
Theses and Dissertations

Spring 2013

Predictive analysis of coronary plaque morphology and composition on a one year timescale

Richard Wesley Downe
University of Iowa

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Electrical and Computer Engineering Commons](#)

Copyright © 2013 Richard Wesley Downe

This dissertation is available at Iowa Research Online: <https://ir.uiowa.edu/etd/2479>

Recommended Citation

Downe, Richard Wesley. "Predictive analysis of coronary plaque morphology and composition on a one year timescale." PhD (Doctor of Philosophy) thesis, University of Iowa, 2013.

<https://doi.org/10.17077/etd.ynig7tzu>

Follow this and additional works at: <https://ir.uiowa.edu/etd>



Part of the [Electrical and Computer Engineering Commons](#)

PREDICTIVE ANALYSIS OF CORONARY PLAQUE MORPHOLOGY AND
COMPOSITION ON A ONE YEAR TIMESCALE

by

Richard Wesley Downe

An Abstract

Of a thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy degree
in Electrical and Computer Engineering
in the Graduate College of
The University of Iowa

May 2013

Thesis Supervisor: Professor Milan Sonka

ABSTRACT

Coronary artery disease is a leading cause of death in the Western world. Symptoms present only late in the progression of the disease, limiting treatment options; moreover, the inability to biopsy arterial tissue in a living patient makes it difficult to study the pathology effectively. 89 patients were imaged twice at a one year interval using x-ray angiography (the traditional modality for assessment of arterial stenosis) and intravascular ultrasound (IVUS), which yields a detailed image of the structure of the vessel wall. 32 of these 89 patients were made available for analysis in this study.

The Volcano Corp. IVUS acquisition systems include software that provides a *virtual histology* (VH) characterization of plaque composition that provides information otherwise only available by biopsy. Using a geometric reconstruction method described in previous work, a full working model of wall shear stresses (WSS) produced by blood flow and vessel wall composition is created. Using these, the morphologic structural information gleaned from the 3-D reconstruction, and some additional composite indices, combined with demographic information collected at enrollment and serum biomarkers collected from each patient during imaging, a detailed portrait of each patient's disease state is created, with the objective of predicting disease evolution over a 1 year timescale.

We have, in the course of this study, accomplished the following 5 aims towards the goal of predicting localized changes in disease state on a 1 year timescale: Aim 1: Develop and validate a method of compensating for rotational motion of the catheter within the vessel and its effect upon the 3-D orientation of the reconstruction. Aim 2: Develop and validate a method of registering the reconstructed vessels that permits identification of a point-to-point correspondence on all quantitative indices. Aim 3: Successfully reconstruct, register, and analyze image sets for each of as many patients as possible for analysis. Aim 4: Identify statistically significant indices in the data

suitable for use as features in a classifier. Aim 5: Construct and assess performance of a classification system that can draw useful conclusions about the 1-year progression of the arterial pathology in a patient not used in the training set.

Aim 2 was a complete success. Branches were reliably present in the IVUS data in sufficient quantities to facilitate reliable identification of the overlap and the requisite spatial transformation required to map points from one pullback onto another.

Aim 1 was much more problematic. While a method was developed which showed promise, the image acquisition protocol did not provide for orientation of the angiograms with an eye towards bifurcation identification. With neither an analytic model, nor reasonable fiducials, the method developed could only be tested on a small subset of the data, limiting both our confidence in its validation, as well as its usability in this study. It is hoped that the method can be refined and used in any subsequent study, given proper planning during the acquisition of the images, and that in turn the spatial reliability of the reconstructions can be improved beyond what is possible today.

Regarding aim 3, 32 patients were ultimately processed completely.

Aims 4 and 5 were completed successfully. Meaningful correlations were identified in the data, and the results illustrate that while we were by no means able to obtain perfect classification, we were able to handily beat a both a random, and a maximum likelihood classifier.

Abstract Approved: _____

Thesis Supervisor

Title and Department

Date

PREDICTIVE ANALYSIS OF CORONARY PLAQUE MORPHOLOGY AND
COMPOSITION ON A ONE YEAR TIMESCALE

by

Richard Wesley Downe

A thesis submitted in partial fulfillment of the
requirements for the Doctor of Philosophy degree
in Electrical and Computer Engineering
in the Graduate College of
The University of Iowa

May 2013

Thesis Supervisor: Professor Milan Sonka

Graduate College
The University of Iowa
Iowa City, Iowa

CERTIFICATE OF APPROVAL

PH.D. THESIS

This is to certify that the Ph.D. thesis of

Richard Wesley Downe

has been approved by the Examining Committee for the thesis requirement for the Doctor of Philosophy degree in Electrical and Computer Engineering at the May 2013 graduation.

Thesis Committee: _____

Milan Sonka, Thesis Supervisor

Andreas Wahle

Reinhard R. Beichel

Joseph Reinhardt

Mona Garvin

TABLE OF CONTENTS

LIST OF TABLES	iv
LIST OF FIGURES	v
LIST OF ALGORITHMS	vii
CHAPTER	
1 OVERVIEW	1
1.1 Background	1
1.2 Project History	3
1.3 Aims	4
1.3.1 New work supporting aims	4
1.4 Analysis Pipeline Description	6
1.5 Patient Data Description	9
1.5.1 Optical Gating	11
1.6 Similar Work	13
2 BIFURCATION MODELING	14
2.1 Methods	14
2.1.1 Trajectory Estimation	17
2.1.2 Diameter Estimation	17
2.1.3 Mesh Merging	20
2.2 Validation	24
3 REGISTRATION	29
3.1 Baseline/Followup Registration	30
3.1.1 Methods	30
3.1.2 Experimental Methods	34
3.1.3 Results	36
3.2 Catheter Twist Correction	39
3.2.1 Phantom Study	39
3.2.2 Methods	43
3.2.3 Experimental Methods	46
3.2.4 Results	50
4 DATA ANALYSIS	56
4.1 Data Model	56
4.2 Quantitative Indices	57
4.2.1 Methods	57

4.2.2	List of Indices	66
5	CLASSIFICATION SYSTEM	69
5.1	Preprocessing	70
5.1.1	Exclusion Marking	70
5.1.2	Globbering	72
5.1.3	Imputation of Missing Values	73
5.1.4	Output Binning	74
5.1.5	Resampling	77
5.2	Feature Selection	78
5.2.1	Methods	78
5.2.2	Experimental Methods	80
5.2.3	Results	80
5.3	Classifier	83
5.3.1	Algorithm overview	84
5.3.2	Methods	86
5.3.3	Experimental Methods	87
5.3.4	Results	90
6	DISCUSSION	101
6.1	Interpretation of results	101
6.2	Sources of uncertainty	103
6.3	Future lines of inquiry	105
6.3.1	Acquisition protocol	105
6.3.2	3-D reconstruction	106
6.3.3	TCFA analysis	107
7	CONCLUSION	109
	REFERENCES	112
	APPENDIX: SCHEMA DIAGRAM	119

LIST OF TABLES

Table

3.1	Registration error statistics	36
3.2	Detwisting error statistics.	51
4.1	List of indices used in classifier	68
5.1	Classifier output variables	70
5.2	Glob sizes used in preprocessing.	72
5.3	Feature selection summary	82
5.4	Classifier styles	84
5.5	Grid describing class performance	87
5.6	Statistical power measures per output feature	89
5.7	Average in-class accuracy	91
5.8	Confusion matrices for 5-way output	93
5.9	Best performers, 5 class output	94
5.10	Best performers, 3 class output	97
5.11	Confusion matrices for 3-way output	98
5.12	Best performers, 5 class output, with SFFS	99
5.13	Confusion matrices for 5-way output, with SFFS	100

LIST OF FIGURES

Figure

1.1	Data processing flowchart	7
1.2	Explanation of vessel anatomy from the point of view of IVUS.	9
1.3	Comparison of motion blur versus ECG gating.	12
2.1	Interface for manual refinement in branch capture program.	16
2.2	Output of 2-D mean-shift segmentation used to estimate branch diameter.	17
2.3	Branch junctions, after different merging procedures.	18
2.4	Positioning and radius error in branch identification.	27
2.5	Illustration of branch impact on wall shear stress	28
3.1	Before/after for baseline/followup registration.	31
3.2	Density curves for lumen-lamina border registration error.	37
3.3	Density curves for media-adventitia border registration error.	38
3.4	Illustration of twisting artifact	40
3.5	Phantoms used to verify catheter twisting artifact.	41
3.6	Images from phantom study.	42
3.7	Software interface used to identify branches in angiograms.	44
3.8	Application used for 3-D reconstructions	46
3.9	Curvature index error, measured as a function of rotational misalignment.	48
3.10	Density plot for curvature index error taken over the entire data set.	49
3.11	Curvature index density plot	52
3.12	Correction angle shown alongside reliability weight, baseline reconstruction.	54
3.13	Correction angle shown alongside reliability weight, followup reconstruction.	55
4.1	Wall shear stress, baseline versus followup	62

4.2	Phantom image used to validate connected component labeling.	65
5.1	Tracing discrepancy between baseline and followup.	70
5.2	Output feature density plots.	76
5.3	Flowchart describing feature selection procedure.	79
5.4	Inter-feature correlations for all inputs and outputs	81
5.5	Overview of classifier performance	92
5.6	Inter-patient variability, box and whisker plots.	95
6.1	Slight border mismatch in the VH application due to lack of synchronization.104	
A.1	<i>Root</i> portion of the database schema hierarchy.	120
A.2	<i>IVUS</i> portion of the database schema hierarchy.	121
A.3	<i>Angiogram</i> portion of the database schema hierarchy.	122
A.4	<i>Fusion</i> portion of the database schema hierarchy.	123

LIST OF ALGORITHMS

Algorithm

1	Procedure for removing endcaps from triangulated surface mesh.	24
2	Procedure for identifying control points for registration spline.	34
3	Procedure for estimating localized twisting from bifurcations.	45
4	J48 pseudocode	85
5	Random forests pseudocode	85

CHAPTER 1 OVERVIEW

1.1 Background

At its heart this work is motivated by our efforts to develop and validate a method to characterize the progression of plaque formations as related to atherosclerotic coronary artery disease (CAD) in-vivo.

Coronary atherosclerosis starts at a young age and advances until occlusive coronary heart disease causes chronic symptoms such as angina, or the rupture of vulnerable plaque causes acute symptoms, up to and including myocardial infarction or sudden death. Clinical prognosis for patients with coronary atherosclerosis is directly related to the progression of the disease, thereby making the methods for monitoring disease progression and regression crucial in both the detection and the treatment of the disease. Emerging evidence of the initiation and progression of coronary atherosclerosis reveals a complex process where genetic factors interact with inflammatory, thrombotic, and other systemic factors, as well as vascular mechanics and hemodynamics.

Modalities currently in use for assessing coronary artery disease in vivo include x-ray angiography (the oldest tool for assessing stenosis) [1], intravascular ultrasound (IVUS) [2], computed tomography (CT) [3], and optical coherence tomography (OCT) [4]. CT provides the most geometrically/spatially accurate portrait of the vasculature by virtue of being a true 3-D imaging modality, however, it lacks the spatial resolution of IVUS and OCT, and like traditional x-ray angiography, is poorly suited to imaging the vessel wall itself, and therefore is useful primarily for assessing stenosis, but not for analyzing specific lesions. OCT has the highest spatial resolution of any modality, and by virtue of its speed of acquisition, avoids the need to gate the images against the electrocardiogram (ECG) signal. However, it lacks the depth of field of IVUS. Both OCT and IVUS may be combined with x-ray angiography using the 3-D

reconstruction technique described in [5], and used in this project, to obtain a 3-D model of the vessel.

While x-ray angiography provides a less detailed picture of interior structures than IVUS, due primarily to the lack of radio-opacity of vascular soft tissue, it is crucial for the 3-D reconstruction required for accurate morphologic and hemodynamic analysis. Wahle et al. have previously developed methods for geometrically correct reconstruction of IVUS [5], as well as analysis of coronary hemodynamics [6],[7],[8].

A combination of either IVUS or OCT, and x-ray angiography, was required for this study, because of the need to examine morphologic features of the plaque and their impact upon vessel hemodynamics. IVUS was chosen because of its greater depth of field, and, more importantly, because of the possibility of including a *virtual histology* (VH) analysis, wherein the composition of the plaque is inferred from the ultrasound signal. This technique was developed in Nair [9], and is peculiar to the proprietary software bundled with the acquisition system used for the IVUS, as well as an off line version of the software supplied to us by the vendor (Volcano Corp., San Diego, CA). This type of analysis is unique to IVUS because it relies on an analysis of the raw ultrasound signal, and provides a compelling case for use of the modality in this type of study.

X-ray angiography is often used to assist in localization of the catheter within the artery, but this is not required for pure clinical IVUS. Our acquisition protocol as developed at the beginning of the study (see [10]) does call for concurrent angiograms and IVUS in order to facilitate an accurate 3-D reconstruction; the protocol was not always rigorously followed by the collaborating clinicians, which in some extreme cases required data to be excluded from the study (see section 1.5).

In this work, we demonstrate a body of work that uses a variety of techniques to extract useful quantitative indices (section 4.2) from angiographic and IVUS image sets, and then trains a classification system on this data in order to make useful

inferences about the progression of plaque structure and composition, and therefore the way the disease evolves over the course of a year. Biomarkers, demographics, and genetics are also used when available.

1.2 Project History

Previous work on this project has focused on the development of image segmentation techniques and the development of an integrated suite of tools for tracking the joint analysis of angiograms and IVUS images. Olszewski [11] integrated the older generation of tools on which much of the software in this project is based into a coherent suite whose functions included

- Importation of angiograms and IVUS data sets from DICOM as individual raster frames.
- De-warping of angiograms acquired using the older CRT-style image intensifier (described by Prause [12]).
- Segmentation of both angiograms (using dynamic programming, first on the catheter, and then on the vessel borders [12]) and IVUS data sets (using the optimal surface detection graph search approach developed by Li [13] and the cost function described by Olszewski in [14],[8]).
- Fusion of angiograms with the IVUS images, by way of the techniques described in [5], to produce both VRML scenes and structured grid meshes suitable for downstream processing (e.g., computation of morphologic and hemodynamic features).
- Computation of morphologic indices to describe the vessel and the plaque structures in 3 dimensions, described by Medina in [15],[16],[17].

Newer angiography systems (all those used in this study) use a more modern flat panel image intensifier grid, which requires no dewarping. Pixel spacing values are

computed from fiducial points supplied in the image.

A new cost function was developed by this author [18] to better cope with the IVUS data used in this study (the previous cost function was developed for 40 MHz catheters supplied by Boston Scientific; this study relied upon 20 MHz catheters supplied by Volcano Corp.).

1.3 Aims

The aims of this thesis were 5-fold:

1. Develop and validate a method of compensating for rotational motion of the catheter within the vessel and its effect upon the 3-D orientation of the reconstruction. (Section 3.2)
2. Develop and validate a method of registering the reconstructed vessels that permits identification of a point-to-point correspondence on all quantitative indices. (Section 3.1)
3. Successfully reconstruct, register, and analyze image sets for each of as many patients as possible (not all can be included due to errors at acquisition time in some patients) for analysis. (Chapters 2, 3, 4)
4. Identify statistically significant indices in the data suitable for use as features in a classifier. (Section 5.2)
5. Construct and assess performance of a classification system that can draw useful conclusions about the 1-year progression of the arterial pathology in a patient not used in the training set. (Section 5.3)

1.3.1 New work supporting aims

Computational fluid dynamic (CFD) analysis provides a powerful tool for the analysis of atherosclerotic plaque development in the coronary artery *in vivo* by allowing

correlation of plaque development over time with vessel wall shear stress (WSS), opening the door for prediction of disease progression at specific sites within the vessel, as anticipated in [8].

Most existing work on the effects of WSS on plaque development in the coronary artery either omits the effects of bifurcations entirely ([6],[7],[19],[20],[21],[8],[22]) or uses an overly simplistic model of them ([23],[24]). (Soulis [23] bases the model purely on angiograms. While this captures much of the important information about the bifurcations, it omits much of the important morphological detail available through IVUS. Samady [24] is much more thorough, though the details of their method for attaching the branches is omitted from their publication.)

The dearth of existing publications treating coronary bifurcations in WSS studies owes to the difficulty of segmenting branches in IVUS: coronary branches are spatially small, often appearing in only 3 or 4 frames during a pullback (assuming the frames are sampled on the ECG R-wave, the branch would appear in 3-4 successive R-peak frames, corresponding to an axially-represented diameter of roughly 2 mm for a 0.5 mm/s pullback speed). If the trajectory of the branch leads away from the main vessel, they rapidly escape from the field of view (typically about 5 mm radially). If the ECG signal during the pullback was imperfect, proper segmentation can further be complicated by spatial displacement of the branch from one gated frame to the next. These factors limit the performance of most conventional IVUS segmentation methods.

In chapter 2 we outline a novel method for identifying and modeling the bifurcations and incorporating them into the mesh used for WSS analysis. This work was previously published in [25], but is reprinted here because it is comparatively new, it is essential to much of the subsequent work described in this study, and the method for mesh construction changed after [25] was submitted for publication.

1.4 Analysis Pipeline Description

Figure 1.1 depicts the full array of tasks required to fully analyze a patient's data for use in this study. Blue blocks indicate tasks that were fully developed by prior authors, and which were performed in this study by software that was already written, without modification. Green blocks indicate tasks that had been previously developed, but for which substantial modifications were made to the software used to perform the task, in order to accommodate substantial changes to the task itself or the workflow in general. Red blocks reflect wholly new work. The majority of this paper will be devoted to covering these blocks. Demographic and biomarker data are depicted as a yellow ellipse, reflecting the fact that they are slurped into the database described in section 4.1, but otherwise used as-is.

The first step is to import the angio and IVUS data. For the IVUS data, the goal is to import the entire pullback. For the series of acquisition systems used in this study, the data comes in the form of a DICOM file containing the IVUS data, and a black-box file containing the radio frequency (RF) data from the ultrasound. This second file is used to retrieve the timestamps at which the acquisition system detected r-wave peaks in the ECG of the patient during acquisition. The import tool uses known information about the DICOM frame-rate, combined with the timestamps in the RF data file, to attempt to locate all frames corresponding to r-wave peaks.

This subset of frames are then imported into our system, along with the maximum number of trailing frames that can be taken from between all pairs of r-wave peaks – in theory, chronicling the full cardiac cycle for the course of the entire pullback, leaving the door open to 4-D analysis in the future.

The 3-D reconstruction method was developed with true biplane x-ray angiography in mind. The current data set was not collected with true biplane angiography, due to a lack of the required equipment at the hospital where the data was collected. To simulate the geometry of biplane angiography, at import time a pair of projections

Figure 1.1: Data processing flowchart. Red blocks are new work, blue blocks are untouched preexisting work, green are old software that was extensively modified. Dotted lines indicate alternative paths.



that were as close as possible to 90 degrees apart was identified. For the reconstruction to be feasible, the distance between the projections must be between 60 and 120 degrees. Thus at the hospital, the physician injects the patient with contrast dye, collects the images from 1 projection, then rotates the beam source about the gantry, and repeats the process, injecting a new dose of contrast and collecting another set of images.

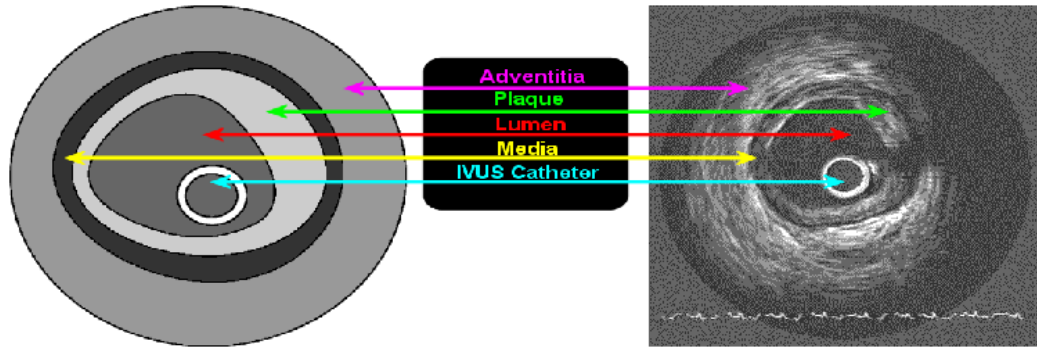
A single cardiac cycle is imported in the case of the angiograms. The objective is to begin the import at an r-wave peak where the vessel lumen is flush with contrast dye, but the catheter is still visible as well. If successful, this means a) the angiograms are synchronized to the IVUS in terms of the position in the cardiac cycle and b) the segmentation of the angiogram will be trivial, as 2-D dynamic programming is extremely successful in segmenting the catheter and vessel borders in these cases, requiring no manual correction if both are clearly visible.

