that a signal would be reflected while passing through the space. To generate more realistic images, scaled normally distributed random noise was introduced to the

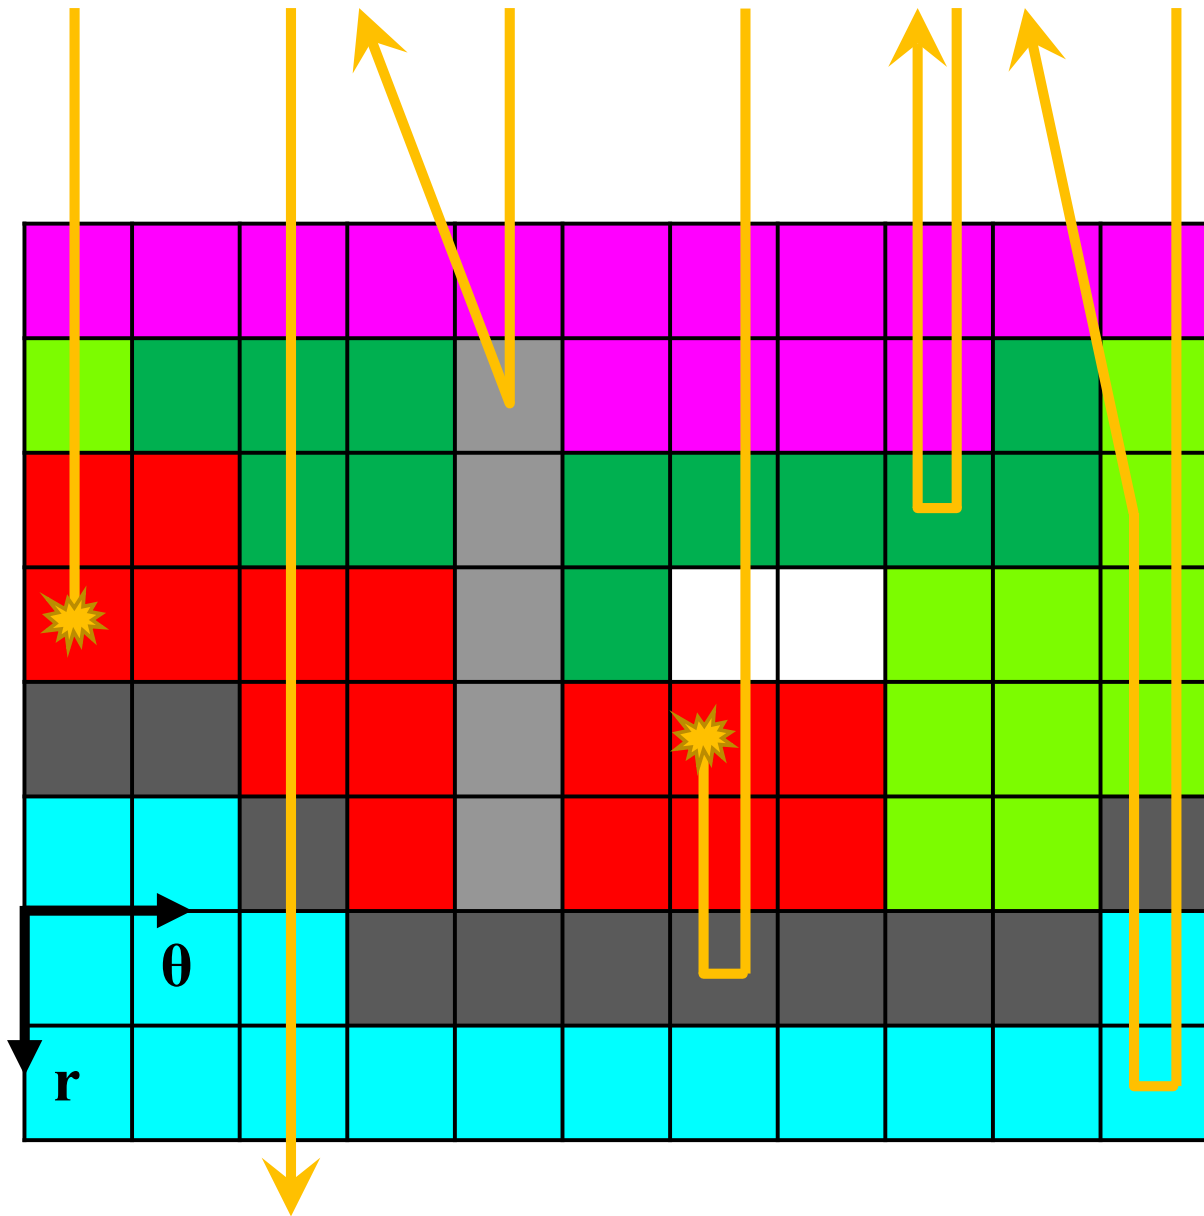


Figure 5.1. Schematic of the physics-based method used in the synthetic image generation method. By tracing the path of a “ray” through discretized tissue elements, an estimate of the probability that the signal is reflected and returns to the imaging sensor for a given element (and thus the strength, or brightness, of the signal) could be made. Each material class (shown here in various colors) had different properties, which dictated the probability that a signal was reflected or attenuated (or scattered) within a discrete region of that material. In each element, a signal could be reflected back towards the receiver, attenuated (absorbed), scattered (redirected away from the receiver), or pass through unaffected; various such outcomes are shown. However, rather than a single ray of light/signal with a single outcome (as shown), outcomes were considered in probabilistic terms for a large starting population of signal emitted into the column (A-line). Element size is exaggerated here for clarity; individual pixels served as the discrete material elements.

maps of p_{ref} and p_{att} (imitating in part the effect of discrete scatters distributed throughout the tissue). Consequently, the probability that signal passes through the tissue is the remainder:

$$p_{pass}(r, \theta, \mu) = 1 - p_{att}(r, \theta, \mu) - p_{ref}(r, \theta, \mu). \quad (69)$$

(Note that an equivalent interpretation is that the probabilities, p , indicate the fraction of a signal that experiences the given outcome.) The resulting strength of the signal reflected at each pixel, attenuated in both directions of travel, formed the reconstructed, synthetic image (I_s), which can be expressed:

$$I_s(r, \theta, \mu) = B \times p_{ref}(r, \theta, \mu) \times \prod_{\rho=0}^{r-1} p_{pass}(\rho, \theta, \mu)^2, \text{ where} \quad (70)$$

B is a scalar brightness multiplier to achieve a maximum brightness equal to that observed in the training dataset. Note that the return of signal undergoing multiple internal reflections within the wall, which increases the complexity of the model exponentially, was not considered.

The image, as the real one, was generated in initial raw, uncompressed format. For optimization purposes, the compressed grayscale image format of the actual (i.e. real) images (I_{compr}) therefore had to be transform back to their initial uncompressed format (I_a) [218]:

$$I_a(r, \theta) = \left(\frac{I_{compr}(r, \theta)}{255} \right)^4. \quad (71)$$

The optimal material optical properties, μ , were identified by employing a nonlinear programming solver. The standard interior-point approach to constrained minimization found the minimum of a constrained nonlinear multivariable function to identify the model parameters (μ^*) that minimized mean absolute error (L1 loss) for all pixels indices (r, θ) across all N_{train} images in the training dataset:

$$\mu^* = \underset{\mu}{\operatorname{argmin}} \frac{1}{N_{train}} \sum_{n=1}^{N_{train}} \mathcal{L}(I_s(\mu), I_a), \text{ where} \quad (72)$$

$$\mathcal{L}(I_R(\mu), I_a) = \mathcal{L}_{L1}(\mu) = \sum |I_s(r, \theta, \mu) - I_a(r, \theta)|. \quad (73)$$

To reduce degrees of freedom while imposing domain knowledge of the modality, the optical properties of shadow were fixed such that its probability of attenuation/scatter was 100%.

Because the physical elements which produce the catheter and guidewire artifact were not encoded in the tissue map, imaging artifacts, if desired, had to be added in a secondary step. The location and arc length of the guidewire shadow of the vessel wall was used to infer the position

of the guidewire relative to the imaging catheter, and this information was used to introduce the high-intensity intra-luminal crescent-shaped guidewire reflection artifact. The concentric rings of the imaging catheter artifact in the center of the image were added using the artifact observed in the training dataset as a simple template.

Regression Model Formulation

In the basic pixel-by-pixel regression approach, pixel-wise grayscale intensity values were predicted based on a set of curated features. Using standard supervised ML regression techniques, regression models were trained to predict grayscale values for individual pixels, thereby “generating” images in a piece-wise manner. In this case, as for the physics-based model, the only aforementioned goal that can be readily expressed in the formulation of this model is (1), minimizing the intensity difference between corresponding pixels of the real and generated images.

Feature selection and representation presented the primary challenge, and required consideration of the underlying physics governing the imaging to ensure all relevant information was made available to the model (an undesirable feature, for reasons described earlier, which motivates the use of alternative ML methods). However, unlike the physics-based formulation, the processes did not need to be explicitly modeled, providing some added freedom and flexibility. Considering the physical processes, described above, encoded data included pixel tissue class assignment, radial depth, and the total number of proximal pixels (i.e. discretized tissue elements between the catheter and the given tissue element) belonging to each of 8 pixel classes (calcium, fibrous, lipid, mixed, healthy or media, perivascular, obscured/shadowed, and lumen), altogether resulting in 10 features. Pixel tissue class assignment was represented as a one-hot encoded vector: $y = [y_1, \dots, y_K]^T$, where $y_k = 1$ if pixel i belongs to class k , and $K = 9$ is the total number of encoded pixel classes (the 8 classes noted above and another for the central catheter artifact region). All other features (including target grayscale intensity values) were standardized:

$$x_{(f)}^{(i)} = \frac{x_{(f)}^{(i)} - \bar{x}_{(f)}}{\sigma_{(f)}}, \quad (74)$$

where the subscript f indicates the feature index, $\bar{x}_{(f)}$ is the feature value mean:

$$\bar{x}_{(f)} = \frac{1}{n} \sum_{i=1}^n x_{(f)}^{(i)}, \quad (75)$$

and $\sigma_{(f)}$ is the feature value variance:

$$\sigma_{(f)} = \sqrt{\frac{1}{n} \sum_{i=1}^n \left(x_{(f)}^{(i)} - \bar{x}_{(f)} \right)^2}. \quad (76)$$

Recognizing that non-linear interaction between features may play an important role, feature spaces were constructed using 1st- and 2nd-order polynomial bases.

Standard supervised ML regression techniques were used to train regression models to predict grayscale values for individual pixels, thereby “generating” images in a piece-wise manner. Each pixel was associated with a set of values $(x^{(i)})$ corresponding to the aforementioned curated features and calculated during a pre-processing step, and the intensity difference between corresponding pixels of the generated and real images was minimized. Squared error was implemented as the loss function:

$$\mathcal{L} \left(I_s(r^{(i)}, \theta^{(i)}), I_a(r^{(i)}, \theta^{(i)}) \right) = \left(I_s(r^{(i)}, \theta^{(i)}) - I_a(r^{(i)}, \theta^{(i)}) \right)^2 \quad (77)$$

where \mathcal{L} is the loss function, $I(r^{(i)}, \theta^{(i)})$ is the grayscale value at discrete polar coordinate/pixel index $(r^{(i)}, \theta^{(i)})$, and subscripts s and a indicate synthetic and actual images, respectively. A linear hypothesis class h was considered for transformed feature set $\varphi(x)$:

$$I_s(r, \theta) = h(x; \Theta, \Theta_0) = \Theta^T \varphi(x) + \Theta_0 \quad (78)$$

where Θ and Θ_0 are the fit model parameters. Regression was treated as an optimization problem to find parameters that minimize mean squared error (with optimal parameters indicated by an asterisk, *):

$$\Theta^*, \Theta_0^* = \underset{\Theta, \Theta_0}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \mathcal{L} \left(I_s(r^{(i)}, \theta^{(i)}), I_a(r^{(i)}, \theta^{(i)}) \right) + \lambda \|\Theta\|^2. \quad (79)$$

Here, λ is a trade-off parameter providing weight to the regularizing term, which was added to prevent overfitting and reduce estimation error. Following best practices, stochastic gradient descent (SGD) with minimum step size η was performed to identify these optimal model parameters:

$$\eta(t) = \frac{0.01}{\sqrt{t+1}}. \quad (80)$$

Conditional Generative Adversarial Network Model Formulation

In the cGAN image generation method, full grayscale images were generated by the generative network based on the full input class label map (Figures 5.2 and 5.3). Neither physical

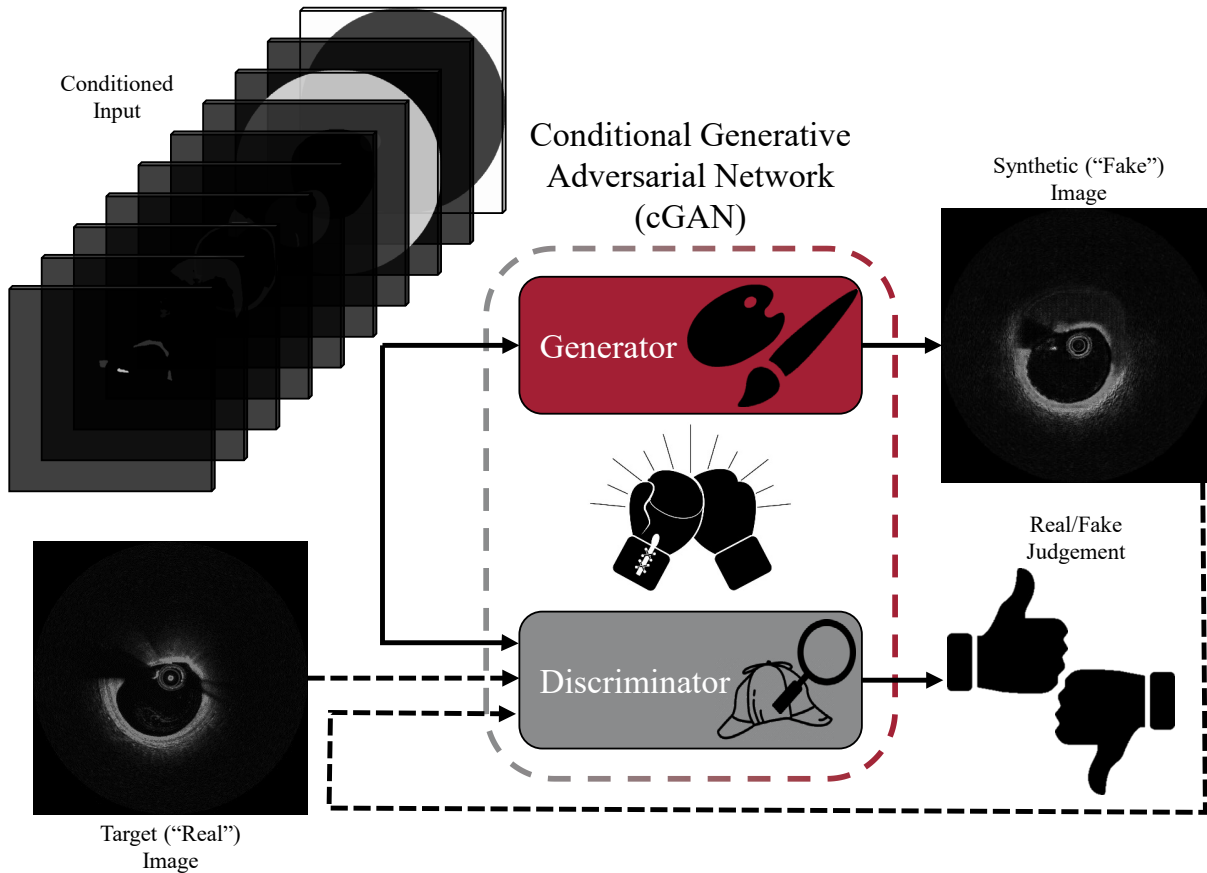


Figure 5.2. Schematic structure of the cGAN used in the synthetic image generation method. The network was comprised of a generator and a discriminator. The generator was provided with a conditioned input encoding the distribution of various tissue types and other key features of the image, and generated a corresponding image. The discriminator was provided with the same conditioned input as well as an image, and judges whether the image is real or fake (given the known tissue distribution and image features encoded in the conditioned input). During training, the generator is penalized based on direct comparison between the synthetic and target image and on the discriminator’s ability to identify its synthetic images as fake; the discriminator is penalized for falsely identifying the real target as fake or the synthetic image as real.

process modeling nor physics-informed feature design was required, providing the greatest freedom and flexibility of the three approaches described here. Furthermore, in this case, both goals (1) and (3)—minimizing the intensity difference between corresponding pixels of the real and generated images and the ability to determine whether or not the generated image is real—can be readily expressed in the formulation of the model.

In the formulation of the network that follows, note that notation differs here somewhat from that used in the physics-based and regression model formulations as described above. Here, x represents the conditioned input (the tissue class label map), y represents the target (actual grayscale image I_a) image, and z represents a random noise vector. Furthermore, G represents the

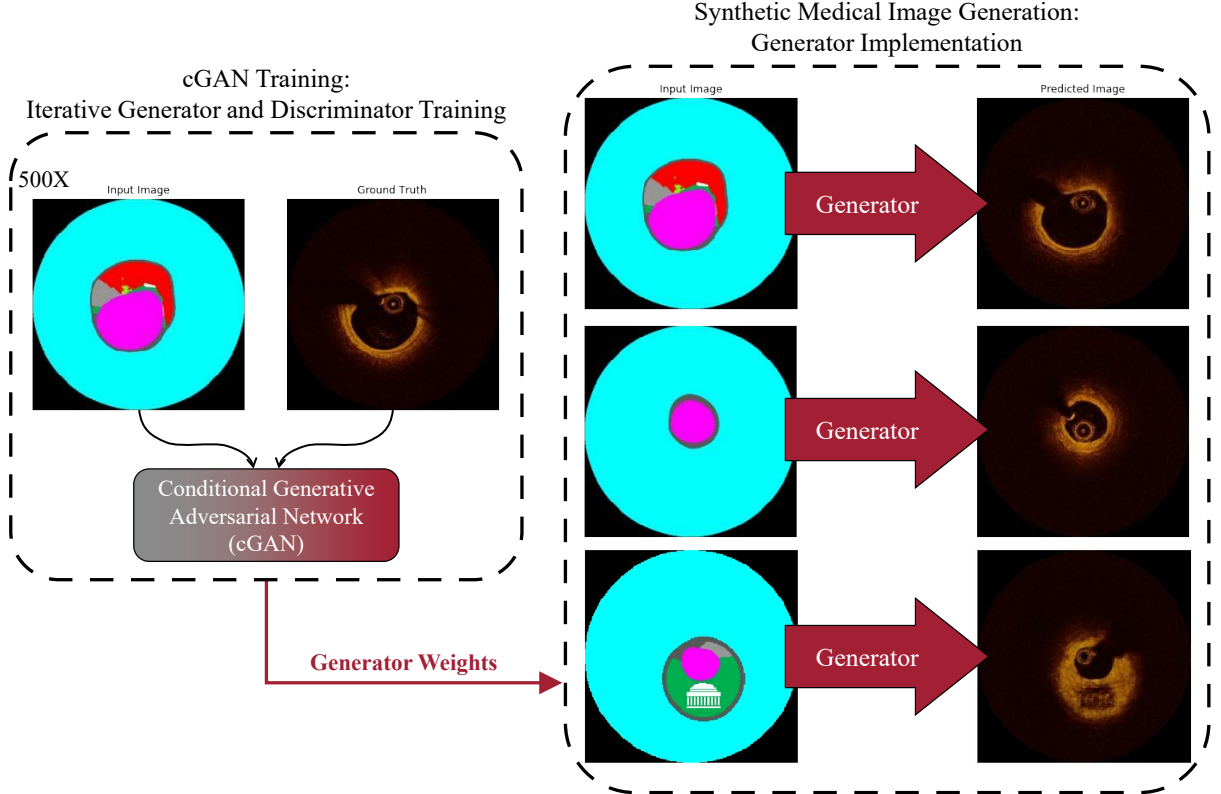


Figure 5.3. Workflow of the synthetic image generation method, illustrating the training of the conditional generative adversarial network and application of the resulting generator. Classification maps utilized as input images to train the conditional generative adversarial network were generated through image classification (e.g. by CNN or VH) and prepared for cGAN input by subsequent image processing. Here, the OCT generator is shown, though an additional generator was trained to generate IVUS images. Conditioned input of classification maps are shown here as color-coded images to represented multidimensional one-hot encoded arrays.

generation network (which produces synthetic image I_s) and D represents the discriminator network.

The overall loss of a cGAN can be expressed as the combined binary cross-entropy loss of the discriminator D when evaluating the true image (y) and the output image from generator G , given the same conditioned input x :

$$\mathcal{L}_{cGAN}(G, D) = \log D(x, y) + \log(1 - D(x, G(x, z))). \quad (81)$$

As an adversarial network, generator G is optimized to minimize this loss, while discriminator D is optimized to maximize this loss. The generator is further constrained, as in previous approaches to image generation [286], to concurrently minimize mean absolute error (i.e. L1 loss):

$$\mathcal{L}(G, D) = \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G), \text{ where} \quad (82)$$

$$\mathcal{L}_{L1}(G) = \sum |y - G(x, z)|. \quad (83)$$

Generator network weights were updated to minimize this weighted combination of mean absolute error (L1 loss) and error resulting from the discriminator’s ability to differentiate the generated and target images.

Summation was taken over the entirety of the image on a pixel-by-pixel basis, and hyperparameter λ was again a trade-off parameter, here providing weight to the L1 loss term. L1 loss has been found to encourage the generator to match the low frequency features of target images, but alone this loss function typically produces blurry images [286]. The cGAN loss term has been found to encourage the generator to match the high frequency features of target images, such as texture and style, but may not enforce global consistency and structure [286]. These loss terms are therefore combined to achieve sharp images with high-level structure.

During training, discriminator and generator loss, \mathcal{L}_D and \mathcal{L}_G , respectively, are considered separately for the purposes of weight optimization and back-propagation:

$$\mathcal{L}_D(G, D) = -\left(\log D(x, y) + \log\left(1 - D(x, G(x, z))\right)\right); \text{ and} \quad (84)$$

$$\mathcal{L}_G(G, D) = \log\left(1 - D(x, G(x, z))\right) + \lambda\mathcal{L}_{L1}(G). \quad (85)$$

The logarithm function prevents the loss from being dominated by outlier contributions, facilitating the consideration of overall trends. It can be seen that these component loss functions result in the same optimal points for the discriminator and generator as the overall loss function \mathcal{L} (Equation 82). In particular, the optimal generator (i.e. generator with optimal weights; G^*) is the one that minimizes both \mathcal{L} and \mathcal{L}_G (given an optimal discriminator):

$$G^* = \underset{G}{\operatorname{argmin}} \max_D \mathcal{L}(G, D) = \underset{G}{\operatorname{argmin}} \max_D \mathcal{L}_G(G, D). \quad (86)$$

Feature selection and representation were fairly straightforward in comparison with the physics-based and regression models, as specific features of relevance need not be identified *a priori*—one of the benefits of neural networks. However, choices regarding data structure did leverage knowledge of the imaging modality and imaged subject matter. For example, by utilizing grayscale rather than colored representations, the network required only one output channel, rather than three (corresponding to the red, green, and blue channels of colored images)—and their corresponding learned series of weighted connections. To facilitate use of standard activation

functions and normalize the operating range, target images were linearly scaled from 8-bit integer representation in the range $[0, 255]$ to the consolidated range $[-1, +1]$:

$$y = \frac{y_0}{255} \times 2 - 1. \quad (87)$$

Reducing the complexity of the problem in this way, and consequently the dimensionality of the output, yielded immediate improvements in the speed of training progression (both in number of iterations and time to execute each iteration), and may have contributed to reduction of some artifacts that were observed in preliminary systems which utilized colored images.

Furthermore, unlike image-to-image translation, which accepts input with numeric representation [286], the tissue class label map was represented with one-hot encoding of the discrete features. Each of the K channels represented a class, such that the information was encoded as a $500 \times 500 \times K$ array (resized to base-2 friendly $512 \times 512 \times K$ for network processing) with each channel representing a feature and exactly one channel activated for any given pixel (Figure 5.4). The K classes that were represented required some consideration and heuristic exploration. Initial classes included the 7 tissue types considered above (calcium, fibrous, lipid, mixed, healthy or media, perivascular, and obscured/shadowed), as well as lumen and image frame (the black region filling the square image in which the round data-inclusive image—converted from acquired polar coordinates—is set). After observing that the network struggled with the localization of the catheter artifact that appears in the center of every OCT image, an additional feature representing this artifact’s location was added, thus adding at 10th layer ($K = 10$; Figure 5.4).

The architecture of the implemented cGAN reflected that of a model investigated as a general-purpose image-to-image translation solution by Isola *et al.* [286], with key modifications to better suit the application. The generator encoder-decoder network (Figure 5.5) progressively downsampled then upsampled the input (thereby forming a “U” shape) with recurring convolution,




Figure 5.4. Feature representation for a single image prepared for input to the cGAN, using one-hot encoding with $K = 10$ features. Binary maps show, from left to right, (A) calcium, (B) fibrous tissue, (C) mixed tissue, (D) lipid tissue, (E) media and non-pathological tissue, (F) shadow, (G) lumen, (H) perivascular tissue, (I) catheter artifact, and (J) image frame. Channels were stacked in 3D (out of the plane of the page; $500 \times 500 \times 10$), but are shown here in concatenated (2D) form for viewing and convenience.




Figure 5.5. U-net architecture used to generate synthetic intravascular images. Blue boxes indicate multi-channel feature maps resulting from convolution, normalization, and activation operations; map sizes of hidden layers are symmetric. Skip connections (black arrows) concatenate channels of equal size.




Figure 5.6. CNN architecture of the discriminator used in adversarial training of the U-net to generate synthetic intravascular images. Blue boxes indicate multi-channel feature maps resulting from convolution, normalization, and activation operations. A grayscale image (real or fake) and tissue map (one-hot encoded) are provided as input; a scaled confidence map is output. During training (concurrent to the generator), the discriminator is penalized by binary cross-entropy loss for incorrectly designating a classifier-generated (fake) image as authentic or a ground truth (real) image as inauthentic (Figure 5.2).

batch normalization, and rectified linear unit (ReLU) activation operations. Several blocks of the generator encoder were subject to stochastic dropout. Additionally, skip connections concatenated channels in layers of the encoder with channels in a subsequent, equally-sized layers of the decoder, thereby generating a U-net [285]. The discriminator network (Figure 5.6) constituted a convolutional neural network (CNN) which progressively downsampled the conditioned input and accompanying image with recurring convolution, batch normalization, and ReLU activations.

Detailed descriptions of the network architectures and illustrative diagrams can be found in groundbreaking papers by Ronneberger *et al.* [285] and Isola *et al.* [286]. The network implemented here deviated somewhat from those in the papers, particularly in the number of input and output channels, as well as the number of upsampling and downsampling layers, in order to accommodate the selected data structures and input size. Otherwise, the hidden layers reflect those implemented by Isola *et al.* [286]. Note the similarity of the network to that used in semantic segmentation to characterize images (Chapter 3), which was, in fact, an adaptation of the cGAN network presented here.

Training was performed for 50 epochs using SGD with adaptive step size η selected by the Adam method, which accounts for objective function gradient and momentum to promote timely convergence. Whereas other image generation approaches often apply random jittering and mirroring to the training dataset to induce randomness and avoid overfitting, it was expected that jittering would be counterproductive in this application, as the center of the image, which always includes a predictable imaging artifact, should remain fixed. However, the rotational orientation of an image around that center point is arbitrary (dependent upon the uncontrolled initial radial position of the catheter at the time the acquisition is initiated—a general characteristic of intravascular imaging), so random rotation was instead implemented to introduce randomness (in addition to random mirroring, which was retained). Therefore, target images and conditioned inputs were subjected to random rotation (in 90° increments, to avoid interpolation-driven loss of information) and mirroring to avoid, or at least delay and reduce, overfitting of the data.

5.2.2 Demonstration, Refinement, and Assessment of Image Generation Approaches

To demonstrate, refine, and assess the various image generation approaches, synthetic image generate was implemented using clinical datasets of characterized intravascular images.

Dataset

Two image datasets, one including OCT images and the other including IVUS images, were utilized in this study. Both datasets were described in detail in Chapters 2 and 3 (and previous publications [157], [160], [161]). All three generation methods were trained to synthesize OCT images and used the former dataset. Only the cGAN method was trained to synthesize IVUS images and used the latter dataset.

As described previously, intravascular OCT and IVUS provide cross-sectional images of a blood vessel wall. In diseases such as atherosclerosis, aberrations and lesions in the lumen (central blood passageway) and wall can be visualized, identified, and delineated. Such features include stenosis (narrowing of the lumen), devices (such as stents and vascular scaffolds), and plaque (heterogeneous fatty deposits; hallmark of atherosclerosis) [18]. While manual annotation of various plaque components is possible and the resulting information can be useful in quantifying the disease state, this process is time consuming, requires significant training, and is subject to high variability between experts. Some IVUS systems are equipped with virtual histology (VH) capabilities, which utilize supplementary radiofrequency to automatically label tissue, albeit with relatively low resolution, but no such capabilities exist in established OCT systems. Motivated by these challenges, discussed in Chapter 3, deep learning methods and neural networks have been applied to fully segment vessel walls imaged with OCT into standard tissue classes associated with atherosclerosis—calcium, lipid tissue, fibrous tissue, mixed tissue, non-pathological tissue or media, and no visible tissue (i.e. guidewire shadow artifact) [160], [349]. Similar methods have been applied to fully segment grayscale IVUS images into the comparable tissue classes—dense calcium, necrotic core, fibrous tissue, fibrofatty tissue, and non-pathological tissue or media [161]. The resulting dataset offers cross-sectional tissue class label maps of diseased coronary arteries along the full tomographic series (Figure 5.7). This progress in image segmentation within the medical imaging field has assumed a critical role in enabling growth in the field of medical image generation.

The first dataset consisted of 751 OCT cross-sectional images of coronary arteries and their associated tissue classification maps (Figure 5.7). These medical images were drawn from pullbacks performed in 7 distinct patients in the course of their treatment; each patient underwent invasive angiography and were found to have one or more coronary stenosis of intermediate or greater severity. OCT imaging was performed in accordance with standard clinical acquisition procedures [18], as described for the same imaging dataset utilized in work described in Chapters 2 and 3 [157], [160]. The acquired images were 500×500 pixels in size. OCT images were fully segmented into standard tissue classes associated with atherosclerosis using the CNN-based method described in Chapter 3 [160].




Figure 5.7. A representative image pair illustrating an original cross-sectional OCT image of a coronary artery (*left*) and the associated map of the plaque classification (*right*). White: calcium; Red: lipid tissue; Dark green: fibrous tissue; Light gray: shadow (no visible tissue/catheter artifact); Dark gray: media and non-pathological tissue; Magenta: lumen; Cyan: perivascular tissue; Black: outer frame. An additional class (mixed tissue; light green), does not appear in this cross-section but is present in others.

The second dataset consisted of 553 IVUS cross-sectional images of coronary arteries and their associated tissue classification maps. These medical images were drawn from pullbacks performed in 8 distinct patients in the course of their treatment. VH-IVUS imaging was performed in accordance with standard clinical acquisition procedures [19], as described for the same imaging dataset utilized in work described in Chapter 3 [161]. The acquired images were 648×648 pixels in size. IVUS images were fully segmented into standard tissue classes associated with atherosclerosis by VH [235].

The encoded classes varied somewhat by modality, as reflected in the tissue characterization work discussed earlier (Chapter 3). In OCT, the six segmented classes were: calcium, lipid tissue, fibrous tissue, mixed tissue, non-pathological tissue or media, and no visible tissue. In IVUS, the five classes were: dense calcium, necrotic core, fibrous tissue, fibrofatty tissue, and non-pathological tissue or media. Of note, classified plaque types were unequally represented in the dataset (Tables 5.1 and 5.2). The label maps generated by the classification network and VH system were post-processed to label the vessel interior (lumen), exterior (perivascular tissue), catheter artifact, and black outer frame (Figure 5.7). (Intravascular image data are acquired in polar coordinates but displayed in Cartesian coordinates, so the circular image is framed to complete the

Table 5.1. Relative content of tissue classes within the classification maps of the OCT dataset

	Tissue Class					
	Calcium	Lipid Tissue	Fibrous Tissue	Mixed Tissue	Media/Non-Pathological	No Visible Tissue
Portion of Vessel Wall	2.07%	19.51%	29.98%	4.78%	36.94%	6.73%
Portion of Total Image*	0.09%	0.83%	1.27%	0.20%	1.56%	0.29%

*Balance of image (95.76%) comprised of lumen, perivascular tissue, and outer “frame”

Table 5.2. Relative content of tissue classes within the classification maps of the IVUS dataset

	Tissue Class				
	Dense Calcium	Necrotic Core	Fibrous Tissue	Fibrofatty Tissue	Media/Non-Pathological
Portion of Vessel Wall	2.75%	7.62%	37.70%	19.17%	32.76%
Portion of Total Image*	0.25%	0.69%	3.39%	1.73%	2.95%

*Balance of image (90.99%) comprised of lumen, perivascular tissue, and outer “frame”

rectangular representation.) Locations of these additional classes were all deterministically implied by the combined geometry of the original tissue classes and the image dimensions.

A portion of each set was reserved exclusively for testing and validation. Of the image-map pairs, 500 from each set were randomly selected for use in training and the remaining pairs were reserved for testing and validation. For the regression model, pixels were pooled for the frames allocated to each subset and randomly selected for use.

Evaluation Process

Parametric analysis was performed to evaluate the impact of, and sensitivity to, various hyperparameters, and to select the best values for the ML methods. Refinement, assessment, and comparison between methods was performed for OCT images. The hyperparameters which resulted in the best (most desirable) performance were selected for subsequent use. For the regression model, regularizer weight (λ ; Equation 79), feature polynomial base order, and maximum number of training iterations were varied, and changes in cross-validation root mean squared error (RMSE) and execution time were compared. Images generated from the resulting models were qualitative compared (Figure 5.10). For the cGAN model, the relative weight (λ ; Equations 82 and 85) of the generator loss function components—L1 loss (mean absolute error) and GAN loss (differentiability of generated images from real images by the discriminator)—were varied, and the resulting changes in the various discriminator and generator loss functions were quantified. Again, images generated from the resulting models were qualitative compared. The

cGAN model was subsequently extended and applied to generate IVUS images, which were qualitatively assessed.

5.2.3 Results and Analysis

The three methods—explicit physics modeling, feature-based regression, and conditional adversarial generation—differed substantially in quantitative and qualitative performance, differing in training time, execution speed, and synthetic image fidelity. Representative images generated with physics-based and ML approaches, utilizing the tissue class label map and emulating the real image shown in Figure 5.7, are displayed side-by-side in Figure 5.8.




Figure 5.8. Representative actual (*top left*) and synthetic OCT images generated through simplified physics-based simulation (i.e. ray tracing; *top right*), regression modeling (*bottom right*) and cGAN synthesis (*bottom left*), corresponding to the same underlying tissue label map (i.e. structure; Figure 5.7). The three synthetic images illustrate some distinct similarities and differences in brightness, texture, and realism.

Physics-Based Image Generation and Impact of Noise

The physics-based image generation method employing a simple model of the OCT imaging modality successfully generated synthetic images from the provided tissue class label maps. Optical property (μ) optimization (i.e. training; Equation 72) took 2.4 ± 0.1 hours to complete. Subsequent image generation required just 0.56 ± 0.06 seconds per frame, though the computation time was dominated by polar-Cartesian conversion (0.49 ± 0.02 seconds per frame); the method was otherwise highly scalable and rapid.




Figure 5.9. The physics-based method captured general features of the true images it aimed to recapitulate. Relative to the real image (*top left*), the generated images appeared somewhat posterized and cartoonish due to their deterministic relationship with the provided label map (*top right*). This was particularly true without added noise (*bottom left*); the introduction of noise to the attenuation and reflection probability maps introduced a semblance of texture, enhancing the realism of the image (*bottom right*).

A resulting image generated in this manner is shown in Figure 5.8 (top right). The image is notably quite bright, with high levels of contrast. Also of note, the physics-based generated image demonstrates a posterized, cartoonish appearance arising from the image's deterministic relationship with the provided label map, though the introduction of artificial noise somewhat disrupted this artificial appearance and introduced a semblance of texture, enhancing the realism of the image (Figure 5.9).

Parametric Study of Regression Model Hyperparameters

Training the regression model was a relatively fast endeavor. A single iteration took 63.2 seconds, but 10^6 iterations took just 151.9 seconds, highlighting a highly non-linear scaling relationship (Figure 5.10). For the model, 1st-order polynomial basis feature representation offered far better performance than 2nd-order polynomial basis not only in execution time—which is expected, as the number of transformed features was increased from 18 to 189—but also in cross-validation RMSE. Therefore, 1st-order polynomial basis feature representation was implemented in subsequent testing and validation. RMSE was found to increase monotonically with regularizer λ , which indicated that any decrease in RMSE attributable to greater generalizability arising from the regularization was more than offset by increases in loss arising from the repression of hypothesis complexity (structural error); the regularizer did have some positive effect in a narrow range for the exponentially higher dimension (and worse performing) 2nd-order polynomial basis model. Nevertheless, a marginal degree of regularization ($\lambda=10^{-1}$) was retained as a precautionary measure. With these hyperparameters, results converged within 10^5 - 10^6 iterations (Figure 5.10).




Figure 5.10. Parametric analysis quantifying the impact of hyperparameters on regression model performance. Model performance is quantified in terms of cross-validation RMSE and execution time. *Top:* Regularizer loss term weight (λ) and feature polynomial base order were varied and results were evaluated after 10^5 iterations of training. *Middle:* Detailed view showing a subset of values for the 1st-order polynomial basis from the first plot (*top*). *Bottom:* Convergence was assessed over increasing training iterations for the 1st-order basis with $\lambda = 10^{-1}$.

Interestingly, though not surprisingly, results of the regression model share several features with those generated by a simple physics-based model. This is largely expected, as the features selected for representation in the regression model input data were deliberately selected with consideration of the same governing physical imaging process and principles. Due to the inherently deterministic relationship with the provided tissue label map arising from a linear relationship among basic features, images appeared similarly posterized and cartoonish—even more so than the physics-based method due to the absence of introduced noise. However, a departure of the physics-based image from the regression-generated image was the brightness of each, with the former being generally much brighter than the latter. This is likely because expected exponential decay of the light signal with depth was explicitly modeled by the physics-based method, but could not be achieved by the simple linear regression model. As a consequence of the limited dynamic range which could be achieved, the optimal model which minimized the global error for generating the images—which, due to limited penetration depth of the signal, are dominated by poorly-visualized tissues with dark appearance—was one with broadly dark grayscale values. As rapid drop-off and large intensity differences between adjacent tissue was not feasible, this regression approach may only be valid in scenarios where imaged tissues possess qualities resulting in constant or linearly diminishing signal intensities, which precludes most pathological vascular tissue (particularly lipid).

cGAN Loss Contributions and Progression with Training

Although hyperparameter λ (L1 loss weight; Equation 85) only modified the generator loss function, it was found to play a significant role in the training progression and performance of both the generator and the discriminator of the cGAN. Regardless of this weight value, the general trend during training was for both generator and discriminator loss magnitudes to initially decrease rapidly, then subsequently continue to increase (Figures 5.11 and 5.12). When observing training and outcome with respect to weight value, including the edge cases in which only L1 loss or GAN loss were used (but not a weighted combination of both), interesting trends arose. (Note that the “ $\lambda = \infty$ ” case indicates that only the λ -weighted L1 loss term was considered and the GAN loss was neglected, not that the actual applied weight was infinitely large.) L1 loss quickly allowed the generator to defeat the discriminator (incur low GAN loss), but the generator subsequently becomes progressively less successful as training continued and the discriminator developed further. Indeed, when L1 loss was used alone, the discriminator was, over time, able to achieve

nearly 0 loss (i.e. correctly identify every generated image as fake and every real image as such). Furthermore, while heavily weighting or exclusively using L1 loss allowed the generator to rapidly achieve low L1 loss (as expected, because this was the objective which the generator was being optimized to achieve), GAN loss alone was still able to achieve similarly low L1 loss values nearly as quickly and in a sustained manner. Finally, it can be observed that generated loss and real loss—discriminator loss resulting from misclassification of generated and real input images, respectively—follow each other closely as training progressed. This is not unexpected, as convincing output from the generator makes it more likely that the discriminator misclassifies both real and generated images (since they are alike), while unconvincing output from the generator makes the discriminator’s task easy, such that it is able to successfully distinguish one from the other consistently and incurs little loss on either front (akin to a high-margin problem). As noted earlier, a successful image generation system is therefore one in which generator loss is low but discriminator loss is high.

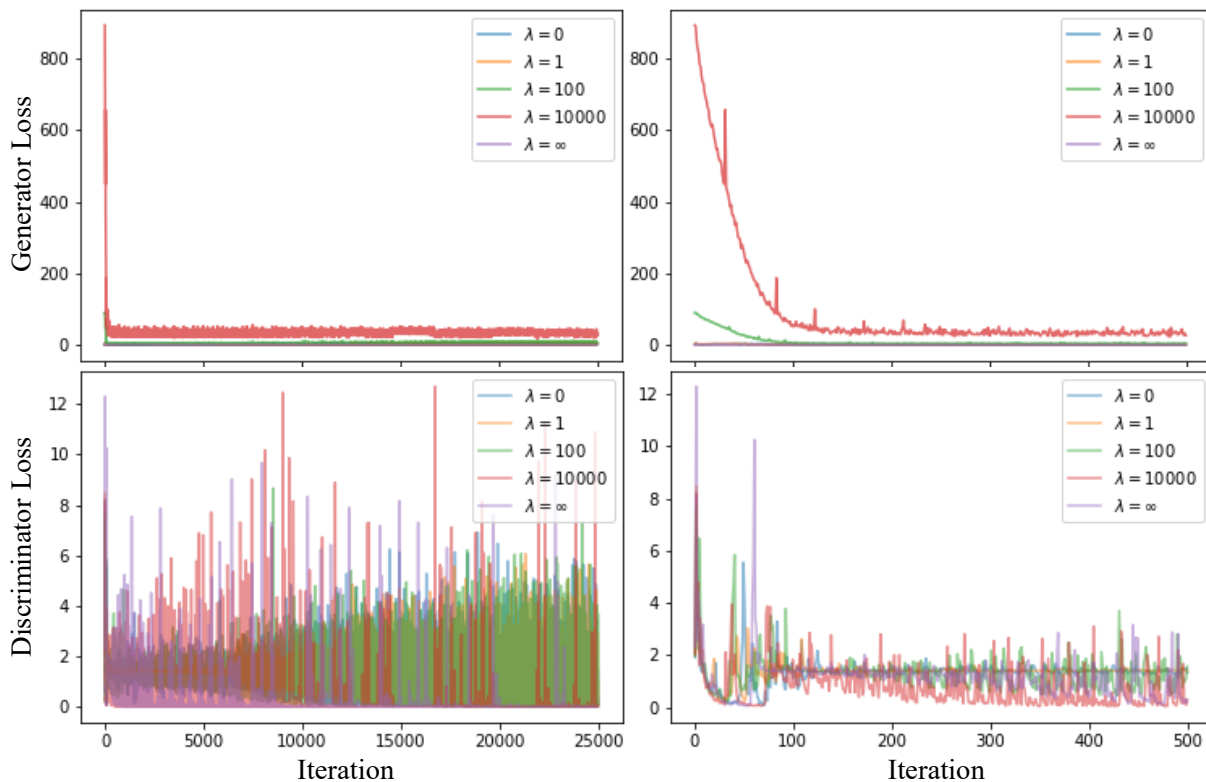


Figure 5.11. Progression of cGAN generator and discriminator losses over 50 epochs (*left*), with a close-up view of the 1st epoch (*right*), for various values of hyperparameter λ (L1 loss weight). Both generator and discriminator loss magnitudes initially decrease rapidly, then subsequently continue to increase. (This phenomenon is present but masked in the generator loss plots due to the exceptionally high initial loss values in networks with large λ values and the resulting scaling.)

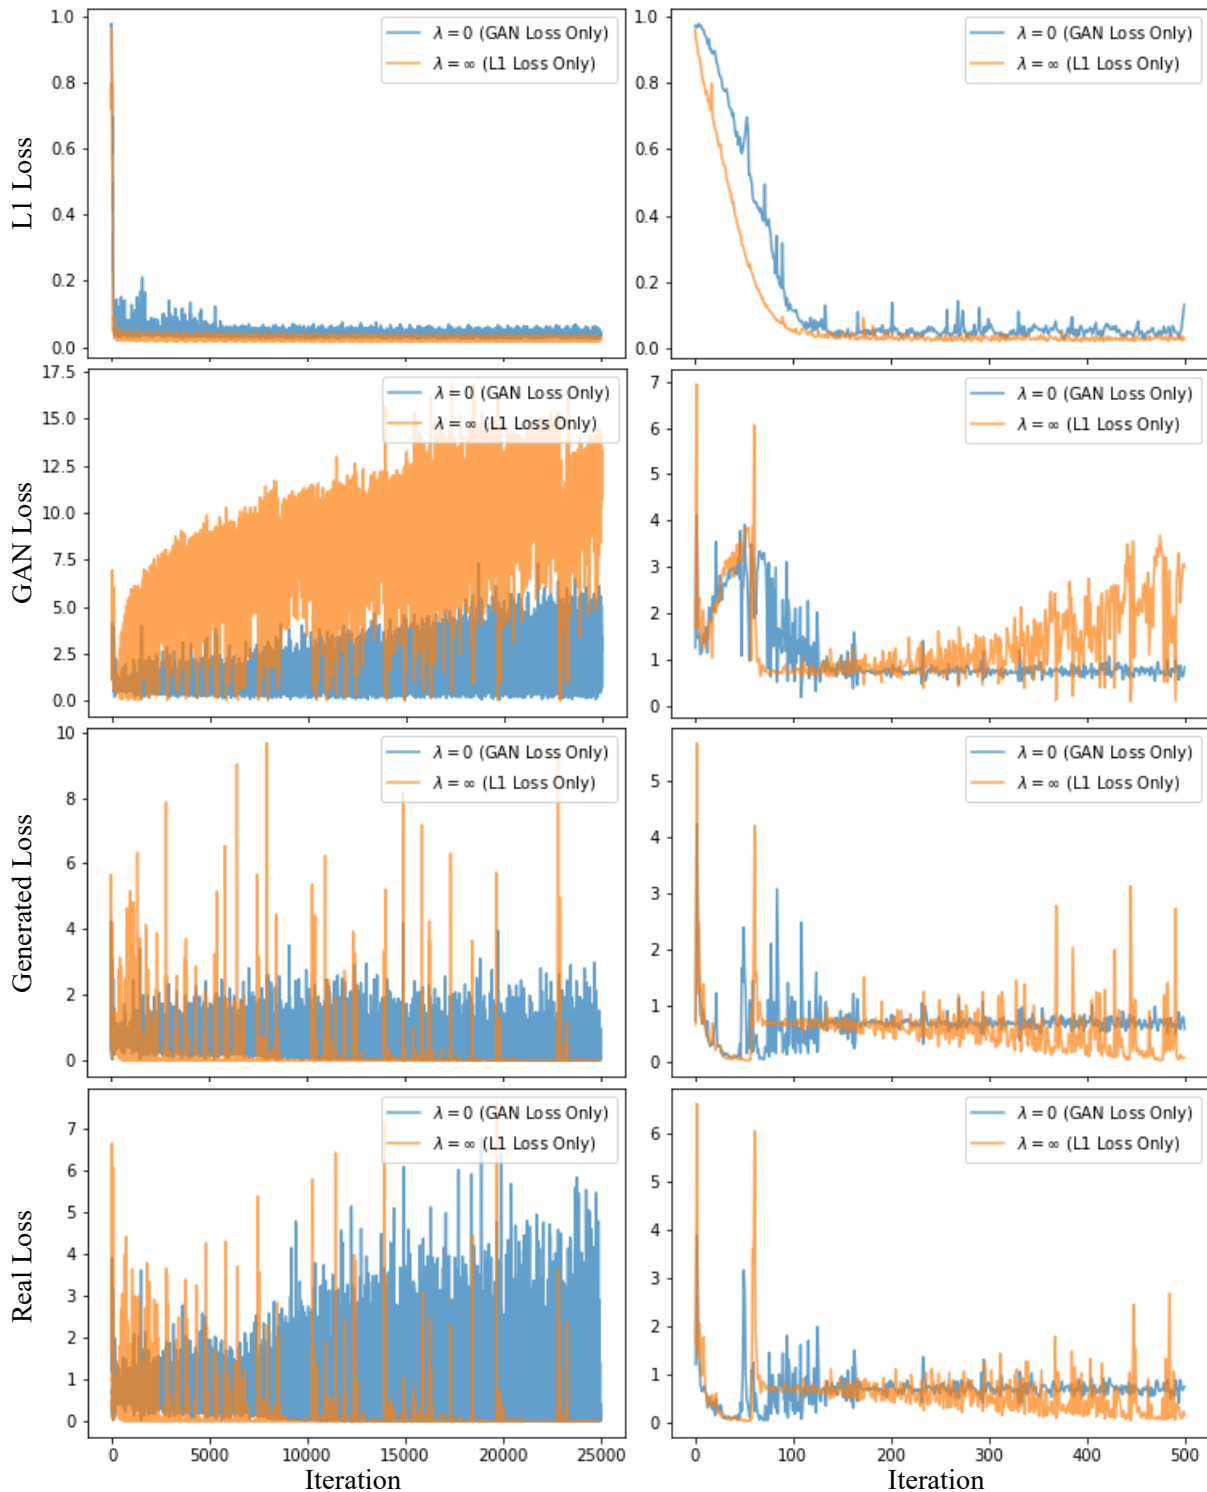


Figure 5.12. Progression of cGAN loss components over 50 epochs (*left*), with a close-up view of the 1st epoch (*right*), for extreme values of hyperparameter λ (i.e. when only GAN or L1 loss, but not both, are considered by the generator). Components of loss for the generator (L1 and GAN loss; Equation 85) and the discriminator (Generated and Real; Equation 84) coevolved. When the generator weight updates considered only GAN loss ($\lambda = 0$; ability to fool the discriminator), L1 loss dropped marginally more slowly but achieves low values while constraining GAN loss. When considering only L1 loss ($\lambda = \infty$; deviation of grayscale values), L1 loss dropped quickly and remained low, but GAN loss (inability to fool the discriminator) increased quickly and continuously.