The IVUS is automatically segmented using an application developed by this group prior to the current generation of the project, using a cost function published in [18]. The segmentation always requires manual correction, and a team of cardiology residents was employed to perform the manual tracing to bring the segmented borders up to an acceptable standard for publication. Only the endiastolic phase (the subset of frames corresponding to the r-wave peaks) is segmented.

Figure 1.2 illustrates the anatomy from the point of view of the IVUS segmentation. The inner border that gets segmented represents the *lumen/intima* boundary, where the intima is the innermost layer of cells of the vessel wall (1 cell thick), and the lumen is the interior of the vessel where the blood flows. The outer border represents the *media/adventitia* interface, where the media is the vessel smooth muscle, and the adventitia is the surrounding layer of connective tissue. In a healthy vessel, there is no plaque; when plaque is present, the segmentation includes it in the *media* region.

The angiograms are automatically segmented using an application developed by

Figure 1.2: Explanation of vessel anatomy from the point of view of IVUS.



this group in prior years. If the lumen and the catheter are clearly visible, no manual correction is required. If either is missing, it must be manually traced, at times by estimating from watching the video of the image loop from the raw DICOM.

During the 3-D reconstruction, the borders are sampled at 72 points per contour, where a *contour* refers to the border from 1 frame of the IVUS data set. Thus, throughout the rest of this paper, but particularly in section 4.2, 72 *wedges* are referenced. Where this occurs, a wedge corresponds to all data gleaned from the portion of the vessel residing between a pair of border points in the 3-D reconstruction in a given frame; it is called a wedge because when viewed as subdivisions taken along the perimeter of the border, extending from the edge to the center, they have a wedge shape.

From this point, the VH procedure is described in section 4.2.1.2, the bifurcation modeling and finite element (FE) meshing in chapter 2, the registration in section 3.1, the angiogram branch identification and its impact upon the 3-D reconstruction in section 3.2, and the feature selection and classifier in chapter 5.

1.5 Patient Data Description

The data used in this study was taken from the HEAVEN study [26]. 89 patients were originally enrolled in the HEAVEN study; all signed statements of informed

consent. Because the objective of the study was to identify clinically significant features before the onset of acute symptoms, the study was restricted to patients with stable angina.

Of the 89 enrolled, we received data for 67 patients; of those, we received only partial data for 22, yielding 45 patients for whom we received complete baseline and followup data sets. Repeated efforts were made to secure the missing pieces from the collaborating physician. Of the 45 for which complete data were received, one was excluded because there was no catheter visible in any of the angiographic projections, and therefore there was no way to plausibly localize the 3-D reconstruction. In 2 of the cases, there was no meaningful overlap between the baseline and followup portions of the data, making the type of analysis performed in this study impossible. A further 10 patients had gating issues that excluded them. This yielded a final total of 32 patients that we were able to include in the study for analysis.

The pullbacks were gated against the R-peaks of the ECG, which correspond to the end-diastolic phase of the cardiac cycle. With the Volcano IVUS acquisition system used in this study, each time the system detects an R-peak it stores a time stamp in the file that contains the raw RF feed from the ultrasound; this can then be correlated with the DICOM B-mode video feed to select phase-stationary image frames with which to compose an image volume.

If the patient's heartbeat was sufficiently erratic during acquisition, or if the ECG signal was merely noisy (e.g., due to a bad connection/cable), these timestamps may be incorrect, or they may be too inconsistent for a constant slice thickness to be assumed, and saw-tooth artifacts may be introduced to the resultant image volume. There is active research into localizing the end-diastolic phase from the images themselves as described by [27],[28]. The method of [27], as adapted by us, is described below in section 1.5.1; we considered the method to be effective, however, concerns about the availability of VH led us to omit those patients in spite of the improved

gating.

1.5.1 Optical Gating

Gatta’s technique [27] of using motion blur is simple, but effective. The equation as described in Gatta’s paper is reprinted in equation 1.1; our adapted version is shown in equation 1.2. $B(I)$ refers to the *blurring* value as a function of image intensity. E refers to the expectation. r is the radial coordinate and θ the angular coordinate, in a log-polar representation of the image data. δ refers to the gradient, in the case of equation 1.1, the change in intensity taken with respect to change in radius.


In the case of our version, $C(I)$ refers to the centroid – our version seeks out the rotational information, by tracking the centroid of the image intensity, and using only the θ component thereof, rather than tracking vertical shift in the frame; we found that on the data sets in question this provided pullbacks that were smoother and more consistent.

The output of both versions is subsequently processed with a high-pass infinite impulse response (IIR) Butterworth filter. We stuck with the example provided in [27], after experimenting with a few adjustments. The filter has 10 poles, with $f_c = 0.5$, and f_s set to the frame rate of the IVUS pullback.

Figures 1.3a and 1.3b illustrate the difference for 1 pullback in which the ECG was particularly bad. The number of frames in the two sequences is approximately the same; the time past between the first and last frame in figure 1.3a (ECG gated) is 54 seconds; the time past between first and last frame in figure 1.3b is 2 minutes, 41 seconds, indicating a factor of 3 in terms of the amount of distance and time covered in the same number of frames.

We can see that the pullback “looks” much more correct, in terms of being less jagged, and following the shape that might be expected of a phase-stationary IVUS pullback; the difference in time covered suggests that the average frame spacing of approximately 10.9 frames per trigger in the motion blur gated case, versus 3.75

Figure 1.3: Comparison of motion blur versus ECG gating.



frames per trigger in the ECG gated case. The former would roughly correspond to a brisk, but not unheard of, resting heart rate of approximately 164 beats per minute; the latter to a completely impossible 479 beats per minute. While both almost certainly have spurious gating events, one much more closely maps to the likely underlying physiological reality.

$$B(I) = -E\left(\left|\frac{\delta I(r, \theta)}{\delta r}\right|\right) \quad (1.1)$$

$$C(I) = \frac{\sum_{\forall r, \theta} \theta I(r, \theta)}{\sum_{\forall r, \theta} I(r, \theta)} \quad (1.2)$$

1.6 Similar Work

The closest existing work to the current study available in the literature was performed by Samady [24], albeit in a more limited scope. They demonstrate a positive correlation between change in plaque burden and low WSS and a negative correlation between change in plaque burden and high WSS ($P < 0.001$). Moreover, they showed that the plaque composition as depicted by Volcano's VH algorithm moves towards a considerably greater fraction of necrotic core in the high WSS patients. This was computed over a 6 month window for baseline/followup study.

Our work looks at a one year window, and attempts to take a much more general view of the pathology. The principal improvement that our work brings to the table is the development of a classifier that predicts the plaque evolution in a patient at a specific location based upon the morphologic, VH, and hemodynamic indices, as well as serum biomarkers and patient demographics.

CHAPTER 2 BIFURCATION MODELING

Bifurcations are used as the landmarks in the registration described in chapter 3. As such they are indispensable to the rest of the work presented here. They are also important for the computation of the wall shear stress terms used in training and evaluating the classifier in chapter 5.

An approach to automatically segmenting the branches using region growing techniques was explored; however, it was concluded that the nature of coronary bifurcations in IVUS makes them poorly suited to automated segmentation: they are small, often appearing in only 3 or 4 frames during a pullback (assuming the frames are sampled on the ECG R-wave, the branch would appear in 3-4 successive R-peak frames, corresponding to an axially-represented diameter of roughly 2 mm for a 0.5 mm/s pullback speed). If the trajectory of the branch leads away from the main vessel, they rapidly escape from the field of view (typically about 5 mm).

Moreover, the anisotropic nature of IVUS means that branches that lie normal to the main vessel in the image will be represented by an extremely small number of voxels, providing too little information for a reliable segmentation. The automated segmentation yielded segmented regions that could have been used as registration landmarks, but would have been wholly unsuitable for CFD modeling. Because we are not presently investigating the properties of the branches themselves, a decision was made to model them with a cylindrical shape that can be guaranteed to be suitable for CFD analysis; this method is presented here.

2.1 Methods

To facilitate robust modeling of the bifurcating vascular segments, we have developed a hybrid user-guided approach to branch identification (automatic segmentation from few control points) and generalized-cylinder representation. Segmentation of the main vessel (both lumen-intima surface and media-adventitia surface) is performed


prior to the branch marking, using procedures discussed in [8] and [18]. This maintains accuracy in the main vessel, while permitting analysis of the newly modeled branches' impact upon flow conditions in the vessel.

The method for marking the branches can be summarized as follows: a user indicates the approximate trajectory of the bifurcation by placing control points for a B-spline along the branch. An automated segmentation is computed in 2-D at each control point normal to the curve computed from said points, and the diameter at that location is estimated, based on the assumption that the segmented region through which the curve passes corresponds to the branch. The control points are adjusted based upon the automated segmentation, yielding a curved surface diffeomorphic to a cylinder. The user can then refine the diameter and the position if the algorithm failed.

The entire process takes no more than 5 minutes; less than a minute for the automated estimate, and perhaps another 2 or 3 if significant adjustment is required or there's some ambiguity making the manual adjustment difficult. A more detailed description follows, and a screenshot of the interface is shown in figure 2.1. The yellow dot is the proximal seed point of the branch, and the two white dots are each control points for the spline curve; the cyan border reflects the current position of the branch.

1. A branch starting point (corresponding to the vessel branch junction), a distal branch point, and one or more intermediate control points are interactively identified by the user.
2. A uniform cubic B-spline is computed over the range of identified points to yield a 3-D space curve that follows the branch's trajectory.
3. For each control point, the image is resampled at the nearest location on the spline curve, such that a 2-D plane normal to the spline curve is produced, using linear interpolation (Fig. 2.2a).


Figure 2.1: Interface for manual refinement in branch capture program.




4. Mean-shift segmentation [29] is performed on the 2-D normal plane (Fig. 2.2b).
5. The area of the branch region identified in the segmentation and the first order image moments are used to compute the radius and centroid of the branch at the position of the control point.
6. The control point positions are adjusted to match the computed centroid points and the spline is recomputed.
7. A mesh is constructed around the spline with a cylindrical shape, parallel to the spline at each position along the curve, using the radius computed from the 2-D segmentation.
8. The branch radius is manually adjusted, if necessary, using the mouse wheel.
9. The trajectory of the branch is manually adjusted, if necessary, by dragging the control points. No new control points may be added, nor existing ones removed, once the automated procedure has run.

Figure 2.2: Output of 2-D mean-shift segmentation used to estimate branch diameter.

(a) Raw 2-D plane as resampled from the initial control-point based spline.



(b) Same 2-D plane with the mean-shift segmentation output overlaid.



2.1.1 Trajectory Estimation

B-splines [30] are a well known and established method for curve fitting and interpolation, defined as piecewise polynomials, such that each disjoint piece is a smooth function on the interval defined, and the spline as a whole is generally C^2 smooth (or twice differentiable) where the pieces meet.


We have employed a generic uniform cubic B-spline blending function defined as a 4×4 matrix, utilizing user-defined control points as input.

2.1.2 Diameter Estimation


Using the spline curve, the torsion is computed along with the Frenet frame at each point along the spline [5]. This information is used to compute a 2-D image plane, $P(u, v)$, normal to the spline at each point. The u coordinate is resampled along the binormal, and the v coordinate is resampled along the normal. This guarantees that the 2-D image plane is perfectly aligned with our trajectory estimation, and thus well

Figure 2.3: Branch junctions, after different merging procedures.


(a) Branch/artery junction with naive triangle stripping.



(b) Junction as constructed using Nef polyhedra.



(c) Nef junction after Delaunay refinement.



suitied for estimating the diameter (Fig. 2.2a).

The diameter estimation is performed using a 2-D segmentation in the plane sampled from the spline. Prior to segmentation, the entire data set is smoothed with a $3 \times 3 \times 3$ spherical kernel. The smoothed image is then sampled at the nearest spline point to each control point along the spline, in a plane normal to the spline. Because the z axis has much lower resolution than the x and y (for the data studied, the z resolution is on the order of 0.5 mm/voxel, exactly matching this value at 60 beats per minute (bpm), whereas the x and y directions are 0.026 mm/voxel), we must

interpolate along the z axis in order to achieve an isotropic result.

Equation (2.1) shows the bi-linear interpolation function used; a is the fraction portion of the floating point representation of z , and is therefore a measure of how far between two points on the z axis the current $f(u, v)$ is. In equation (2.2), f_0 represents the top left corner of the projection plane, as determined by measuring the distance to the edges of the 3-D volume along \hat{N} and \hat{B} .

$$P(u, v) = (1 - a)I(f_x, f_y, f_z) + aI(f_x, f_y, f_z + 1) \quad (2.1)$$

$$f(u, v) = f_0 + u\hat{N} + v\hat{B} \quad (2.2)$$

The segmentation method is the synergistic approach described in [29], which combines the recursive mean-shift based density estimator described by [31] with a confidence map computed using the edge detection model described by [32]. The mean-shift procedure, as implemented (the reference implementation published by the authors of [29], [32], and [31] was used in this study), treats the pixel intensity values and the two spatial dimensions as its feature space, yielding a 3-D effective space for the clustering. The confidence map is computed from the image gradients computed at each pixel by examining the orientation of the gradient in the context of its neighbors.

The gradient strength is combined with the confidence measure computed from the gradient orientations into a weighting, and this weighting steers the fusing of modes computed in the mean-shift procedure into regions defined in the final segmentation, seen in fig. 2.2b. The implementation used has a small number of parameters that can be set by the user. These are: edge detection kernel size (set to 2), minimum region size (set to 20 pixels), spatial bandwidth of the segmentation, which defines the search window during the mean-shift (set to 7), and the color bandwidth, which

defines the range of gray levels to be considered during the procedure. This last value is initialized to a default value of 20, and then recursively reduced by 2 if the area identified as the branch occupies more than 40% of the total area of the resampled plane, until a result that takes up less than 40% of the input plane is found.

This last step prevents an overzealous segmentation from always slurping up the entire input plane, but it also means that if the search area is too small, it may fail to converge.

The area of the branch is computed by summing the pixels matching the segmentation label at the control point, and the radius is estimated by assuming a circular cylindrical shape as given in eq. (2.3)

$$r = \sqrt{\frac{\sum_{u,v} T(u,v)}{\pi}}, \quad (2.3)$$

where $T(u,v)$ is a threshold function defined in Eq. (2.4), $L(u,v)$ represents the segmentation label, and u and v are the transposed coordinates in the 2-D normal plane.

$$T(u,v) = \begin{cases} 1 & \text{if } L(u,v) = L(u_{seed}, v_{seed}) , \\ 0 & \text{else} \end{cases} \quad (2.4)$$

2.1.3 Mesh Merging

We initially developed an approach to mesh merging that involved deleting the overlapping points from the respective surfaces (e.g., the main vessel and the branch), and then using a plane fitted to the junction with least-squares estimation to triangulate the junction. The result of this procedure is depicted in fig. 2.3a. This proved

inadequate and problematic: the junction was too abrupt for truly robust CFD analysis, as abrupt junctions result in pockets of turbulence unrepresentative of the underlying reality. Moreover, because the two meshes were invariably sampled at different resolutions, constructing a junction that was manifold (e.g., no holes, and no self-intersecting triangles) was extremely difficult to do consistently.

A modified approach making extensive use of the CGAL (Computational Geometry Algorithms Library) library [33] was devised, with a two step process that eliminates the topological inconsistencies that stymied the initial effort. The list of steps performed for the mesh merging is outlined below, and details follow after.

1. Branch mesh points are transformed into physical coordinate space using the intermediate Frenet-Serret geometry model derived from code based upon [5].
2. 10-diameter flow extensions and end-caps are added to each of the branches and to the main vessel. The flow extensions are required to ensure fully developed flow in the CFD analysis; the end-caps are required to make the surface manifold for the remeshing procedure that follows.
3. The branches are merged with the vessel into a single surface using CGAL's [33] *NEF_Polyhedron class*, based on Walter Nef's [34] work with boolean half-spaces.
4. The resulting surface is remeshed using a Delaunay refinement procedure described in [35].
5. The end-caps are removed from the surface. The details of this procedure are described in algorithm 1.
6. The final surface is topologically sufficient for use in the remainder of the CFD pipeline, as described in section 4.2.1.3.

First, all surfaces, including the main vessel, and each of the branches, is extended with 10-diameter flow extensions (branches are extended at outlets only because the branch inlets will be merged into the final surface). It is standard practice in CFD analysis to add flow extensions to guarantee fully developed, convergent behavior within the region being modeled. 10-diameters is a commonly used length that is widely agreed to be sufficient, and is the convention that was used in our case. Prior to remeshing, each of the surfaces has a form known as a *structured grid*, which means that a surface consists of a finite number of *contours* and *points per contour*. (In the case of our work, the number of *points per contour* was always equal to 72, but that is by convention and not by necessity).

To add the extensions, a normal vector is taken from the last pair of points in the centerline of each sub-surface. This normal vector is then rescaled to be equal in length to the diameter at the outlet in question. Then 11 additional contours are added, each displaced from the last contour at the outlet by a multiple of the rescaled vector. The 11th contour is added to make it easy to remove the endcaps later while still keeping a 10 diameter extension.

End caps are added by adding a centerline point coplanar to the last contour, which then forms 72 triangles in concert with the 72 points of the last contour. During this and the last step, 2 data structures are created and cached to permit the removal of the end caps after remeshing: first, a circumscribing cuboid is constructed around each flow extension, and then an oriented plane is stored at the position of the 10th diameter extension. A coordinate transform object is also created, which shifts the axes of the mesh to where the centerline of the flow extension is coincident with the x-axis, thereby permitting the cuboid to be computed as an axis-aligned bounding box.

After end caps and flow extensions are added, the meshes are merged, using CGAL's Nef polyhedron class, which builds upon Walter Nef's [34] work on Boolean

halfspaces, and was fully developed for work with 3-D polyhedral meshes by Hachenberger and Kettner [36]. The algorithm performs a logical union of the volume of the main vessel mesh with those of the bifurcations, yielding a closed, manifold surface free of intersections (though neither smoothness nor triangle quality is guaranteed). The result of this procedure can be seen in fig. 2.3b.

This initial surface is then used to initialize a second procedure, described in [35]. The technique in question uses an iterative Delaunay refinement driven by an implicit surface definition to produce a smooth mesh whose vertices have valence of 6 wherever possible, guaranteeing a well-behaved topology suitable for CFD analysis. Because the branched coronary artery cannot be defined analytically, we construct the implicit surface definition around iterative intersection of axis aligned bounding boxes, where the bounding boxes are constructed over various subsets of points taken from the initializing polyhedron produced in the first step. The final output of the smoothed, resampled junction can be seen in fig. 2.3c.

Before the final surface is written out, the end caps are removed, so that the location of inlets and outlets is unambiguous to the downstream pieces of the pipeline that dispatch the CFD solver. The stored 3-D coordinate transformation is used to map each point in the surface into an alternate frame in which each of the bounding cuboids lies along the x-axis, centered at the origin. This in turn makes it trivial to determine if a given point a) lies within this cuboid and b) lies on the side of the plane where the end cap is located, or the side where the rest of the flow extension is located. If the point lies on the *negative* side of the plane, but within the cuboid, it, and any triangles in which it is contained, are removed from the surface. This procedure is described in algorithm 1; the function $MAP()$ refers to the coordinate transform mentioned above; $AABB$ refers to the axis aligned circumscribing cuboid for a given flow extension, and $BELOW()$ tests to see if a point is on the negative side of the boundary plane.

Algorithm 1 Procedure for removing endcaps from triangulated surface mesh.

```

for all POINT do
  for all FLOW_EXTENSION do
     $MAPPED \leftarrow MAP(POINT, CENTERLINE(FLOW\_EXTENSION))$ 
    if  $MAPPED \in AABB(FLOW\_EXTENSION)$  then
      if  $BELOW(MAPPED, BOUNDING\_PLANE)$  then
         $DELETE(POINT)$ 
         $NEXT(POINT)$ 
      end if
    end if
  end for
end for

```

2.2 Validation

Suitability of the method for marking/identifying bifurcations for purposes of downstream analysis was assessed by comparing two separate users of the software to each other and to the initial estimate computed algorithmically. While some bifurcations are difficult to mark, the unambiguous cases were typically very close, both between users and between the users and the automated estimate. Moreover, user training can further improve the results, as it is possible to make a much more accurate estimate of the branch trajectory once one is familiar with the quirks of the interface (In particular, setting too many spline control points can result in contorted trajectories that are very difficult to straighten. If a bifurcation is roughly normal to the catheter path within the IVUS field of view, the spline should be limited to one intermediate control point to allow for a straight trajectory to be modeled.).

The quantitative results identified in [25] can be seen in figure 2.4. The validation was computed on 9 different bifurcations identified for purposes of assessing the quality of the method when [25] was originally published. For each of figures 2.4a-2.4c, the metric in question is computed on each of the 9 branches. Figure 2.4a computes the variation in centerline position between 3 expert observers used to trace the bifurcations; figure 2.4b compares the centerline position between the observers and

the initial raw estimate of the algorithm as described in section 2.1; figure 2.4c shows the variation in the radius compared both between tracers and against the raw initial estimate. The results for each branch are averaged across observers; error bars show the range of variation in plots 2.4a and 2.4b.