Images generated from the resulting models were qualitative compared as well (Figure 3.14). As expected, when the L1 loss was heavily weighted (e.g. $\lambda = 10^4$) or used exclusively ($\lambda = \infty$), the image reflected the general intensity pattern but lacked local, high-frequency features such as texture and were smooth and blurry. (This expectation is discussed in Section 5.2.1’s subsection on “Conditional Generative Adversarial Network Model Formulation.”) However, when L1 loss was neglected entirely ($\lambda = 0$, considering only GAN loss), global structure was somewhat less consistent, with some abnormal splotches and non-uniformities present. Overall, the loss weight hyperparameter (λ) was found to significantly affect the training progression and performance of both the generator and the discriminator, and a loss function with a small L1 loss contribution (weight) produced the best performing generator (Figures 5.11, 5.12, and 3.14). The importance of the combination of both loss contributions suggests that a reductionist approach is insufficient. Replicating local structure, as evaluated by the discriminator, without enforcing global coherence and structure was insufficient, just as a focus on average intensity matching without integration of high-frequency features and texture failed to achieve desired outcomes.

Time to execute each of the 50 training epochs was largely consistent for and between each weighted loss function (Figure 5.13). Average times per epoch were 61.3 ± 0.4 , 64.4 ± 0.5 , 62.9 ± 0.6 , 61.3 ± 0.7 , and 57.5 ± 0.5 seconds, for the $\lambda = 0, 1, 10^2, 10^4$, and ∞ cases, respectively. These measured times represent a maximum deviation of 12% between cases, with standard deviations for each of less than 1% of their mean time. Complete training time (over 50 epochs) took 51.3 ± 2.0 minutes.




Figure 5.13. Time to execute each of the 50 training epochs for given weighted loss functions. The initial epoch included several logistical steps beyond the basic training, and iterations 20 and 40 were designated as checkpoints to save progress, so these three epochs exceeded normal training times. Excluding these values, average times per epoch were 61.3 ± 0.4 , 64.4 ± 0.5 , 62.9 ± 0.6 , 61.3 ± 0.7 , and 57.5 ± 0.5 seconds, respectively.

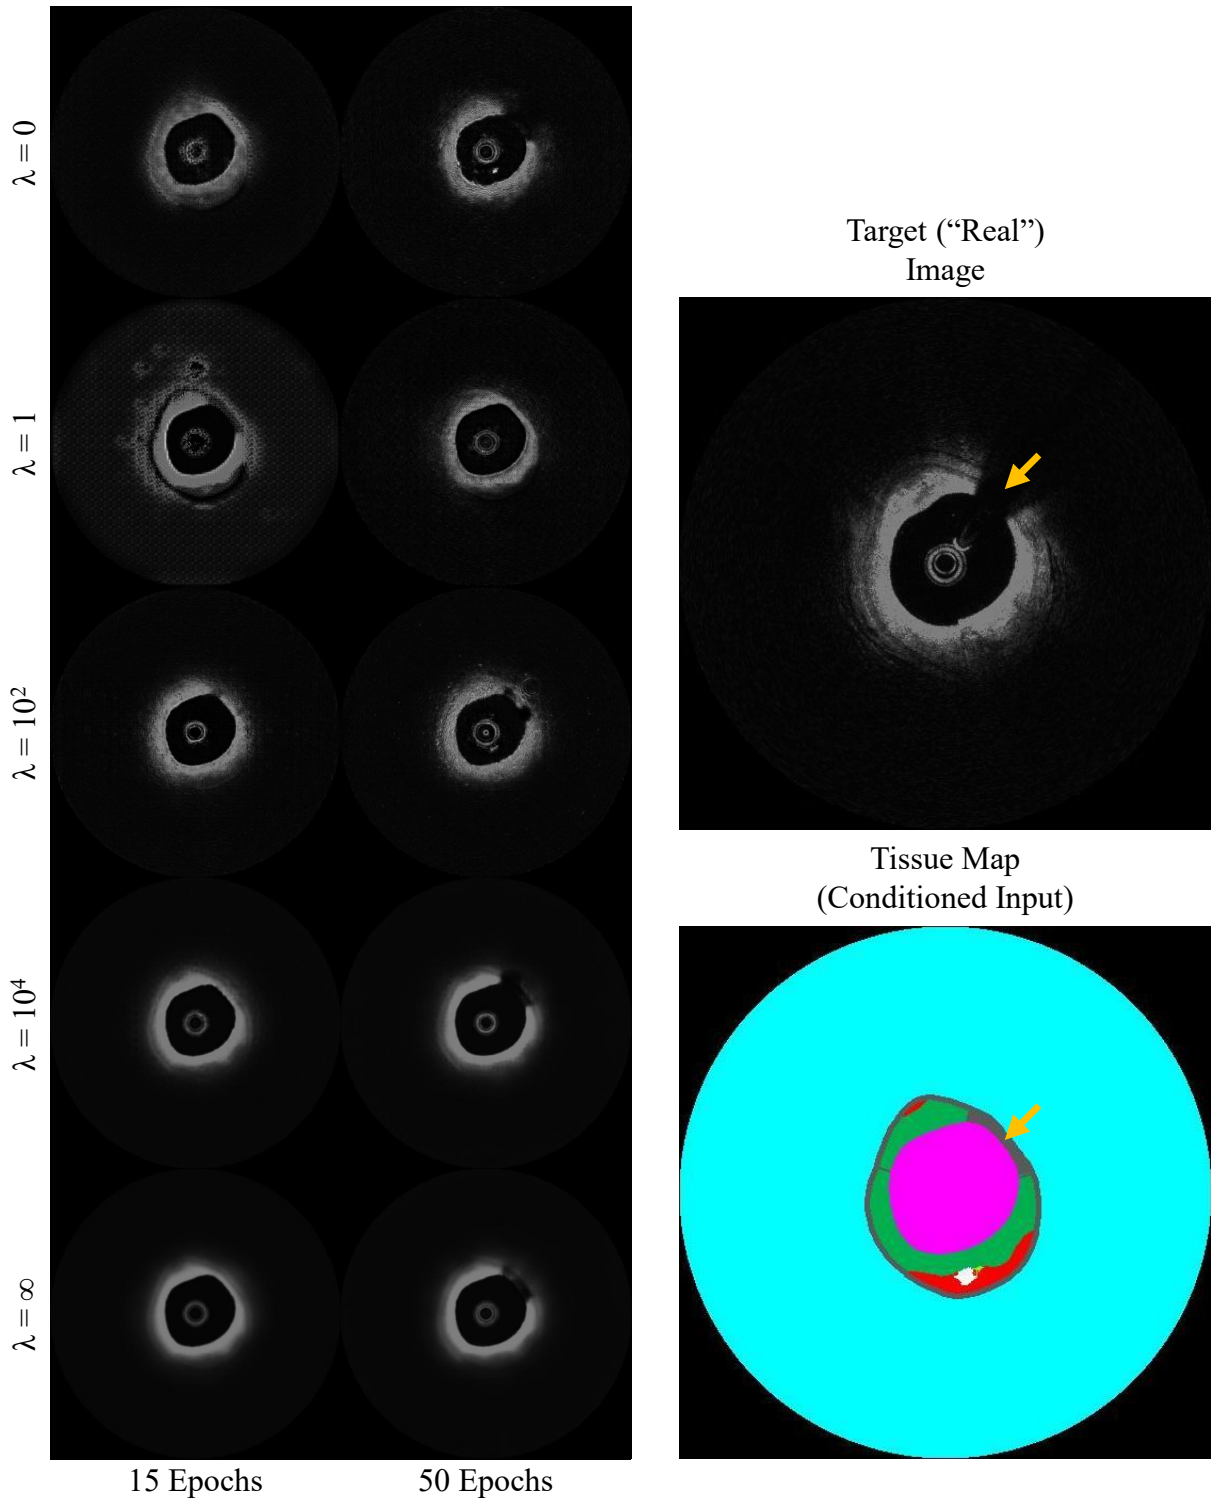


Figure 3.14. OCT images generated with various hyperparameters governing the balance of error contribution (λ) after 15 and 50 epochs of training. When neglecting adversarial loss ($\lambda = \infty$), generated images reflected general intensity pattern but lacked local, high-frequency features such as texture and were blurry. When considering only adversarial loss ($\lambda = 0$), global structure was less consistent, with abnormalities and non-uniformities present. Signs of overfitting emerged after extensive training, as illustrated by the inclusion of the shadow artifact visible in the true image but not encoded in the conditioned input (yellow arrow); such signs had not emerged by the end of the 15th training epoch. Conditioned input shown here as color-coded representation of multidimensional array.

Some signs of overfitting emerged after extended training regardless of which weighted loss function was used. While each OCT-derived class label map should have included exactly one region of no visible tissue (i.e. shadow), the classification algorithm used to generate the images (described in Chapter 3 [160]) was fallible; in the particular true image replicated in Figure 3.14, the shadow occurred in a non-pathological segment of the wall, so no shadow artifact was segmented, and consequently, this class label did not appear in the map used to generate the images shown. Therefore, the dark, shadowed region indicative of this class should not necessarily appear in the generated image, as it did in the real one. However, after several epochs, such a feature did appear in some models. While the model may have learned some more fundamental rules about the location and occurrence of such features independent of its class label, it is probable that this instead constituted an example of overfitting as the model learned to replicate the target image from nearby images which weren't excluded from training. This observation reinforced the need to segregate training and testing data by patient (rather than pooling patient data prior to set division) and to closely monitor validation progression during training to identify signs of degrading performance with increased training. Failing to do so would be expected to lead to poor performance when applied to new datasets. Equivalent to memorizing solutions to an equation rather than learning to solve them, through overfitting the model may learn to recognize specific label maps and output the associate image rather than learning underlying structure-appearance relationships allowing for the construction of entirely new, unfamiliar images.

Building upon lessons learned from comprehensive assessment of image synthesis in OCT, the cGAN method of synthetic image generation was applied to IVUS. Both adversarial loss and mean absolute error (L1 loss) were considered to achieve a balance between general, globally consistent intensity pattern and local, high-frequency features (e.g. texture). As opposed to the CNN characterization output, VH served as a highly-detailed, reliable, and convenient source of conditioned input for cGAN training to synthesize corresponding grayscale IVUS images. Early attempts revealed that the network sometimes struggled to learn the proper placement of the scale bars embedded in the images (which is unsurprising given the spatial invariance and loss of relative spatial information inherent to the convolutional process). Therefore, in addition to the various tissue classes, perivascular tissue, lumen, image frame, and catheter artifact features encoded in OCT image generation (Figure 5.4), the scale bars imprinted in the IVUS images were included as an additional one-hot encoded channel. With these tweaks to the optimal strategy learned for OCT,

IVUS images were generated with a cGAN network that recapitulated the true images with excellent fidelity (Figure 5.15).




Figure 5.15. IVUS images generated with a cGAN network recapitulate the true images with excellent fidelity. Building upon lessons learned from comprehensive assessment of image synthesis in OCT, both adversarial and L1 loss were considered ($\lambda = 1$); images reflected general intensity pattern and demonstrated local, high-frequency features (e.g. texture). Conditioned input derived from the VH-IVUS data (*top right*) is shown as a concatenated ribbon along the top of the image (*top*). In addition to the various tissue classes, perivascular tissue, lumen, image frame, and catheter artifact features encoded in OCT image generation, the scale bars imprinted in the IVUS images were included as an additional one-hot encoded channel. The true (target) grayscale image (*top left*) is very similar to the cGAN output after 15 (*bottom left*) and 50 (*bottom right*) epochs of training.

5.3. Applications in Cardiology

In the previous section, physics-based and AI systems were developed to generate intravascular images based upon (patho)physiological tissue morphology. Models could likely be generated for any desired imaging system or acquisition parameters given sufficient training data. Exciting applications of these models in ancillary systems are envisioned that empower clinicians, who retain full autonomy and authority in the interpretation and decision-making processes.

The above-described systems achieve competences that otherwise either require significant artistic and specialized medical training or simply could not be achieved by a human. They offer capabilities that may assist clinicians in making diagnoses, determining intervention success, evaluating patient state, and more, as well as capabilities that assist medical device developers in optimizing and/or training new computational tools. Because tasks performed by these systems either cannot be completed by a human, or require an immense amount of time and high degree of training, skill, and expertise, there are no viable, currently available alternative options to the above-described systems. Previously, one technical hurdle has been the lack of training data to develop the ML component of the above-described systems; this has only recently been solved by work described in this thesis (Chapters 3). As such, these systems were not feasible, and therefore not developed, previously. However, with further work in the field of tissue classification quickly growing, becoming more accurate and more widely available, opportunities to refine and expand application of synthetic image generation models becomes increasingly feasible. The absence of precedent invites speculation and inspired vision for this budding technology.

One field in which synthetic imaging and virtual medicine offers tremendous unrealized promise is cardiology. Idealized and image-based models of diseased coronary arteries are already used to better understand the pathological mechanisms of lesion progression and vulnerability [103]. ML models have also blossomed in the field of cardiology, though with few applications translated to the clinic [114]. Fully realized concepts of “computational cardiology” [67] remain encumbered by barriers to integration with physical medicine, stranding years of work in modeling and image analysis across a chasm from the clinical bedside. New approaches are therefore sorely needed to bridge the gap; synthetic imaging may serve this role bridging role to merge physical and virtual medicine.

The full implications of such technology are still yet to be considered, but some potential applications are described here. The framework suggested for the bridging of physical and virtual medicine alludes to several possible applications of synthetic medical imaging in (interventional) cardiology. The ability of the methods described in Section 5.2 to generate synthetic medical images of tissue or systems which may or may not have a basis in a physical, existing tissue or systems enables several potential uses in the medical imaging and other medical or healthcare technology fields. While these systems more generally utilize tissue classification maps as intermediaries to achieve various tasks requiring full or partial image synthesis, the applications described below specifically address systems which would use ML and classified plaque maps (as specified in the technical description of Section 5.2) to construct whole or partial intravascular cross-sectional images of blood vessels. In this section, an exploration of several specific applications is presented with some preliminary results and proof-of-concept demonstrations.

5.3.1 Modality Conversion

Synthetic imaging from a plaque map intermediary facilitates the ability to select the modality in which image-encoded information is expressed. Indeed, once the underlying tissue architecture is established, images portraying physical structure can be synthesized for any applicable image modality. As such, images can be translated between modality renderings through this shared commonality, facilitate image sharing for consult or transfer of care when modality standards, preferences, or specialties differ, or in co-registering image series acquired with varying modalities in the course of longitudinal patient care or trials.

The presented method has been applied to convert from VH-IVUS to OCT images via a plaque map intermediary (Figure 5.16). VH-IVUS, as well as some validated ML methods using grayscale IVUS (Chapter 3), can generate tissue/plaque maps of arterial wall cross-sections. These maps may then be used to generate corresponding OCT images of the same wall segment. The same methods have also been applied to generate grayscale IVUS images from tissue/plaque maps obtained from OCT tissue characterization routines (Chapter 3) in the same manner.



Figure 5.16. Synthetic image generation enables “multimodal imaging” through translation between IVUS and OCT. Synthetic OCT images were generated from VH data associated with IVUS frames, which provided plaque label maps that were fed to the trained OCT generator. As such, the synthetic OCT images represent the exact “co-registered” image, allowing for direct comparison of the synthetic OCT image to the real IVUS one it emulates.

This application allows for imaging datasets acquired with different modalities, potentially during different patient visits or at different medical centers, to be directly compared, and also for a clinician to review the conveyed morphological data in the visualization and representation of choice. As such, this system might be useful for integration into (electronic) medical record systems by allowing for “conversion” of arterial imaging data among various representations/modalities. It may also aide in clinical study and research data management, as it could loosen limitations regarding equipment and expertise for clinical center inclusion (e.g. if a center only has IVUS capabilities, but a study is seeking and utilizing OCT data to monitor clinical endpoints).

Modality conversion may assist in the co-registration of image sequences of inconsistent modalities. Co-registration is the identification (and alignment) of corresponding frames and relative orientation between multiple image acquisition sequences. Given two image sequences of the same vessel segment (or some common section thereof) acquired with different modalities, one image sequence could be converted from one to the other such that both are visualized as if they were acquired by the same modality. The two sequences could then be more readily be co-registered, providing a one-to-one matching of the frames of one series to those in the another.

However, the system need not be constrained to replicating standard modalities exactly; by training the AI to replicate contrast-enhanced images or providing the system with additional information to integrate, fundamental limitations of a modality can be transcended.

5.3.2 Enhancing Images

Synthetic image generation can address shortcomings in medical imaging by ameliorating challenges in image quality, interpretability, coherence, completeness, and granularity (Figure 5.17). This can be achieved in part either through the fusion of imaging data or the employment of computational sub-routines.

One exciting prospect is the ability to select not only the modality in which image-encoded information is expressed, but also which enhancements to apply to standard modality rendering. Indeed, the goal need not be exact replication, replete with flaws and shortcomings. For example, by training the AI network to replicate images enhanced to aid in interpretation and use of the images [65], inherent limitations of a modality can be addressed without subsequent post-processing. Similarly, artifacts can be excluded and obscured regions manifested.

Image generation technology offers advanced capabilities and improved imaging, for example by completing an image with an obscured region (Figure 5.18). Such areas include the shadowed region present in intravascular OCT images due to the presence of a guidewire. In the sample implementation shown (Figure 5.18), expected, likely, or nominal classifications of obscured pixels could be determined from supplemental or contextual information (e.g. through region growing and/or iterative mode filtering). A complete synthetic image from the resulting classification map was generated. This image could stand alone, or obscured pixels of the original, incomplete image could be replaced with the corresponding pixels of the generated image. (Additional ML techniques for completing partial images, not requiring tissue classification map precursors, may be applicable here as well, though the use of tissue maps to condition generation

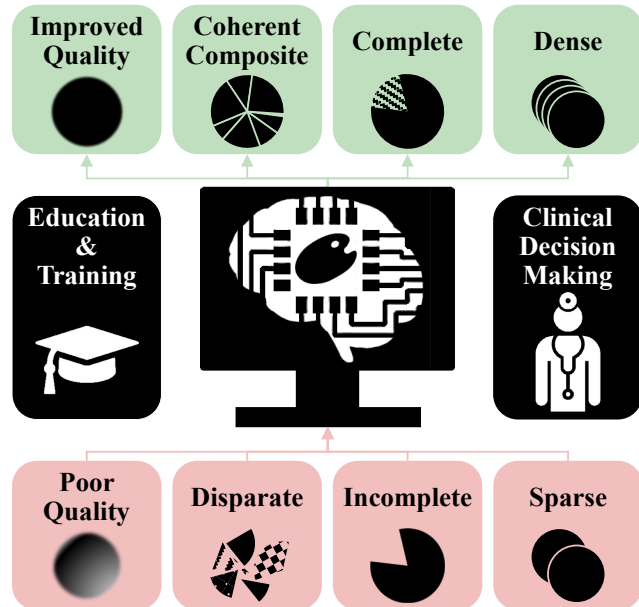


Figure 5.17. Synthetic image generation addresses shortcomings in medical imaging. Artificial intelligence can ameliorate challenges in image quality, interpretability, coherence, completeness, and granularity, offering promising applications in clinical decision-making and education.

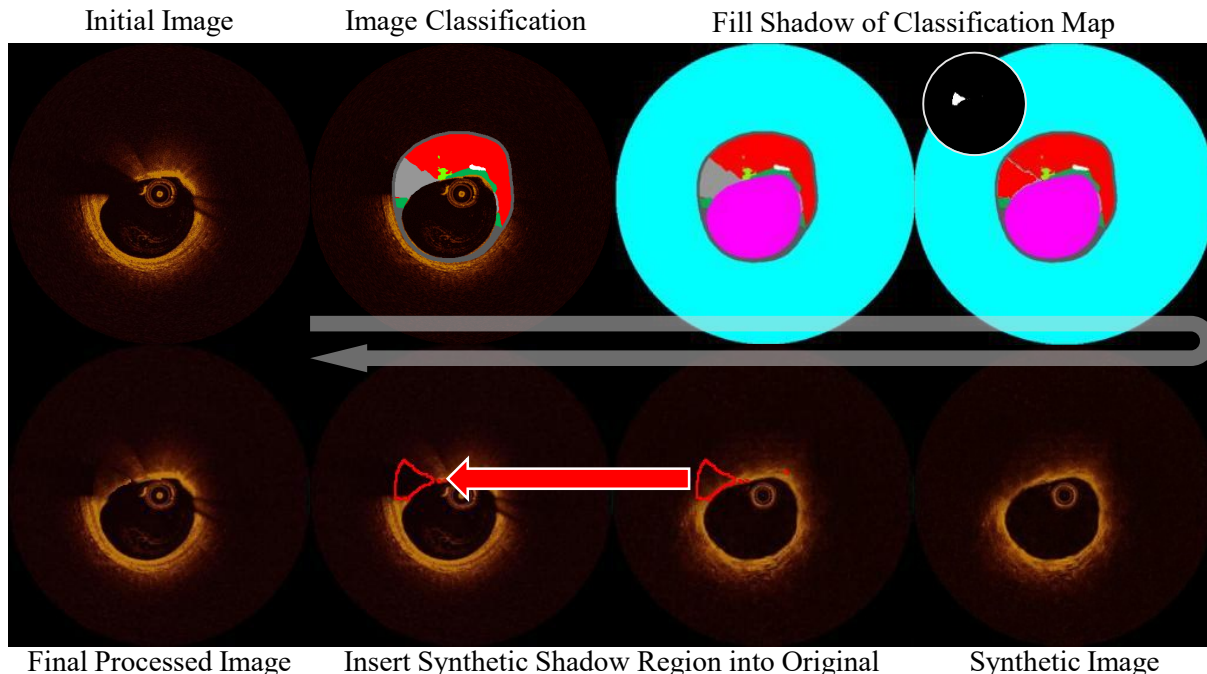


Figure 5.18. One possible approach for an application to improve incomplete or partially obscured images. An image is classified (e.g. by a CNN), from which a classification map and shadow mask is prepared through image processing. The shadowed region of the corresponding classification map is filled (e.g. through region-growing), and a synthetic image is generated from the resulting conditioned input. The synthetic image may be used independently, or those pixels corresponding to the shadow mask may be transferred to the original image.

of the missing segment enables the imposition of continuity and integration of *a priori* and domain knowledge.)

In addition to addressing regions obscured by imaging artifacts, ambiguities resulting from limitations of the modality itself can be clarified using *a priori* and contextual information. The platform can generate synthetic intravascular images intuitively visualizing augmented morphological information for clinician interpretation. For example, information on the vessel geometry and morphology extracted from the initial OCT image, such as those described in Chapters 2 and 3, can be reported to the AI system. One illustrative set of images is shown in Figure 5.19. An initial series of acquired images was automatically processed and analyzed to ascertain information on the vessel morphology, and this enriched set of information was in turn used to generate enhanced images. In this case, automated image processing algorithms developed in this thesis characterized the frame and leveraged contextual information from adjacent frames and physical continuity constraints to enrich the underlying anatomy beyond that immediately visualized by the lone OCT frame. Consequently, while OCT has limited penetration depth in the presence of fatty plaques and regions obscured by guidewire shadows, the synthetic OCT images

could be generated a fully-visible wall region and no such obfuscating artifact. Whereas limited penetration depth and image artifacts of the initial images do not allow for complete assessment of plaque burden or composition, remodeling, or vessel size—information critical for guiding treatment decisions and planning (e.g. pharmacotherapy prescription or lesion preparation and stent sizing for intervention)—this information is presented for clinician assessment in the subtly enhanced, synthetic set of images. Synthetic OCT imaging could therefore potentially rival IVUS imaging in far field visualization while retaining advantages of OCT imaging in the near field, overcoming the main limitation of this modality, with implications for both diagnostic and clinical decision-making purposes.

In addition to leveraging contextual and domain knowledge through computational sub-routines enhancing the information extracted from within a single image acquisition sequence, the platform can increase visible information and decrease ambiguity by integrating data from multiple sources. The platform transforms a digital representation of tissue anatomy (specified underlying physical structure of diseased vascular morphology), typically derived from stacked clinical images, into corresponding synthetic images for clinician interpretation. However, one benefit of




Figure 5.19. Synthetic image generation addresses shortcomings in medical imaging. Images are synthetically generated based on anatomy, determined using any combination of data sources and modalities. Relative to the unprocessed acquired OCT image (here the sole source of underlying anatomy), the generated synthetic image shows the full vessel wall with enhanced feature discrimination, with no arc concealed by guidewire shadow (*) or far-field feature (e.g. vessel border) obscured due to limited penetration depth and signal dropout (†). Note that the corresponding IVUS frame was considered in verifying the geometry and morphology extracted from the OCT frame, but was not actually used to generate the “output” image in any way.

this transformation is that knowledge of the tissue architecture underlying the images need not be derived from a single source, but can be enhanced by integrating geometric and morphological information from various imaging acquisitions, modalities, processing and analysis algorithms, or patient test results, thereby unifying all relevant available information into a single coherent representation (Figure 5.19). For example, the labeled map of atherosclerotic plaques and adjacent tissue can be derived from manual annotation, virtual histology, automated image characterization, or any combination thereof.

A specific form of this integration is the fusion of multimodal imaging. Due to the tradeoffs typical of imaging modalities, hybrid and multimodal technologies (combining modalities) are being explored to take advantage of distinct strengths and negate shortcomings [21], [26], [32]. However, preliminary efforts to overlay or fuse the various acquired images are typically unintuitive and difficult to interpret [202], [210], [211], [214], [232]. The framework developed here suggests an entirely alternative approach whereby structural information—the key insights offered by imaging—are first extracted from each available modality, then combined into a single digital representation of physical tissue structure from which a single image (or whichever desired modality) is generated, negating the need to contrive an entirely new visual representation that otherwise results from fusion of images from different modalities.

Finally, acquired images can be recreated with decreased ambiguity and noise. Because the generation method is dependent only on provided anatomical (conditioned) input, the system is unlikely to learn and encode acquisition artifacts and irregularities, and as such does not inject them into the synthetically generated images. Figure 5.19 shows how synthetically generated images can enhance views of diseased vessel geometry and morphology from modalities which suffer from shadowing artifact and undesirable signal attenuation. The resulting image shows the entire vessel wall in familiar and interpretable format without gaps, greater discrimination of lesions, and fewer artifacts. Note that the guidewire shadow typically present in OCT images was electively excluded from the synthetic image (Figure 5.19), and artifacts arising from insufficient blood flush were not present. Whereas guidewire shadow obscures a significant arc of the original acquired image and highly-attenuating lesion obfuscates abluminal tissue, corresponding segments of the generated synthetic image clearly illustrate vessel wall structures.

5.3.3 Abnormality Identification

Visible abnormality identification could assist clinicians making diagnoses and providing clinical care. In this application, a system could generate a synthetic image based upon the expected or classified tissue/plaque map, and this image would be compared to an actual image of the same physical segment. Direct comparison between the synthetic image (generated based upon standard presentation of the patient's given morphology) and the actual/real image (reflecting the actual presentation of the patient's morphology), could highlight unexpected findings and potential abnormalities warranting attention and further review.

The discriminator, trained in tandem with the generator as part of the GAN to distinguish real or fake images based upon the provided plaque map, may also offer utility. The discriminator may aid in distinguishing between imaging artifact and physiological signal by flagging artifact-affected image segments as "fake." It may also assist in detecting possible instances of image tampering, a concern briefly addressed in Section 5.4.

5.3.4 Model Training and Method Evaluation

Image data can be generated for model training or method evaluation in research and development. Similar to the following application, images can be generated with a plurality of various plaques in known locations within a displayed arterial cross-section. Such a system allows for a practically infinite pool of inherently-labeled images; the method offers an unlimited source of sample images of nearly any conceivable physical scenario. The images and labels may thereafter be provided as input to a computational or numerical model (such as a ML model) to train, test, validate, and/or otherwise optimize the model or image processing, quantification, and analysis method, technique, or tool. Such derived image data could contribute to an emerging paradigm of computer-simulated clinical trials for regulatory evaluation in which image analysis systems are trained or validated on synthetic images [341]; the U.S. Food and Drug Administration is already exploring the use of synthetic images generated from physics-based models in computer-simulated (virtual) clinical trials for regulatory evaluation [341].

In addition to employing the generator to synthesize images on which to test new methods, the trained discriminator could also be employed in quantifying confidence in classification or characterization of images. For example, the characterization map resulting from a new characterization method could be supplied as conditional input to the discriminator along with the

real image. The output of the discriminator would reflect the accuracy of the characterization method, since the actual image should always be discriminated as “real” given an accurate corresponding characterization map (conditional input). Alternatively, the system could generate a synthetic image based upon the tissue/plaque maps determined by the new characterization process. The synthetic image and actual image—upon which the characterization process was performed—could then be compared and quantified (either through direct comparison of images or by the discriminator).

5.3.5 Education and Training

Before reaching the clinic, initial application may well be in education and training. Just as synthetic images can be used to train and test computational algorithms, AI-powered medical image generation offers a unique tool to facilitate human education and training, particularly for medical professionals. Synthesizing medical images from morphological templates offers a limitless source of physiology-based images without the challenges of appropriate patient data acquisition and management, including ethical concerns, acquisition risk, and procedural cost. Nearly any conceivable physical scenario—even rare phenotypes, contraindicative conditions, or mild disease states which would not justify invasive imaging—could be rendered for illustrative educational purposes. Moreover, inherent label maps provide keys for identification tasks. Learners can therefore be presented with images of one or more relevant modalities depicting any underlying pathology to reinforce structure-appearance and inter-modality relationships. In the course of, for example, medical student training, professional certification or licensing exams, or continuing education courses, images could be generated and presented to a trainee, who would then be tasked with identifying the type and location of plaques, or performing diagnostic tasks.

5.3.6 Representing Virtual Models through Clinical Images

The ability to synthetically generate clinical images presents rich opportunities to integrate patient-specific virtual phantoms into the image-based workflows and decision-making processes of personalized medicine [67]. Computational models—virtual and *in silico* representations of physical systems—not only enable generation of different projections and views, but, through simulation and modeling, visualization of the system at different time points and conditions. This application is central to the paradigm proposed herein to surmount the schism between physical and virtual medicine (Figure 5.20). The suggested process is as follows: clinicians collect data,

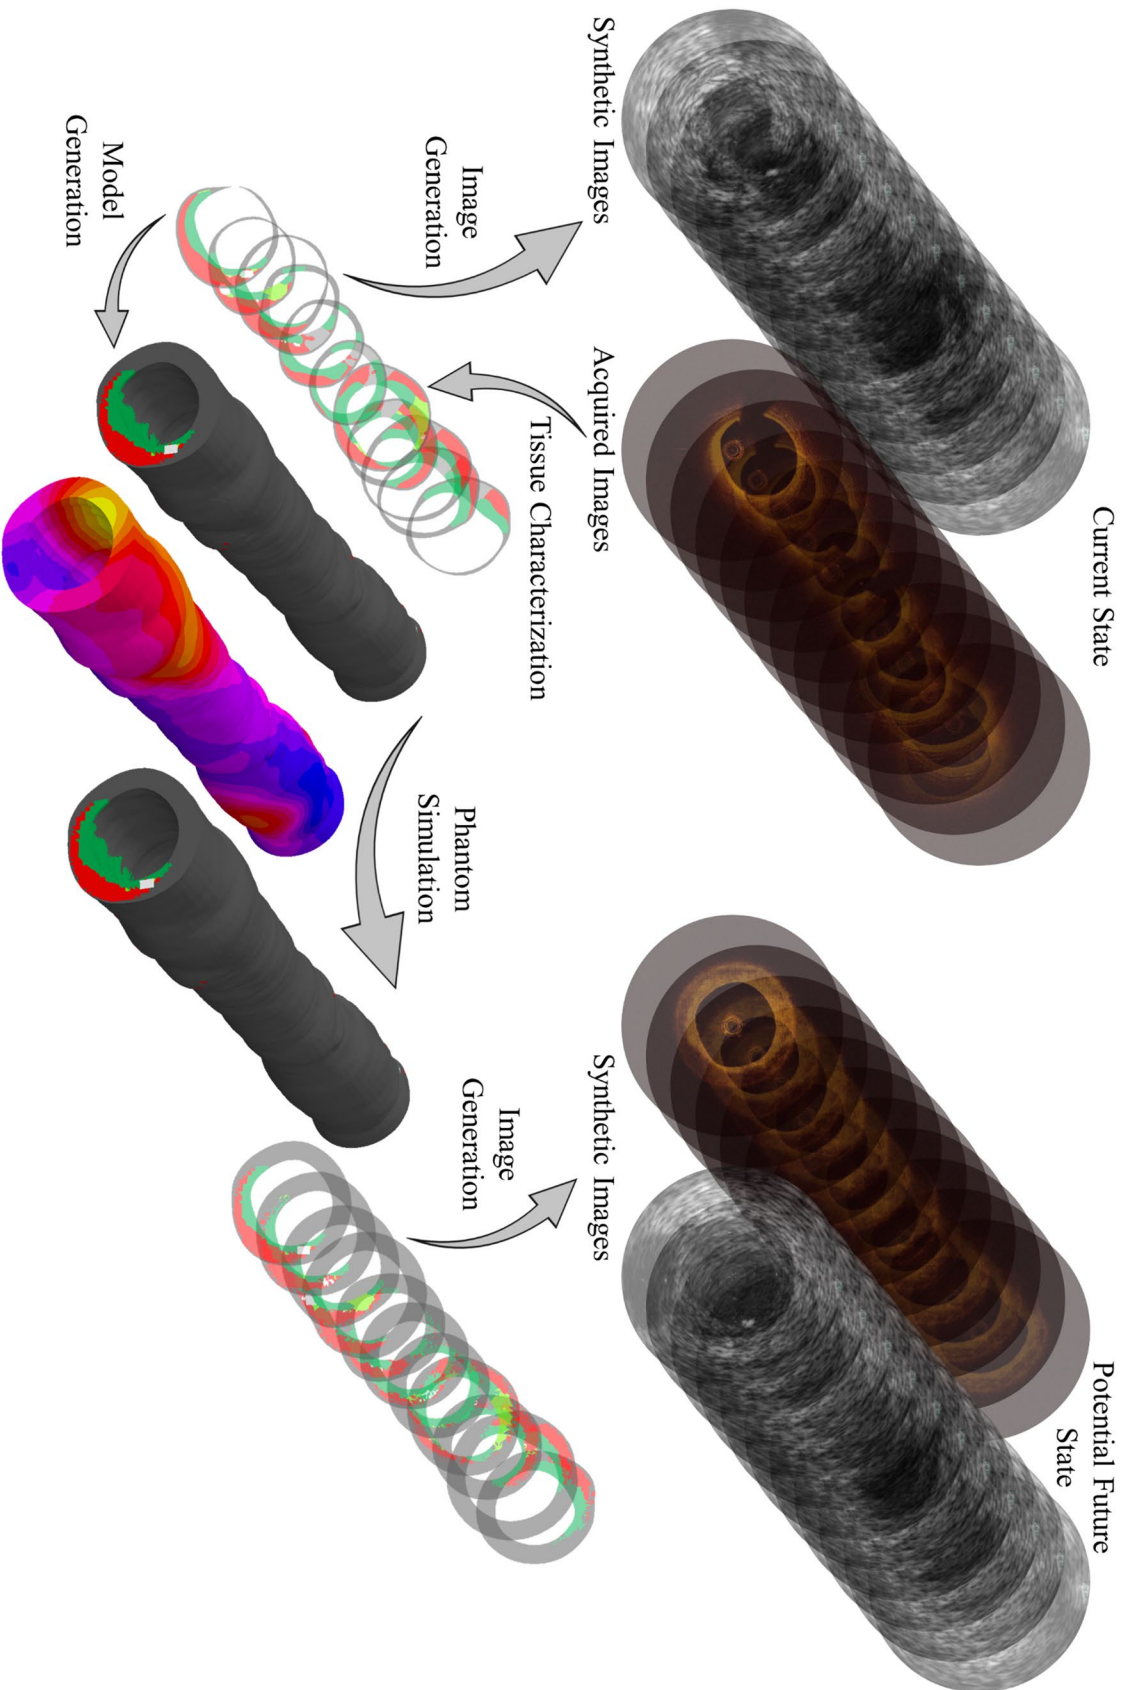


Figure 5.20. To surmount the schism between physical and virtual medicine, clinicians collect data, including images or other anatomical and morphological data, from a patient. These data are used to generate patient-specific computational models. Various simulated tests and procedures are performed on the virtual patient. Results of the simulations are transformed back into clinical data, including images, to enable seamless integration and assessment by the care team.

including images or other anatomical and morphological data, from a patient; these data are used to generate patient-specific computational models; various simulated tests and procedures are performed on the virtual patient; and finally, results of the simulations are transformed back into clinical data, including images, to enable seamless integration and assessment by the care team. Outcomes could inform decision-making and guide procedure execution for the patient, or could be used to validate and assess outcomes of tests and procedures that are undertaken.

This framework for clinician-driven personalized medicine is enabled, rather than governed, by ML. The role of ML, specifically deep learning neural networks, here is twofold (Figure 5.20): classification to segment images (to generate detailed patient-specific models) and tissue map-to-image translation to generate clinically-relevant images from model-derived (or image-derived) morphological maps. By introducing novel application of these approaches, formerly-disparate functions spanning several fields can be unified and integrated.

The proposed framework has been applied to a cardiology case of coronary artery disease to demonstrate the broader paradigm in which physical and virtual medicine form a fluid continuum (Figure 5.20). Intravascular OCT data was acquired from a patient and characterized to segment geometric and morphological features following the process introduced in Chapters 2 and 3. This information was subsequently transformed to an alternative modality (IVUS) image series, and also formed the basis of a computational finite element model (with 5 neo-Hookean material classes) following the procedure introduced in Chapter 4. A simple simulation of lumen pressurization was performed in the heterogeneous vessel model. Images of the final simulated state were then generated in multiple modalities for review by a clinician, along with standard clinical metrics extracted from the simulation results.

More inventive applications of generating synthetic images in reflecting expected appearance of underlying morphology may serve to identify in real time the outcome of, or analyze *post hoc* system conditions during, intervention or testing by comparing acquired images to those generated from the outcomes of various simulations. For example, a characterized image would be used to construct a model system, the model system could be perturbed (e.g. simulated deployment of interventional medical device), and synthetic images for that modified/perturbed system would be generated. Several iterations with simulated interventions could be carried out with various parameters (e.g. device placement, settings, tissue characteristics, etc.). Resulting

images offered to clinicians would provide a glimpse ahead of time into what a successful—or unsuccessful—intervention for a given patient would look like, allowing for easy evaluation once the real intervention is undertaken and imaging acquired. Determining possible image manifestations of clinical outcomes could allow rapid and direct verification of outcome state in the clinical setting, enabling faster assessment and decision-making in the operating theater.

5.4. Implications & Future Directions

Building upon and unifying my progress in determining vascular geometry from intravascular imaging, extracting vascular morphology from intravascular imaging extends, and constructing lesion-specific models of diseased arteries, my work in generating synthetic intravascular images offers some immediate, and many speculative and forward-looking, implications.

Some interesting insights are derived from the adequacy of the physics-based and ML image synthesis approaches. The simplified physics-based approach yielded reasonable though generally unconvincing images that were easily distinguished from real OCT images, particularly without added noise. This observation highlights the prevalence of stochastic distributions and discrete scatters in living systems and their prominence in dictating the results of imaging processes—by extension the limitations of bulk or continuum assumptions in physical descriptions of imaging mechanisms. A simple regression model proved inadequate for generating realistic images, as it couldn't account for complex interactions beyond the features explicitly encoded. Yet, the model was provided the same basic descriptive information and metrics as that made available to the physics-based method (namely depth, proximal/adluminal tissue composition, and tissue class of the present pixel), but with worse outcomes. The primary reason this ML approach likely failed was its inability to generate exponential changes in signal with depth on account of its restriction to linear combinations of variables. This observation highlights the highly non-linear nature of OCT imaging and the insufficiency of linear approximations or linearity assumptions in utilizing this modality. The results of the physics-based and regression models further suggests that the uncompressed format of the images (rather than the compressed or clipped 0-to-255 format amenable to 8-bit base 2 binary encoding) may yield richer information for analysis and should be utilized when possible. cGANs present a promising approach to generate intravascular images and don't require a deep understanding of the complex imaging physics. However, adversarial training,

rather than minimizing grayscale deviation alone, was critical, further highlighting the role of inhomogeneity in generating image texture. The networks were, further, very sensitive to hyperparameters and data representation strategies. As for many deep learning approach (like those described in Chapter 3), this finding once again underscores the importance of deliberate encoding strategies informed by domain knowledge, the importance of comprehensive hyperparameter evaluation, and the need for design decisions to be guided by end application-driven goals; the intended use of these images should always remain a foremost consideration.

Further implications lie in the array of potential applications, particularly in cardiology. While image generation has been implemented in numerous fields and with various applications, the training data necessary to implement this method in OCT has only recently become available through the advancements presented in this thesis. The uniqueness of the body of work in straddling imaging, modeling, and mechanics also yields different opportunities and perspectives than previous advancements in the image generation academic community—particularly its application to interactive patient phantoms and the associated bridging of physical and virtual medicine. Image generation could soon enable researchers in the cardiovascular space to represent hypothetical patient states or conditions, including those resulting from virtual or simulated interventions and procedures.

Application of AI to synthesize image data remains an underexplored paradigm for AI in medicine, mandating transparency in noting current uncertainties and limitations. Further work is needed to precisely determine the extent and context of resolution loss or gain, as well as to quantify reliability and robustness. Performance must be further assessed and validated in any proposed applications before use in clinical workflows informing diagnosis or treatment. Systems will likely fail in extenuating scenarios which violate standard morphology-appearance relationships present in training datasets; devices or pathologies not represented in the dataset would not be reliably rendered.

As for any new technology, particularly those impacting human health, it is important to consider, insofar as possible, unintended consequences or malicious abuses of the technology. Concerning the former, concerns have been raised regarding the application and potential misuse of ML-based decision support systems [136]. In particular, there is risk that such systems may not be subjected to the same standards and scrutiny as usual changes in practice, and could be adopted

without robust evidence of clinical benefit and improved outcomes [136]. As such, efforts should be made to ensure that the presented framework and its constituent components are well validated and positively impact patient care in demonstrable and relevant ways. Regarding the latter issue of malicious abuse, clinicians should be aware that this paradigm can be employed by malicious actors to manipulate and tamper with patient data with dire consequence. The ability to tamper with medical records through application of deep learning image generation techniques, and consequently manipulate clinical diagnosis and treatment planning, has been demonstrated [154]. In a disconcerting exhibition, researchers attacked a hospital's image archive system (in a controlled manner not impacting actual patients) and modified images to insert or remove cancerous nodules, deceiving multiple radiologists (and an AI screening system) who missed or issued false diagnoses [154]. As this work demonstrated, clinics must continue to take steps to ensure the security of their networks and systems, particularly as the role of virtual medicine becomes more integral to, and integrated with, that of physical medicine.

The full clinical, scientific, academic, and commercial applications remain unclear. However, the potential for such technology could be profound, and it will be important for various stakeholders—technologists and clinicians foremost among them—to shepherd in its development and use responsibly to achieve its full benefit.

5.5. Conclusion

In this chapter, I have addressed a new way to bridge the chasm between virtual and physical medicine with synthetic imaging. Several uniquely-applied methods of synthetic image generation for intravascular OCT and IVUS were described, assessed, and compared, and several proposed clinical applications of such technology in cardiology were presented.

Synthetic images were generated by models trained in the familiar visual vernacular of vascular imaging. More than just a new way of presenting old processing methods for ease of interpretation, the work advances the art of image processing. The developed systems generated subtly enhanced intravascular images based upon (patho)physiological tissue morphology, which could be extracted from existing images, computational models, or entirely contrived. Comparable systems could likely be generated for any desired imaging system or acquisition parameters given sufficient training data. Exciting—but as yet unproven—applications in ancillary systems are

envisioned that empower cardiologists, who retain full autonomy and authority in the interpretation and decision-making processes.

In a time when seeing is no longer believing, the visually-driven field of cardiology is amenable and susceptible to AI-driven synthetic media generation technologies. The full implications and ramifications of such technology should be considered by the field; to constructively utilize and benefit from these advancements—and realize the tremendous opportunities—technical and medical stakeholders must work together to determine its useful and appropriate application.

Chapter 6

Conclusion: Biomechanical Assessment & Modeling of Diseased Coronary Arteries

Despite great progress in the past decades [2], cardiovascular disease, and coronary artery disease in particular, remains a persistent devastating and prevalent menace to health and wellbeing globally [1]. While biomechanical forces are known to play a driving role in the natural history of atherosclerosis [9], the nuanced yet profound impact of patient- and lesion-specific biomechanics in disease presentation, course, and treatment are not fully understood, appreciated, or accounted for in clinical care. These shortcomings result, in part, from a lack of quantitative approaches for assessing and modeling diseased arteries. Development of such methods have been inhibited by fundamental gaps in knowledge at the intersection of imaging, pathological anatomy, and biomechanical structural behavior. Through the application of image processing, artificial intelligence (AI), computational modeling, and numerical methods, new windows have been opened into the biomechanical state of atherosclerotic arteries, providing expanded awareness of the mechanical context of the disease. Through the insights described in this thesis, greater information can be extracted, augmented, and made accessible from data currently collected in the clinical setting. This information, currently unavailable to clinicians and researchers, could be used to monitor mechanical status of a lesion over time or in response to treatment, more quantitatively and reliably classify disease states, optimize intervention strategies, and assess tradeoffs between alternative treatment approaches. As such, troves of existing data could be leveraged, through use of tools and methods developed in the course of this thesis, to identify heretofore unrecognized trends and associations that may offer mechanistic insight into disease development, progression, and treatment response, ultimately leading to improved personalized patient care in an emerging era of computational cardiology [67].

AIArtificial Intelligence
FEFinite Element
IVUS ...Intravascular Ultrasound
OCTOptical Coherence Tomography
2DTwo-Dimensional
3DThree-Dimensional

6.1. Summary of Work & Contributions

The contributions of this thesis have focused on several complementary aspects of cardiovascular biomechanics, including the determination, quantification, and assessment of three-dimensional (3D) geometric and spatial parameters, structural morphology, constitutive material properties, and visualization, among others. Further work in collaboration with clinical partners, led by Ben-Assa *et al.*—beyond the immediate scope of this thesis—further elucidated the importance of cardiovascular biomechanics and the role of vascular stiffness and impedance on outcomes and impact of transcatheter aortic valve replacement in treating aortic stenosis [350], [351]. The work presented in this thesis, however, has considered and leveraged knowledge of diseased coronary arteries, with particular attention given to intravascular imaging.

6.1.1 Introduction to Biomechanical Assessment & Modeling of Diseased Arteries

Chapter 1 provided a comprehensive overview of the unique convergence of cardiovascular pathology, clinical practice, medical imaging, computational modeling, and artificial intelligence, at the intersection of which the work of this thesis resides.

6.1.2 Extracting Vessel Geometry

Chapter 2 presented a breakthrough method of extracting vessel geometry which overcomes current limitations of intravascular imaging [155]–[158], as well as a method and illuminating analysis of disparate geometric dataset alignment [159]. Human-like performance in annotating the outer border of diseased vessels with limited visibility in optical coherence tomography (OCT) was achieved only when consideration evolved from two-dimensional (2D) to 3D context, and when continuity in 3D space was leveraged. Success of an equilibrated mechanical lumped-parameter model—an anisotropic linear-elastic spring mesh—demonstrated that anatomic geometry—and consequently plaque distribution—adheres to a consistent and predictable pattern which limits variation. Applications of the novel method of smooth surface fitting extend beyond image processing of intravascular images, and offers unique, flexible abilities in handling variably-sparse and questionably-reliable data in any 3D context. Furthermore, a simple yet effective method to align vascular imaging datasets on the basis of luminal geometry was described and validated. Success in aligning vessel segments longitudinally without major features, such as bifurcations, and acquisitions acquired before and after lumen-altering interventions demonstrated that naturally occurring, subtle variations in lumen area represent unique, stable, and robust

features which can be matched across acquisitions and modalities. However, challenges in rotationally aligning acquisitions without major features (e.g. bifurcations) demonstrated that lumen eccentricity is not similarly unique and stable, so not a robust and reliable signal.

6.1.3 Determining and Characterizing Vessel Morphology

Chapter 3 described the development and critical assessment of a unified framework to automatically determine and characterize vessel wall morphology and indwelling devices in OCT [160] and intravascular ultrasound (IVUS) [161] images. Studies highlighted the importance of enriching such methods with domain knowledge of vascular geometry and imaging [161] and carefully considering the curation of training data [162]. Characterization of vascular tissue and indwelling devices was achieved in OCT and IVUS through the application of AI in specified regions of an image. Defining a region of interest on the basis of vessel geometry and knowledge of plaque distribution derived from past *ex vivo* studies proved critical in constraining characterization outputs to physiologically valid solutions and achieving accurate results. Plaque classes which shared compositional features were most commonly confused by automated characterization methods, just as they are by expert readers. However, neural networks still performed remarkably well across applications, even in cases with substantial variability. This was particularly true when contextual information on population prevalence was incorporated by commensurately weighting errors or balancing datasets during training of the system. Still, the work demonstrated that a single universal or general model is rarely the best choice. Rather, application-specific models optimized for specific use cases are typically preferable.

6.1.4 Computational Modeling and Simulation of Patient-Specific Vessels

Chapter 4 described use of clinical OCT data to construct highly detailed patient-specific finite element (FE) models of an *in situ* diseased artery, which enabled exciting new research and potential future clinical applications. One key application was demonstrated in the chapter through a novel approach to estimate lesion-specific biomechanical tissue characteristics, which was validated *in silico* and *in vitro*, highlighting potential for future *in vivo* application. Two methods of FE model construction were developed, and as for neural network models, the FE models demonstrated tradeoffs and preferred use cases. Whereas coarse meshes allow for morphological detail to be more readily retained and yielded useable displacement estimates, sharp edges compromised stress distribution estimates in some regions. Refined meshes, with plaque feature

geometries constraining mesh generation, resulting in smooth contours and stress distributions not contaminated by non-physiological distortions. However, retaining high levels of detail was more difficult. Inverse FE analysis studies with the coarse models allowed for the recovery of constitutive mechanical properties, as demonstrated for heterogeneous vessels *in silico* and homogeneous vessels phantoms *in vitro*. By imposing some degree of posterizing on the basis of characterized plaque morphology, and leveraging displacement among tissue feature interfaces, the approach may offer greater amenability to clinical use scenarios than alternatives which instead reside at the extremes of simplification—full displacement field measurement with no simplifying binning assumptions, or extensive displacement simplification considering only maximum and minimum diameter and therefore requiring multiple acquisitions at constant pressures.

6.1.5 Bridging the Chasm between Virtual and Physical Medicine with Synthetic Imaging

Chapter 5 presented a paradigm to surmount the schism between physical and virtual medicine through the generation of synthetic medical images. By learning the unspoken nomenclature of images in conveying geometric and morphological information, an unconventional approach to image processing was developed to convey such information in familiar visual vernacular from any digitized source of choice. Techniques developed to construct such images were described and compared, and potential applications and opportunities highlighted, forming the basis for the unification and integration of physical and virtual medicine. A fundamental approach explicitly modeling simplified imaging physics yielded interpretable but not entirely realistic images. However, the addition of noise to represent the stochastic distribution of discrete scatterers did improve realism and texture of generated images. A regression model considering the same feature as the physics-based model proved insufficient in cases with highly-attenuating plaques, in which exponential dropout in signal is a dominant feature which cannot be replicated by linear relationships. An alternative approach, whereby an AI system learned underlying relationships between system morphology and appearance in varying image modalities without explicit modeling of imaging physics, was quite successful in generating both OCT and IVUS images but sensitive to model hyperparameters. In particular, tradeoff between grayscale deviation loss and adversarial discriminator loss was immensely consequential, with both serving an important role in achieving sharp, textured, and globally cohesive images. Several exciting applications of this technology were described, with some executed in promising proof-of-concept demonstrations.