These results are considered good enough, given the intended use of the modeled bifurcations. The inter-observer positioning error is an indication of significant uncertainties in identifying bifurcations in IVUS; they nonetheless were typically less than 0.5 mm, which indicates a high degree of similarity in trajectory, even if the exact location varies. Likewise, the variation in radius estimates was typically on the order of 0.1 mm, which is approximately 5 pixels for the image frames used in this study, and again a reasonable uncertainty value. The average radius estimate for marked branches in our dataset is 0.85 mm, which is considerably larger.

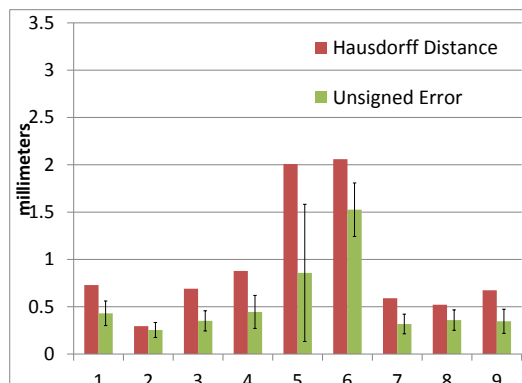
Branches 5 and 6 had considerably greater inter-observer variability than the others, and branch 6 also had by far the greatest automated segmentation error. These branches were the most difficult to segment, because they had trajectories that were almost normal to the main vessel. From an automated segmentation standpoint, this is extremely difficult because it means that the size of the image plane resampled (section 2.1.2) is extremely small; more significantly, the anisotropic nature of IVUS comes heavily into play: a bad automated segmentation can easily cause the trajectory estimate to be off by a few millimeters. Fortunately, this is a rare occurrence.

Figure 2.5d illustrates how the WSS in a single vessel changes from the inclusion of a single branch. Figure 2.5b shows the WSS with the branch omitted; figure 2.5c shows the WSS with the branch included (the bright dot at the branch's location is not actually WSS, but rather an annotation to make the site of the branch visually clear). Figure 2.5a shows the surface the WSS runs in figure 2.5 were based on. The drop in WSS after the inclusion of the branch (fig 2.5d is unambiguous, and makes the case that the effects of the branches are important to consider in any evaluation


of coronary WSS values.

Figure 2.4: Positioning and radius error in branch identification.

(a) Inter-observer positioning error



(b) Segmentation error



(c) Radius error

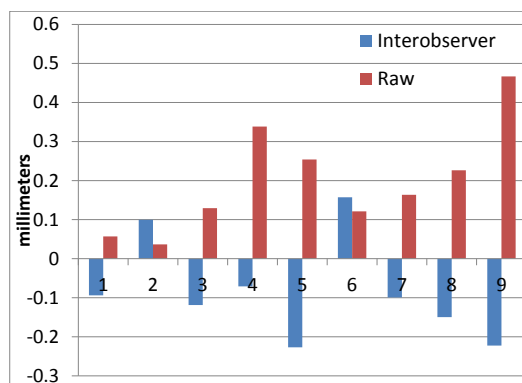
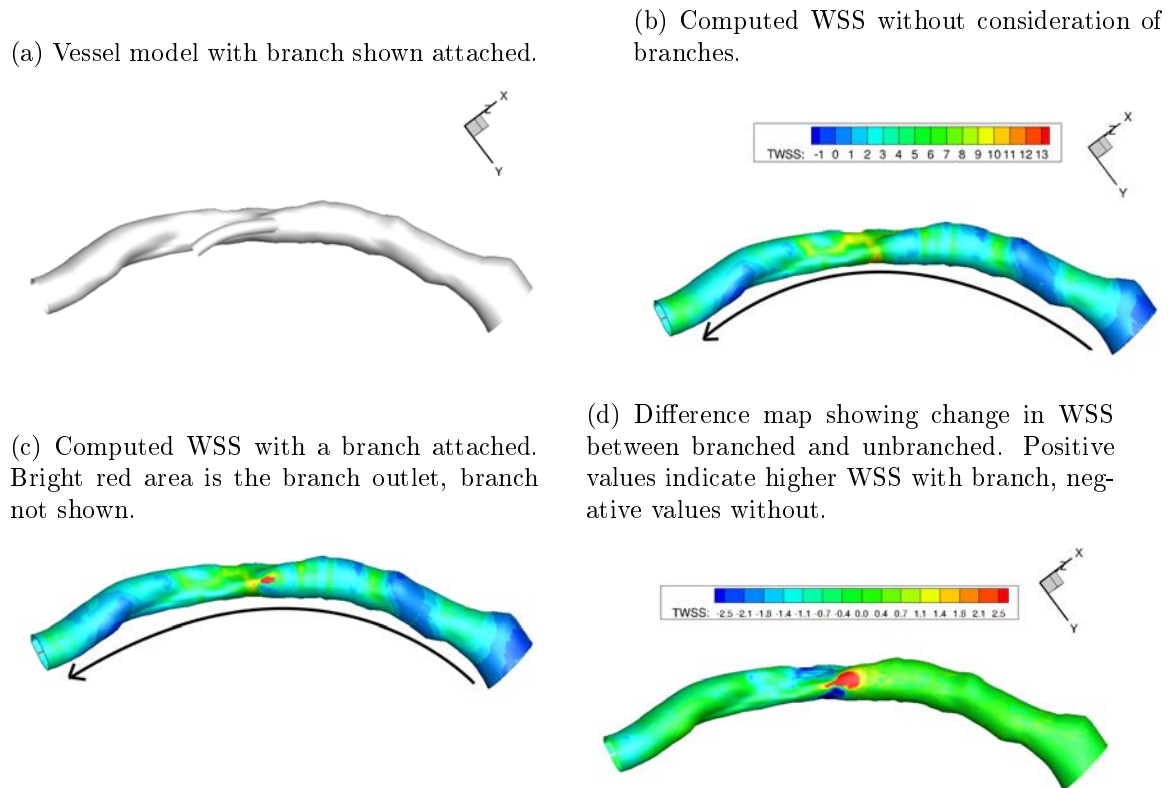


Figure 2.5: WSS for a right coronary artery, shown with and without considering the flow effect of a branching vessel. Notice reduced WSS (light blue in 2.5d) at the distal end of the vessel. Arrow depicts the direction of blood flow.



CHAPTER 3 REGISTRATION

The registration portion of the analysis consists of two parts: the most obvious and important part consists of mapping the baseline and followup portions of the data into a common space so that meaningful comparisons can be made between given regions of the vessel. Our method for this portion of the registration is detailed in section 3.1. The method developed uses a piecewise transform to scale the coordinates along the z-axis while applying a B-spline generated rotational shift to compensate for rotational variation between the pullbacks.

Were the 3-D modeling not a part of this project – if the analysis was limited to virtual histology and measures of plaque thickness – this would be sufficient. Section 3.2, and in particular section 3.2.1, detail a nondeterministic variation in the catheter orientation that sometimes occurs during a pullback. The phantom study was performed to investigate the nature of the twisting, and to see if analytical model could be developed, permitting it to be corrected for automatically. This was found not to be the case.

To deal with this effect, a method was developed that uses the appearance of branches in the angiograms as fiducials onto which the IVUS branches can be mapped. If the branches can be clearly identified in the IVUS, and in both planes of the angiograms, locations in 3-D where the vessel has twisted beyond what should be expected from the model of the catheter’s torsion can be detected and unwound.

The method we developed to address this problem required branches to be clearly visible in the angiograms to be useful. Because this was seldom the case in this data set, we performed an initial validation study of the method on a minority of patients for whom branches were visible in the x-ray projections, but did not employ it in the creation of the 3-D reconstructions used for the analysis discussed in chapters 4 and 5. We present the method anyway as a possible means of improving future studies,

provided the inclusion of branches as fiducial points is addressed at acquisition time for the angiograms.

3.1 Baseline/Followup Registration

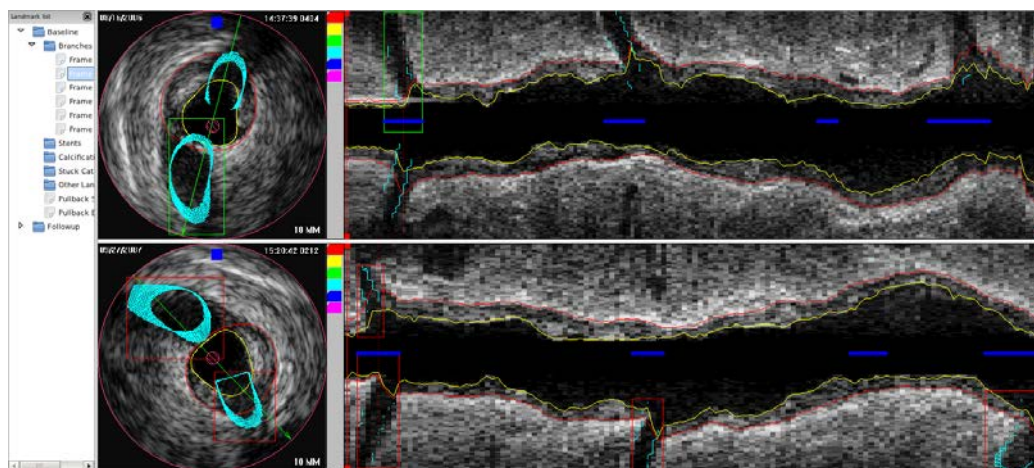
3.1.1 Methods

If the methods described in section 3.2 are essential to ensuring accurate wall shear stress values and 3-D morphologic indices, the actual registration of the baseline and followup pullbacks is in many ways much simpler: because the indices are all sampled in the IVUS space of cylindrical coordinates, corresponding to an axial (z) value and an angular (θ) value, a rubber sheet transform similar to an affine transform in the cylindrical IVUS space is sufficient to align the indices we compute for the purpose of running a classifier on the data.

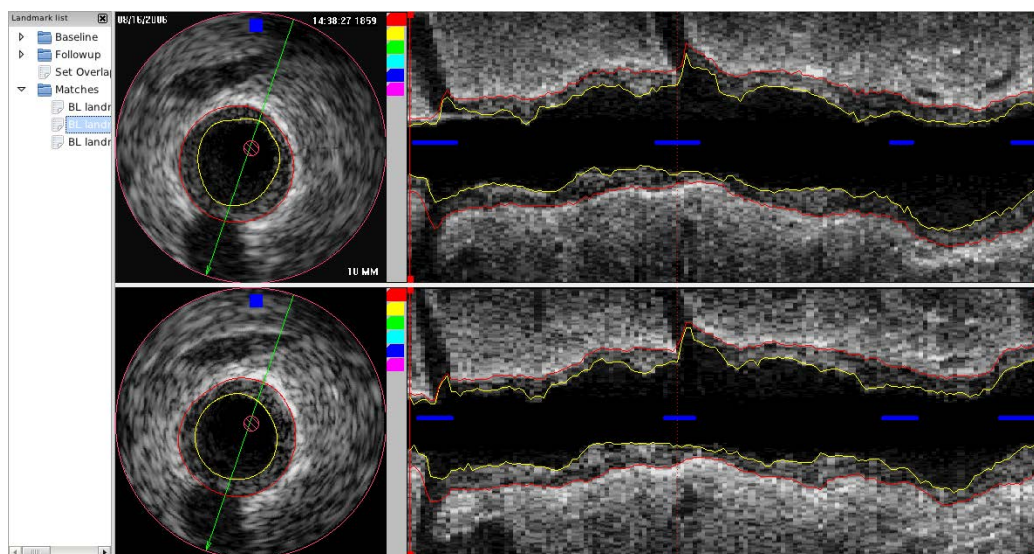
However, if there exists more than 1 branch, a single affine transform is inadequate to represent multiple directions of rotational distortion, and attempting to construct one piecewise presents a significant problem in the form of dependencies between the pieces: the edge of one must align with the middle of the next and so on. (Audette et al. [37] demonstrated the robustness of a piecewise affine technique in considerably more complex image registrations involving otolaryngologic surgery with CT guidance; however, they stitched chunks of image together end-to-end, and had no overlap between pieces of transforms.) Thus the transformation has been broken into two components, computed separately: an axial scaling transform, computed along the length of the pullback, with a scaling term computed separately for each pair of landmarks, and a rotational transform, computed using a Catmull-Rom B-spline with a control point set at each pair of landmarks. While thin-plate splines have generally become the standard technique for image registration in cases like these, the method used was chosen because it was important to treat rotation and axial scaling separately, as one is driven by catheter motion and/or initial position, and the other by catheter speed.

Figure 3.1: Before/after for baseline/followup registration.

(a) Tool used to capture landmark correspondence. Green boxes match a branch selected from the tree control on the left. Red boxes represent all available branches in the opposite pullback. Software attempts to infer “baseline” and “followup” from the database described in section 4.1; else the order is arbitrary.



(b) Patient depicted in figure 3.1a, after registration transform is applied to followup image. Timestamp information is gone from the lower image because only the portion inside the field of view (outer red ring) is resampled.



For clarity, we have adopted the following convention: lowercase z and θ refer to the actual coordinates (in the case of equations 3.3 and 3.4, the input and output values for the mapping), and upper case Z and Θ refer to functions, taking a landmark

as an argument and returning the corresponding coordinate of that landmark.

At each point, given a pair of input coordinates from the baseline pullback, the followup equivalents are computed, treating the baseline as the *master*. Throughout this section, the abbreviation *BL* refers to the baseline coordinate space, and the abbreviation *FU* refers to the followup coordinate space. LM_x and LM_{x+1} are meant to describe the previous and next marked corresponding landmarks. The math is described in equations 3.3 and 3.4. The angular components from the B-spline function are indexed to the length of the followup pullback, so that for a pair of baseline coordinates, the equivalent followup pair is found by first computing the adjusted z coordinate, and then retrieving the corresponding θ from an array of values computed from the B-spline, with 1 rotational value computed for each followup frame. Terms $Z_{BL}(LM_x)$ and $Z_{FU}(LM_x)$ correspond to the baseline and followup z coordinates respectively of the x th landmark in the registration. $\Theta()$ is a function of z , thus $\Theta_{FU}(z_{FU})$ gets the value of the B-spline interpolation function at the followup z coordinate z_{FU} .

In every registration, $LM_0 = (z_0, \Theta(Z_{FU}(LM_1)))$, for every pullback, where z_0 is the frame number where the overlap was identified to begin, and $\Theta(Z_{FU}(LM_1))$ is the angular correction for the first matching pair of branches. This ensures that the registration examines only regions of the pullbacks that overlap, and it gives the θ coordinate a plausible initialization – the first matched pair of branches is the earliest information available regarding the angular error between the two pullbacks.

$$\delta_{BL}(x) = Z_{BL}(LM_{x+1}) - Z_{BL}(LM_x) \quad (3.1)$$

$$\delta_{FU}(x) = Z_{FU}(LM_{x+1}) - Z_{FU}(LM_x) \quad (3.2)$$

$$z_{FU} = (z_{BL} - Z_{BL}(LM_x)) \times \frac{\delta_{FU}(x)}{\delta_{BL}(x)} + Z_{FU}(LM_x) \quad (3.3)$$

$$\theta_{FU} = \theta_{BL} + \Theta(z_{FU}) \quad (3.4)$$

The control points for the B-spline are selected by first finding the centroid of the axis-aligned bounding box surrounding each marked branch, then computing the angle between the centroids (followup is subtracted from baseline). From this angle, a control point value is selected based on the criteria of minimizing the change in angle from the previous control point (*i.e.* given the circular wrapping nature of angles, if 4 different representations of an angle all transform the followup branch centroid to the baseline's position, the correct angle should be the one that requires the minimal amount of movement from the previous control point location). The procedure is described in detail in algorithm 2.

The values *BL_CENTROID* and *FU_CENTROID* refer respectively to the matched baseline and followup centroid positions for the branch being used as a landmark.

The software tool used for this step (figure 3.1a) builds upon the same framework as the tool used for identifying the branches in IVUS. For the moment, only bifurcations are used as anchors, because they provide the most information and are also the only landmarks that provide two forms of localization (other landmarks are identified in terms of the rough location by the user, but the identification is coarser. The door is open to the inclusion of others should they be necessary, but they are for now identified only for other purposes, *i.e.* excluding stents and heavy calcifications from analysis – described in section 5.1.1 – or compensating for a stuck catheter). Upon

Algorithm 2 Procedure for identifying control points for registration spline.

```

CTRLPTS  $\leftarrow$   $\emptyset$ 
for all MARKED_PAIR do
   $\theta_{BL} \leftarrow ATAN2(BL\_CENTROID_Y, BL\_CENTROID_X)$ 
   $\theta_{FU} \leftarrow ATAN2(FU\_CENTROID_Y, FU\_CENTROID_X)$ 
   $\Delta_0 \leftarrow \theta_{BL} - \theta_{FU}$ 
   $\Delta_1 \leftarrow \Delta_0 + 360$ 
   $\Delta_2 \leftarrow \Delta_0 - 360$ 
   $\Theta(LM_{Next}) \leftarrow MIN(\forall(\Delta - \Theta_{Last}(CTRLPTS)))$ 
   $Z(LM_{Next}) \leftarrow FU\_CENTROID_Z$ 
  APPEND(CTRLPTS, LMNext)
end for

```

load, a list of branches for both baseline and followup is displayed. When the user selects one, the program navigates to that branch, and then highlights all branches in the opposite pullback as cursors which can be selected to indicate a match.

Once a sufficient number of overlapping branches have been marked, the correspondences are committed to the relational database. (If 2 branches overlap somewhat along the z-axis, it's unnecessary to mark both.) Figure 3.1b shows the appearance of the image after the registration transform is applied, in contrast to 3.1a, which shows the interface pre-transformation. The correspondence is excellent, and the differences remaining owe, presumably, primarily to actual physiological differences of the sort we seek to measure.

The end result is that for a given location in the baseline image set, the quantitative indices computed (morphologic, virtual histology, wall shear stress) can all be retrieved trivially from the followup data set, making it possible to do detailed data mining, and subsequently build the aforementioned classifier.

3.1.2 Experimental Methods

This method was validated across all data sets used in the study, as described in section 1.5.

Figures 3.1a, 3.1b provide a visual, qualitative assessment of the efficacy of the

registration method described in section 3.1.1. Each pair of pullbacks was visually assessed to confirm that the method was working correctly; indeed, during the development of the software this provided the essential clues when things went awry, most often because the incorrect rotational direction was selected at some point, causing the registration to introduce a full, 360-degree twist rather than the correct adjustment.

However, because this registration is performed entirely within the IVUS space, it makes sense that it should be evaluated principally in this space as well. The visual examination was performed by examining the image warping and border displacement that was performed with the computed transformation, and the ultimate metrics of the accuracy are built upon a comparison of the IVUS borders between baseline and followup post registration. Zero error can reasonably be considered impossible, even with a perfect registration, as a result of the fact that the pullbacks are one year apart, and therefore some physiological change is inevitable – however, we seek to demonstrate a successful minimization of that error.

Three assessments of the difference are taken in section 3.1.3; Z_{BL} and θ_{BL} in the descriptions below correspond to their definitions in equations 3.3, 3.4:

1. The registered difference. (Eq. 3.5)
2. The unregistered difference, beginning at the first frame, and continuing until one of the pullbacks ends. (Eq. 3.6)
3. The difference taken after registering only along the z-axis, illustrating the portion of the error that results from rotational misalignment alone. (Eq. 3.7)

Table 3.1: Registration error statistics

(a) Lumen-lamina border (mm)				
Comparison	Max	Mean	MSE	SD
Registered	2.8753	-0.0081	0.1830	0.4278
Unregistered	3.9531	0.0388	0.7696	0.8764
Z-axis only	3.4883	-0.0106	0.6558	0.8097

(b) Media-adventitia border (mm)				
Comparison	Max	Mean	MSE	SD
Registered	2.9063	-0.0238	0.1881	0.4330
Unregistered	3.4716	0.0835	0.7869	0.8832
Z-axis only	3.4338	-0.0215	0.6651	0.8153

$$\Delta = f_{BL}(z_{BL}, \theta_{BL}) - f_{FU}(z_{FU}, \theta_{FU}) \quad (3.5)$$

$$\Delta = f_{BL}(z_{FU}, \theta_{FU}) - f_{FU}(z_{FU}, \theta_{FU}), \forall z_{FU} < Length(BL) \quad (3.6)$$

$$\Delta = f_{BL}(z_{BL}, \theta_{FU}) - f_{FU}(z_{FU}, \theta_{FU}) \quad (3.7)$$


For each of these three categories of comparison, we have tabulated the max, mean, and mean squared error, as well as the standard deviation, and plotted density curves for all 3 categories of results as well.