6.1.6 Conclusion

Chapter 6 now summarizes the work of the thesis, concluding the dissertation with a look forward at an envisioned future of cardiovascular medicine in the context of new insights into biomechanical assessment and modeling of diseased coronary arteries.

6.2. A Vision and Roadmap for the Future of Cardiovascular Medicine

Building upon the central pillars of the thesis, a personal optimistic vision for the clinical application of such technology in cardiology—and the concomitant unification and integration of physical and virtual medicine—has begun to take shape (Figure 6.1).

In this section, a brief vignette illustrating envisioned applications of the approaches developed in this thesis within the field of cardiology is first presented, followed by a cursory overview of future work needed to address current obstacles to realizing this vision.

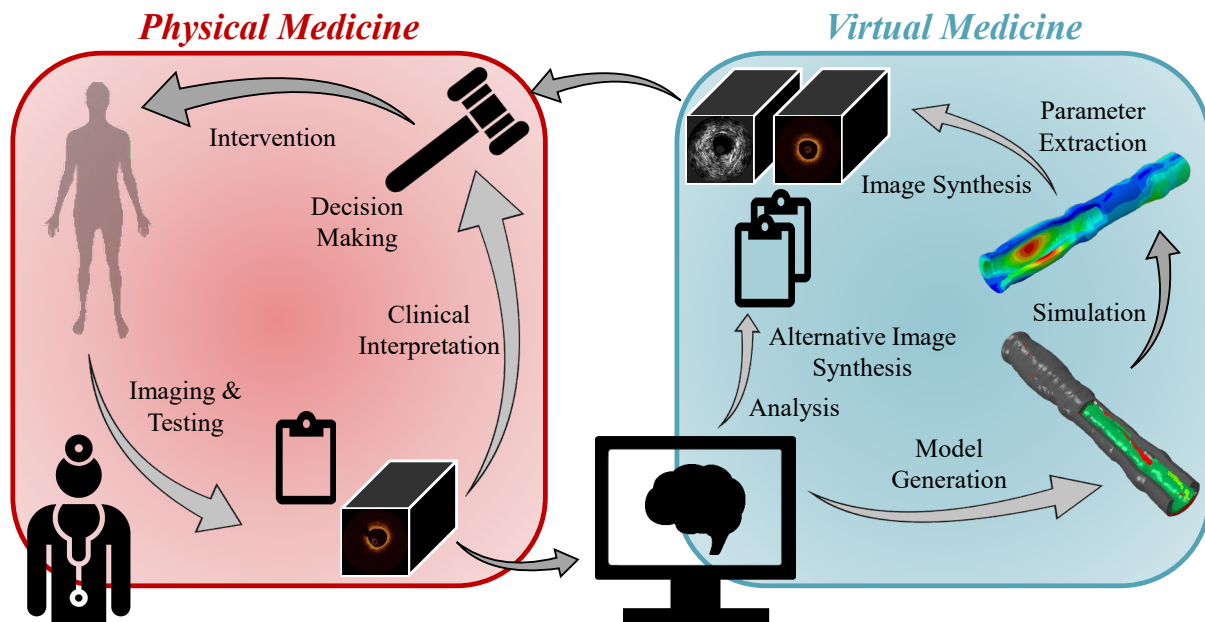


Figure 6.1. The proposed vision for the future of cardiovascular medicine is one in which physical and virtual medicine forms a continuum. The envisioned flow of information is as follows: clinicians collect data, including images or other test results, from a patient; anatomical (geometric) and morphological information is automatically extracted by algorithmic processing routines; this set of information is both distilled into reported quantitative metrics and used to generate patient-specific computational models; various simulated tests and procedures are performed on the virtual patient; and finally, results of the analyses and simulations are transformed back into clinical data, including images and standard metrics, to enable seamless integration and assessment by the care team. Outcomes could inform decision-making and guide procedure execution for the patient, or could be used to validate and assess outcomes of tests and procedures that are undertaken.

6.2.1. Vision for the Future of Integrated Physical-Virtual Medicine: A Vignette

A patient is admitted to the hospital with acute chest pain. Following vital sign measurement and preliminary screening with electrocardiography and blood biomarkers assays, doctors decide to proceed with minimally-invasive catheterization. Angiographic imaging of the coronary artery tree is performed to identify the likely cause of the symptoms and assess the severity and the condition. Spotting a stenosed artery which was flagged by integrated software for closer expert examination, interventional cardiologists pass catheters into the suspected culprit vessel to perform fractional flow reserve pressure measurement and two consecutive sets of intravascular OCT imaging, acquiring detailed pressure differential and morphological information on the lesion. Automated enhancement and augmentation of the data—leveraging previous observations from the angiographic imaging as well as algorithmic routines automatically processing the images—improves image clarity and integrates additional information for consideration. Key metrics of geometry and morphology are automatically computed and reported to the clinicians alongside the images.

Throughout this evaluation process, collected data has been assimilated into a virtual patient model, which is further refined with each subsequent input. A 3D model of the coronary tree lumina was generated from the angiographic image, and contrast dispersion rates were automatically computed to provide rough estimates of relative pressures across the branches. Direct invasive pressure measurements update the pressure values in the model and refine the boundary conditions. Vessel wall geometry and morphology in the segment imaged by OCT is subsequently integrated following lumen-based co-registration of the angiographic and OCT acquisitions. Material characterization is performed through inverse methods utilizing the measured pressure waveforms and two OCT pullbacks (which have been similarly co-registered).

Once sufficient information has been collected, the model undergoes several simulations. Risk assessments and failure analyses are performed, disease progression is prognosticated, and virtual interventions are evaluated to assess procedural risk and outcomes. Clinical metrics likely to result from various possible courses of treatment are reported alongside image sequences from corresponding future states.

The clinicians review the treatment options and likely outcomes of each. A consult is requested from a colleague. This clinician has in recent weeks been honing image interpretation

skills in OCT through a continuing education program using interactive learning exercises pairing images—generated from a known tissue distribution—with their corresponding morphologies and pathologies. However, primary experience and comfort remains in reading IVUS images, so the acquired and generated sequences are toggled from one modality to the other for easier interpretation.

Active intervention through stent deployment is pursued. The lesion preparation, stent design, device placement, and procedure parameters identified through simulations as yielding the lowest risk and highest probability of favorable outcome are selected. Following device deployment, another sequence of intravascular imaging is acquired and compared to images generated from the simulated outcome to confirm successful execution of the procedure. Deviation between the two sets highlights a region of sub-optimal state, and post-dilation of the stent is performed. Quantitative comparison of the post-procedural and synthetic image sequences also allows for the determination of some tissue constitutive properties under supraphysiological loading, which refine the virtual model for future use.

The patient returns several weeks later for a follow-up appointment. New measurements reveal the patient's longitudinal physiological, pathological, and pharmacological response. Deviations of model predictions from this observation are minimized, further refining and optimizing the patient phantom, which is used to simulate the vessel further forward in time and identify areas of concern to monitor. Reports of clinical metrics and imaging of future states are generated for the clinician, who uses them to determine the duration and aggression of medical treatment and frequency of follow-up. The future state images are also used as a concrete aid in discussion with the patient, enabling sounder understanding of the disease, ongoing risk, and reasons for treatment, thereby empowering the patient while promoting treatment adherence.

6.2.2. Next Steps and Future Work

The current state of the field is far from the envisioned vignette. Avenues for fruitful future work are addressed in each of the four preceding chapters in the context of work presented therein, with specific recommendations offered for areas of improvement and exploration. Overarching themes reveal a broad map to guide next steps in addressing current shortcomings which inhibit full realization of benefits from assessment and modeling of diseased arteries.

A primary focus of future work towards the presented vision is more extensive and robust validation. Whereas methods developed in this thesis for extracting vascular geometry and morphology were validated against expert annotations in intravascular images, more comprehensive validation should be carried out with histological studies. Development of improved phantoms with realistic optical properties could also provide a useful benchmarking tool. Similarly, improved phantoms with known heterogeneous properties should extend *in vitro* validation of the mechanical characterization approach of the inverse FE workflow. Here, as well, ultimate validation should be carried out with excised tissue, characterized through direct material testing procedures. Validation of synthetic images, which has so far been purely qualitative and subjective, should be executed more objectively by quantitatively assessing their interpretation by interventional cardiologists. In all cases, standardized datasets associated with reliable “ground truth” data would prove tremendous benefit to the field, enable greater contribution by a wider range of technical experts, allow the best methods and approaches to be more readily identified, and accelerate progress in advancing the state-of-the-art.

Related to standardized datasets, more data will need to be collected to develop generalizable methods and robustly validate technologies. First, direct measurement or improved estimation methods to determine luminal pressuring during image acquisition will be critical in advancing and demonstrating material characterization strategies. Methods to manipulate and exert some level of control over those pressures (e.g. through flush injection), while not strictly necessary, would provide even greater flexibility and power in the measurement of mechanical properties. Beyond pressure data, longitudinal data (repeated image acquisitions at various times) and associated clinical outcomes will be instrumental in realizing the full vision outlined above, which includes simulation of disease progression and long term response to treatments. Unfortunately, such datasets are largely lacking in the field. Furthermore, compartmentalization of data and institutions results in smaller and less representative datasets being available to developers of technology. As a result, models may be built around datasets which are not representative of the clinical population in which they are intended to be applied. As shown throughout this thesis, and as supported by many other observations, such biased data can lead to non-optimal performance in the “real world,” and unequal outcomes. Future work should expand available datasets to improve the robustness of methods and validity of validation, and deliberately

assess how performance of methods varies with respect to the clinical population in which it is applied.

Beyond more comprehensive and target population-aligned datasets, other assumptions will also need to be assessed. Modeling assumptions implemented in the construction of patient-specific computational models require comprehensive examination and quantification to ensure the validity of their output. Assumptions enshrined in the integration of contextual and domain knowledge, shown to generally provide tremendous benefit, must also be carefully considered to ensure misconceptions and unsubstantiated dogma are not propagated. For example, the tissue characterization method presented in this thesis segmented non-pathological tissue and media on the basis of standard thresholds set by previous histology studies and clinical standards. However, by using virtual histology IVUS as the benchmark, which itself utilizes the same assumptions, the validity of such thresholds, as well as the grouping of non-pathological and media tissue, was perpetuated without *a priori* justification. Such assumptions should be carefully examined and justified prior to widespread application.

Future work will need to accelerate the execution speed of many of the processes and analyses devised in this thesis. While the vision laid out above calls for nearly instantaneous processing and return of results for real time informing of decision-making, several workflows are currently far too slow for such application. Optimization of both hardware and software in image processing and generation, more strategic application of time-consuming processes, and greater automation of model construction are all needed.

One avenue of investigation alluded to throughout the thesis but not directly addressed by the dissertation research is the interaction of arterial biomechanical state, interventions, and devices. Incorporating these features and effects into models will enable the investigation of relationships key to treatment planning and risk profiling in the manner suggested in the vignette. Additional properties and constitutive parameters will need to be measured and modeled—those governing pharmacodynamics and pharmacokinetics, injury, foreign body, and inflammatory responses, tissue and structural remodeling, and other multifactorial behaviors and processes which play key roles in disease progression, treatment response, and disease regression. Furthermore, the geometry-morphology relationships, and the visualization of key features in

imaging, may be modified by such devices in interventional scenarios; investigations should evaluate when and how these relationships hold or are invalidated.

Finally, and perhaps most importantly, clinical needs must be comprehensively examined and understood. Scenarios in which additional insights may or would not impact clinical course, diminishing benefits or model accuracy—and consequently complexity—and the relative value of sensitivity, specificity, positive predictive value, and negative predictive value in various applications and scenarios must be honestly assessed. Doing so is imperative to ensuring the proper models are advanced and evaluation metrics optimized to yield clinical value.

6.3. In Closing

The achievements of this thesis have opened new windows into the biomechanical state of atherosclerotic arteries, and provided expanded awareness of the mechanical context of the disease. Through the work described in this thesis, greater insight is extracted from data currently collected in the clinical setting. With further refinement and validation, the methods and approaches will empower clinicians and researchers with enhanced information on diseased vessel constitution and biomechanical properties. Leveraging this information, which could previously only be derived through destructive *in vitro* testing, arterial modeling and clinical assessment will allow clinicians to make more informed choices regarding their course of treatment and researchers to pursue new insights into the pathobiology of atherosclerosis and the impact of interventions. These insights, currently unavailable to clinicians and researchers, could be used to monitor mechanical status of a lesion over time or in response to treatment, more quantitatively and reliably classify disease states, optimize intervention strategies, and assess tradeoffs between alternative treatment approaches. As such, troves of existing data could be leveraged, through use of tools and methods developed in the course of this thesis, to identify heretofore unrecognized trends and associations that may offer mechanistic insight into disease development, progression, and treatment response, ultimately leading to improve patient care.

With profound awe for the exquisite complexity of the cardiovascular system—and the body as a whole—and appreciation for predecessors who have advanced knowledge thereof, it is my profound hope that the insights and advancements of this thesis contribute to continued progress in diminishing the detrimental effects of cardiovascular disease and atherosclerosis.

References⁶

- [1] S. S. Virani *et al.*, “Heart Disease and Stroke Statistics—2020 Update: A Report From the American Heart Association,” *Circulation*, vol. 141, no. 9, pp. E139–E596, Mar. 2020.
- [2] E. G. Nabel and E. Braunwald, “A Tale of Coronary Artery Disease and Myocardial Infarction,” *N. Engl. J. Med.*, vol. 366, no. 1, pp. 54–63, Jan. 2012.
- [3] G. A. Holzapfel and C. J. Cyron, “Preface of the special issue on mathematical and computational modeling in biomechanics,” *ZAMM - J. Appl. Math. Mech. / Zeitschrift für Angew. Math. und Mech.*, vol. 98, no. 12, pp. 2044–2046, Dec. 2018.
- [4] W. Harvey, *Exercitatio anatomica de motu cordis et sanguinis in animalibus*. Frankfurt, 1628.
- [5] L. S. Lilly, Ed., *Pathophysiology of Heart Disease: A Collaborative Project of Medical Students and Faculty*, 5th ed. Lippincott Williams & Wilkins, 2011.
- [6] G. J. Tortora and M. T. Nielsen, *Principles of Human Anatomy, 13th Edition*. Wiley, 2013.
- [7] P. Libby, *The Vascular Biology of Atherosclerosis*, Ninth Edit. Elsevier Inc., 2011.
- [8] B. F. Waller *et al.*, “Anatomy, histology, and pathology of coronary arteries: A review relevant to new interventional and imaging techniques-Part I,” *Clin. Cardiol.*, vol. 15, no. 6, pp. 451–457, Jun. 1992.
- [9] A. J. Brown, Z. Teng, P. C. Evans, J. H. Gillard, H. Samady, and M. R. Bennett, “Role of biomechanical forces in the natural history of coronary atherosclerosis,” *Nat. Rev. Cardiol.*, vol. 13, no. 4, pp. 210–220, Apr. 2016.
- [10] B. Boon, “Leonardo da Vinci on atherosclerosis and the function of the sinuses of Valsalva,” *Netherlands Hear. J.*, vol. 17, no. 12, pp. 496–9, Sep. 2010.
- [11] C. H. Parry, *An Inquiry into the Symptoms and Causes of the Syncope Anginosa, Commonly Called Angina Pectoria; Illustrated by Dissections*. Bath: R. Cruttwell, 1799.
- [12] W. Heberden, *Commentaries on the history and cure of diseases*. Boston: Wells and Lilly, 1818.
- [13] W. Osler, “The Lumleian Lectures on Angina Pectoris,” *Lancet*, vol. 175, no. 4519, pp. 973–977, Apr. 1910.
- [14] J. L. Fleg *et al.*, “Detection of High-Risk Atherosclerotic Plaque: Report of the NHLBI Working Group on Current Status and Future Directions,” *JACC Cardiovasc. Imaging*, vol. 5, no. 9, pp. 941–955, Sep. 2012.
- [15] S. Celi, N. Martini, L. E. Pastormerlo, V. Positano, and S. Berti, “Multimodality Imaging for Interventional Cardiology,” *Curr. Pharm. Des.*, vol. 23, no. 22, pp. 3285–3300, Sep. 2017.
- [16] G. S. Mintz, “Clinical utility of intravascular imaging and physiology in coronary artery disease,” *Journal of the American College of Cardiology*, vol. 64, no. 2, pp. 207–222, 2014.

⁶ Referenced articles of which I am an author are shown in bold, with position in author list underlined

- [17] T. W. Johnson *et al.*, “Clinical use of intracoronary imaging. Part 2: acute coronary syndromes, ambiguous coronary angiography findings, and guiding interventional decision-making: an expert consensus document of the European Association of Percutaneous Cardiovascular Intervent,” *Eur. Heart J.*, vol. 40, no. 31, pp. 2566–2584, Aug. 2019.
- [18] G. J. Tearney *et al.*, “Consensus Standards for Acquisition, Measurement, and Reporting of Intravascular Optical Coherence Tomography Studies,” *J. Am. Coll. Cardiol.*, vol. 59, no. 12, pp. 1058–1072, Mar. 2012.
- [19] G. S. Mintz *et al.*, “American College of Cardiology clinical expert consensus document on standards for acquisition, measurement and reporting of intravascular ultrasound studies (IVUS),” *J. Am. Coll. Cardiol.*, vol. 37, no. 5, pp. 1478–1492, Apr. 2001.
- [20] E. Arbustini, M. Urtis, and F. Prati, “OCT/atherectomy/pathology studies open new perspectives for in vivo characterization of plaque composition,” *Int. J. Cardiol.*, vol. 284, pp. 14–15, Jun. 2019.
- [21] D. Vancraeynest, A. Pasquet, V. Roelants, B. L. Gerber, and J. L. J. Vanoverschelde, “Imaging the vulnerable plaque,” *J. Am. Coll. Cardiol.*, vol. 57, no. 20, pp. 1961–1979, May 2011.
- [22] F. Prati *et al.*, “In vivo vulnerability grading system of plaques causing acute coronary syndromes: An intravascular imaging study,” *Int. J. Cardiol.*, vol. 269, no. 5, pp. 350–355, Oct. 2018.
- [23] E. Romagnoli, L. Gatto, and F. Prati, “The CLIMA study: assessing the risk of myocardial infarction with a new anatomical score,” *Eur. Hear. J. Suppl.*, vol. 21, no. Supplement B, pp. B80–B83, Mar. 2019.
- [24] M. I. Papafaklis, M. C. Mavrogiannis, and P. H. Stone, “Identifying the progression of coronary artery disease: prediction of cardiac events,” *Contin. Cardiol. Educ.*, vol. 2, no. 2, pp. 105–114, Jun. 2016.
- [25] L. Räber *et al.*, “Clinical use of intracoronary imaging. Part 1: guidance and optimization of coronary interventions. An expert consensus document of the European Association of Percutaneous Cardiovascular Interventions,” *Eur. Heart J.*, vol. 39, no. 35, pp. 3281–3300, Sep. 2018.
- [26] K. C. Koskinas, G. J. Ughi, S. Windecker, G. J. Tearney, and L. Räber, “Intracoronary imaging of coronary atherosclerosis: validation for diagnosis, prognosis and treatment,” *Eur. Heart J.*, vol. 37, no. 6, pp. 524–535, Feb. 2016.
- [27] Y. Ueki, T. Otsuka, K. Hibi, and L. Räber, “The Value of Intracoronary Imaging and Coronary Physiology When Treating Calcified Lesions,” *Interv. Cardiol. Rev.*, vol. 14, no. 3, pp. 164–168, Nov. 2019.
- [28] G. S. Mintz, “Intravascular Imaging, Stent Implantation, and the Elephant in the Room,” *JACC Cardiovasc. Interv.*, vol. 10, no. 24, pp. 2499–2501, Dec. 2017.
- [29] T. F. Lüscher, “Assessing coronary plaques: non-invasive and intracoronary imaging and haemodynamic measurements,” *Eur. Heart J.*, vol. 39, no. 35, pp. 3265–3268, Sep. 2018.
- [30] J. M. Hodgson, “How Quickly We Forget,” *JACC Cardiovasc. Imaging*, vol. 11, no. 1, pp. 124–126, Jan. 2018.
- [31] L. Shaw and Y. Chandrashekar, “Progress in Cardiovascular Imaging,” *JACC Cardiovasc. Imaging*, vol. 12, no. 12, pp. 2589–2610, Dec. 2019.
- [32] P. Donnelly *et al.*, “Multimodality imaging atlas of coronary atherosclerosis,” *JACC Cardiovasc. Imaging*, vol. 3, no. 8, pp. 876–880, 2010.

- [33] D. Huang *et al.*, “Optical coherence tomography,” *Science*, vol. 254, no. 5035, pp. 1178–81, 1991.
- [34] G. J. Tearney *et al.*, “In Vivo Endoscopic Optical Biopsy with Optical Coherence Tomography,” *Science*, vol. 276, no. 5321, pp. 2037–2039, Jun. 1997.
- [35] G. J. Tearney *et al.*, “Porcine coronary imaging in vivo by optical coherence tomography,” *Acta Cardiol.*, vol. 55, no. 4, pp. 233–237, Aug. 2000.
- [36] E. A. Swanson and J. G. Fujimoto, “The ecosystem that powered the translation of OCT from fundamental research to clinical and commercial impact [Invited],” *Biomed. Opt. Express*, vol. 8, no. 3, pp. 1638–1664, Mar. 2017.
- [37] F. Prati *et al.*, “Expert review document part 2: methodology, terminology and clinical applications of optical coherence tomography for the assessment of interventional procedures,” *Eur. Heart J.*, vol. 33, no. 20, pp. 2513–2520, Oct. 2012.
- [38] S. Takarada *et al.*, “Advantage of next-generation frequency-domain optical coherence tomography compared with conventional time-domain system in the assessment of coronary lesion,” *Catheter. Cardiovasc. Interv.*, vol. 75, no. 2, pp. 202–206, Feb. 2010.
- [39] F. Prati *et al.*, “Expert review document on methodology, terminology, and clinical applications of optical coherence tomography: physical principles, methodology of image acquisition, and clinical application for assessment of coronary arteries and atherosclerosis,” *Eur. Heart J.*, vol. 31, no. 4, pp. 401–415, Feb. 2010.
- [40] T. Kubo *et al.*, “OCT compared with IVUS in a coronary lesion assessment: The OPUS-CLASS study,” *JACC Cardiovasc. Imaging*, vol. 6, no. 10, pp. 1095–1104, 2013.
- [41] R. Virmani, A. P. Burke, F. D. Kolodgie, and A. Farb, “Pathology of the thin-cap fibroatheroma: A type of vulnerable plaque,” *Journal of Interventional Cardiology*, vol. 16, no. 3. John Wiley & Sons, Ltd (10.1111), pp. 267–272, 01-Jun-2003.
- [42] P. R. Moreno, “The high-risk thin-cap fibroatheroma: A new kid on the block,” *Circ. Cardiovasc. Interv.*, vol. 2, no. 6, pp. 500–502, 2009.
- [43] G. Van Soest *et al.*, “Pitfalls in plaque characterization by OCT: Image artifacts in native coronary arteries,” *JACC: Cardiovascular Imaging*, vol. 4, no. 7. JACC: Cardiovascular Imaging, pp. 810–813, 01-Jul-2011.
- [44] M. J. Suter *et al.*, “Optimizing flushing parameters in intracoronary optical coherence tomography: an in vivo swine study,” *Int. J. Cardiovasc. Imaging*, vol. 31, no. 6, pp. 1097–1106, Aug. 2015.
- [45] F. Imola *et al.*, “Safety and feasibility of frequency domain optical coherence tomography to guide decision making in percutaneous coronary intervention,” *EuroIntervention*, vol. 6, no. 5, pp. 575–581, Nov. 2010.
- [46] F. Prati, M. Cera, V. Ramazzotti, F. Imola, R. Giudice, and M. Albertucci, “Safety and feasibility of a new non-occlusive technique for facilitated intracoronary optical coherence tomography (OCT) acquisition in various clinical and anatomical scenarios,” *EuroIntervention*, vol. 3, no. 3, pp. 365–370, Nov. 2007.
- [47] R. Shah, “Optical coherence tomography-guided PCI,” *Lancet*, vol. 389, no. 10079, p. 1607, Apr. 2017.
- [48] K. Shimamura *et al.*, “Intracoronary pressure increase due to contrast injection for optical coherence tomography imaging,” *J. Cardiol.*, vol. 75, no. 3, pp. 296–301, Mar. 2020.