3.1.3 Results

Figures 3.2 and 3.3 depict density curves over all the data sets for difference values described in equations 3.5-3.7 in the borders for registered, unregistered, and z-axis-only registered pairs of pullbacks respectively. Tables 3.1a and 3.1b depict mean, max, MSE, and SD values for the same in tabular format.

Figures 3.2 and 3.3 depict Gaussian kernel density estimates computed using the *R* software package. The bandwidth value depicted in the plot is computed by the

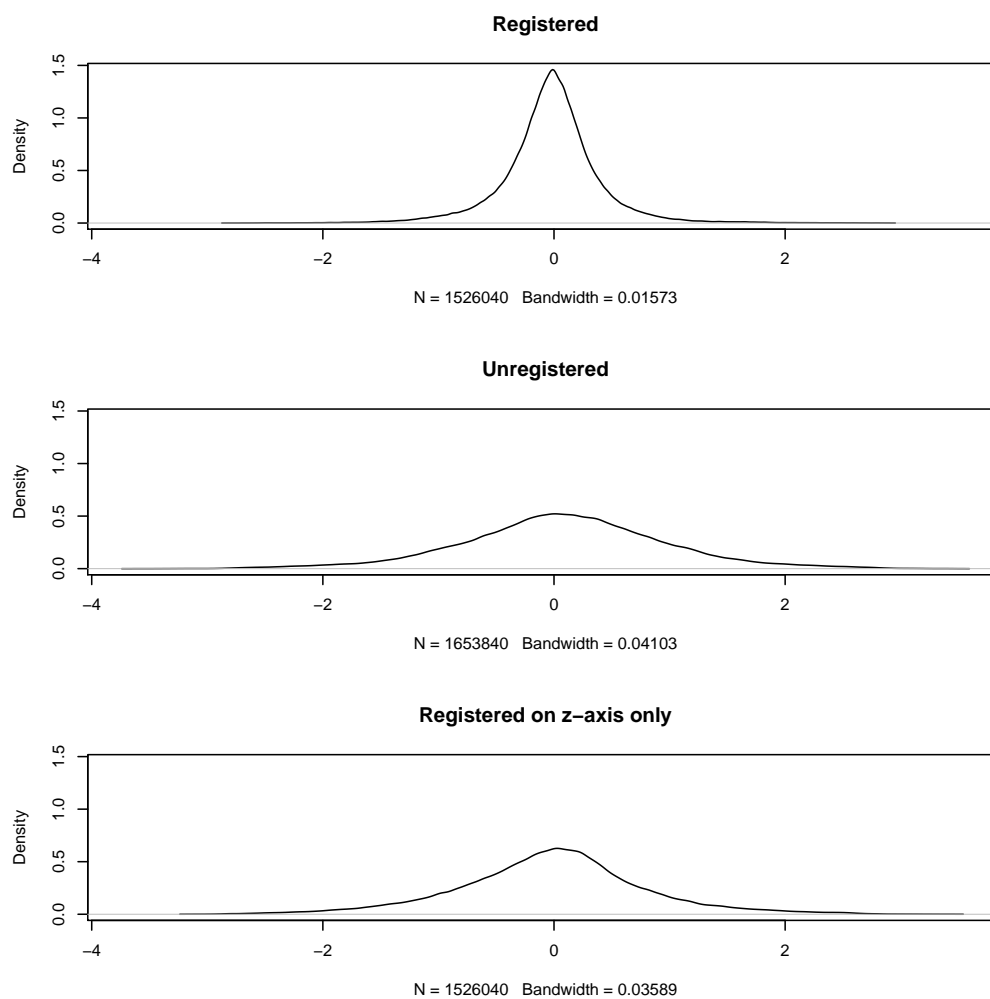
Figure 3.2: Density curves for lumen-lamina border registration error.



kernel density estimator, based on the standard deviation of the input data, and is used for smoothing the output curve. The formula for computing the estimate was first published in [38] as “Silverman’s rule of thumb,” and is described in equation 3.8. SD is standard deviation, IQR is interquartile range, and n is sample size. Given the density of the points, these curves are functionally equivalent to continuously-valued histogram plots.

$$BW = 0.9 \times \min(SD, IQR \times 1.34 \times n^{-1/5}) \quad (3.8)$$

Figure 3.3: Density curves for media-adventitia border registration error.



Figures 3.2 and 3.3, and table 3.1 show quite clearly that the discrepancies between pullbacks are visibly significantly reduced by the registration. A paired t-test was performed for all each of registered/unregistered and registered/z-only-registered for each border. In both cases, the p -value was found to be less than 0.001, indicating the changes seen are statistically significant. Given the fact that the registered difference reflects differences measured over a year in which disease is expected to evolve, we consider this registration approach to perform satisfactorily, and to be a success.

3.2 Catheter Twist Correction

IVUS catheters can be broadly divided into solid-state and rotational-transducer designs. The rotational design requires a sheath to insulate and stabilize the rotating core; the solid state transducer does not require the sheath and thus is generally used without it because of cost savings, and because it makes the assembly narrower. The rotational design produces a known phase-shift artifact in the presence of friction points which was examined and discussed in [5]. A twisting artifact previously unseen in the rotational designs was observed in this study, in which solid-state catheters were used (figure 3.4). The twisting is best described as a corkscrew artifact, which appeared physiologically implausible, given contemporary understanding of plaque growth and development.


A phantom study (section 3.2.1) confirmed that this twisting is beyond what a purely geometric model would predict. This means that what necessarily must be treated as non-deterministic twisting can occur, which adversely affects the accuracy of the geometric fusion, and subsequently the CFD analysis.

We have undertaken to address this artifact by adjusting the orientation of the IVUS frames along the reconstructed centerline in accordance with a mapping between IVUS branch locations and branch locations estimated from the angiograms. Unfortunately, the branches are seldom visible in the angiograms with sufficient clarity for us to have confidence in the projected location; thus validation of this work can be considered cursory at best, and this technique was not employed on the data that was fed into the classification system of chapter 5.

3.2.1 Phantom Study

After observing the twisting in patient data as shown in figure 3.4, a phantom study was undertaken to confirm the presence of an artifact that could not reasonably be corrected analytically. The phantoms shown in figure 3.5a were created to model a variety curved vessel shapes. One of them was designed as a straight pipe as a control

Figure 3.4: Plaque thickness projected onto 3-D reconstruction with twisting artifact. Measure of thickness scales from dark blue (very thin) to red (very thick). Isolated red square near center of the image is the application's cursor for getting detailed information at a given location; translucent yellow shadow corresponds to the media/adventitia border.



segment; the others were intended to roughly match the curvature of a real coronary artery. The fixture depicted in figure 3.5b was created to add some torsional stress comparable to the portion of the circulatory system the catheter must pass through prior to reaching the coronary artery.

There are two alternate paths through which the catheter can be threaded into the phantom shown: either directly through the length of PVC pipe at the far left of figure 3.5b, in which case there is very little bending or torsional stress on the incipient catheter, or through the clear tygon shown descending through the vertical length of PVC, which induces much more bending, in theory comparable to what is present in the aortic arch, and therefore suitable for inducing the twisting observed in figure 3.4. The vessel itself was modeled with 1/4 inch diameter weather stripping purchased at a hardware store, selected for both its diameter, which is comparable to that of a coronary arterial lumen, and for its distinctive D-shape, which makes it

Figure 3.5: Phantoms used to verify catheter twisting artifact.

(a) Photo showing all trajectories used in phantom study.




(b) Photo showing phantom apparatus fully assembled.



easy to visually detect any twisting during the pullback.

The plastic enclosure depicted in figure 3.5b was filled with water to simulate blood. After that, the catheter was threaded into the phantom, while the x-ray machine was used to locate the catheter and the transducer within the phantom. During this procedure, a continuous feed from the x-ray was shown on the screen. Snapshots were taken at specific locations, for later off-line analysis. One such snapshot can be

Figure 3.6: Images from phantom study.



seen in figure 3.6a. The catheter is difficult to see, but is visible as a thin curved wire, snaking all the way up to the top anchor post. Once the position of the transducer was set (generally at the furthest point within the phantom that still left the transducer completely enclosed within the D-shaped weather stripping), an automatic pullback commenced, with periodic snapshots recorded on the x-ray machine in order to localize specific twisting events within the phantom. An image frame taken from the IVUS data is depicted in figure 3.6b. The sharp D-shape of the weather stripping is clearly visible in this figure. The catheter is positioned all the way in the tight notched edge of the tubing in this frame.

The phantom confirmed the twisting of the catheter during the pullback, and also that the twisting was not a function of any easily measurable parameter: even in the phantom with a straight trajectory, a 120 degree twist was observed in the IVUS image data over the course of the pullback, thus confirming the need for the methodology outlined in section 3.2.2.

3.2.2 Methods

The process of angio-IVUS fusion described in [5] uses a Frenet-Serret model of the catheter to estimate the position of the IVUS frames relative to each other, and the torsional motion that takes place between them. This model is described in detail in [39], and is accurate as long as the catheter behaves in the way that a rigid model would suggest. Sheathed catheters typically do this reasonably well; the unsheathed catheters made by Volcano Corp., used in this study, do not. We are presently unaware of any means of analytically predicting this behavior, and thus we have explored a method of correcting it through the use of anchors and tie points.

An illustration of the unnatural twisting can be seen in figure 3.4. The figure depicts plaque thickness values, which should be correlated with vessel curvature, but here take on a spiral shape that suggests the catheter orientation is erroneous.


The only reliably available anchors present in both the angiograms and the IVUS images are the bifurcations described in section 2. The method used for angiographic segmentation relies both on contrast dye and on a catheter with high x-ray opacity, so that traditional dynamic-programming 2 dimensional graph search techniques can be applied reliably. When the angiogram meets both these criteria, the segmentation proceeds with minimal error and virtually no need for manual intervention. When one or the other is lacking, a tedious procedure of observing the video of the raw angiogram and then manually inferring the correct positions of the borders and catheter is required. Such pullbacks are often the ones with the greatest uncertainty with regards to the absolute orientation, thus even in the absence of twisting, this method should be able to improve the quality of the reconstruction.

Segmenting bifurcations is difficult because there is no catheter, they are frequently narrow enough to limit the usefulness of the contrast dye, and, more significantly, they are often badly occluded in the image by the vessel of interest, or suffer from foreshortening artifacts due to the fact that the projection planes were chosen

with respect to the main vessel, not the bifurcations. For this method to be made useful, the image acquisition protocol would have to stipulate that the projections used capture a minimal number of usable branches.

“Segmenting” the bifurcations in the angiograms as such is infeasible. We instead have devised a method that relies upon identifying two points along each bifurcation in each plane (“proximal” and “distal” respectively), ideally matched in both planes. To compensate for the inevitable ambiguity of identifying peripheral vessels in an angiogram, epipolar projection is used to assist in matching points between the two projections, with the transformation initialized based upon the parameters of the original vessel segmentation. Figure 3.7 illustrates the interface. Depicted are the angiograms of one of the patients in the study for whom one of the branches is clearly visible in both projections. The yellow squares are the cursor used to identify the branch locations; the yellow lines are the epipolar projection of the cursor from the other pane, which in this case provide visual reassurance that the same branch is being marked in both.

Figure 3.7: Software interface used to identify branches in angiograms.



In this interface, the green lines denote the previously identified borders of the main vessel. The red line is the segmented catheter path. The red square is the estimated proximal position of the bifurcation, or the location of the junction based upon its previously identified location from the IVUS as marked with the method of section 2, and the yellow square denotes the estimated distal end of the branch. The software locks the position of the proximal location to the segmented catheter, whereas the user is free to move the distal point anywhere in the image. The straight yellow line is the epipolar projection of the distal point from the opposite pane. The software does not force the user to place the cursor along this line, but merely draws it as a guide, to provide a valuable clue as to which vessel-like object in the image is most likely the bifurcation of interest.

Algorithm 3 Procedure for estimating localized twisting from bifurcations.

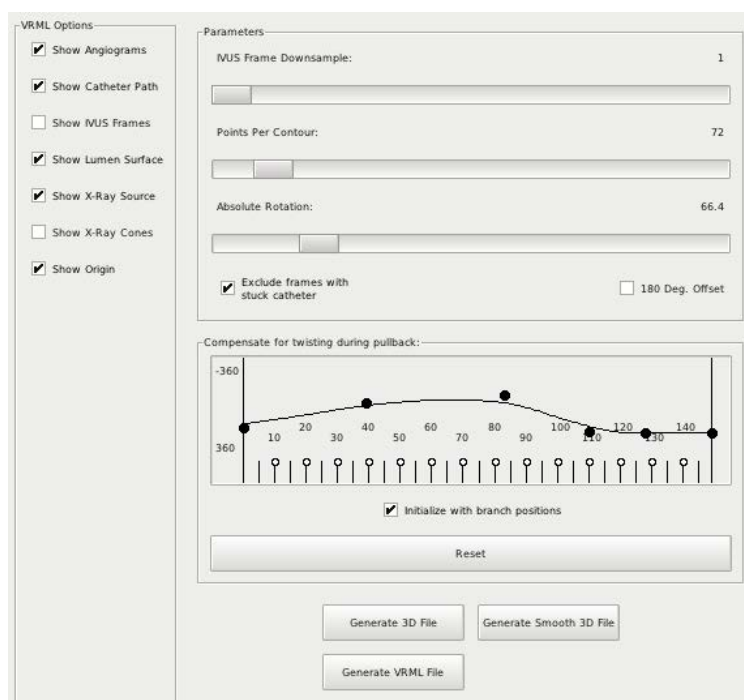
```

CTRLPTS  $\leftarrow$   $\emptyset$ 
for all BRANCH do
  Z  $\leftarrow$  DISTAL_FRAME(BRANCH)
  DISTAL_3D  $\leftarrow$  RECONSTRUCT_PT(DISTAL_2Da, DISTAL_2Db)
  DISTAL_VEC  $\leftarrow$  DISTAL_3D - CATH_PT[Z]
  PROJECTIONu  $\leftarrow$  DISTAL_VEC · SVEC[Z]
  PROJECTIONv  $\leftarrow$  DISTAL_VEC · TVEC[Z]
   $\theta_{angio}$   $\leftarrow$  atan2(PROJECTIONv, PROJECTIONu)
  NEW_CTRLPTu  $\leftarrow$  Z
  NEW_CTRLPTv  $\leftarrow$   $\theta_{IVUS}$  -  $\theta_{angio}$ 
  NEXT(CTRLPTS)  $\leftarrow$  NEW_CTRLPT
end for

```

Traditionally, the IVUS frames have been oriented along the catheter path according only to the Frenet-Serret torsion result as described in [5] and [39]. Once bifurcations have been identified in the angiograms, however, an additional step is inserted, prior to the computation of the absolute orientation of the vessel. The application that is used to perform the 3-D reconstruction is depicted in figure 3.8;

Figure 3.8: Application used for 3-D reconstructions. Spline curve for detwisting is visible towards the bottom of the panel.



the spline curve estimated for reversing the twisting, along with the control points corresponding to marked branches, is depicted in the lower half of the panel.

First, a list of branches that have been identified in both the IVUS and in the angiograms is compiled. The procedure described in algorithm 3 is run on the set of branches to produce a set of control points for a spline curve. This spline curve defines the angular deflections of each frame along the IVUS pullback that are to be performed prior to the 3-D reconstruction; a spline curve is used to guarantee smooth progression of twisting over the course of the pullback that is minimally disruptive to the ultimate 3-D model.

3.2.3 Experimental Methods

The *curvature index* is a morphologic parameter defined in [17], and reprinted in equation 3.10, that measures the projection of the curvature of the vessel centerline

onto the lumen surface at each sampled point in the 3-D reconstruction. Whereas the point-to-point registration is validated using differences in the IVUS borders prior to 3-D reconstruction, the curvature index provides a unique window into a more complete picture of the accuracy of registration which takes into account the 3-D uncertainty we attempt to address in this section.

\vec{T} and \vec{N} are the tangent and normal vectors with respect to the surface; κ_j is the curvature value at a given frame; $\bar{\kappa}_n$ is the mean curvature taken over a 23 point window at frame n . The magnitude of $\bar{\kappa}_n$ is multiplied by the cosine of the angle between the tangent and normal surface vectors to get the curvature index value ($\dot{\kappa}_n$) at a given frame.

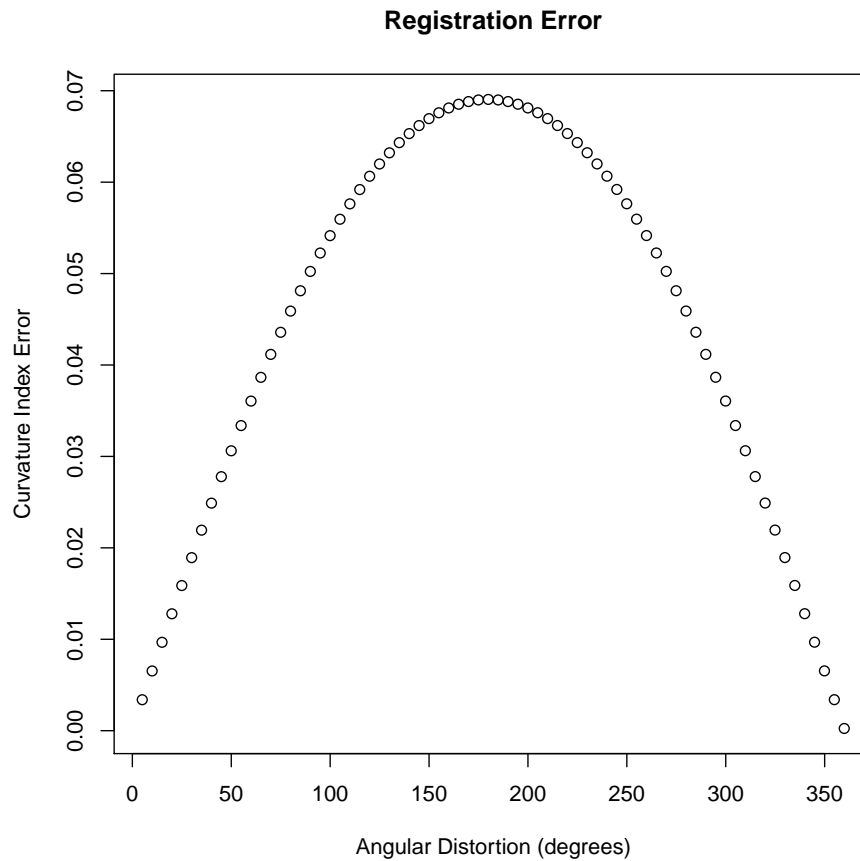
$$\bar{\kappa}_n = \frac{\sum_{j=n-11}^{n+11} \kappa_j}{23} \quad (3.9)$$

$$\dot{\kappa}_n = |\bar{\kappa}_n| \times \frac{\vec{T} \cdot \vec{N}}{|\vec{T}| |\vec{N}|} \quad (3.10)$$

Figure 3.10 illustrates the density function of the curvature index for registered and unregistered data sets. While the bandwidth is narrowest in the registered case, the extrema are far further from the center of the distribution. This is because in cases where uncompensated variation in catheter orientation, or uncertainty in absolute orientation, the baseline and followup are at times seen to bend in opposite directions. Nonetheless, the sharpness of the peak can be taken as an indication of the correctness of the 3-D alignment of a pair of pullbacks.

Any change in the orientation of specific IVUS frames will have a significant impact upon the absolute orientation as computed by the methods of [39]. If the frames are reoriented correctly, they should improve the accuracy of the final result – the goal is not to change the overall approach, but rather to remove a previously unaccounted for source of error, and permit the absolute orientation detection algorithm to work

Figure 3.9: Curvature index error, measured as a function of rotational misalignment.