- [49] G. M. Stiel, A. Lattermann, K.-P. Schaps, M. Siebes, and C. A. Nienaber, “Effect of contrast injection on coronary artery pressure and dimension during angiography,” in *Computers in Cardiology 1995*, 1995, pp. 157–160.
- [50] I.-K. Jang *et al.*, “Visualization of coronary atherosclerotic plaques in patients using optical coherence tomography: comparison with intravascular ultrasound,” *J. Am. Coll. Cardiol.*, vol. 39, no. 4, pp. 604–609, Feb. 2002.
- [51] H. G. Bezerra *et al.*, “Optical coherence tomography versus intravascular ultrasound to evaluate coronary artery disease and percutaneous coronary intervention,” *JACC Cardiovasc. Interv.*, vol. 6, no. 3, pp. 228–236, 2013.
- [52] F. Abnoui *et al.*, “Variability in quantitative and qualitative analysis of intravascular ultrasound and frequency domain optical coherence tomography,” *Catheter. Cardiovasc. Interv.*, vol. 82, no. 3, pp. E192–E199, Sep. 2013.
- [53] J. N. van der Sijde *et al.*, “Safety of optical coherence tomography in daily practice: a comparison with intravascular ultrasound,” *Eur. Hear. J. – Cardiovasc. Imaging*, vol. 18, no. 4, p. jew037, Mar. 2016.
- [54] T. Yamaguchi *et al.*, “Safety and Feasibility of an Intravascular Optical Coherence Tomography Image Wire System in the Clinical Setting,” *Am. J. Cardiol.*, vol. 101, no. 5, pp. 562–567, Mar. 2008.
- [55] F. Alfonso *et al.*, “Combined use of optical coherence tomography and intravascular ultrasound imaging in patients undergoing coronary interventions for stent thrombosis,” *Heart*, vol. 98, no. 16, pp. 1213–1220, Aug. 2012.
- [56] N. Gonzalo *et al.*, “Second-Generation Optical Coherence Tomography in Clinical Practice. High-Speed Data Acquisition is Highly Reproducible in Patients Undergoing Percutaneous Coronary Intervention,” *Rev. Española Cardiol. (English Ed.)*, vol. 63, no. 8, pp. 893–903, Jan. 2010.
- [57] M. Habara *et al.*, “In vivo tissue characterization of human atherosclerotic plaques by optical coherence tomography: A directional coronary atherectomy study with histopathologic confirmation,” *Int. J. Cardiol.*, vol. 268, no. 5, pp. 1–10, Oct. 2018.
- [58] Y. Liu *et al.*, “Comparison of longitudinal geometric measurement in human coronary arteries between frequency-domain optical coherence tomography and intravascular ultrasound,” *Int. J. Cardiovasc. Imaging*, vol. 30, pp. 271–277, 2014.
- [59] N. Meneveau *et al.*, “Optical Coherence Tomography to Optimize Results of Percutaneous Coronary Intervention in Patients with Non–ST-Elevation Acute Coronary Syndrome,” *Circulation*, vol. 134, no. 13, pp. 906–917, Sep. 2016.
- [60] F. Burzotta *et al.*, “Fractional Flow Reserve or Optical Coherence Tomography to Guide Management of Angiographically Intermediate Coronary Stenosis,” *JACC Cardiovasc. Interv.*, vol. 13, no. 1, pp. 49–58, Jan. 2020.
- [61] F. Prati *et al.*, “Clinical Impact of OCT Findings During PCI,” *JACC Cardiovasc. Imaging*, vol. 8, no. 11, pp. 1297–1305, Nov. 2015.
- [62] T. Kubo *et al.*, “Feasibility of Optical Coronary Tomography in Quantitative Measurement of Coronary Arteries With Lipid-Rich Plaque,” *Circ. J.*, vol. 79, no. 3, pp. 600–606, 2015.

- [63] K. Sakata, M. Kawashiri, and K. Hayashi, “Expansion of the Clinical Application of Optical Coherence Tomography to Percutaneous Coronary Intervention and Assessment of the Instability of Coronary Atherosclerosis,” *Circ. J.*, vol. 79, no. 3, pp. 513–514, 2015.
- [64] K. Otsuka *et al.*, “Intravascular Polarimetry in Patients With Coronary Artery Disease,” *JACC Cardiovasc. Imaging*, vol. 13, no. 3, pp. 790–801, Mar. 2020.
- [65] E. Gerbaud *et al.*, “Plaque burden can be assessed using intravascular optical coherence tomography and a dedicated automated processing algorithm: a comparison study with intravascular ultrasound,” *Eur. Hear. J. - Cardiovasc. Imaging*, vol. 21, no. 6, pp. 640–652, Jul. 2019.
- [66] B. E. Bouma, M. Villiger, K. Otsuka, and W.-Y. Oh, “Intravascular optical coherence tomography [Invited],” *Biomed. Opt. Express*, vol. 8, no. 5, p. 2660, May 2017.
- [67] L. S. Athanasiou, F. R. Nezami, and E. R. Edelman, “Computational Cardiology,” *IEEE J. Biomed. Heal. Informatics*, vol. 23, no. 1, pp. 4–11, Jan. 2019.
- [68] S. Morlacchi and F. Migliavacca, “Modeling Stented Coronary Arteries: Where We are, Where to Go,” *Ann. Biomed. Eng.*, vol. 41, no. 7, pp. 1428–1444, Jul. 2013.
- [69] C. A. Taylor and D. A. Steinman, “Image-Based Modeling of Blood Flow and Vessel Wall Dynamics: Applications, Methods and Future Directions,” *Ann. Biomed. Eng.*, vol. 38, no. 3, pp. 1188–1203, Mar. 2010.
- [70] C. Chiastra, S. Migliori, F. Burzotta, G. Dubini, and F. Migliavacca, “Patient-Specific Modeling of Stented Coronary Arteries Reconstructed from Optical Coherence Tomography: Towards a Widespread Clinical Use of Fluid Dynamics Analyses,” *J. Cardiovasc. Transl. Res.*, vol. 11, no. 2, pp. 156–172, Apr. 2018.
- [71] M. J. Turner, R. W. Clough, H. C. Martin, and L. J. Topp, “Stiffness and Deflection Analysis of Complex Structures,” *J. Aeronaut. Sci.*, vol. 23, no. 9, pp. 805–823, Sep. 1956.
- [72] G. A. Holzapfel, J. J. Mulvihill, E. M. Cunnane, and M. T. Walsh, “Computational approaches for analyzing the mechanics of atherosclerotic plaques: A review,” *J. Biomech.*, vol. 47, no. 4, pp. 859–869, 2014.
- [73] A. I. Sakellarios *et al.*, “Available Computational Techniques to Model Atherosclerotic Plaque Progression Implementing a Multi-Level Approach,” in *Computational Biomechanics for Medicine*, K. Miller and P. Nielsen, Eds. Cham: Springer International Publishing, 2017, pp. 39–55.
- [74] R. T. Lee, A. J. Grodzinsky, E. H. Frank, R. D. Kamm, and F. J. Schoen, “Structure-dependent dynamic mechanical behavior of fibrous caps from human atherosclerotic plaques,” *Circulation*, vol. 83, no. 5, pp. 1764–70, May 1991.
- [75] H. M. Loree, R. D. Kamm, R. G. Stringfellow, and R. T. Lee, “Effects of fibrous cap thickness on peak circumferential stress in model atherosclerotic vessels,” *Circ. Res.*, vol. 71, no. 4, pp. 850–858, 1992.
- [76] A. C. Akyildiz *et al.*, “Intima heterogeneity in stress assessment of atherosclerotic plaques,” *Interface Focus*, vol. 8, no. 20170008, Feb. 2018.
- [77] C. Conway, J. P. McGarry, E. R. Edelman, and P. E. McHugh, “Numerical Simulation of Stent Angioplasty with Predilation: An Investigation into Lesion Constitutive Representation and Calcification Influence,” *Ann. Biomed. Eng.*, vol. 45, no. 9, pp. 2244–2252, Sep. 2017.

- [78] E. L. Boland, J. A. Grogan, C. Conway, and P. E. McHugh, “Computer Simulation of the Mechanical Behaviour of Implanted Biodegradable Stents in a Remodelling Artery,” *J. Miner. Met. Mater. Soc.*, vol. 34, no. 4, 2016.
- [79] K. D. Everett *et al.*, “Structural Mechanics Predictions Relating to Clinical Coronary Stent Fracture in a 5 Year Period in FDA MAUDE Database,” *Ann. Biomed. Eng.*, vol. 44, no. 2, pp. 391–403, 2016.
- [80] S. Sankaran, H. J. Kim, G. Choi, and C. A. Taylor, “Uncertainty quantification in coronary blood flow simulations: Impact of geometry, boundary conditions and blood viscosity,” *J. Biomech.*, vol. 49, no. 12, pp. 2540–2547, Aug. 2016.
- [81] M. Malvè, A. García, J. Ohayon, and M. A. Martínez, “Unsteady blood flow and mass transfer of a human left coronary artery bifurcation: FSI vs. CFD,” *Int. Commun. Heat Mass Transf.*, vol. 39, no. 6, pp. 745–751, Jul. 2012.
- [82] L. Cardoso, A. Kelly-Arnold, N. Maldonado, D. Laudier, and S. Weinbaum, “Effect of tissue properties, shape and orientation of microcalcifications on vulnerable cap stability using different hyperelastic constitutive models,” *J. Biomech.*, vol. 47, no. 4, pp. 870–877, Mar. 2014.
- [83] J. Ohayon *et al.*, “Influence of residual stress/strain on the biomechanical stability of vulnerable coronary plaques: potential impact for evaluating the risk of plaque rupture,” *Am. J. Physiol. Circ. Physiol.*, vol. 293, no. 3, pp. H1987–H1996, Sep. 2007.
- [84] C. Polindara, T. Waffenschmidt, and A. Menzel, “Simulation of balloon angioplasty in residually stressed blood vessels—Application of a gradient-enhanced fibre damage model,” *J. Biomech.*, vol. 49, no. 12, pp. 2341–2348, Aug. 2016.
- [85] J. Bols, J. Degroote, B. Trachet, B. Verhegghe, P. Segers, and J. Vierendeels, “A computational method to assess the in vivo stresses and unloaded configuration of patient-specific blood vessels,” *J. Comput. Appl. Math.*, vol. 246, pp. 10–17, Jul. 2013.
- [86] P. D. Morris, J. Iqbal, C. Chiastra, W. Wu, F. Migliavacca, and J. P. Gunn, “Simultaneous kissing stents to treat unprotected left main stem coronary artery bifurcation disease; stent expansion, vessel injury, hemodynamics, tissue healing, restenosis, and repeat revascularization,” *Catheter. Cardiovasc. Interv.*, vol. 00, pp. 1–12, Apr. 2018.
- [87] A. P. Antoniadis *et al.*, “Biomechanical Modeling to Improve Coronary Artery Bifurcation Stenting: Expert Review Document on Techniques and Clinical Implementation,” *JACC Cardiovasc. Interv.*, vol. 8, no. 10, pp. 1281–1296, Aug. 2015.
- [88] F. Iannaccone *et al.*, “Impact of plaque type and side branch geometry on side branch compromise after provisional stent implantation: A simulation study,” *EuroIntervention*, vol. 13, pp. e236–e245, 2017.
- [89] C. Rogers, D. Y. Tseng, J. C. Squire, and E. R. Edelman, “Balloon-Artery Interactions During Stent Placement: A Finite Element Analysis Approach to Pressure, Compliance, and Stent Design as Contributors to Vascular Injury,” *Circ. Res.*, vol. 84, no. 4, pp. 378–383, Mar. 1999.
- [90] G. A. Holzapfel, M. Stadler, and T. C. Gasser, “Changes in the Mechanical Environment of Stenotic Arteries During Interaction With Stents: Computational Assessment of Parametric Stent Designs,” *J. Biomech. Eng.*, vol. 127, no. 1, p. 166, Mar. 2005.
- [91] G. S. Karanasiou *et al.*, “In Silico Assessment of the effects of Material on Stent Deployment,” in *2017 IEEE 17th International Conference on Bioinformatics and Bioengineering (BIBE)*, 2017, pp. 462–467.

- [92] G. E. Ragkousis, N. Curzen, and N. W. Bressloff, “Multi-objective optimisation of stent dilation strategy in a patient-specific coronary artery via computational and surrogate modelling,” *J. Biomech.*, vol. 49, no. 2, pp. 205–215, Jan. 2016.
- [93] S. Morlacchi, G. Pennati, L. Petrini, G. Dubini, and F. Migliavacca, “Influence of plaque calcifications on coronary stent fracture: A numerical fatigue life analysis including cardiac wall movement,” *J. Biomech.*, vol. 47, no. 4, pp. 899–907, 2014.
- [94] E. Yamamoto *et al.*, “Low Endothelial Shear Stress Predicts Evolution to High-Risk Coronary Plaque Phenotype in the Future,” *Circ. Cardiovasc. Interv.*, vol. 10, no. 8, p. e005455, Aug. 2017.
- [95] V. Thondapu, C. V. Bourantas, N. Foin, I. K. Jang, P. W. Serruys, and P. Barlis, “Basic science for the clinician: Biomechanical stress in coronary atherosclerosis: Emerging insights from computational modelling,” *Eur. Heart J.*, vol. 38, no. 2, pp. 81–92, 2017.
- [96] V. Thondapu *et al.*, “Endothelial shear stress 5 years after implantation of a coronary bioresorbable scaffold,” *Eur. Heart J.*, vol. 39, no. 18, pp. 1602–1609, May 2018.
- [97] S. Migliori *et al.*, “A Patient-Specific Study Investigating the Relation between Coronary Hemodynamics and Neo-Intimal Thickening after Bifurcation Stenting with a Polymeric Bioresorbable Scaffold,” *Appl. Sci.*, vol. 8, no. 9, p. 1510, Sep. 2018.
- [98] S.-J. Kang *et al.*, “Plaque structural stress assessed by virtual histology-intravascular ultrasound predicts dynamic changes in phenotype and composition of untreated coronary artery lesions,” *Atherosclerosis*, vol. 254, pp. 85–92, Nov. 2016.
- [99] L. H. Timmins, M. W. Miller, F. J. Clubb, and J. E. Moore, “Increased artery wall stress post-stenting leads to greater intimal thickening,” *Lab. Investig.*, vol. 91, no. 6, pp. 955–967, Jun. 2011.
- [100] G. Finet, J. Ohayon, and G. Rioufol, “Biomechanical interaction between cap thickness, lipid core composition and blood pressure in vulnerable coronary plaque: impact on stability or instability,” *Coron. Artery Dis.*, vol. 15, no. 1, pp. 13–20, 2004.
- [101] L. Cardoso and S. Weinbaum, “Changing views of the biomechanics of vulnerable plaque rupture: A review,” *Ann. Biomed. Eng.*, vol. 42, no. 2, pp. 415–431, Feb. 2014.
- [102] J. Ohayon *et al.*, “Biomechanics of Atherosclerotic Coronary Plaque: Site, Stability and In Vivo Elasticity Modeling,” *Ann. Biomed. Eng.*, vol. 42, no. 2, pp. 269–279, Feb. 2014.
- [103] D. Tang *et al.*, “Image-based modeling for better understanding and assessment of atherosclerotic plaque progression and vulnerability: Data, modeling, validation, uncertainty and predictions,” *J. Biomech.*, vol. 47, no. 4, pp. 834–846, Mar. 2014.
- [104] C. Stefanadis, C. Antoniou, D. Tsiachris, and P. Pietri, “Coronary Atherosclerotic Vulnerable Plaque: Current Perspectives,” *J. Am. Heart Assoc.*, vol. 6, no. 3, p. e005543, Mar. 2017.
- [105] R. A. Gray and P. Pathmanathan, “Patient-Specific Cardiovascular Computational Modeling: Diversity of Personalization and Challenges,” *J. Cardiovasc. Transl. Res.*, vol. 11, no. 2, pp. 80–88, Apr. 2018.
- [106] C. A. Taylor, T. A. Fonte, and J. K. Min, “Computational Fluid Dynamics Applied to Cardiac Computed Tomography for Noninvasive Quantification of Fractional Flow Reserve,” *J. Am. Coll. Cardiol.*, vol. 61, no. 22, pp. 2233–2241, Jun. 2013.
- [107] D. Bluestein *et al.*, “Influence of microcalcifications on vulnerable plaque mechanics using FSI modeling,” *J. Biomech.*, vol. 41, no. 5, pp. 1111–1118, Jan. 2008.

- [108] H. Chen and G. S. Kassab, “Microstructure-based biomechanics of coronary arteries in health and disease,” *J. Biomech.*, vol. 49, no. 12, pp. 2548–2559, 2016.
- [109] A. C. Akyildiz, L. Speelman, and F. J. H. Gijzen, “Mechanical properties of human atherosclerotic intima tissue,” *J. Biomech.*, vol. 47, no. 4, pp. 773–783, 2014.
- [110] B. F. Kennedy, K. M. Kennedy, and D. D. Sampson, “A Review of Optical Coherence Elastography: Fundamentals, Techniques and Prospects,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 2, pp. 272–288, Mar. 2014.
- [111] S. K. Nadkarni *et al.*, “Characterization of Atherosclerotic Plaques by Laser Speckle Imaging,” *Circulation*, vol. 112, no. 6, pp. 885–892, Aug. 2005.
- [112] R. C. Chan *et al.*, “OCT-based arterial elastography: robust estimation exploiting tissue biomechanics,” *Opt. Express*, vol. 12, no. 19, pp. 4558–4572, Sep. 2004.
- [113] R. L. Winslow, N. Trayanova, D. Geman, and M. I. Miller, “Computational Medicine: Translating Models to Clinical Care,” *Sci. Transl. Med.*, vol. 4, no. 158, pp. 158rv11–158rv11, Oct. 2012.
- [114] P. Bizopoulos and D. Koutsouris, “Deep Learning in Cardiology,” *IEEE Rev. Biomed. Eng.*, vol. 12, pp. 168–193, 2019.
- [115] K. W. Johnson *et al.*, “Artificial Intelligence in Cardiology,” *J. Am. Coll. Cardiol.*, vol. 71, no. 23, pp. 2668–2679, Jun. 2018.
- [116] K. Shameer, K. W. Johnson, B. S. Glicksberg, J. T. Dudley, and P. P. Sengupta, “Machine learning in cardiovascular medicine: are we there yet?,” *Heart*, vol. 104, no. 14, pp. 1156–1164, Jul. 2018.
- [117] D. Tsay and C. Patterson, “From Machine Learning to Artificial Intelligence Applications in Cardiac Care,” *Circulation*, vol. 138, no. 22, pp. 2569–2575, Nov. 2018.
- [118] P. Sardar, J. D. Abbott, A. Kundu, H. D. Aronow, J. F. Granada, and J. Giri, “Impact of Artificial Intelligence on Interventional Cardiology: From Decision-Making Aid to Advanced Interventional Procedure Assistance,” *JACC Cardiovasc. Interv.*, vol. 12, no. 14, pp. 1293–1303, Jul. 2019.
- [119] C. Krittanawong, H. Zhang, Z. Wang, M. Aydar, and T. Kitai, “Artificial Intelligence in Precision Cardiovascular Medicine,” *J. Am. Coll. Cardiol.*, vol. 69, no. 21, pp. 2657–2664, May 2017.
- [120] J. P. Singh, “It Is Time for Us to Get Artificially Intelligent!,” *JACC Clin. Electrophysiol.*, vol. 5, no. 2, pp. 263–265, Feb. 2019.
- [121] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [122] A. Rajkomar, J. Dean, and I. Kohane, “Machine Learning in Medicine,” *N. Engl. J. Med.*, vol. 380, no. 14, pp. 1347–1358, Apr. 2019.
- [123] A. L. Fogel and J. C. Kvedar, “Artificial intelligence powers digital medicine,” *npj Digit. Med.*, vol. 1, no. 1, p. 5, Dec. 2018.
- [124] V. H. Buch, I. Ahmed, and M. Maruthappu, “Artificial intelligence in medicine: current trends and future possibilities,” *Br. J. Gen. Pract.*, vol. 68, no. 668, pp. 143–144, Mar. 2018.
- [125] C. D. Naylor, “On the Prospects for a (Deep) Learning Health Care System,” *JAMA*, vol. 320, no. 11, p. 1099, Sep. 2018.

- [126] R. C. Deo, “Machine Learning in Medicine,” *Circulation*, vol. 132, no. 20, pp. 1920–1930, Nov. 2015.
- [127] Z. Obermeyer and E. J. Emanuel, “Predicting the Future — Big Data, Machine Learning, and Clinical Medicine,” *N. Engl. J. Med.*, vol. 375, no. 13, pp. 1216–1219, Sep. 2016.
- [128] L. Zhu and W. J. Zheng, “Informatics, Data Science, and Artificial Intelligence,” *JAMA*, vol. 320, no. 11, p. 1103, Sep. 2018.
- [129] G. Hinton, “Deep Learning—A Technology With the Potential to Transform Health Care,” *JAMA*, vol. 320, no. 11, p. 1101, Sep. 2018.
- [130] N. Kagiya, S. Shrestha, P. D. Farjo, and P. P. Sengupta, “Artificial Intelligence: Practical Primer for Clinical Research in Cardiovascular Disease,” *J. Am. Heart Assoc.*, vol. 8, no. 17, p. e012788, Sep. 2019.
- [131] W. S. Weintraub, “Role of Big Data in Cardiovascular Research,” *J. Am. Heart Assoc.*, vol. 8, no. 14, p. e012791, Jul. 2019.
- [132] W. S. Weintraub, A. C. Fahed, and J. S. Rumsfeld, “Translational Medicine in the Era of Big Data and Machine Learning,” *Circ. Res.*, vol. 123, no. 11, pp. 1202–1204, Nov. 2018.
- [133] J. H. Chen and S. M. Asch, “Machine Learning and Prediction in Medicine — Beyond the Peak of Inflated Expectations,” *N. Engl. J. Med.*, vol. 376, no. 26, pp. 2507–2509, Jun. 2017.
- [134] W. W. Stead, “Clinical Implications and Challenges of Artificial Intelligence and Deep Learning,” *JAMA*, vol. 320, no. 11, p. 1107, Sep. 2018.
- [135] B.-K. Koo and H. Samady, “Strap In for the Artificial Intelligence Revolution in Interventional Cardiology,” *JACC Cardiovasc. Interv.*, vol. 12, no. 14, pp. 1325–1327, Jul. 2019.
- [136] F. Cabitza, R. Rasoini, and G. F. Gensini, “Unintended Consequences of Machine Learning in Medicine,” *JAMA*, vol. 318, no. 6, p. 517, Aug. 2017.
- [137] T. Modine and P. Overtchouk, “Machine Learning Is No Magic: A Plea for Critical Appraisal During Periods of Hype,” *JACC Cardiovasc. Interv.*, vol. 12, no. 14, pp. 1339–1341, Jul. 2019.
- [138] F. Rodriguez, D. Scheinker, and R. A. Harrington, “Promise and Perils of Big Data and Artificial Intelligence in Clinical Medicine and Biomedical Research,” *Circ. Res.*, vol. 123, no. 12, pp. 1282–1284, Dec. 2018.
- [139] M. van Assen and L. J. Cornelissen, “Artificial Intelligence: From Scientific Curiosity to Clinical Precocity?,” *JACC Cardiovasc. Imaging*, vol. 13, no. 5, pp. 1172–1174, May 2020.
- [140] R. Puri and S. R. Kapadia, “Intelligently Advancing Cardiovascular Health Care,” *JACC Cardiovasc. Interv.*, vol. 12, no. 20, pp. 2102–2105, Oct. 2019.
- [141] R. J. Westcott and J. E. Tchong, “Artificial Intelligence and Machine Learning in Cardiology,” *JACC Cardiovasc. Interv.*, vol. 12, no. 14, pp. 1312–1314, Jul. 2019.
- [142] R. Tomsett *et al.*, “Rapid Trust Calibration through Interpretable and Uncertainty-Aware AI,” *Patterns*, vol. 1, no. 4, p. 100049, Jul. 2020.
- [143] C. M. Stultz, “The Advent of Clinically Useful Deep Learning,” *JACC Clin. Electrophysiol.*, vol. 5, no. 5, pp. 587–589, May 2019.

- [144] H. Greenspan, B. van Ginneken, and R. M. Summers, “Guest Editorial Deep Learning in Medical Imaging: Overview and Future Promise of an Exciting New Technique,” *IEEE Trans. Med. Imaging*, vol. 35, no. 5, pp. 1153–1159, May 2016.
- [145] G. Litjens *et al.*, “A survey on deep learning in medical image analysis,” *Med. Image Anal.*, vol. 42, pp. 60–88, Dec. 2017.
- [146] D. Shen, G. Wu, and H.-I. Suk, “Deep Learning in Medical Image Analysis,” *Annu. Rev. Biomed. Eng.*, vol. 19, no. 1, pp. 221–248, Jun. 2017.
- [147] L. Carin and M. J. Pencina, “On Deep Learning for Medical Image Analysis,” *JAMA*, vol. 320, no. 11, p. 1192, Sep. 2018.
- [148] B. M. Kuehn, “Cardiac Imaging on the Cusp of an Artificial Intelligence Revolution,” *Circulation*, vol. 141, no. 15, pp. 1266–1267, Apr. 2020.
- [149] G. Litjens *et al.*, “State-of-the-Art Deep Learning in Cardiovascular Image Analysis,” *JACC Cardiovasc. Imaging*, vol. 12, no. 8, pp. 1549–1565, Aug. 2019.
- [150] D. Dey *et al.*, “Artificial Intelligence in Cardiovascular Imaging,” *J. Am. Coll. Cardiol.*, vol. 73, no. 11, pp. 1317–1335, Mar. 2019.
- [151] M. Henglin, G. Stein, P. V. Hushcha, J. Snoek, A. B. Wiltchko, and S. Cheng, “Machine Learning Approaches in Cardiovascular Imaging,” *Circ. Cardiovasc. Imaging*, vol. 10, no. 10, Oct. 2017.
- [152] N. K. Sabharwal, “Could Deep Learning Change Our Working Lives?,” *JACC Cardiovasc. Imaging*, vol. 11, no. 11, pp. 1664–1665, Nov. 2018.
- [153] X. Yi, E. Walia, and P. Babyn, “Generative adversarial network in medical imaging: A review,” *Med. Image Anal.*, vol. 58, p. 101552, Dec. 2019.
- [154] Y. Mirsky, T. Mahler, I. Shelef, and Y. Elovici, “CT-GAN: Malicious tampering of 3D medical imagery using deep learning,” in *Proceedings of the 28th USENIX Security Symposium*, 2019, pp. 461–478.
- [155] **M. L. Olender *et al.***, “Estimating the internal elastic membrane cross-sectional area of coronary arteries autonomously using optical coherence tomography images,” in *2017 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, 2017, pp. 109–112.
- [156] **M. L. Olender *et al.***, “Algoritmo computacional para detectar el borde interno de la capa media de arterias coronarias ateromatosas mediante tomografía de coherencia óptica,” in *SEC 2017: El Congreso de las Enfermedades Cardiovasculares*, 2017, p. 70.
- [157] **M. L. Olender**, L. S. Athanasiou, J. M. de la Torre Hernández, E. Ben-Assa, F. R. Nezami, and E. R. Edelman, “A Mechanical Approach for Smooth Surface Fitting to Delineate Vessel Walls in Optical Coherence Tomography Images,” *IEEE Trans. Med. Imaging*, vol. 38, no. 6, pp. 1384–1397, Jun. 2019.
- [158] **M. L. Olender**, L. S. Athanasiou, J. M. de la Torre Hernandez, E. Ben-Assa, and E. R. Edelman, “Simultaneous Multi-Surface Fitting for Vessel Wall Layer Delineation,” in *2019 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI)*, 2019, pp. 1–4.
- [159] A. Karmakar, **M. L. Olender**, *et al.*, “Detailed Investigation of Lumen-Based Tomographic Co-Registration,” in *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, 2020, pp. 1038-1042.