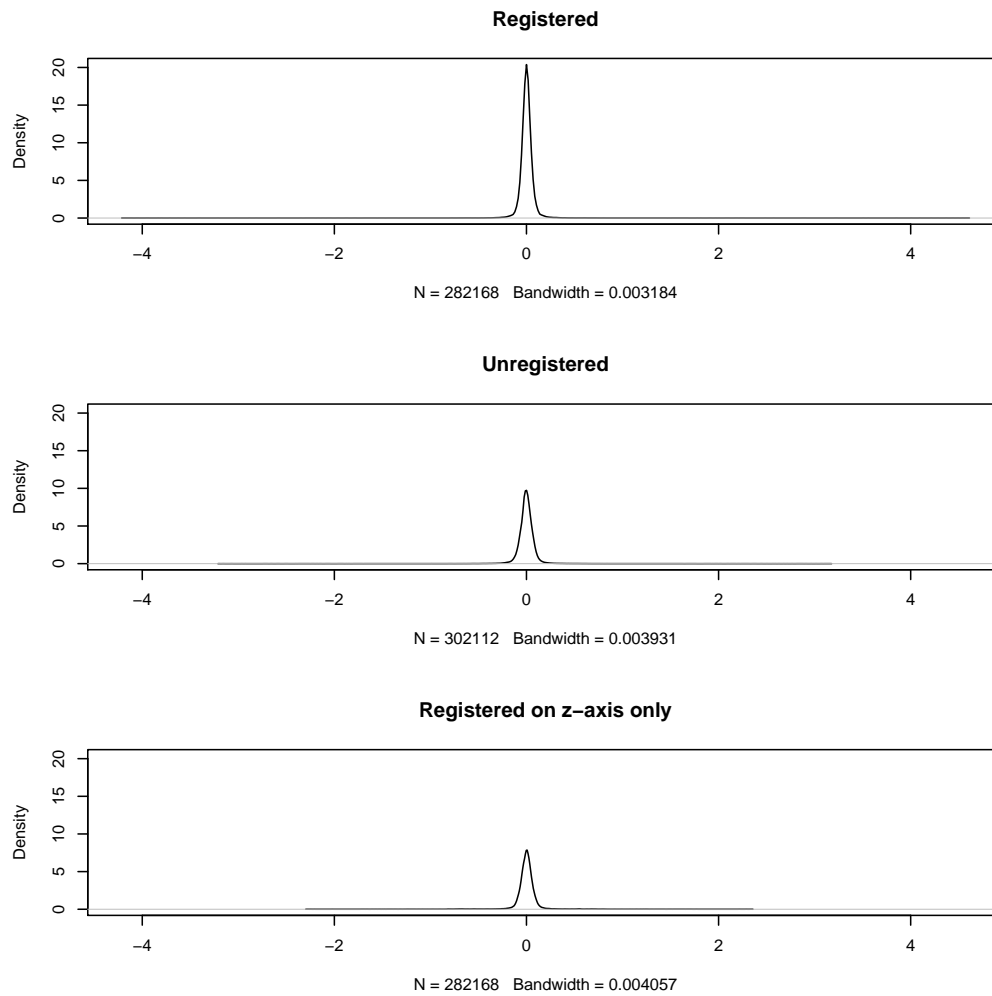


correctly.

Because the reliability weight is based in part on how far out of center the catheter is in a given frame (and therefore how likely it is that the difference angle can be meaningfully computed), it's unlikely that the reliability weight will change significantly. Rather, if the approach described in section 3.2.2 works correctly, we should expect that anywhere the reliability weight is reasonably high, the difference angle should be consistent, *i.e.* there should not be significant fluctuations among difference angle values perceived as reliable.

Because it is closely coupled with the 3-D reconstruction, the curvature index has limited use in validating the mapping described in section 3.1. However, it provides

Figure 3.10: Density plot for curvature index error taken over the entire data set.



a useful metric in assessing the effectiveness of the detwisting procedure described in section 3.2.2. Figure 3.9 illustrates the average difference in curvature index, computed over all pullbacks, for a given angle of displacement (e.g., each pullback is compared to a rotated version of itself over all points). The shape of the curve illustrates that the error is strictly a function of rotational misalignment in this limited case.

Figure 3.10 also shows a plot of curvature index presented in the same way as figures 3.2 and 3.3 were in section 3.1.3, with *registered* versus *unregistered* versus

z-only registration over all patients. It is less persuasive than figures 3.2 and 3.3 because the uncertainties in the 3-D reconstruction and any potential misalignment directly impact the result seen there.

Difficulty identifying bifurcations in the available angiograms has limited our ability to employ this technique in this study. However, to illustrate the effects of its use, we have chosen 6 baseline-followup pairs for demonstration. For these 6 pairs, we have created a table showing each of the absolute orientation and RMS positioning error from the 3-D reconstruction both with and without the detwisting respectively, and the mean and standard deviation of the registration error for curvature index for the pair (tables 3.2a, 3.2b). CI refers to curvature index. RMS error and absolute rotation are in degrees; CI is unitless. We also plot the density curves for the registered detwisted curvature index against that of the unregistered, *z-only* registration, and registered non-detwisted data for the 6 pair subset in question (figure 3.11).

Figures 3.12, 3.13 also show the reliability weight and correction angle values plotted over the entire pullback for both pullbacks for the patient that inspired the development of this method (see figure 3.4) for both regular and detwisted versions. In the course of the reconstruction, the estimated correction angles at each position along the pullback are averaged with weights corresponding to their reliability weight value, to determine the final global absolute orientation value, depicted as the green horizontal line on each plot. Reliability weight values are low if the segmentation borders at a particular location are centered and circular, because at those locations there is very little to differentiate different possible orientations. A reconstruction can be considered to be of high quality if the estimated correction angle values are all very similar at all locations where reliability weight is high.

3.2.4 Results

Figure 3.11 shows density curves for the detwisted and nondetwisted subsets used in this section; it also includes the unregistered versions for comparison, illustrat-

Table 3.2: Detwisting error statistics.

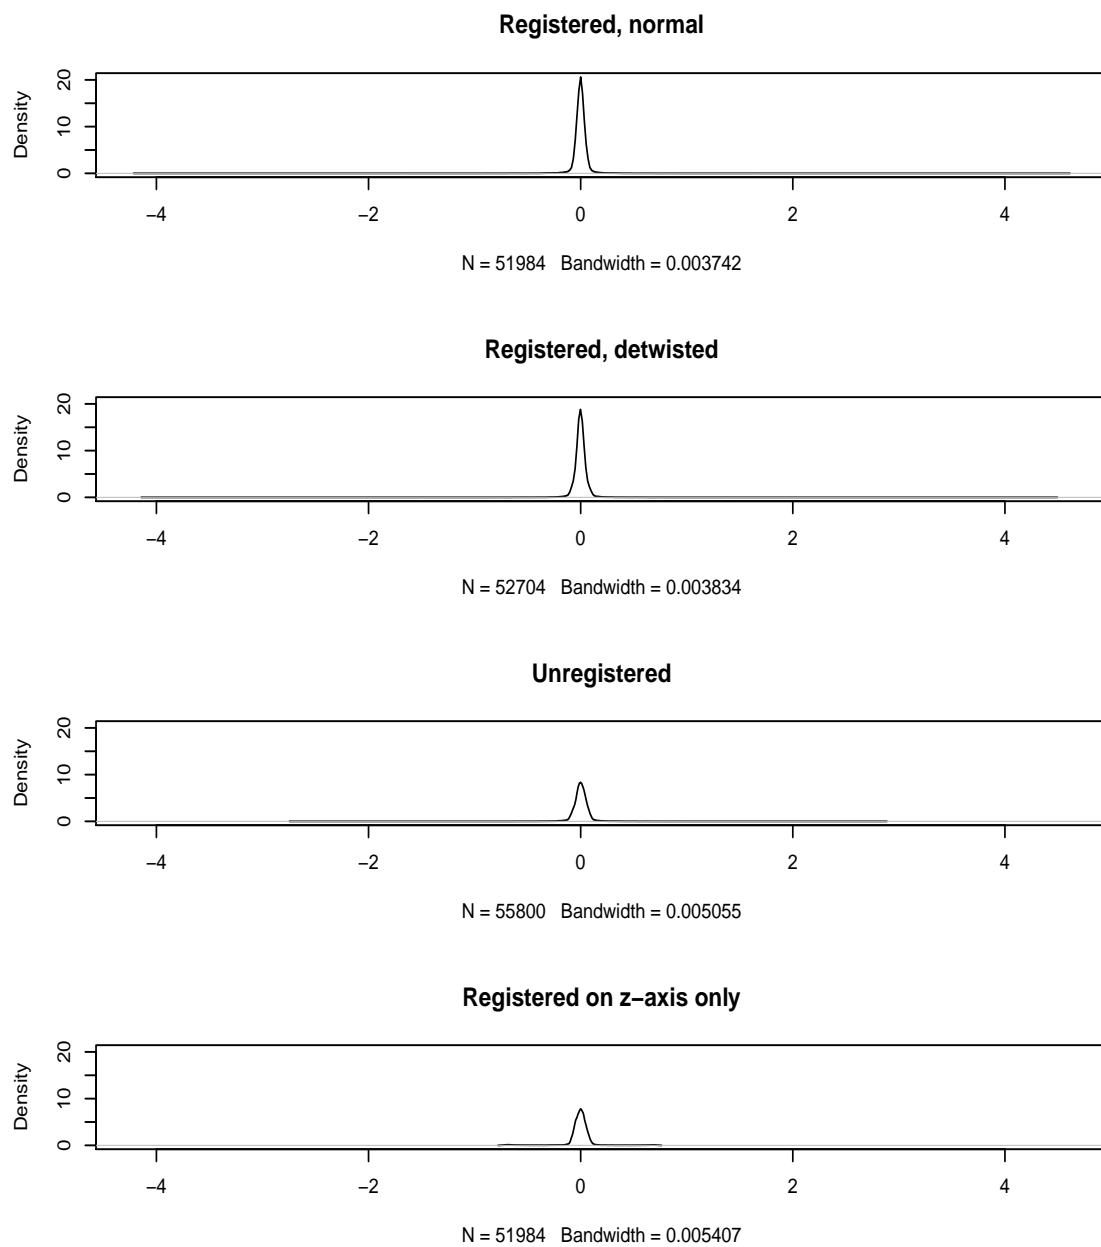
(a) Detwisted pullbacks.

Patient	Study	Abs Rot	RMSerr	CI Mean	CI SD
prague17	Baseline	36.69	83.47	-4.25E-04	3.82E-01
	Followup	-71.23	52.41		
prague28	Baseline	64.67	60.42	-3.25E-03	2.25E-01
	Followup	-71.36	37.42		
prague56	Baseline	-38.19	74.38	-2.85E-03	1.74E-01
	Followup	69.37	43.14		
prague69	Baseline	242.4	55.17	-2.42E-03	1.56E-01
	Followup	104.3	64.07		
prague100	Baseline	178.7	66.89	-1.93E-03	1.39E-01
	Followup	-15.18	64.55		
Prague18	Baseline	95.15	9.656	-1.96E-03	1.33E-01
	Followup	-74.2	43.77		

(b) Non-detwisted pullbacks.

Patient	Study	Abs Rot	RMSerr	CI Mean	CI SD
prague17	Baseline	34.31	89.71	1.37E-03	3.97E-01
	Followup	188	67.71		
prague28	Baseline	169.7	68.26	-1.38E-03	2.34E-01
	Followup	-39.83	36.53		
prague56	Baseline	2.974	80.49	-1.90E-03	1.81E-01
	Followup	74.71	41.01		
prague69	Baseline	11.91	70.59	-1.41E-03	1.61E-01
	Followup	151	50.36		
prague100	Baseline	30.53	84.5	-1.14E-03	1.43E-01
	Followup	-11.29	59.69		
Prague18	Baseline	24.85	66.4	-1.36E-03	1.39E-01
	Followup	194.1	82.12		

Figure 3.11: Curvature index density plot for detwisted versus nondetwisted reconstructions.



ing that both the detwisted and the nondetwisted versions perform better than the unregistered versions.

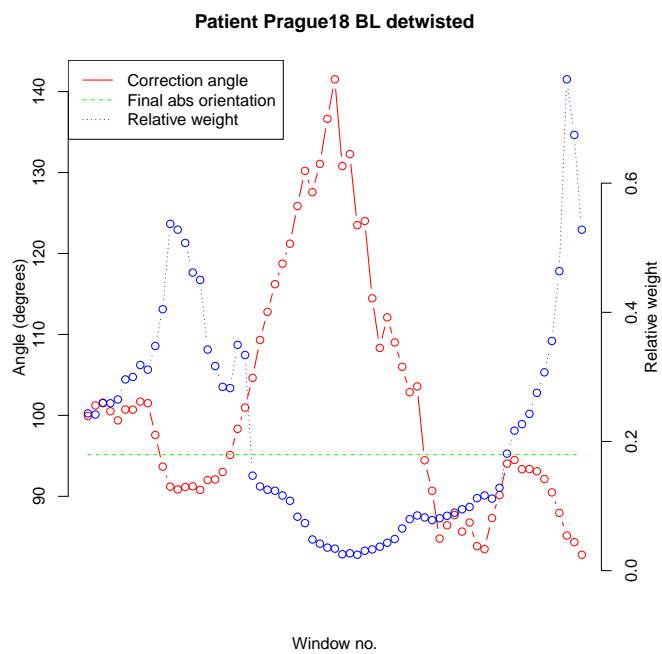
Table 3.2a shows the curvature index and reconstruction error statistics for the detwisted versions of the 6 pullbacks; table 3.2b shows the same for equivalent non-detwisted versions of the pullbacks in question. The RMS error reflects the 3-D reconstruction algorithm's confidence in its positioning of the pullback in space, and of the vessel borders with respect to the centerline. It is a useful, but not infallible, metric of the correctness of the reconstruction. The mean value over all the detwisted pullbacks is 54.6; the mean value over all non-detwisted pullbacks the mean RMS error is 66.4. A paired t test comparing the two columns gives a p - value of 0.13.

Between this relatively high p - value, and the fact that the sharpness of the curvature index density plot (figure 3.11) is reduced, the overall result is somewhat inconclusive. Some pullbacks saw a noticeable reduction in error, others much less so.

Figures 3.12, 3.13 show reconstruction error statistics juxtaposed with the final absolute orientation value for each for the detwisted and non-detwisted versions of the baseline and followup for patient #18, which is where the twisting artifact displayed in figure 3.4 was first observed. As can be inferred from the description of these plots at the end of section 3.2.3, because these are all computed on actual patients in-vivo, they are more a measure of confidence in the geometric model than of measurable error with regard to objective physiological reality. In both the baseline (fig 3.12) and followup (fig 3.13), this consistency has been improved, which indicates that we accomplished our goals with the detwisting procedure for this particular patient.

Figure 3.12: Correction angle shown alongside reliability weight, baseline reconstruction.

(a)



(b)






Figure 3.13: Correction angle shown alongside reliability weight, followup reconstruction.

(a)



(b)



CHAPTER 4 DATA ANALYSIS

4.1 Data Model

Much of the structure of the data hierarchy reflects the long history of the project; nonetheless, efforts have been made to modernize in order to improve the efficacy of the analysis and data mining stages of the project. A variety of XML files and implicit relationships have been merged into a relational database schema. The database has been fully normalized to *third normal form* (3NF), which loosely means that every column in a table contains information that uniquely applies to a single row in the table. 3NF is a common software engineering best practice, which makes it easier to enforce data consistency, and optimize search speed. Cached views are used to maintain read performance, given the heavy read-bias of the database in this situation. The advantages of this configuration include:

1. One unified repository for metadata eliminates versioning and permits sharing across multiple users.
2. Flexibility with regard to change management: if a new column is added to a table, existing software and data continue to work; moreover, because all records are contained in the altered table, a sensible value can be added to all existing rows when the column is added. This was much more difficult to do reliably with the XML based system.
3. Searching the patient metadata becomes vastly easier, because searching is intrinsic to the nature of a relational database, whereas previously it required parsing many XML files and traversing one-off data structures.
4. Because searching is easier, identifying relationships among the data, and grouping patients based on metadata criteria for downstream analysis becomes trivial.

Figures A.1-A.4, in appendix 7, illustrate the schema.

Many of the relationships in the schema embody the “has-a” relationship. In summary, a “pullback” consists of one or more angiographic and IVUS data sets. The pullback itself should actually consist of a single data acquisition; multiple angio and IVUS imports allow for some experimentation with gating and synchronization, which are on some level an inexact science in this case (this owes to the fact that any errors, both human and equipment related, in the cath lab, can make identification of the R-peaks difficult in both mediums). Likewise, multiple segmentations of each in the data model permits comparison of automated and manual methods, as well as permitting a human tracer to save multiple versions in the case of uncertainty.

4.2 Quantitative Indices

4.2.1 Methods

The 72 *wedges* described in section 1.4 will be referred to in this section also as *elements*. VH (sec. 4.2.1.2) and WSS (sec. 4.2.1.3) indices were computed on an element-by-element basis, with exactly 1 value for each wedge situated between a pair of points in the 3-D reconstruction. Most morphologic indices were as well, with the exceptions being curvature (a metric of the centerline, and therefore with meaning only on a slice-by-slice basis), and eccentricity, which is the ratio of the thickest plaque element in a slice to the thinnest plaque element in a slice, and therefore also has meaning only at the slice level. Thin cap fibroatheroma (TCFA, sec. 4.2.1.4) was also computed only on a slice-by-slice basis.

4.2.1.1 Morphology

The morphologic indices include cross sectional area, surface area, and volume computations of the lumen, plaque, and full vessel for each element. In a straight cylinder, this would be redundant, but the volume and surface area measurements reflect the shape of the 3-D reconstruction. Volume values are computed using the polytope method described in [40]. [16] evaluated 3 competing methods for imple-

mentation of the volume computation and the polytope method was consistently the most accurate.

Eccentricity, defined as P_{Tmax}/P_{Tmin} , where P_{Tmax} is the maximum plaque thickness for a slice, and P_{Tmin} is the minimum plaque thickness for a slice. Eccentricity is, therefore, defined as a “per slice” metric.

Curvature is defined per slice, while curvature index, used in 3.1.2 to validate the registration, is the projection of the curvature onto the lumen surface at each point, and as such is defined for each of the 72 elements per slice. The computation of both of these indices is described in [17].

4.2.1.2 *Virtual Histology*

Virtual histology indices are computed by a 3rd party proprietary software package provided by the acquisition system vendor (Volcano Corp). The software is based on the work of [9]. It uses the raw RF signal from the ultrasound to infer plaque composition, and classifies each voxel as belonging to one of fibrous, fibrous fatty, necrotic core, or dense calcium. A thin “media tissue” layer is provided for at the edge of the EEL border.

According to Volcano’s representatives, fibrous fatty translates into the softer “lipid pool” plaques; fibrous is the firmer, more heavily fibrotic plaque. Because dense calcium often effects a partial or complete reflection, plaque classification behind a large calcified block is suspect [41],[42]; nonetheless, the classification supplied by the VH software is the best information we have regarding plaque composition.

The VH in this study is computed using an older version of Volcano’s software which is no longer sold by them but which is still in fairly wide clinical circulation. This program provides no intrinsic hooks for interaction with third party software, thus some creativity was required to properly make use of the data. The company supplied us with sample code for reading the timestamp offsets in their RF data file; no other details of the format were provided, thus we could not analyze the RF, but

could retrieve the gating from the RF file and use the embedded timestamps to load our segmentation into their program.

The procedure involved in acquiring the VH data from the Volcano software package for these data sets is as follows:

1. Import images from DICOM using timestamps from RF data file.
2. Segment IVUS images, refining the segmentation manually to produce an acceptable border.
3. Run the Volcano software package.
4. Import the required data set.
5. Step through the Volcano workflow until the Volcano software package's segmentation procedure has completed. This step is required to ensure that a variety of files whose format we have no documentation of and which we do not have the ability to manipulate are in the required state for computation of the VH.
6. Resample our segmentation to the spacing used by Volcano. (Our software uses a radial format, centered at 0, with 360 points per frame, per border, with units in pixels. Volcano's uses 512 points, scaled from 0..1, with the image center coinciding to 0.5.)
7. Overwrite the Volcano borders with the resampled version of our segmentation.
8. Reload the data set in the Volcano program. Step through the workflow all the way to completion (because a segmentation is detected, it will not attempt to resegment, but will merely reload the borders, and then proceed to the VH analysis).

9. Import the result of the VH analysis into our patient data repository; annotate the database to reflect the presence of an up to date VH analysis.

We compute percentages of each plaque type both for each slice, and for the 72 elements composing each slice; further composite indices as described in 4.2.1.4 are computed using the individual voxel classifications along with their adjacency and proximity to the lumen.

4.2.1.3 Wall Shear Stress

For computation of the WSS, first a surface mesh is generated, using the procedure outlined in 2.1.3.

The triangulation procedure is wrapped in a python loadable module, which connects in a software pipeline with the meshing facilities of the Vascular Modeling Toolkit (VMTK), [43]. After the triangulation is completed and the endcaps are stripped, the surface, along with the centroids of each of the inlets/outlets, is passed downstream. The procedure has been fully automated, to be entirely turnkey such that a job is dispatched for a particular reconstruction, and the results appear in the database upon completion.

1. Create surface/identify boundary centerpoints (procedure from section 2.1.3).
2. Build centerline tree with radius values for main vessel and all branches (VMTK).
3. Create volume mesh, using centerline and radius information to adaptively vary tetrahedra size, thereby optimizing running time and accuracy jointly (VMTK).
4. Get distance from each boundary point to boundaries identified in the volume meshing. Use this information to correctly determine the flow inlet (python script developed by us).