- [160] L. S. Athanasiou, M. L. Olender, J. M. de la Torre Hernandez, E. Ben-Assa, and E. R. Edelman, “A deep learning approach to classify atherosclerosis using intracoronary optical coherence tomography,” in *Medical Imaging 2019: Computer-Aided Diagnosis*, 2019, vol. 10950, p. 22.
- [161] M. L. Olender, L. S. Athanasiou, L. K. Michalis, D. I. Fotiadis, and E. R. Edelman, “A Domain Enriched Deep Learning Approach to Classify Atherosclerosis Using Intravascular Ultrasound Imaging,” *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1210–1220, Oct. 2020.
- [162] A. Gowrishankar, L. S. Athanasiou, M. L. Olender, and E. R. Edelman, “Neural Network Training Data Profoundly Impacts Texture-Based Intravascular Image Segmentation,” in *2019 IEEE 19th International Conference on Bioinformatics and Bioengineering (BIBE)*, 2019, pp. 989–993.
- [163] K. Fujii *et al.*, “Expert consensus statement for quantitative measurement and morphological assessment of optical coherence tomography,” *Cardiovasc. Interv. Ther.*, vol. 35, no. 1, pp. 13–18, Jan. 2020.
- [164] B. F. Waller, “The eccentric coronary atherosclerotic plaque: Morphologic observations and clinical relevance,” *Clin. Cardiol.*, vol. 12, no. 1, pp. 14–20, Jan. 1989.
- [165] B. F. Waller, C. M. Orr, J. D. Slack, C. A. Pinkerton, J. Van Tassel, and T. Peters, “Anatomy, histology, and pathology of coronary arteries: A review relevant to new interventional and imaging techniques-Part II,” *Clin. Cardiol.*, vol. 15, no. 7, pp. 535–540, Jul. 1992.
- [166] B. F. Waller, C. M. Orr, J. D. Slack, C. A. Pinkerton, J. Van Tassel, and T. Peters, “Anatomy, histology, and pathology of coronary arteries: A review relevant to new interventional and imaging techniques-Part III,” *Clin. Cardiol.*, vol. 15, no. 8, pp. 607–615, Aug. 1992.
- [167] B. F. Waller, C. M. Orr, J. D. Slack, C. A. Pinkerton, J. Van Tassel, and T. Peters, “Anatomy, histology, and pathology of coronary arteries: A review relevant to new interventional and imaging techniques-Part IV,” *Clin. Cardiol.*, vol. 15, no. 9, pp. 675–687, Sep. 1992.
- [168] L. S. Athanasiou *et al.*, “3D reconstruction of coronary arteries using Frequency Domain Optical Coherence Tomography images and biplane angiography,” in *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 2012, vol. 2012, pp. 2647–2650.
- [169] M. I. Papafaklis *et al.*, “Anatomically correct three-dimensional coronary artery reconstruction using frequency domain optical coherence tomographic and angiographic data: head-to-head comparison with intravascular ultrasound for endothelial shear stress assessment in humans,” *EuroIntervention*, vol. 11, no. 4, pp. 407–415, Aug. 2015.
- [170] C. Chiastra *et al.*, “Reconstruction of stented coronary arteries from optical coherence tomography images: feasibility, validation, and repeatability of a segmentation method,” *PLoS One*, vol. In press, no. 6, pp. 1–23, 2017.
- [171] S. Migliori *et al.*, “A framework for computational fluid dynamic analyses of patient-specific stented coronary arteries from optical coherence tomography images,” *Med. Eng. Phys.*, vol. 47, pp. 105–116, 2017.
- [172] C. Chiastra *et al.*, “Computational replication of the patient-specific stenting procedure for coronary artery bifurcations: From OCT and CT imaging to structural and hemodynamics analyses,” *J. Biomech.*, vol. 49, no. 11, pp. 2102–2111, Jul. 2016.

- [173] L. S. Athanasiou *et al.*, “Methodology for fully automated segmentation and plaque characterization in intracoronary optical coherence tomography images,” *J. Biomed. Opt.*, vol. 19, no. 2, p. 026009, Feb. 2014.
- [174] S. Celi and S. Berti, “In-vivo segmentation and quantification of coronary lesions by optical coherence tomography images for a lesion type definition and stenosis grading,” *Med. Image Anal.*, vol. 18, no. 7, pp. 1157–1168, Oct. 2014.
- [175] L. S. Athanasiou *et al.*, “Currently available methodologies for the processing of intravascular ultrasound and optical coherence tomography images,” *Expert Rev. Cardiovasc. Ther.*, vol. 12, no. 7, pp. 885–900, Jul. 2014.
- [176] Feng Zhao, Xianghua Xie, and M. Roach, “Computer Vision Techniques for Transcatheter Intervention,” *IEEE J. Transl. Eng. Heal. Med.*, vol. 3, no. June, pp. 1–31, 2015.
- [177] H. S. Nam, C.-S. Kim, J. J. Lee, J. W. Song, J. W. Kim, and H. Yoo, “Automated detection of vessel lumen and stent struts in intravascular optical coherence tomography to evaluate stent apposition and neointimal coverage,” *Med. Phys.*, vol. 43, no. 4, pp. 1662–1675, 2016.
- [178] M. Bologna *et al.*, “Automatic segmentation of optical coherence tomography pullbacks of coronary arteries treated with bioresorbable vascular scaffolds: Application to hemodynamics modeling,” *PLoS One*, vol. 14, no. 3, p. e0213603, Mar. 2019.
- [179] J. M. Amrute, L. S. Athanasiou, F. Rikhtegar, J. M. de la Torre Hernández, T. G. Camarero, and E. R. Edelman, “Polymeric endovascular strut and lumen detection algorithm for intracoronary optical coherence tomography images,” *J. Biomed. Opt.*, vol. 23, no. 3, pp. 1–14, Mar. 2018.
- [180] S. Balocco *et al.*, “Standardized evaluation methodology and reference database for evaluating IVUS image segmentation,” *Comput. Med. Imaging Graph.*, vol. 38, no. 2, pp. 70–90, Mar. 2014.
- [181] A. Hoogendoorn *et al.*, “OCT-measured plaque free wall angle is indicative for plaque burden: overcoming the main limitation of OCT?,” *Int. J. Cardiovasc. Imaging*, vol. 32, pp. 1477–1481, 2016.
- [182] G. Zahnd *et al.*, “Contour segmentation of the intima, media, and adventitia layers in intracoronary OCT images: application to fully automatic detection of healthy wall regions,” *Int. J. Comput. Assist. Radiol. Surg.*, vol. 12, no. 11, pp. 1923–1936, Nov. 2017.
- [183] Z. Chen *et al.*, “Quantitative 3D Analysis of Coronary Wall Morphology in Heart Transplant Patients: OCT-Assessed Cardiac Allograft Vasculopathy Progression,” *Med. Image Anal.*, vol. 50, pp. 95–105, Dec. 2018.
- [184] W. A. Catterall, *Webvision: The Organization of the Retina and Visual System*, vol. 16, no. 1. University of Utah Health Sciences Center, 1995.
- [185] H. Ishikawa, D. M. Stein, G. Wollstein, S. Beaton, J. G. Fujimoto, and J. S. Schuman, “Macular Segmentation with Optical Coherence Tomography,” *Investig. Ophthalmology Vis. Sci.*, vol. 46, no. 6, p. 2012, Jun. 2005.
- [186] D. Cabrera Fernández, H. M. Salinas, and C. a. Puliafito, “Automated detection of retinal layer structures on optical coherence tomography images,” *Opt. Express*, vol. 13, no. 25, p. 10200, Dec. 2005.
- [187] A. M. Bagci, M. Shahidi, R. Ansari, M. Blair, N. Paul, and R. Zelkha, “Thickness Profiles of Retinal Layers by Optical Coherence Tomography Image Segmentation,” *Am. J. Ophthalmol.*, vol. 146, no. 5, pp. 679–687, 2008.

- [188] A. Mishra, A. Wong, K. Bizheva, and D. A. Clausi, “Intra-retinal layer segmentation in optical coherence tomography images,” *Opt. Express*, vol. 17, no. 26, p. 23719, Dec. 2009.
- [189] J. Duan, C. Tench, I. Gottlob, F. Proudlock, and L. Bai, “Optical Coherence Tomography Image Segmentation,” *Int. Conf. Image Process.*, pp. 4278–4282, 2015.
- [190] M. K. Garvin, M. D. Abramoff, W. Xiaodong, S. R. Russell, T. L. Burns, and M. Sonka, “Automated 3-D Intraretinal Layer Segmentation of Macular Spectral-Domain Optical Coherence Tomography Images,” *IEEE Trans. Med. Imaging*, vol. 28, no. 9, pp. 1436–1447, Sep. 2009.
- [191] R. Kafieh, H. Rabbani, M. D. Abramoff, and M. Sonka, “Intra-retinal layer segmentation of 3D optical coherence tomography using coarse grained diffusion map,” *Med. Image Anal.*, vol. 17, no. 8, pp. 907–928, Dec. 2013.
- [192] I. Amidror, “Scattered data interpolation methods for electronic imaging systems: a survey,” *J. Electron. Imaging*, vol. 11, no. 2, p. 157, Apr. 2002.
- [193] L. L. Schumaker, “Fitting surfaces to scattered data,” *Approx. Theory II*, pp. 203–68, 1976.
- [194] V. Weiss, L. Andor, G. Renner, and T. Várady, “Advanced surface fitting techniques,” *Comput. Aided Geom. Des.*, vol. 19, no. 1, pp. 19–42, Jan. 2002.
- [195] W. S. Cleveland, “Robust Locally Weighted Regression and Smoothing Scatterplots,” *J. Am. Stat. Assoc.*, vol. 74, no. 368, p. 829, Dec. 1979.
- [196] E. Keeve, S. Girod, R. Kikinis, and B. Girod, “Deformable modeling of facial tissue for craniofacial surgery simulation,” *Comput. Aided Surg.*, vol. 3, no. 5, pp. 228–238, Jan. 1998.
- [197] P. Meseure and C. Chaillou, “Deformable body simulation with adaptive subdivision and cuttings,” *Proc. WSCG*, vol. 97, pp. 361–370, 1997.
- [198] K. Tanaka, “Statistical-mechanical approach to image processing,” *J. Phys. A. Math. Gen.*, vol. 35, no. 37, pp. R81–R150, Sep. 2002.
- [199] E. Shlofmitz and A. Jeremias, “Precision percutaneous coronary intervention: Is optical coherence tomography co-registration the future?,” *Catheter. Cardiovasc. Interv.*, vol. 92, no. 1, pp. 38–39, Jul. 2018.
- [200] D. M. Leistner *et al.*, “Real-time optical coherence tomography coregistration with angiography in percutaneous coronary intervention-impact on physician decision-making: The OPTICO-integration study,” *Catheter. Cardiovasc. Interv.*, vol. 92, no. 1, pp. 30–37, Jul. 2018.
- [201] M. Michail *et al.*, “Intravascular multimodality imaging: feasibility and role in the evaluation of coronary plaque pathology,” *Eur. Hear. J. - Cardiovasc. Imaging*, vol. 18, no. 6, pp. 613–620, Jun. 2017.
- [202] L. Räber *et al.*, “Offline fusion of co-registered intravascular ultrasound and frequency domain optical coherence tomography images for the analysis of human atherosclerotic plaques,” *EuroIntervention*, vol. 8, no. 1, pp. 98–108, May 2012.
- [203] J. N. van der Sijde *et al.*, “The OPTIS Integrated System: real-time, co-registration of angiography and optical coherence tomography,” *EuroIntervention*, vol. 12, no. 7, pp. 855–860, Sep. 2016.
- [204] S. Tu, N. R. Holm, G. Koning, Z. Huang, and J. H. C. Reiber, “Fusion of 3D QCA and IVUS/OCT,” *Int. J. Cardiovasc. Imaging*, vol. 27, no. 2, pp. 197–207, Feb. 2011.

- [205] D. Rotger *et al.*, “Corresponding IVUS and angiogram image data,” in *Computers in Cardiology 2001. Vol.28 (Cat. No.01CH37287)*, 2001, pp. 273–276.
- [206] L. H. Timmins *et al.*, “Framework to Co-register Longitudinal Virtual Histology-Intravascular Ultrasound Data in the Circumferential Direction,” *IEEE Trans. Med. Imaging*, vol. 32, no. 11, pp. 1989–1996, Nov. 2013.
- [207] M. Alberti, S. Balocco, X. Carrillo, J. Mauri, and P. Radeva, “Automatic Non-rigid Temporal Alignment of Intravascular Ultrasound Sequences: Method and Quantitative Validation,” *Ultrasound Med. Biol.*, vol. 39, no. 9, pp. 1698–1712, Sep. 2013.
- [208] L. Zhang *et al.*, “Simultaneous Registration of Location and Orientation in Intravascular Ultrasound Pullbacks Pairs Via 3D Graph-Based Optimization,” *IEEE Trans. Med. Imaging*, vol. 34, no. 12, pp. 2550–2561, Dec. 2015.
- [209] D. S. Molony, L. H. Timmins, E. Rasoul-Arzumly, H. Samady, and D. P. Giddens, “Evaluation of a Framework for the Co-registration of Intravascular Ultrasound and Optical Coherence Tomography Coronary Artery Pullbacks,” *J. Biomech.*, 2016.
- [210] C. Huang *et al.*, “A New Framework for the Integrative Analytics of Intravascular Ultrasound and Optical Coherence Tomography Images,” *IEEE Access*, vol. 6, pp. 36408–36419, May 2018.
- [211] O. Pauly, G. Unal, G. Slabaugh, S. Carlier, and T. Fang, “Semi-automatic matching of OCT and IVUS images for image fusion,” in *Medical Imaging 2008: Image Processing*, 2008, vol. 6914, p. 69142N.
- [212] W. R. Crum, T. Hartkens, and D. L. G. Hill, “Non-rigid image registration: theory and practice,” *Br. J. Radiol.*, vol. 77, no. suppl_2, pp. S140–S153, Dec. 2004.
- [213] L. S. Athanasiou *et al.*, “Computerized methodology for micro-CT and histological data inflation using an IVUS based translation map,” *Comput. Biol. Med.*, vol. 65, pp. 168–176, Oct. 2015.
- [214] G. Unal, S. Lankton, S. Carlier, G. Slabaugh, and Y. Chen, “Fusion of IVUS and OCT Through Semi-Automatic Registration,” *Proc. CVII-MICCAI*, pp. 162–170, 2006.
- [215] N. Otsu, “A Threshold Selection Method from Gray-Level Histograms,” *IEEE Trans. Syst. Man. Cybern.*, vol. 9, no. 1, pp. 62–66, 1979.
- [216] C. Tomasi and R. Manduchi, “Bilateral Filtering for Gray and Color Images,” *Int. Conf. Comput. Vis.*, pp. 839–846, 1998.
- [217] J. Canny, “A Computational Approach to Edge Detection,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, Nov. 1986.
- [218] M. J. A. Girard, N. G. Strouthidis, C. R. Ethier, and J. M. Mari, “Shadow Removal and Contrast Enhancement in Optical Coherence Tomography Images of the Human Optic Nerve Head,” *Investig. Ophthalmology Vis. Sci.*, vol. 52, no. 10, p. 7738, Sep. 2011.
- [219] A. Maity, A. Pattanaik, S. Sagnika, and S. Pani, “A Comparative Study on Approaches to Speckle Noise Reduction in Images,” in *2015 International Conference on Computational Intelligence and Networks*, 2015, pp. 148–155.
- [220] I. Sobel and G. Feldman, “A 3x3 isotropic gradient operator for image processing,” in *Stanford Artificial Project*, 1968, pp. 271–272.
- [221] R. Hooke, *De Potentia Restitutiva, or of Spring: Explaining the Power of Springing Bodies*. London: The Royal Society, 1678.

- [222] G. S. Mintz, K. M. Kent, A. D. Pichard, L. F. Satler, J. J. Popma, and M. B. Leon, "Contribution of Inadequate Arterial Remodeling to the Development of Focal Coronary Artery Stenoses: An Intravascular Ultrasound Study," *Circulation*, vol. 95, no. 7, pp. 1791–1798, Apr. 1997.
- [223] P. Schoenhagen, K. M. Ziada, S. R. Kapadia, T. D. Crowe, S. E. Nissen, and E. M. Tuzcu, "Extent and Direction of Arterial Remodeling in Stable Versus Unstable Coronary Syndromes: An Intravascular Ultrasound Study," *Circulation*, vol. 101, no. 6, pp. 598–603, Feb. 2000.
- [224] M.-K. Hong *et al.*, "Intravascular ultrasound assessment of patterns of arterial remodeling in the absence of significant reference segment plaque burden in patients with coronary artery disease," *J. Am. Coll. Cardiol.*, vol. 42, no. 5, pp. 806–810, Sep. 2003.
- [225] L. S. Athanasiou *et al.*, "Optimized Computer-Aided Segmentation and Three-Dimensional Reconstruction Using Intracoronary Optical Coherence Tomography," *IEEE J. Biomed. Heal. Informatics*, vol. 22, no. 4, pp. 1168–1176, Jul. 2018.
- [226] A. M. Kok, L. Speelman, R. Virmani, A. F. W. van der Steen, F. J. H. Gijzen, and J. J. Wentzel, "Peak cap stress calculations in coronary atherosclerotic plaques with an incomplete necrotic core geometry," *Biomed. Eng. Online*, vol. 15, no. 1, p. 48, Dec. 2016.
- [227] G. Lamouche *et al.*, "Review of tissue simulating phantoms with controllable optical, mechanical and structural properties for use in optical coherence tomography," *Biomed. Opt. Express*, vol. 3, no. 6, p. 1381, Jun. 2012.
- [228] C. V. Bourantas *et al.*, "Angiographic derived endothelial shear stress: a new predictor of atherosclerotic disease progression," *Eur. Hear. J. - Cardiovasc. Imaging*, Jul. 2018.
- [229] P. Libby, "Inflammation in atherosclerosis," *Nature*, vol. 420, no. 6917, pp. 868–874, Dec. 2002.
- [230] J. F. Bentzon, F. Otsuka, R. Virmani, and E. Falk, "Mechanisms of Plaque Formation and Rupture," *Circ. Res.*, vol. 114, no. 12, pp. 1852–1866, Jun. 2014.
- [231] S. Koganti, T. Kotecha, and R. D. Rakhit, "Choice of Intracoronary Imaging: When to Use Intravascular Ultrasound or Optical Coherence Tomography," *Interv. Cardiol. Rev.*, vol. 11, no. 1, pp. 11–16, May 2016.
- [232] C. V. Bourantas *et al.*, "Hybrid intravascular imaging: current applications and prospective potential in the study of coronary atherosclerosis," *J. Am. Coll. Cardiol.*, vol. 61, no. 13, pp. 1369–78, Apr. 2013.
- [233] D. G. Vince *et al.*, "Automated coronary plaque characterization with intravascular ultrasound backscatter: In vivo and ex vivo validation," *J. Acoust. Soc. Am.*, vol. 119, no. 5, pp. 3256–3256, May 2006.
- [234] H. M. Garcia-Garcia, M. A. Costa, and P. W. Serruys, "Imaging of coronary atherosclerosis: intravascular ultrasound," *Eur. Heart J.*, vol. 31, no. 20, pp. 2456–2469, Oct. 2010.
- [235] A. Konig and V. Klauss, "Virtual histology," *Heart*, vol. 93, no. 8, pp. 977–982, Aug. 2007.
- [236] A. P. Burke, F. D. Kolodgie, A. Farb, D. Weber, and R. Virmani, "Morphological Predictors of Arterial Remodeling in Coronary Atherosclerosis," *Circulation*, vol. 105, no. 3, pp. 297–303, Jan. 2002.
- [237] H. G. Bezerra, M. A. Costa, G. Guagliumi, A. M. Rollins, and D. I. Simon, "Intracoronary Optical Coherence Tomography: A Comprehensive Review: Clinical and Research Applications," *JACC Cardiovasc. Interv.*, vol. 2, no. 11, pp. 1035–1046, Nov. 2009.

- [238] J. Hou *et al.*, “Comparison of Intensive Versus Moderate Lipid-Lowering Therapy on Fibrous Cap and Atheroma Volume of Coronary Lipid-Rich Plaque Using Serial Optical Coherence Tomography and Intravascular Ultrasound Imaging,” *Am. J. Cardiol.*, vol. 117, no. 5, pp. 800–806, Mar. 2016.
- [239] L. Xing *et al.*, “Clinical Significance of Lipid-Rich Plaque Detected by Optical Coherence Tomography,” *J. Am. Coll. Cardiol.*, vol. 69, no. 20, pp. 2502–2513, May 2017.
- [240] A. Maehara, M. Matsumura, Z. A. Ali, G. S. Mintz, and G. W. Stone, “IVUS-Guided Versus OCT-Guided Coronary Stent Implantation,” *JACC Cardiovasc. Imaging*, vol. 10, no. 12, pp. 1487–1503, Dec. 2017.
- [241] M. V Madhavan, M. Tarigopula, G. S. Mintz, A. Maehara, G. W. Stone, and P. Généreux, “Coronary artery calcification: Pathogenesis and prognostic implications,” *J. Am. Coll. Cardiol.*, vol. 63, no. 17, pp. 1703–1714, 2014.
- [242] Y. Kobayashi *et al.*, “Impact of Target Lesion Coronary Calcification on Stent Expansion,” *Circ. J.*, vol. 78, no. 9, pp. 2209–2214, Sep. 2014.
- [243] A. Fujino *et al.*, “A new optical coherence tomography-based calcium scoring system to predict stent underexpansion,” *EuroIntervention*, vol. 13, no. 18, pp. 2182–2189, Apr. 2018.
- [244] R. Hoffmann, “Treatment of calcified coronary lesions with Palmaz–Schatz stents An intravascular ultrasound study,” *Eur. Heart J.*, vol. 19, no. 8, pp. 1224–1231, Aug. 1998.
- [245] A. R. Tzafriri *et al.*, “Calcified plaque modification alters local drug delivery in the treatment of peripheral atherosclerosis,” *J. Control. Release*, vol. 264, pp. 203–210, Oct. 2017.
- [246] G. S. Mintz, “Intravascular imaging of coronary calcification and its clinical implications,” *JACC Cardiovasc. Imaging*, vol. 8, no. 4, pp. 461–471, 2015.
- [247] A. J. Brown *et al.*, “Plaque Structural Stress Estimations Improve Prediction of Future Major Adverse Cardiovascular Events after Intracoronary Imaging,” *Circ. Cardiovasc. Imaging*, vol. 9, no. 6, p. e004172, Jun. 2016.
- [248] A. C. Akyildiz *et al.*, “Effects of intima stiffness and plaque morphology on peak cap stress,” *Biomed. Eng. Online*, vol. 10, no. 1, p. 25, Apr. 2011.
- [249] E. Gerbaud *et al.*, “Multi-laboratory inter-institute reproducibility study of IVOCT and IVUS assessments using published consensus document definitions,” *Eur. Hear. J. – Cardiovasc. Imaging*, vol. 17, no. 7, pp. 756–764, Jul. 2016.
- [250] M. E. Plissiti, D. I. Fotiadis, L. K. Michalis, and G. E. Bozios, “An Automated Method for Lumen and Media–Adventitia Border Detection in a Sequence of IVUS Frames,” *IEEE Trans. Inf. Technol. Biomed.*, vol. 8, no. 2, pp. 131–141, Jun. 2004.
- [251] A. Moshfegh, A. Javadzadegan, M. Mohammadi, L. Ravipudi, S. Cheng, and R. Martins, “Development of an innovative technology to segment luminal borders of intravascular ultrasound image sequences in a fully automated manner,” *Comput. Biol. Med.*, vol. 108, pp. 111–121, May 2019.
- [252] Y. Wang, C. Qiu, J. Jiang, and S. Xia, “Detecting the Media-adventitia Border in Intravascular Ultrasound Images through a Classification-based Approach,” *Ultrason. Imaging*, vol. 41, no. 2, pp. 78–93, Mar. 2019.