5. Generate a batch file script for running commercial solver *Fluent*, with initialization and boundary conditions properly configured (python script developed by us).
6. Generate a final mesh output file compatible with *Fluent*, so the job can be launched automatically (VMTK).
7. Launch fluent; wait for it to terminate successfully.
8. Postprocess: interpolate the *Fluent* output back onto the original structured grid of the vessel from the original 3-D reconstruction.
9. Compute the component of the WSS that lies tangential to the vessel wall, to be used as a classification feature ([19],[20],[21]).
10. Write out a final file in the patient directory hierarchy; add an entry to the relational database for the WSS data for easy retrieval.

The output of the solver is interpolated back onto the 72 points per slice from the original 3-D reconstruction, and then the tangential component of WSS is projected onto the lumen surface. This final value, termed TWSS from here forward, is the terminal value of the WSS indices as we use them. The details of this computation are described in [19],[20],[21].


Figure 4.1 shows the final TWSS output projected onto the vessel wall for both the baseline and followup pullbacks for one of the patients. The coloring is bracketed at -5 Pa (blue) to 8 Pa (red). Most of the values lie between those two extrema at small values (yellow/green). The TWSS values are not computed for the bifurcations, thus they are uniformly green by virtue of having TWSS values of uniformly zero. The blue ring near the proximal end of the baseline (fig. 4.1a) indicates the proximal point of the registration, corresponding to the end of the followup image. The registration for this patient is distally aligned at frame 0 for both pullbacks, thus there is no

Figure 4.1: Wall shear stress, both pullbacks for 1 patient. Proximal end is top-right, distal end is bottom-left, for both.

(a) Baseline



(b) Followup



distal blue ring. The red directly inside the branch is a reflection of the fact that WSS within branches is relabeled to infinity in post processing in order to make the branch locations easy to visually identify.

In both images, a sharp drop in WSS immediately downstream of the middle branch can be seen, reflecting one of the more obvious impacts of the branch's presence on the WSS values. The fact that in spite of having the coordinate frame aligned between the two, the two vessels are clearly positioned slightly differently, and the discrepancy grows between the distal and proximal ends, illustrates the challenges and uncertainties involved with the 3-D reconstruction.

4.2.1.4 TCFA

A *thin cap fibroatheroma*, or TCFA, is a specific type of lesion known to have high risk of rupture and thrombosis, and therefore a high degree of implication in mortality caused by sudden cardiac arrest. Composite quantitative indices (considering VH and morphologic parameters jointly) based on the AHA's definition of a TCFA are also computed as features to be used in the classification system described in chapter 5.

From [44]: "Early in atherogenesis, the nascent fatty streak evolves into a transitional lesion, also known as the preatheroma. Death of macrophages, hypoxia, extravasation of erythrocytes, and free cholesterol replace lipid pools to form necrotic cores. Disease progresses through necrotic core expansion by active digestion of collagen, reducing the rim of fibrous tissue separating the core from the lumen. This advanced lesion is the hallmark of human atherosclerosis, and the American Heart Association termed it the fibroatheroma. When macrophage-derived collagenolytic activity ruptures the cap of the fibroatheroma, the highly thrombogenic core is exposed to circulating blood and triggers thrombosis [45]. In fact, when quantified by ocular micrometry, all ruptured plaques are characterized by a thin fibrous cap, ($<65 \mu m$ in thickness), a large necrotic core, and increased macrophage infiltration [46]. Adding these 3 characteristics to the American Heart Association definition of

fibroatheroma led to the distinction of TCFA as a separate entity [46].”

The principal difficulty with identifying TCFA lesions with IVUS is the axial resolution of the modality itself: depending on the frequency used for imaging, spatial resolution typically ranges between 100-200 μm . Given that the *thin cap* in question in the TCFA is defined to be $<65 \mu m$, the cap itself is by definition invisible. (OCT does have sufficient spatial resolution to see the thin cap, but lacks the VH plaque characterization capability of IVUS. Rathore et al. [4] have done work to combine the two, as an alternate means of getting around the limitation.) Due to this limitation, and based on the understanding that necrotic core plaque type will never abut the lumen directly, a VH-TCFA definition was created as follows [44]:

- Necrotic core $\geq 10\%$ total plaque area.
- Total plaque burden $\geq 40\%$.
- Necrotic core in direct contact with the lumen.


This roughly matches the definition used by [47] in their effort to correlate TCFA with Framingham score. Their results were promising, if preliminary, and suggest that there is value in searching for these lesions in IVUS images. Working from this accepted definition as a starting point, and terming it *definition A*, we have attempted to refine and improve the selectivity and specificity of this definition with regards to positive identification of TCFA candidate lesions in the IVUS image volume. 4 classes of definitions, which can be further specified by constraining the portion of the atheroma examined, have been defined, as below.

(A) Same as the definition above.

(B) Same as above, but necrotic core behind a calcium shadow is excluded from consideration, in keeping with the observations of [41].

- (C) Only count necrotic core that is part of a connected component that abuts the lumen. Compute connected component labeling in 2-D, with an 8-neighborhood.
- (D) Same as (C), but using 10-neighborhood (*e.g.*, the logical OR of 6-neighborhood and 8-neighborhood, chosen because of the significant anisotropy of IVUS voxel spacing).
- (E) Same as (A), except that all necrotic core further than 1 mm from the lumen-lamina border is ignored.

Figure 4.2: Phantom image used to validate connected component labeling.



Connected component labeling used in definitions C and D is performed using a depth first search, with metadata attached to the vertices of the graph in order to permit a useful extraction of the underlying image locations and plaque types at the end of the graph search. Figure 4.2 provides an example of the phantom images

that were generated to validate the connected component labeling and the subsequent TCFA classification based on the labeling. The procedure was run on synthetically generated phantom images similar to figure 4.2. Both the numbers of connected components detected, and the TCFA status (e.g., true or false) was verified to match the expected values; in each case the result was correct.

4.2.2 List of Indices

In addition to the indices described in section 4.2.1, a variety of demographic data, serum biomarkers (collected at the time of the imaging, therefore potentially differing between baseline and followup) and genetic markers are compiled for each patient to permit division of the patients into subgroups as appropriate and to assist in finding meaningful correlations.

The overall list of indices is based on those outlined in [10]; the indices used in this study are listed in table 4.1. *Slice indices* refers to indices computed once for each IVUS frame. *Element indices* are computed once for each of the 72 points per IVUS frame computed in the 3-D reconstruction. Per patient and per pullback refer to demographic and chemical biomarkers that apply more globally. Many of the slice indices are aggregate (either sum or average) versions of the element indices. The TCFA indices are either evaluated to “true” or “false” per frame, and thus apply only to the slice indices. Eccentricity, as a ratio of max plaque thickness to min plaque thickness, also can apply only to a slice.

The indices that measure a volume or a percentage can be aggregated over any arbitrary region of the vessel as is useful and convenient. This is in fact done to some extent in section 5.1.2. The TCFA indices can be aggregated into what might be thought of as either a percentage or a confidence value, depending on perspective, where each “true” is represented with a 1 and each “false” as a 0, and the mean computed.

Some indices were not actually used in training and testing the classifier. These

are displayed in ~~strikeout~~ font. These include the alternative TCFA definitions (only definition E was actually used, based on a parallel study that found it to be the most discriminative), and the volume indices, which, based upon uncertainties involved in the 3-D reconstruction, were determined to be redundant with respect to the cross sectional area (CSA) indices.

Table 4.1: List of indices used in classifier

Slice Indices:	Element Indices:	Per Patient Indices:	Per Pullback Indices:
% Necrotic Core	% Necrotic Core	Sex	Cholesterol Total
% Fibrous Fatty	% Fibrous Fatty	Age	Cholesterol HDL
% Fibrous	% Fibrous	Smoking History	Cholesterol LDL
% Dense Calcium	% Dense Calcium	Current Smoker	Triglycerides
Volume Necrotic Core [mm ³]	Volume Necrotic Core [mm ³]	Diabetes	Cholesterol APOA
Volume Fibrous Fatty [mm ³]	Volume Fibrous Fatty [mm ³]	Hypertension	Cholesterol APOB
Volume Fibrous [mm ³]	Volume Fibrous [mm ³]	Hyperlipidemia	Marker IL6
Volume Dense Calcium [mm ³]	Volume Dense Calcium [mm ³]	Family History CAD	Marker VCAM
Mean TWSS [Pa]	TWSS [Pa]	Previous MI	Marker TNAlfa
TWSS Std Dev [Pa]	Lumen CSA [mm ²]	HO Polymorphism	Marker CD40
Percent Stenosis	Plaque CSA [mm ²]	ENOS Polymorphism	Marker CRP
Mean Plaque Thickness [mm]	Lumen Volume [mm ³]	Beta Blockers	Diastolic BP [mmHg]
Lumen CSA [mm ²]	Plaque Volume [mm ³]	ACE Inhibitors	Systolic BP [mmHg]
Plaque CSA [mm ²]	Plaque Thickness [mm]	Statins	
Eccentricity	Curvature Index	Race	
TCFA-Def A	Lumen Surface Area [mm ²]	Stable Disease	
TCFA-Def B	Vessel Surface Area [mm ²]	Stable Angina	
TCFA-Def C		Unstable Angina	
TCFA-Def D			
TCFA-Def E			

CHAPTER 5 CLASSIFICATION SYSTEM

The Weka [48] software suite was chosen for the classification task based on its rich library of classifiers and preprocessing tools, and its reputation and ease of use. A preliminary evaluation of the data found decision tree algorithms to perform the best consistently, when compared against backpropagation neural networks (using Weka's multilayer perceptron implementation), support vector machines (SVM), and a naive Bayesian classifier.

Because of the limited separability of the data, and in order to maximize performance of the classifier, a large variety of classifier designs were benchmarked against each other in an exhaustive search. The parameters varied were feature selection (absence/presence, and style, described in section 5.2), resampling rate (described in section 5.1), classifier style (random forests were evaluated on their own, J48 trees in conjunction with various meta-classifiers), and what will heretofore be referred to as *glob-size*. Glob-size is described in detail in section 5.1.2, but refers to the practice of aggregating samples into larger chunks as a means of data smoothing and noise reduction.

Each classifier was trained on the input data listed in table 4.1. Outputs were as specified in table 5.1. Per-patient indices are displayed in ~~strikeout~~ type to indicate that they were dropped from the classifier. The reason for this is that there were too few examples of any of these incidents to meaningfully classify. There were no incidences of death, hospitalization, or infarction, or any other clinical event deemed significant by the physician who collected the data.

The output values were quantized based on criteria described in section 5.1.4; they were also analyzed using a simple binary $+/-$ measure as an alternative standard. TCFA was assessed using both definition *A*, which is the most traditional standard, and definition *E*, which was found to be one of the most descriptive in a related, but

Table 5.1: Classifier output variables

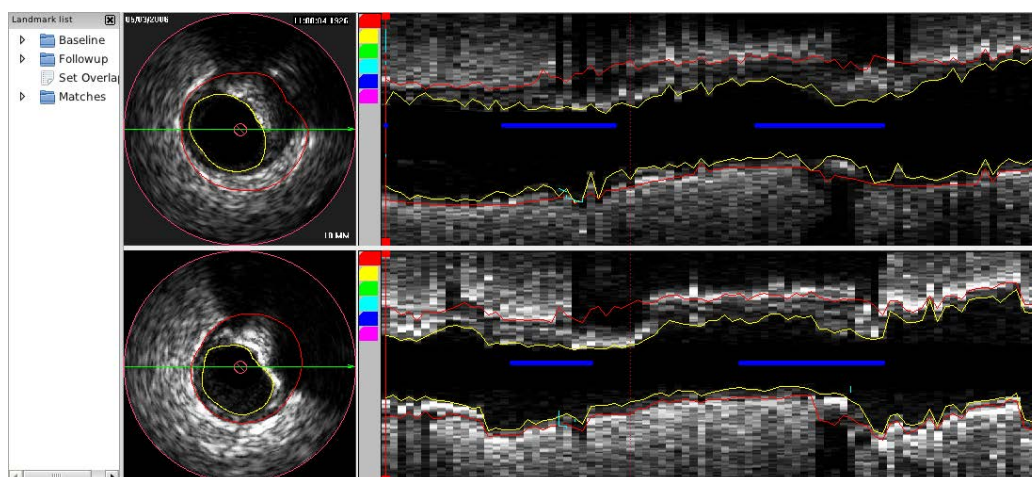
Region Output:	Global Patient Output:
Lumen Cross Sectional Area Change [++ / + / 0 / - / --]	Myocardial Infarction [Y/N]
Avg. Plaque Thickness Change [++ / + / 0 / - / --]	Cardiac Hospitalization [Y/N]
Fibrous % Change [++ / + / 0 / - / --]	Revascularization [Y/N]
Fibrous Fatty % Change [++ / + / 0 / - / --]	Clinical Event in 12 months [Y/N]
Necrotic Core % Change [++ / + / 0 / - / --]	
Dense Calcium % Change [++ / + / 0 / - / --]	
TCFA Confidence Value Change [++ / + / 0 / - / --]	

yet unpublished study, and was found to be present in a large minority of frames in this data set.

5.1 Preprocessing

5.1.1 Exclusion Marking

Figure 5.1: Tracing discrepancy between baseline and followup.



IVUS is notoriously difficult to segment automatically, and even among expert tracers, there is often disagreement on more heavily diseased frames. In particular, stents and calcifications reflect the ultrasound signal, resulting in shadows that occlude arterial features. While the inner (lumen/lamina interface) border may be

clearly visible, the outer (media/adventitia interface) border is impossible to discern. This in turn means that even if the same expert tracer evaluates both the baseline and the followup, there is a high probability of inaccuracy with regards to plaque thickness and any other metrics that depend upon it. Thus we label regions of the pullbacks where stents are present, and mark calcifications that produce shadowing sufficiently severe as to make meaningful tracing of the outer border impossible.

The interface used is completely manual, but simple and fast to use. The marking is performed using the same tool that is used to segment bifurcations. A dialog box for exclusion marking can be opened at the click of a button. The cursor is placed at the start and end frame, or, alternately, the user can type in frame numbers, and a button is clicked to enter the exclusion into the database. Previously marked exclusions for the currently opened IVUS segmentation are stored in a drop-down list in the dialog box used to mark them, and they can be modified or deleted from this interface. When an exclusion is marked, a color-coded bar appears in the longitudinal view, indicating the type of exclusion. (All landmarks, including branches, are marked in this way. The blue bars in figure 5.1 identify the locations of branches; calcifications are marked in yellow, and stents in white.)

Figure 5.1 illustrates the difference in tracings for the same heavily calcified location in the baseline and followup pullbacks for one patient. While it's clear the traced borders do not match, more problematic is the fact that one can plainly see that there is no discernible feature past the heavy calcification that can be fairly regarded as the border. The media/adventitia interface is completely concealed within the shadow, making accurate adjudication of the border impossible. Due to the fact that such locations make it impossible to meaningfully predict changes in plaque thickness – given that we have no confidence in our estimate of the border location – we exclude them from the training and testing data.

Table 5.2: Glob sizes used in preprocessing.

ag	sg
18	1
36	3
72	5
9	
3	

5.1.2 Globbing

While the registration performs well, there are nonetheless uncertainties present. Moreover, in cases where the gating is choppy, either due to cardiac arrhythmias or to poor signal, there is a high degree of spatial uncertainty within the pullback, never mind in its mapping to another. One important way of compensating for this while still making good predictions about changes in disease is to average or aggregate features spatially.

To perform the *globbing*, each data set was first viewed in terms of the resolution of the 3-D reconstruction: 72 wedges, or elements, as discussed in sections 1.4 and 4.2, along a ring, 1 ring per frame. A wedge therefore constitutes a 5 degree ($5 = 360/72$) angular subset along a ring, or slice, where a slice is defined as all the data corresponding to a single IVUS frame.

Globs were then constructed by aggregating in the angular direction (along the ring) and in the axial direction (across multiple frames). The sizes used can be seen listed in table 5.2, where *ag* indicates the size of the chunk around the ring (number of 5 degree wedges concatenated to form a larger wedge), and *sg* indicates the number of slices chunked together. Thus each glob covers a specific volume determined by sum of the number of angular wedges and longitudinal slices concatenated together.

The sizes specified in table 5.2 were combined in an all-to-all fashion, forming 15 total permutations of glob-size for the exploration of the classifier space.

Because the running times for classifier regressions were extremely long, and because in initial tests it never performed significantly better than the 3-wedge variety, the 1-wedge glob size was scrapped, and thus the smallest aggregate in the angular direction was 3, which corresponds to 15 degrees = $3 * (360/72)$.

5.1.3 Imputation of Missing Values

Some value must be present at each location in the training and testing vectors for both learning and classification to work. While patients were excluded from consideration any time image data was missing or unusable, for a variety of patients, some portion of the biomarker or demographic data was missing. In these cases, it was necessary to synthesize replacement values, and it was important that the replacement values not undermine the classification procedure.

Weka's default method of handling missing data is to simply select the median in the case of numeric columns, or the mode for categorical ones. The most recent implementations of random forests include a much more sophisticated method of filling in missing values. It begins by performing a rough replacement, using median and mode values. It then constructs a decision tree to compute proximity among samples, and attempts to infer the missing values based upon this.

This was the method used in this work to infer missing demographic and biomarker data (if missing data was detected in the other features, execution was halted, as it was considered to either be a problem with the data, or an indication of a software bug).

Also, because the granularity of the training and testing samples was at the level of individual locations inside of a vessel, there was the possibility that the random forests imputation method would produce contradictory values for demographic or biomarker information for a given patient. Thus a post-processing pass was used to homogenize the values by taking the mode on a per-patient level, so that the data makes sense in the context in which it was collected. For binary features (*e.g.*,

diabetes, hypertension), the definition of the mode was the same that is in common use: the most frequently occurring value. In the case of numerical features, there exists a wide range of methods in the literature for selection of a value that meets the definition of *mode*; in this case, a kernel density estimate, using a Gaussian kernel, was selected, with the value for all samples corresponding to a given patient taken to be the peak of the kernel density estimate.

In summary, the procedure is:

1. Initialize missing values to median/mode values.
2. Use random forests `rflmpute` function (from the R random forests package) to impute missing values into all samples.
3. Homogenize binary valued features using the most frequently occurring value.
4. Homogenize numerically valued features using kernel density peak (using the R *modeest* package).

5.1.4 Output Binning

To determine the actual quantization levels for the output values ($- - / - / 0 / + / + +$), somewhat arbitrary lines must be drawn to constitute what makes a *major change* versus merely a change, but more importantly, what constitutes a change at all, rather than variation within the noise threshold.

Figure 5.2 shows the distribution of the output samples for the 6 output features, tabulated for all the various glob-sizes. It can be clearly seen that a) all the output features have a slight negative bias and b) the samples for all features are heavily clustered around 0.

The locations of the cuts used to quantize the outputs are described in equation 5.3. The cuts were made somewhat arbitrarily; while we experimented with the width of the zero category, and found we generally got the best results with this configuration

(the mode of variation was principally widening and narrowing the 0 category). While the ++ and -- categories probably could have been expanded at the expense of + and -, given the data set's heavy bias towards 0 we wanted to maximize the + and - minority classes. While narrowing the 0 category would rebalance the data set, given the large uncertainties throughout the process of collecting this data, we feel it is important to be conservative in the labeling, erring on the side of estimating "no change".

The impact of the output binning upon statistical effect size and power is described in section 5.3.3.1.


$$\Delta = v_{max} - v_{min} \quad (5.1)$$

$$\sigma = \frac{\Delta}{9} \quad (5.2)$$


$$F(v) = \begin{cases} 0 & \text{if } v \in \left(-\frac{\sigma}{1.5}, \frac{\sigma}{1.5}\right], \\ + & \text{if } v \in \left(\frac{\sigma}{1.5}, \frac{8\sigma}{3}\right], \\ - & \text{if } v \in \left(-\frac{8\sigma}{3}, -\frac{\sigma}{1.5}\right], \\ ++ & \text{if } v \in \left(\frac{8\sigma}{3}, v_{max}\right], \\ -- & \text{if } v \in \left[v_{min}, -\frac{8\sigma}{3}\right] \end{cases} \quad (5.3)$$

Figure 5.2: Output feature density plots.