- [253] L. Lo Vercio, M. del Fresno, and I. Larrabide, “Lumen-intima and media-adventitia segmentation in IVUS images using supervised classifications of arterial layers and morphological structures,” *Comput. Methods Programs Biomed.*, vol. 177, pp. 113–121, Aug. 2019.
- [254] S. Kim, Y. Jang, B. Jeon, Y. Hong, H. Shim, and H. Chang, “Fully Automatic Segmentation of Coronary Arteries Based on Deep Neural Network in Intravascular Ultrasound Images,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11043 LNCS, 2018, pp. 161–168.
- [255] L. S. Athanasiou, D. I. Fotiadis, and L. K. Michalis, *Atherosclerotic Plaque Characterization Methods Based on Coronary Imaging*. Academic Press, 2017.
- [256] L. S. Athanasiou *et al.*, “Fully automated calcium detection using optical coherence tomography,” in *2013 35th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 2013, vol. 2013, pp. 1430–1433.
- [257] D. A. Borges Oliveira, P. Nicz, C. Campos, P. Lemos, M. M. G. Macedo, and M. A. Gutierrez, “Coronary calcification identification in optical coherence tomography using convolutional neural networks,” in *Medical Imaging 2018: Biomedical Applications in Molecular, Structural, and Functional Imaging*, 2018, vol. 10578, no. 12, p. 69.
- [258] S. He *et al.*, “Convolutional neural network based automatic plaque characterization for intracoronary optical coherence tomography images,” in *Medical Imaging 2018: Image Processing*, 2018, vol. 1057432, no. 2, p. 107.
- [259] C. Xu, J. M. Schmitt, S. G. Carlier, and R. Virmani, “Characterization of atherosclerosis plaques by measuring both backscattering and attenuation coefficients in optical coherence tomography,” *J. Biomed. Opt.*, vol. 13, no. 3, p. 034003, 2008.
- [260] G. van Soest *et al.*, “Atherosclerotic tissue characterization in vivo by optical coherence tomography attenuation imaging,” *J. Biomed. Opt.*, vol. 15, no. 1, p. 011105, 2010.
- [261] C. Kolluru, D. Prabhu, Y. Gharaibeh, H. Bezerra, G. Guagliumi, and D. Wilson, “Deep neural networks for A-line-based plaque classification in coronary intravascular optical coherence tomography images,” *J. Med. Imaging*, vol. 5, no. 04, p. 1, Dec. 2018.
- [262] N. Gessert *et al.*, “Automatic Plaque Detection in IVOCT Pullbacks Using Convolutional Neural Networks,” *IEEE Trans. Med. Imaging*, vol. 38, no. 2, pp. 426–434, Feb. 2019.
- [263] C. Zhang *et al.*, “Convolution Neural Networks and Support Vector Machines for Automatic Segmentation of Intracoronary Optical Coherence Tomography,” *Mol. Cell. Biomech.*, vol. 16, no. 2, pp. 153–161, 2019.
- [264] X. Zhang, C. R. McKay, and M. Sonka, “Tissue characterization in intravascular ultrasound images,” *IEEE Trans. Med. Imaging*, vol. 17, no. 6, pp. 889–899, Dec. 1998.
- [265] E. Brunenberg, O. Pujol, B. ter Haar Romeny, and P. Radeva, “Automatic IVUS segmentation of atherosclerotic plaque with stop & go snake,” *Med. Image Comput. Comput. Assist. Interv.*, vol. 9, no. Pt 2, pp. 9–16, 2006.
- [266] L. S. Athanasiou *et al.*, “A hybrid plaque characterization method using intravascular ultrasound images,” *Technol. Heal. Care*, vol. 21, no. 3, pp. 199–216, Jun. 2013.
- [267] A. Taki, A. Roodaki, O. Pauly, S. K. Setarehdan, G. Unal, and N. Navab, “A new method for characterization of coronary plaque composition via IVUS images,” in *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2009, pp. 787–790.

- [268] A. Taki *et al.*, “A New Approach for Improving Coronary Plaque Component Analysis Based on Intravascular Ultrasound Images,” *Ultrasound Med. Biol.*, vol. 36, no. 8, pp. 1245–1258, Aug. 2010.
- [269] A. Taki, A. Roodaki, S. K. Setarehdan, S. Avansari, G. Unal, and N. Navab, “An IVUS image-based approach for improvement of coronary plaque characterization,” *Comput. Biol. Med.*, vol. 43, no. 4, pp. 268–280, May 2013.
- [270] L. S. Athanasiou *et al.*, “A novel semiautomated atherosclerotic plaque characterization method using grayscale intravascular ultrasound images: Comparison with virtual histology,” *IEEE Trans. Inf. Technol. Biomed.*, vol. 16, no. 3, pp. 391–400, May 2012.
- [271] Y. N. Hwang, J. H. Lee, G. Y. Kim, E. S. Shin, and S. M. Kim, “Characterization of coronary plaque regions in intravascular ultrasound images using a hybrid ensemble classifier,” *Comput. Methods Programs Biomed.*, vol. 153, pp. 83–92, Jan. 2018.
- [272] G. Y. Kim, J. H. Lee, Y. N. Hwang, and S. M. Kim, “A novel intensity-based multi-level classification approach for coronary plaque characterization in intravascular ultrasound images,” *Biomed. Eng. Online*, vol. 17, no. S2, p. 151, Nov. 2018.
- [273] A. Nair, B. D. Kuban, E. M. Tuzcu, P. Schoenhagen, S. E. Nissen, and D. G. Vince, “Coronary Plaque Classification With Intravascular Ultrasound Radiofrequency Data Analysis,” *Circulation*, vol. 106, no. 17, pp. 2200–2206, Oct. 2002.
- [274] K. Nasu *et al.*, “Accuracy of In Vivo Coronary Plaque Morphology Assessment,” *J. Am. Coll. Cardiol.*, vol. 47, no. 12, pp. 2405–2412, Jun. 2006.
- [275] A. Nair, M. P. Margolis, B. D. Kuban, and D. G. Vince, “Automated coronary plaque characterisation with intravascular ultrasound backscatter: ex vivo validation,” *EuroIntervention*, vol. 3, no. 1, pp. 113–20, May 2007.
- [276] J. Van Herck, G. De Meyer, G. Ennekens, P. Van Herck, A. Herman, and C. Vrints, “Validation of in vivo plaque characterisation by virtual histology in a rabbit model of atherosclerosis,” *EuroIntervention*, vol. 5, no. 1, pp. 149–156, May 2009.
- [277] H. M. García-García *et al.*, “Tissue characterisation using intravascular radiofrequency data analysis: recommendations for acquisition, analysis, interpretation and reporting,” *EuroIntervention*, vol. 5, no. 2, pp. 177–189, Jun. 2009.
- [278] S. Balocco, M. González, R. Nanculef, P. Radeva, and G. Thomas, “Calcified Plaque Detection in IVUS Sequences: Preliminary Results Using Convolutional Nets,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11047 LNCS, 2018, pp. 34–42.
- [279] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA: MIT Press, 2016.
- [280] I. Sutskever, J. Martens, G. Dahl, and G. Hinton, “On the importance of initialization and momentum in deep learning,” in *30th International Conference on Machine Learning, ICML 2013*, 2013, no. PART 3, pp. 2176–2184.
- [281] P. K. Siogkas, K. A. Stefanou, L. S. Athanasiou, M. I. Papafaklis, L. K. Michalis, and D. I. Fotiadis, “Art care: A multi-modality coronary 3D reconstruction and hemodynamic status assessment software,” *Technol. Heal. Care*, vol. 26, no. 1, pp. 187–193, Mar. 2018.
- [282] H.-F. Ng, “Automatic thresholding for defect detection,” *Pattern Recognit. Lett.*, vol. 27, no. 14, pp. 1644–1649, Oct. 2006.

- [283] M. Kass, A. Witkin, and D. Terzopoulos, “Snakes: Active contour models,” *Int. J. Comput. Vis.*, vol. 1, no. 4, pp. 321–331, Jan. 1988.
- [284] L. S. Athanasiou *et al.*, “Error propagation in the characterization of atheromatic plaque types based on imaging,” *Comput. Methods Programs Biomed.*, vol. 121, no. 3, pp. 161–174, Oct. 2015.
- [285] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9351, Springer, Cham, 2015, pp. 234–241.
- [286] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-Image Translation with Conditional Adversarial Networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5967–5976.
- [287] I. Goodfellow *et al.*, “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, 2014.
- [288] V. Monga, S. T. Acton, A.-K. Seghouane, A. Munoz-Barrutia, and J. C. Ye, “Editorial: Introduction to the Issue on Domain Enriched Learning for Medical Imaging,” *IEEE J. Sel. Top. Signal Process.*, vol. 14, no. 6, pp. 1068–1071, Oct. 2020.
- [289] R. M. Haralick, K. Shanmugam, and I. Dinstein, “Textural Features for Image Classification,” *IEEE Trans. Syst. Man. Cybern.*, vol. SMC-3, no. 6, pp. 610–621, Nov. 1973.
- [290] K. W. Gossage, T. S. Tkaczyk, J. J. Rodriguez, and J. K. Barton, “Texture analysis of optical coherence tomography images: feasibility for tissue classification,” *J. Biomed. Opt.*, vol. 8, no. 3, p. 570, Sep. 2003.
- [291] G. Beliakov, S. James, and L. Troiano, “Texture recognition by using GLCM and various aggregation functions,” in *2008 IEEE International Conference on Fuzzy Systems (IEEE World Congress on Computational Intelligence)*, 2008, pp. 1472–1476.
- [292] L. Wexler *et al.*, “Coronary Artery Calcification: Pathophysiology, Epidemiology, Imaging Methods, and Clinical Implications,” *Circulation*, vol. 94, no. 5, pp. 1175–1192, Sep. 1996.
- [293] J. C. Teo *et al.*, “Optimization of coronary optical coherence tomography imaging using the attenuation-compensated technique: a validation study,” *Eur. Hear. J. – Cardiovasc. Imaging*, vol. 18, no. 8, p. jew153, Jul. 2016.
- [294] A. de Vecchi *et al.*, “Catheter-Induced Errors in Pressure Measurements in Vessels: An In-Vitro and Numerical Study,” *IEEE Trans. Biomed. Eng.*, vol. 61, no. 6, pp. 1844–1850, Jun. 2014.
- [295] TIMI Study Group, “The Thrombolysis in Myocardial Infarction (TIMI) Trial,” *N. Engl. J. Med.*, vol. 312, no. 14, pp. 932–936, Apr. 1985.
- [296] C. M. Gibson *et al.*, “TIMI Frame Count,” *Circulation*, vol. 93, no. 5, pp. 879–888, Mar. 1996.
- [297] G. A. Holzapfel, G. Sommer, and P. Regitnig, “Anisotropic Mechanical Properties of Tissue Components in Human Atherosclerotic Plaques,” *J. Biomech. Eng.*, vol. 126, no. 5, p. 657, 2004.
- [298] P. D. Richardson, M. J. Davies, and G. V. R. Born, “Influence of Plaque Configuration and Stress Distribution on Fissuring of Coronary Atherosclerotic Plaques,” *Lancet*, vol. 334, no. 8669, pp. 941–944, Oct. 1989.

- [299] G. C. Cheng, H. M. Loree, R. D. Kamm, M. C. Fishbein, and R. T. Lee, “Distribution of Circumferential Stress in Ruptured and Stable Atherosclerotic Lesions A Structural Analysis With Histopathological Correlation,” *Circulation*, vol. 87, pp. 1179–1187, 1993.
- [300] H. Huang, R. Virmani, H. Younis, A. P. Burke, R. D. Kamm, and R. T. Lee, “The Impact of Calcification on the Biomechanical Stability of Atherosclerotic Plaques,” *Circulation*, vol. 103, no. 8, pp. 1051–1056, Feb. 2001.
- [301] R. T. Lee, H. M. Loree, G. C. Cheng, E. H. Lieberman, N. Jaramillo, and F. J. Schoen, “Computational structural analysis based on intravascular ultrasound imaging before in vitro angioplasty: Prediction of plaque fracture locations,” *J. Am. Coll. Cardiol.*, vol. 21, no. 3, pp. 777–782, Mar. 1993.
- [302] A. H. Chau *et al.*, “Mechanical Analysis of Atherosclerotic Plaques Based on Optical Coherence Tomography,” *Ann. Biomed. Eng.*, vol. 32, no. 11, pp. 1494–1503, Nov. 2004.
- [303] Y. Huang, Z. Teng, U. Sadat, M. J. Graves, M. R. Bennett, and J. H. Gillard, “The influence of computational strategy on prediction of mechanical stress in carotid atherosclerotic plaques: Comparison of 2D structure-only, 3D structure-only, one-way and fully coupled fluid-structure interaction analyses,” *J. Biomech.*, vol. 47, no. 6, pp. 1465–1471, Apr. 2014.
- [304] F. J. H. Gijssen and F. Migliavacca, “Plaque mechanics,” *J. Biomech.*, vol. 47, no. 4, pp. 763–764, 2014.
- [305] A. C. Akyildiz *et al.*, “The effects of plaque morphology and material properties on peak cap stress in human coronary arteries,” *Comput. Methods Biomed. Biomed. Engin.*, vol. 19, no. 7, pp. 771–779, May 2016.
- [306] J. Yuan *et al.*, “Influence of material property variability on the mechanical behaviour of carotid atherosclerotic plaques: A 3D fluid-structure interaction analysis,” *Int. j. numer. method. biomed. eng.*, vol. 31, no. 8, pp. 1–12, Aug. 2015.
- [307] L. M. Riou *et al.*, “Effects of mechanical properties and atherosclerotic artery size on biomechanical plaque disruption – Mouse vs. human,” *J. Biomech.*, vol. 47, no. 4, pp. 765–772, Mar. 2014.
- [308] G. J. Tearney and B. E. Bouma, “Atherosclerotic plaque characterization by spatial and temporal speckle pattern analysis,” in *Technical Digest. Summaries of papers presented at the Conference on Lasers and Electro-Optics. Postconference Technical Digest (IEEE Cat. No.01CH37170)*, 2001, pp. 307–308.
- [309] G. J. Tearney and B. E. Bouma, “Atherosclerotic plaque characterization by spatial and temporal speckle pattern analysis,” *Opt. Lett.*, vol. 27, no. 7, p. 533, Apr. 2002.
- [310] E. Maksuti, E. Widman, D. Larsson, M. W. Urban, M. Larsson, and A. Bjällmark, “Arterial Stiffness Estimation by Shear Wave Elastography: Validation in Phantoms with Mechanical Testing,” *Ultrasound Med. Biol.*, vol. 42, no. 1, pp. 308–321, Jan. 2016.
- [311] D. Larsson, J. Roy, T. C. Gasser, M. W. Urban, M. Colarieti-Tosti, and M. Larsson, “An ex-vivo setup for characterization of atherosclerotic plaque using shear wave elastography and micro-computed tomography,” in *2016 IEEE International Ultrasonics Symposium (IUS)*, 2016, vol. 2016-Novem, pp. 1–4.
- [312] D. Marlevi, E. Maksuti, M. W. Urban, R. Winter, and M. Larsson, “Plaque characterization using shear wave elastography—evaluation of differentiability and accuracy using a combined ex vivo and in vitro setup,” *Phys. Med. Biol.*, vol. 63, no. 23, p. 235008, Nov. 2018.

- [313] C. L. de Korte, S. Fekkes, A. J. Nederveen, R. Manniesing, and H. R. H. G. Hansen, "Review: Mechanical Characterization of Carotid Arteries and Atherosclerotic Plaques," *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 63, no. 10, pp. 1613–1623, Oct. 2016.
- [314] C. L. de Korte, H. A. Woutman, A. F. W. van der Steen, G. Pasterkamp, and E. I. Céspedes, "Vascular tissue characterisation with IVUS elastography," *Ultrasonics*, vol. 38, no. 1–8, pp. 387–390, Mar. 2000.
- [315] C. L. de Korte, G. Pasterkamp, A. F. W. van der Steen, H. A. Woutman, and N. Bom, "Characterization of Plaque Components With Intravascular Ultrasound Elastography in Human Femoral and Coronary Arteries In Vitro," *Circulation*, vol. 102, no. 6, pp. 617–623, Aug. 2000.
- [316] J. A. Schaar *et al.*, "Characterizing Vulnerable Plaque Features With Intravascular Elastography," *Circulation*, vol. 108, no. 21, pp. 2636–2641, Nov. 2003.
- [317] Y. Majdoulina *et al.*, "Endovascular Shear Strain Elastography for the Detection and Characterization of the Severity of Atherosclerotic Plaques: In Vitro Validation and In Vivo Evaluation," *Ultrasound Med. Biol.*, vol. 40, no. 5, pp. 890–903, May 2014.
- [318] K. V. Larin and D. D. Sampson, "Optical coherence elastography – OCT at work in tissue biomechanics [Invited]," *Biomed. Opt. Express*, vol. 8, no. 2, p. 1172, Feb. 2017.
- [319] G. van Soest, F. Mastik, N. de Jong, and A. F. W. van der Steen, "Robust intravascular optical coherence elastography by line correlations," *Phys. Med. Biol.*, vol. 52, no. 9, pp. 2445–2458, May 2007.
- [320] A. S. Khalil, R. C. Chan, A. H. Chau, B. E. Bouma, and M. R. Kaazempur-Mofrad, "Tissue Elasticity Estimation with Optical Coherence Elastography: Toward Mechanical Characterization of In Vivo Soft Tissue," *Ann. Biomed. Eng.*, vol. 33, no. 11, pp. 1631–1639, Nov. 2005.
- [321] R. Karimi, T. Zhu, B. E. Bouma, and M. R. Kaazempur-Mofrad, "Estimation of Nonlinear Mechanical Properties of Vascular Tissues via Elastography," *Cardiovasc. Eng.*, vol. 8, no. 4, pp. 191–202, Dec. 2008.
- [322] J. M. Schmitt, "OCT elastography: imaging microscopic deformation and strain of tissue," *Opt. Express*, vol. 3, no. 6, pp. 199–211, Sep. 1998.
- [323] R. K. Wang, Z. Ma, and S. J. Kirkpatrick, "Tissue Doppler optical coherence elastography for real time strain rate and strain mapping of soft tissue," *Appl. Phys. Lett.*, vol. 89, no. 14, p. 144103, Oct. 2006.
- [324] R. K. Wang, S. J. Kirkpatrick, and M. Hinds, "Phase-sensitive optical coherence elastography for mapping tissue microstrains in real time," *Appl. Phys. Lett.*, vol. 90, no. 16, p. 164105, Apr. 2007.
- [325] J. Rogowska, N. A. Patel, J. G. Fujimoto, and M. E. Brezinski, "Optical coherence tomographic elastography technique for measuring deformation and strain of atherosclerotic tissues," *Heart*, vol. 90, no. 5, pp. 556–62, May 2004.
- [326] J. Rogowska, N. Patel, S. Plummer, and M. E. Brezinski, "Quantitative optical coherence tomographic elastography: method for assessing arterial mechanical properties," *Br. J. Radiol.*, vol. 79, no. 945, pp. 707–711, Sep. 2006.
- [327] E. I. Céspedes, C. L. de Korte, and A. F. W. van der Steen, "Intraluminal ultrasonic palpation: assessment of local and cross-sectional tissue stiffness," *Ultrasound Med. Biol.*, vol. 26, no. 3, pp. 385–396, Mar. 2000.

- [328] J. Poree, B. Chayer, G. Soulez, J. Ohayon, and G. Cloutier, “Noninvasive Vascular Modulography Method for Imaging the Local Elasticity of Atherosclerotic Plaques: Simulation and In Vitro Vessel Phantom Study,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 64, no. 12, pp. 1805–1817, Dec. 2017.
- [329] R. Karimi, R. Chan, S. Houser, B. E. Bouma, and M. R. Kaazempur Mofrad, “A Novel Framework for Elastography and Modulus Estimation: Integration of Tissue Mechanics with Imaging,” in *3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano, 2006*, 2006, pp. 602–605.
- [330] L. Dong *et al.*, “Quantitative Compression Optical Coherence Elastography as an Inverse Elasticity Problem,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 22, no. 3, pp. 277–287, May 2016.
- [331] A. C. Akyildiz *et al.*, “A Framework for Local Mechanical Characterization of Atherosclerotic Plaques: Combination of Ultrasound Displacement Imaging and Inverse Finite Element Analysis,” *Ann. Biomed. Eng.*, vol. 44, no. 4, pp. 968–979, Apr. 2016.
- [332] J. Rogowska, N. Patel, J. G. Fujimoto, and M. E. Brezinski, “OCT elastography of the vascular tissue – importance of cross-correlation kernel size,” in *Biomedical Topical Meeting*, 2002, p. PD20.
- [333] C. L. De Korte, E. I. Cespedes, and A. F. W. Van Der Steen, “Influence of catheter position on estimated strain in intravascular elastography,” *IEEE Trans. Ultrason. Ferroelectr. Freq. Control*, vol. 46, no. 3, pp. 616–625, May 1999.
- [334] C. Noble *et al.*, “Patient specific characterization of artery and plaque material properties in peripheral artery disease,” *J. Mech. Behav. Biomed. Mater.*, vol. 101, p. 103453, Jan. 2020.
- [335] B. Narayanan, “Image-based Mechanical Characterisation of Atherosclerotic Plaques - An Inverse Finite Element Approach,” Ecole polytechnique fédérale de Lausanne, 2020.
- [336] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia, “MeshLab: An open-source mesh processing tool,” in *6th Eurographics Italian Chapter Conference 2008 - Proceedings*, 2008, pp. 129–136.
- [337] C. Geuzaine and J.-F. Remacle, “Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities,” *Int. J. Numer. Methods Eng.*, vol. 79, no. 11, pp. 1309–1331, Sep. 2009.
- [338] ASTM International, “D638-14 Standard Test Method for Tensile Properties of Plastics.” ASTM International, West Conshohocken, PA, 2014.
- [339] J. W. Holmes and J. Lumens, “Clinical Applications of Patient-Specific Models: The Case for a Simple Approach,” *J. Cardiovasc. Transl. Res.*, vol. 11, no. 2, pp. 71–79, Apr. 2018.
- [340] P. Libby and G. Pasterkamp, “Requiem for the ‘vulnerable plaque,’” *Eur. Heart J.*, vol. 36, no. 43, pp. 2984–2987, Jul. 2015.
- [341] A. Badano *et al.*, “Evaluation of Digital Breast Tomosynthesis as Replacement of Full-Field Digital Mammography Using an In Silico Imaging Trial,” *JAMA Netw. Open*, vol. 1, no. 7, p. e185474, Nov. 2018.
- [342] W. P. Segars, M. Mahesh, T. J. Beck, E. C. Frey, and B. M. W. Tsui, “Realistic CT simulation using the 4D XCAT phantom,” *Med. Phys.*, vol. 35, no. 8, pp. 3800–3808, Jul. 2008.
- [343] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, “4D XCAT phantom for multimodality imaging research,” *Med. Phys.*, vol. 37, no. 9, pp. 4902–4915, Aug. 2010.
- [344] B. Stimpel *et al.*, “Projection image-to-image translation in hybrid x-ray/MR imaging,” in *Medical Imaging 2019: Image Processing*, 2019, vol. 10949, p. 90.

- [345] W. Yang *et al.*, “Predicting CT Image From MRI Data Through Feature Matching With Learned Nonlinear Local Descriptors,” *IEEE Trans. Med. Imaging*, vol. 37, no. 4, pp. 977–987, Apr. 2018.
- [346] D. Nie *et al.*, “Medical Image Synthesis with Deep Convolutional Adversarial Networks,” *IEEE Trans. Biomed. Eng.*, vol. 65, no. 12, pp. 2720–2730, Dec. 2018.
- [347] H.-C. Shin *et al.*, “Medical Image Synthesis for Data Augmentation and Anonymization Using Generative Adversarial Networks,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 11037 LNCS, 2018, pp. 1–11.
- [348] F. Tom and D. Sheet, “Simulating patho-realistic ultrasound images using deep generative networks with adversarial learning,” in *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, 2018, vol. 2018-April, pp. 1174–1177.
- [349] J. Lee *et al.*, “Automated plaque characterization using deep learning on coronary intravascular optical coherence tomographic images,” *Biomed. Opt. Express*, vol. 10, no. 12, p. 6497, Dec. 2019.
- [350] **E. Ben Assa, J. de la Torre Hernandez, J. Brown, M. Olender, et al.**, “**TCT-299 Pulse Wave Velocity and Aortic Distensibility in Patients with Hypertensive Response Post Transcatheter Aortic Valve Replacement,**” *J. Am. Coll. Cardiol.*, vol. 72, no. 13, p. B123, Sep. 2018.
- [351] **E. Ben-Assa, J. Brown, Z. Keshavarz-Motamed, J. M. de la Torre Hernandez, B. Leiden, M. Olender, et al.**, “**Ventricular stroke work and vascular impedance refine the characterization of patients with aortic stenosis,**” *Sci. Transl. Med.*, vol. 11, no. 509, p. eaaw0181, Sep. 2019.
- [352] M. R. Dweck *et al.*, “Contemporary rationale for non-invasive imaging of adverse coronary plaque features to identify the vulnerable patient: a Position Paper from the European Society of Cardiology Working Group on Atherosclerosis and Vascular Biology and the European Associa,” *Eur. Hear. J. - Cardiovasc. Imaging*, vol. 0, pp. 1–7, Sep. 2020.
- [353] A. Arbab-Zadeh and V. Fuster, “From Detecting the Vulnerable Plaque to Managing the Vulnerable Patient: JACC State-of-the-Art Review,” *Journal of the American College of Cardiology*, vol. 74, no. 12. Elsevier USA, pp. 1582–1593, 24-Sep-2019.