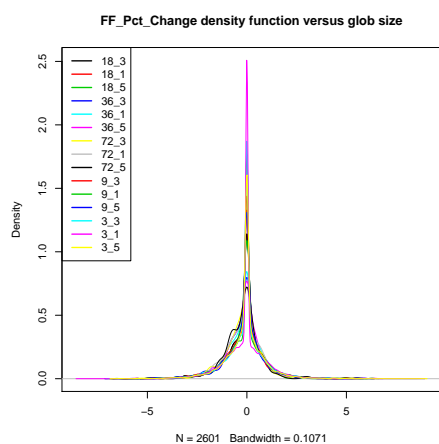
(a)



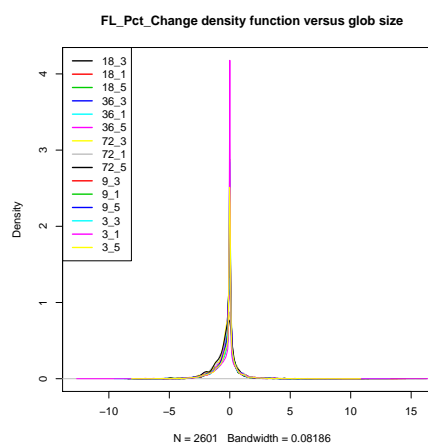
(b)




(c)



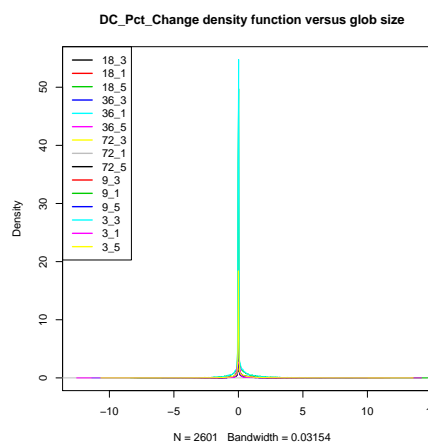
(d)



(e)



(f)



5.1.5 Resampling

Figure 5.2 makes it clear that the data is heavily biased towards a majority class at 0. The choice of glob size has a slight impact upon the location of the mean and the exact spread of the tails, but does little to alleviate the central problem. Because of this, resampling was employed to improve classifier performance. Most results in the literature indicate that undersampling leads to better performance than oversampling ([49]), for most problems. Chawla et al. [50] published an algorithm called SMOTe (Synthetic Minority Oversampling Technique), which attempts to improve the results in oversampling by spatially inferring new minority samples, rather than merely repeating existing samples. In our experiments, this did not work any better than traditional oversampling; this may be a product of the poor separability of this problem.

In our evaluation regressions, we compared for each classifier design a nonresampled data set, and sets resampled, with replacement (same total number of samples, with some fraction of majority samples replaced with repeated minority samples). The Weka resample filter was used. This filter combines undersampling of the majority class with oversampling of the minority class – termed *replacement*, each discarded majority class sample is replaced by a duplicate copy of a minority class sample. The one parameter that is set is called the *bias*, which determines the rate at which the under/oversampling takes place. A bias value of 0 indicates that the output is identical to the input, and a bias value of 1 indicates that the output will have a uniform class distribution.

We ran the filter at rate values of 0, 0.3, 0.6, and 0.9. For most output features, the resampled data sets led to better classification results, though the details of that are discussed fully in section 5.3.4.

5.2 Feature Selection

Feature selection was performed in this study for two reasons: first, it reduces the risk of either overfitting, or training on irrelevant features – thereby improving overall classification performance. The second reason is that it provides an opportunity to glean insights about the relationships in the data beyond mere Pearson correlation coefficients, by illustrating which features most effectively drive classifier performance.

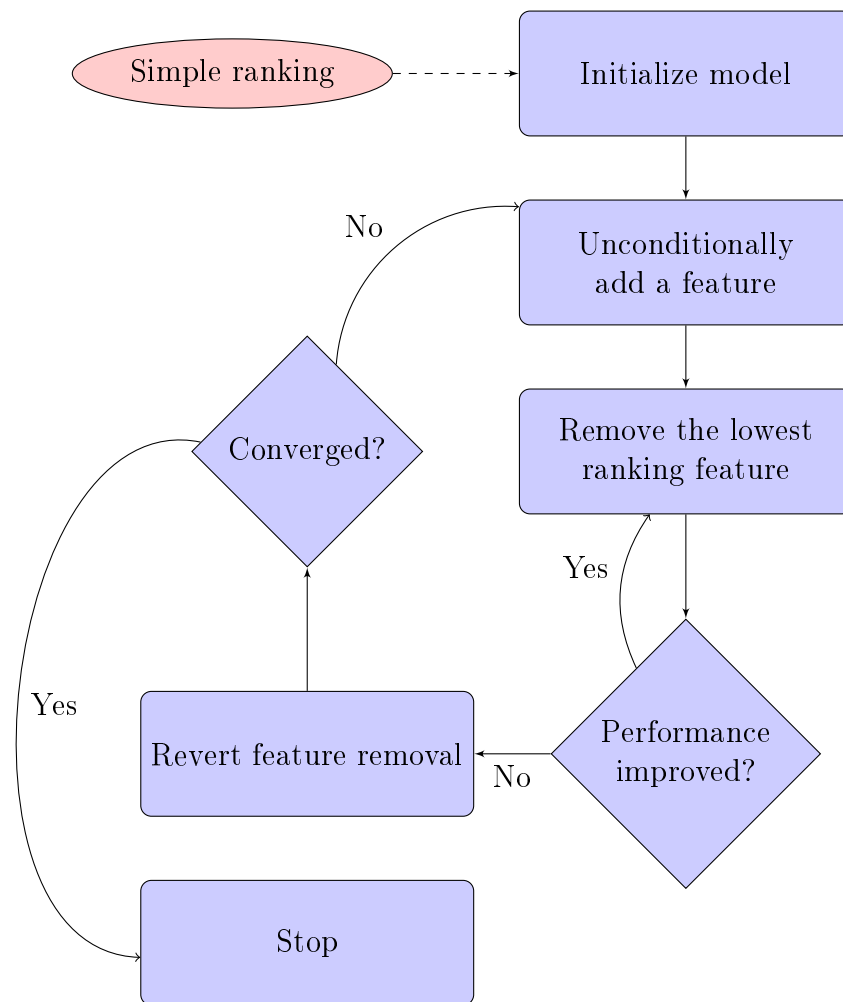
5.2.1 Methods

Pudil [51] demonstrated that robust feature selection cannot be always cleanly separated from classifier design. He outlines an iterative algorithmic procedure for optimizing a feature set against a specific classifier design, wherein features are added and removed until the classifier’s error rate is minimized, known in the literature as *Sequential Floating Forward Search*, hereafter SFFS. It continues to deliver some of the most robust results available, though it has now been expanded into a multitude of variant approaches [52].

In our case, we rely on Weka’s implementation of SFFS, implemented as the *linearForwardSearch* filter, implemented by Gütlein as described in his thesis [53]. This filter is used in conjunction with Weka’s *classifierBasedAttributeSelection* filter, evaluated on a J48 decision tree classifier [54]. (J48 is the open source, Weka implementation of Quinlan’s well known C4.5 algorithm.) The J48 algorithm was used with the confidence threshold for pruning set to 0.25, and the minimum number of instances per leaf set to 2.

The procedure is described in figure 5.3; convergence criteria is based jointly on total number of features selected (or *width*, meaning fraction of available features), and number of iterations performed. The procedure terminates when the selected number of features is within the range of total desired features, and removing additional features yields no improvement.

Figure 5.3: Flowchart describing feature selection procedure.



5.2.2 Experimental Methods

We divided the data set into 32 cross-validation folds, one for each patient. This ensured that generalization was at a per-patient level, and eliminated the risk of training on testing data. A full evaluation run of the classification system was performed prior to performing feature selection, as described in section 5.3.3.

The glob sizes, resampling rates, and classifier types that were most successful in this initial run were each then used in the feature selection run, to explore whether performance could be improved with feature selection. While classifier performance was not output by the feature selection code, the frequency with which each feature appeared in the final subset of the feature selection was stored for each output variable, as seen in table 5.3.

A visualization of inter-feature correlations is shown in figure 5.4 as well. The eccentricity of the ellipse indicates the degree of correlation (hence the main diagonal is all flat discs, as the correlation is uniformly unity), and the direction in which the ellipse is pointed indicates the direction of the correlation. The color varies from red to blue, with red indicating negative correlation, and blue indicating positive correlation. While this is also illustrated with the shape and canting of the ellipses, the color makes it easier to see, given the small size and large number of cells in the matrix.

5.2.3 Results

Table 5.3 shows the frequency with which each feature was selected for each of the output variables. In the case of plaque thickness, it can be seen that each feature was selected at least once; the number of non-zero rows for each output variable were: 36 for plaque thickness change, 35 for lumen CSA change, 24 for fibrous % change, 32 for fibrous-fatty % change, 16 for necrotic core % change, and 30 for dense calcium % change.

Figure 5.4: Inter-feature correlations for all inputs and outputs. Red indicates negative correlation, blue indicates positive.




Table 5.3: Feature selection summary.

	PT	LCSA	F Pct	FF Pct	NC Pct	DC Pct
lumenCSA	14	15	0	3	0	1
plaqueThickness	36	6	5	13	2	5
TWSS	13	5	0	2	0	2
curvatureIndex	23	6	1	3	0	2
eccentricity	13	10	0	2	1	2
FF Percent	28	6	0	3	0	0
FL Percent	20	6	1	13	0	1
NC Percent	13	2	3	2	3	2
DC Percent	20	5	4	3	2	6
TCFA Def E	17	0	3	1	1	2
smokinghistory	5	1	1	1	0	0
currentsmoker	13	2	0	0	0	2
diabetes	10	1	1	2	0	0
hypertension	16	5	3	2	0	5
hyperlipidimia	4	2	1	3	1	3
familyhistory	7	4	0	3	2	2
previousmi	16	2	2	4	1	1
hopolymorphism	3	2	2	2	0	2
enospolymorphism	16	4	2	1	0	3
betablockers	22	5	3	0	1	4
acinhbitors	14	4	5	5	2	1
statins	15	4	1	1	0	1
aggressivetreatment	19	7	2	2	0	1
sex	19	7	3	4	2	1
choltotal	3	6	0	2	0	0
cholhdl	8	2	0	2	0	1
cholldl	5	1	1	1	2	1
triglycerides	7	1	1	1	0	1
markeril6	9	7	0	4	2	0
markervcam	5	1	2	1	2	1
markercrp	9	2	2	3	0	3
cholapoa	12	5	0	8	0	0
cholapob	10	4	0	1	0	2
markericam	9	2	0	0	1	3
markertnfalfa	6	2	3	3	1	1
markercd40	4	1	1	0	0	1

PT=plaque thickness change LCSA=lumen CSA change F=fibrous % change
FF=fibrous fatty % change NC=necrotic core % change DC=dense calcium % change

Figure 5.4 shows a graphical representation of a correlation matrix. Few strong correlations exist between the *per-element* features, such as plaque thickness, lumen

CSA, VH percentages, or the change in any of these – and those that exist are largely between the *change* versions of those indices and the raw input versions thereof. Many of the strong correlations depicted in the chart exist between pairs of biomarkers and/or demographics. This is unsurprising, in the sense that these are among the least noisy indices in the data set, but can still provide for interesting insights.

While few features showed strong correlation with plaque thickness change in figure 5.4, those with the strongest (most elliptical) cells in figure 5.4 were also among the most frequently selected in table 5.3. This pattern breaks down with some of the other outputs, where all features were selected relatively infrequently. This may owe to the fact that the individual test folds, taken at a per-patient level, were overly specific, and led to an over-fitting of the feature selection. (There are folds in the cross validation that, for a given output, contain samples restricted to a single class. This in turn can lead the feature selection to generate a classifier with a single leaf node belonging to the present class, and to conclude that none of the features are necessary).

5.3 Classifier

Decision tree algorithms are known for their good performance when presented with limited separability, and as such they were a natural choice for this particular problem. They do not require normalization of inputs, which obviates an otherwise important piece of preprocessing, can handle both numeric and categorical inputs side-by-side, and are insensitive to missing data, a valuable asset when many biomarkers are only available for an complete majority of patients in the study. They also are typically very fast, performance-wise, in training and testing, which has obvious advantages when there exists a desire to test a variety of configurations.

Experiments focused on the J48/C4.5 model, developed in [54], and the random forests method developed by [55], both of which are well regarded and well studied in the literature, and both of which were available in the Weka software package.

Table 5.4: Classifier styles

Base classifier	Feature selection	Boosting	Resampling rates
J48	Simple Ranking	Adaboost	0,0.3,0.6,0.9
J48	PCA	None	0,0.3,0.6,0.9
J48	None	Adaboost	0,0.3,0.6,0.9
J48	SFFS, Nonzero	Adaboost	0,0.3,0.6,0.9
J48	SFFS, Most selected	Adaboost	0,0.3,0.6,0.9
Random forests	None	None	0,0.3,0.6,0.9
Random forests	SFFS, Nonzero	None	0,0.3,0.6,0.9
Random forests	SFFS, Most selected	None	0,0.3,0.6,0.9

Adaboost [56] was used in a number of the more successful trials.

5.3.1 Algorithm overview

5.3.1.1 J48

J48 is the name of an open source Java implementation of Quinlan's [54] C4.5 algorithm, which is in turn based on his earlier ID3 algorithm. Both are based upon the entropy/information gain inherent in each feature. Pseudocode for the J48 algorithm is shown in algorithm 4. *Base cases* referenced in step 1 refer to situations where either a) all samples belong to a single class, b) none of the features provide any information gain, or c) a previously unseen class is encountered, in which case a node is created higher up the tree.

In this case, information gain refers to the entropy of a feature. Equation 5.4 gives the entropy for a feature f , where n represents the set of all values the feature can take. A feature with more diversity in its representation within the training set will have higher entropy, and therefore its selection will yield greater information gain. Thus the training procedure involves recursively splitting on features presumed to be the most discriminative based upon the variation within them.

The tree is also pruned, though this is not shown in the pseudocode. Based on the configuration used in this work, a feature is pruned if the learner's confidence in

the node is less than 25%. This step is designed to reduce overfitting.

$$E(f) = \sum_{\forall n} f_n(-\log(f_n)) \quad (5.4)$$

Algorithm 4 J48 pseudocode

```

NODES  $\leftarrow$   $\emptyset$ 
Check for base cases
for  $a \in$  ATTRIBUTES do
    Find normalized information gain from splitting on  $a$ 
end for
NODE  $\leftarrow$  NEW(NODE, BEST( $a$ ))
Recurse on sublists from splitting on NODE; add new NODES as children of
NODE
  
```

5.3.1.2 Random forests

Random forests are an ensemble classifier method developed by Breiman [55]. In the course of learning, a large number of decision trees are assembled, which then *vote*, such that the mode of all outputs is selected as the final choice. Pseudocode is given in algorithm 5.

Algorithm 5 Random forests pseudocode

```

M  $\leftarrow$  COUNT(features)
N  $\leftarrow$  COUNT(samples)
for all TREES do
    Take subset of samples of size  $n$ 
    Take subset of features of size  $m$ 
    Grow tree fully, with no pruning
end for
  
```

5.3.1.3 Adaboost

Adaboost is built on the premise that an ensemble of weak classifiers, trained successively, such that the training of later classifiers is biased to emphasize those examples that were the most difficult to train earlier in the process. In general, so long as the weak classifiers perform better than random, adaboost reliably improves the overall classification rate.

5.3.2 Methods

Random forests and J48 trees were assembled in the 8 configurations shown in table 5.4. For an initial run, only the 4 configurations without SFFS were used, to get a baseline of globbers and resampling rates that were effective in classification, before attempting to improve results with feature selection. Each of the 4 was resampled at the rates described in the “resampling rates” column, for a total of 16 classifiers. The J48+PCA combination was actually implemented using a technique called rotation forests, described by [57], in which random subsets of the features are subjected to PCA, the result of which is used for training. In 2 of the configurations, J48 was subjected to Adaboost [56], which considerably improved its performance.

For each of the glob sizes described in table 5.2, each classifier in table 5.4 was run on all 6 output features. At 15 glob size configurations, that yields 15×16 classifiers per output feature, or 240 classifiers per feature. This large number was evaluated as described in section 5.3.3, after which the top performing resampling rates and glob sizes for each output variable were reassembled into the feature selection procedure described in section 5.2.1.

A subsequent regression was performed, using the 4 classifier types from table 5.4, restricted to subsets of the feature set as described in table 5.3. The classifiers labeled *nonzero* were evaluated using all features from the corresponding column in table 5.3 whose value was nonzero; those labeled *most selected* were evaluated using all features that were selected at least half as often as the most frequently selected

Table 5.5: Grid describing class performance.

	--	-	0	+	++
--					
-					
0					
+					
++					

feature.

5.3.3 Experimental Methods

A leave-one-out cross-validation scheme was constructed so that each classifier was trained on a dataset composed of samples from $(n - 1) = 31$ patients, while the test set was composed of samples from the remaining patient. This maximizes our use of our limited test data, while simultaneously ensuring that no overfitting takes place related to training on the test set. A set of box and whisker plots (figs. 5.6a - 5.6f) was composed for each output feature to demonstrate variability in accuracy over the whole set of patients.

Table 5.5 shows a blank confusion matrix with the 5 classes used in the output variables at its margins. Columns represent *detected class*, rows represent *true class*. The diagonal indicates correct in-class classification; a uniform distribution among all cells would indicate a classifier with completely random output. To sift through the large number of classifiers created by the search described in section 5.3.2, a scalar metric for ranking the classifiers was needed. AUC, or area under the curve, which describes the area under the ROC (relative operating characteristic) curve is a popular metric in machine learning because it provides a straightforward, meaningful scalar value by which a classifier can be assessed.

Unfortunately, it's not clearly applicable to multi-class classification. Weka does provide an AUC metric for each class in its output, generated by combining all "other"

classes into the “false” category. Landgrebe [58], Everson [59], and Fieldsend [60] all published varying attempts at extending the AUC to the multiclass case; however, because in this case misclassification costs were purely a function of distance from target class, and the dataset was heavily unbalanced, a simple metric produced by summing the in-class accuracy was used for ranking classifiers. If we call this metric ρ , we save for further consideration all classifiers meeting the criteria described in equation 5.6, where σ_f is the standard deviation computed over all classifier accuracies for a given feature.

$$\rho = \frac{\sum_{\forall x} M_{x,x}}{\sum_{\forall x,y} M_{x,y}} \quad (5.5)$$

$$\rho_{max} - \rho < \sigma_f \quad (5.6)$$

One simple, and useful, interpretation of the metric given in equation 5.5 is that any score > 1 indicates a classifier that performs better than the *zero-R*, or maximum likelihood classifier (that is, classifying all samples to be members of the most common class). Table 5.9 makes it clear that while the results are far from perfect, the best classifier in each case beats a maximum likelihood classifier.

5.3.3.1 Power analysis

Power analysis is a statistical technique that illustrates the sufficiency in a data set for detecting a given change. It serves a purpose complementary to the *p-value*: whereas the p-value provides an indication of whether an observed result was a result of random chance, the statistical power describes the odds of observing a result, should one actually exist. Computation of power requires a p-value, a sample size (n), and an effect size (d). In this case, we set $n = 31$, corresponding to the number of patients in the training set, given our leave-one-out training model. p-value was set

Table 5.6: Statistical power measures per output feature

(a) Plaque thickness change			(b) Lumen CSA change		
Label	d	Power	Label	d	Power
--	0.4095	0.3546	--	0.3901	0.3271
-	0.1259	0.07767	-	0.1224	0.07612
+	0.119	0.07468	+	0.1264	0.07788
++	0.4118	0.3579	++	0.3489	0.272

(c) Fibrous % change			(d) Fibrous-fatty % change		
Label	d	Power	Label	d	Power
--	0.3431	0.2646	--	0.3707	0.3006
-	0.1194	0.07487	-	0.1338	0.08129
+	0.119	0.0747	+	0.1318	0.08037
++	0.3463	0.2687	++	0.3409	0.2618

(e) Necrotic core % change			(f) Dense calcium % change		
Label	d	Power	Label	d	Power
--	0.3419	0.2631	--	0.4207	0.3708
-	0.1261	0.07777	-	0.1307	0.07986
+	0.1264	0.0779	+	0.1376	0.08313
++	0.3479	0.2707	++	0.3903	0.3273

to 0.05, which is a standard cut-off for significance. d was computed using Cohen's formula (5.7), modified to take into account the method by which we binned our output values (5.8), (see section 5.1.4).

This method of computing effect size is justified given that, in this study, the effect size measures changes within a single distribution. The differences in means amongst the classes were all at least 1 standard deviation, which in turn meant that d was always greater than 1, if Cohen's method was used unmodified, and statistical power in turn was always estimated to be unity (or perfect). To give the estimate meaning, we instead measured the distance of the mean of a given class from the mean of the 0 class (signifying no change, whose mean was always definitively 0), as a fraction of the total distribution width.

$$d = \frac{\bar{x}_1 - \bar{x}_2}{s} \quad (5.7)$$

$$d = \frac{\bar{x}_{class}}{x_{max} - x_{min}} \quad (5.8)$$

Statistical power is shown for each of the output variables in tables 5.6a - 5.6f. Because of the small sample size, power is very low for the + and - categories; unfortunately, due to the concentration of samples around 0, there are too few -- and ++ in the data set in general for the classifier to be effectively trained on them in spite of the considerably higher statistical power for detecting them.

5.3.4 Results

Detailed results from each of a) 5-way output, no SFFS, b) 3-way output, no SFFS, and c) 5-way output, using SFFS to inform the subset of features trained on are delineated in sections 5.3.4.1, 5.3.4.2, and 5.3.4.3 respectively. The feature selection had no discernible impact upon the performance of the classifiers. The obvious conclusion would be that while the feature selection successfully identified the

most suitable features in the data set, the decision tree learners used for classification already rather robustly accomplish the same with intrinsic aspects of their algorithms.

On the other hand, table 5.7 demonstrates that accuracy is, on average, improved by restricting the number of classes to 3, and simply merging the tails into the + and – classes, on the grounds that there is insufficient training data to meaningfully identify all five classes.

Table 5.7: Average in-class accuracy.

	ptc	lcsa	f	ff	nc	dc
5way, no fs	0.280	0.320	0.299	0.395	0.421	0.400
3way, no fs	0.445	0.432	0.483	0.529	0.602	0.599
5way, SFFS	0.263	0.278	0.309	0.384	0.389	0.410

ptc = plaque thickness change lcsa = lumen cross sectional area change
f = fibrous plaque % change ff = fibrous-fatty plaque % change
nc = necrotic core plaque % change dc = dense calcium plaque % change

5.3.4.1 5 class, no feature selection

Table 5.9 shows the rankings on each feature after the first regression, prior to SFFS feature selection, was performed. The confusion matrices corresponding to the *best* from each output variable are then shown in tables 5.8a - 5.8f.

A summary of the overall results is also depicted in figure 5.5. The blue bars depict the in-class-accuracy for the best performing classifier configuration for each output feature; the red bars depict additionally the off-by-one percentage, where off-by-one means values misclassified into an adjacent class – so – misclassified as 0, or vice versa. The top of the red bar for each column reflects the sum of the two, depicting the combined in-class and off-by-one accuracy. This combined value is, in general, quite close to one, indicating that most misclassification errors were off by only one class.

Figure 5.5: Overview of classifier performance, scored by feature. Blue is in-class accuracy; red is off-by-one accuracy.




Table 5.8: Confusion matrices for 5-way output

(a) Plaque thickness change (72_3 / 0.0 / random forests)

	--	-	0	+	++
--	0	2	8	2	0
-	0	21	109	5	0
0	0	29	335	20	1
+	0	8	57	21	5
++	0	1	5	0	1

(b) Lumen CSA change (18_3 / 0.9 / adaboosted J48 trees with feature selection)

	--	-	0	+	++
--	7	2	7	0	0
-	10	150	217	22	0
0	6	355	1575	164	7
+	4	31	168	8	0
++	0	4	2	1	0

(c) Fibrous % change (3_1 / 0.9 / adaboosted J48 trees with feature selection)

	--	-	0	+	++
--	0	50	2	3	0
-	4	2210	1858	353	1
0	4	3792	33081	5043	9
+	0	430	2442	705	0
++	1	2	10	5	0

(d) Fibrous-fatty % change (36_3 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	3	1	0	1	0
-	1	87	84	14	0
0	2	131	867	86	1
+	1	18	29	6	0
++	0	0	1	0	0

(e) Necrotic core % change (3_1 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	10	35	0	0	0
-	22	1706	276	51	0
0	1	1836	41080	1579	3
+	0	284	2540	427	1
++	0	15	108	31	0

(f) Dense calcium % change (18_3 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	1	2	0	0	0
-	2	49	10	9	0
0	4	99	2149	158	35
+	0	21	169	18	1
++	0	1	11	1	0

Figure 5.6 illustrates how classification accuracy varied over the data set, as a function of which patient was the hold-out for the testing set. This illustrates the considerable variability from patient to patient – this can also be seen in the results of the feature selection, where many of the patients could be trained successfully with just a small minority of the available features.

Table 5.9: Best performers, 5 class output.

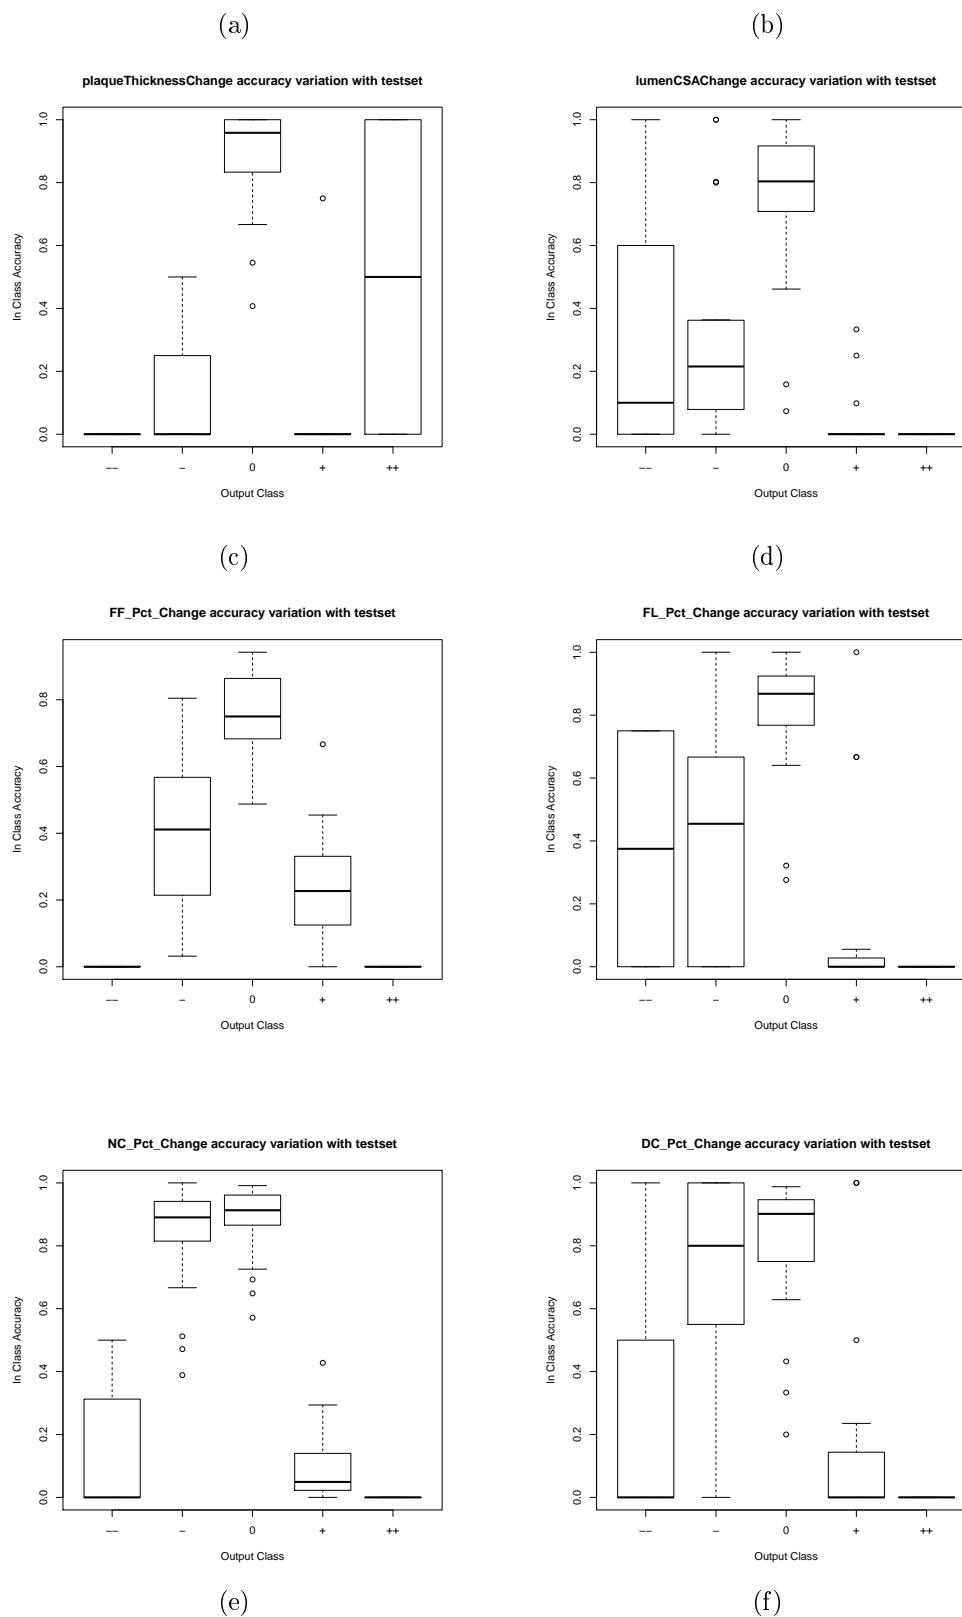
Output Variable	Classifier	Score	
pt change	72 3 randf 0.0	1.399	
	lcsa change	18 3 abwfs 0.9	1.599
lcsa change	9 5 ab 0.3	1.549	
	9 5 abwfs 0.0	1.528	
	9 5 abwfs 0.3	1.503	
	9 1 abwfs 0.3	1.497	
	9 5 ab 0.0	1.496	
	3 5 ab 0.6	1.479	
	9 1 ab 0.9	1.456	
	ff change	3 1 randf 0.9	1.495
		3 1 abwfs 0.9	1.485
		9 1 abwfs 0.9	1.472
		18 3 abwfs 0.9	1.465
		3 3 ab 0.9	1.456
		3 3 abwfs 0.9	1.452
		3 1 ab 0.9	1.45
		3 1 rotf 0.9	1.449
9 5 abwfs 0.9		1.435	
3 3 randf 0.9		1.434	
9 3 ab 0.6		1.429	
3 3 abwfs 0.6		1.427	
9 5 abwfs 0.3		1.422	
9 1 randf 0.9		1.417	
9 3 abwfs 0.6		1.413	
9 1 abwfs 0.3	1.411		
ff change	36 3 abwfs 0.6	1.395	
	9 5 ab 0.9	1.39	
	3 1 abwfs 0.6	1.39	
	3 5 abwfs 0.9	1.388	
	18 5 abwfs 0.3	1.388	
	9 3 randf 0.9	1.383	
	9 1 ab 0.6	1.381	
	3 1 ab 0.6	1.38	
	3 5 abwfs 0.3	1.375	
	fl change	36 3 ab 0.9	1.976
		36 3 abwfs 0.9	1.897
		36 3 ab 0.6	1.838
	nc change	3 1 ab 0.9	2.107
		3 1 abwfs 0.9	1.984
		36 3 abwfs 0.9	1.919
dc change	18 3 ab 0.9	1.998	
	36 3 ab 0.9	1.972	
	36 3 abwfs 0.9	1.952	

Score corresponds to ρ defined in eq. 5.5.

5.3.4.2 3 class, no feature selection

Table 5.10 shows the rankings on each output feature after a regression in which the distribution was cut into only 3 output classes. This is shown for comparison, to illustrate the effect of eliminating the ++ and -- classes, which in general had too little training data to be effectively represented. Confusion matrices follow in tables

Figure 5.6: Inter-patient variability, box and whisker plots.



5.11a-5.11f.

5.3.4.3 5 class, with SFFS

Table 5.12 shows the results of re-running the regression from section 5.3.4.1, restricting features to those selected in the SFFS procedure described in section 5.2. It's worth noting that the feature selection had little impact on the overall accuracy, and that the classifiers most often chosen were versions that retained all features. Confusion matrices are depicted in tables 5.13a-5.13f.

Table 5.10: Best performers, 3 class output

Output Variable	Classifier	Score
pt change	72 5 ab 0.6	1.334
	72 5 rotf 0.9	1.326
	72 3 rotf 0	1.313
	72 3 randf 0.9	1.286
	72 3 abwfs 0.9	1.285
	72 1 abwfs 0.9	1.283
	72 3 randf 0.6	1.265
lcsa change	9 1 abwfs 0.9	1.297
	36 5 randf 0.9	1.292
	36 3 abwfs 0.9	1.284
	9 5 ab 0.9	1.266
	9 1 randf 0.9	1.256
	3 1 rotf 0.9	1.255
	36 5 ab 0.9	1.248
ff change	3 3 abwfs 0.9	1.449
	9 5 abwfs 0.9	1.438
	3 1 ab 0.9	1.432
	3 1 randf 0.9	1.423
	3 1 abwfs 0.9	1.421
	9 1 randf 0.9	1.391
	18 1 randf 0.6	1.387
fl change	3 3 ab 0.9	1.586
	3 3 abwfs 0.9	1.509
	9 5 randf 0.9	1.490
nc change	3 1 ab 0.9	1.807
	9 1 abwfs 0.9	1.805
	9 3 abwfs 0.9	1.800
	3 1 rotf 0.9	1.798
	3 1 abwfs 0.9	1.797
	3 3 abwfs 0.9	1.791
	3 3 ab 0.9	1.775
dc change	3 1 abwfs 0.9	1.796
	18 1 ab 0.9	1.784
	3 1 ab 0.9	1.759
	3 1 rotf 0.9	1.705
	3 1 randf 0.9	1.690
	3 1 ab 0.6	1.685
	18 3 rotf 0.6	1.670

Table 5.11: Confusion matrices for 3-way output

(a) Plaque thickness change (72_5 / 0.6 / adaboosted J48 trees)				(b) Lumen CSA change (9_1 / 0.9 / adaboosted J48 trees with feature selection)			
	-	0	+		-	0	+
-	28	49	4	-	1009	1252	86
0	58	141	31	0	1892	10322	665
+	13	22	21	+	229	1018	88

(c) Fibrous % change (3_3 / 0.9 / adaboosted J48 trees with feature selection)				(d) Fibrous-fatty % change (3_3 / 0.9 / adaboosted J48 trees)			
	-	0	+		-	0	+
-	847	781	105	-	605	394	29
0	1252	11443	1061	0	779	14293	295
+	190	951	168	+	66	310	27

(e) Necrotic core % change (3_1 / 0.9 / adaboosted J48 trees)				(f) Dense calcium % change (3_1 / 0.9 / adaboosted J48 trees with feature selection)			
	-	0	+		-	0	+
-	1648	411	41	-	732	179	17
0	1573	41732	1194	0	859	45651	640
+	264	2856	286	+	132	1720	75

Table 5.12: Best performers, 5 class output, with SFFS

Output Variable	Classifier	Score	
pt change	ab NonZero 0.6 36 5	1.315	
	randf FiftyPercentile 0.9 18 3	1.207	
	ab NonZero 0.9 72 3	1.201	
	randf FiftyPercentile 0.9 18 5	1.192	
	randf FiftyPercentile 0.9 72 5	1.191	
	ab FiftyPercentile 0.9 18 3	1.182	
	randf FiftyPercentile 0.6 18 5	1.179	
	ab FiftyPercentile 0.6 18 3	1.166	
	ab FiftyPercentile 0.9 18 5	1.163	
	randf NonZero 0.9 18 5	1.161	
	ab NonZero 0.6 18 3	1.153	
	ab NonZero 0.9 36 3	1.144	
	lcsa change	ab All 0.9 36 3	1.391
		ab All 0 18 3	1.362
ab NonZero 0.6 36 3		1.335	
ab All 0.6 36 3		1.327	
ab NonZero 0.9 36 3		1.309	
ab All 0.6 18 3		1.308	
ab All 0 36 3		1.303	
ff change		randf All 0.9 9 3	1.547
	randf All 0.9 3 1	1.529	
	ab All 0.9 3 1	1.528	
	randf NonZero 0.9 9 3	1.513	
	ab NonZero 0.9 9 3	1.509	
	ab All 0.9 9 3	1.508	
	ab All 0.9 9 1	1.507	
	ab All 0.9 3 3	1.500	
	ff change	ab FiftyPercentile 0.9 36 1	1.919
		nc change	ab NonZero 0.9 3 5
ab FiftyPercentile 0.9 3 1	1.969		
ab All 0.9 3 1	1.930		
dc change	ab FiftyPercentile 0.9 18 3	2.050	
	ab FiftyPercentile 0.9 18 1	1.949	
	ab All 0.9 3 1	1.905	

Table 5.13: Confusion matrices for 5-way output, with SFFS

(a) Plaque thickness change (All features / 36_5 / 0.6 / adaboosted J48 trees)

	--	-	0	+	++
--	1	1	4	0	0
-	2	40	168	13	0
0	1	91	704	85	2
+	0	5	101	22	0
++	0	0	2	1	0

(b) Lumen CSA change (All features / 36_3 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	8	5	8	4	1
-	37	93	198	53	0
0	11	234	977	206	3
+	8	36	169	40	2
++	4	1	0	1	0

(c) Fibrous % change (All features / 9_3 / 0.9 / random forests)

	--	-	0	+	++
--	4	22	4	1	0
-	11	564	495	81	0
0	3	945	5349	515	0
+	1	111	511	104	0
++	0	3	15	2	0

(d) Fibrous-fatty % change (Top 50% of features / 36_1 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	16	21	0	4	0
-	36	615	255	157	6
0	17	571	3755	465	14
+	1	77	204	60	1
++	0	4	8	2	0

(e) Necrotic core % change (Non-zero subset of features / 3_5 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	3	15	0	0	0
-	11	506	136	35	1
0	6	880	12239	797	8
+	0	148	865	245	1
++	0	4	42	12	0

(f) Dense calcium % change (Top 50% of features / 18_3 / 0.9 / adaboosted J48 trees)

	--	-	0	+	++
--	2	3	0	0	0
-	10	73	24	16	1
0	3	214	3077	436	9
+	2	65	258	102	1
++	0	1	13	6	0

CHAPTER 6 DISCUSSION

6.1 Interpretation of results

A sophisticated workflow for the processing of jointly acquired IVUS and x-ray angiography has been presented here. The bifurcation modeling does not permit analysis of disease in the branches themselves, but efficiently solves 2 other problems: a) analysis of the impact of bifurcations up on WSS, and modeling WSS in heavily branched vessels effectively, and b) creating landmarks suitable for robust and consistent registration even in cases where the borders in the main vessel have enough variability to make shape-based registration techniques unreliable.

The registration consistently maps 2 IVUS pullbacks of a common vessel segment into a common space, so that the data can be analyzed confidently for changes in the physiology. Visual inspection of pairs of pullbacks confirms that, when registered, the pairs of vessel segments match, and the quantitative validation performed in section 3.1.3 shows a statistically significant match.

The database system developed for management of patient data vastly simplified managing a large pool of heterogeneous data. It also enabled increased collaboration among researchers: a data set loaded into the system in Iowa was then available for manual tracing by a cardiology resident in Chicago instantly. The database interfaced seamlessly with a web interface that permitted managing the tracing queues for the various residents working on the manual tracing, to ensure that work was evenly divided, and that different individuals were working on different data sets. It also permitted quick and seamless inventory management: the database could be queried to see which patients were incompletely processed, or missing pieces of data, without requiring us to manually trawl the data store each time.

Feature selection was ineffective at improving classifier results. It does, nonetheless, nicely complement the inter-feature correlations shown in figure 5.4. Average

in-class accuracy in table 5.7 appears to be better for the 3-class variant than for the 5-class, however, the generally poor performance of the -- and ++ classes overall drags down the accuracy figures for the 5-class cases more than it would otherwise.

All outputs had a significant bias just below zero. In the case of the plaque thickness change especially, this can be explained fairly easily with the statin drugs the majority of the patients were on, although it would not explain why all features seemed to have this bias. Nonetheless, the classifiers all proved to be more adept at predicting decreases than at predicting increases. It may be that some careful tuning of the classifier parameters might alleviate this, or it may be that the decreases more often displayed discernible patterns that could be detected.

Input resampling with bias set to 0.9 seemed to make a clear difference in classifier performance. The vast majority of all *winning* classifiers in the study had a 0.9 resampling bias. In fact, the only exceptions were the 3 winning plaque thickness change classifiers: one had a 0 bias, and the other two had 0.6 bias. This variation may mean that resampling had a greater effect for the features other than plaque thickness change.

Adaboosted J48 trees a majority of the best performing classifiers, consistently outperforming random forests – though random forests did appear in a few of the winning classifiers as well – namely plaque thickness change and fibrous plaque % change in the 5 class run.

Plaque thickness change performed the best with larger glob sizes, which would suggest it was the noisiest of the parameters – it consistently did best with *ag* set to 72, which meant that it was best predicted at a per-slice level, and it tended to have either 3 or 5 for *sg*. The other features ran the gamut of glob sizes, which, especially given the stochastic nature of the classifiers involved, would suggest that the other features were relatively insensitive to glob size – although none, other than plaque thickness change ever selected one with *ag* set to 72.

6.2 Sources of uncertainty


This study contains many sources of uncertainty. Successful reduction of them would almost certainly improve classifier performance, by ensuring better alignment of the data. Better quality control on the ECG signals would have eliminated the worst gating issues in the IVUS, and made it easier to synchronize the angiograms with the IVUS temporally. Severe cardiac arrhythmias could still have prevented perfectly smooth images, but smaller defects like these were generally well compensated for in the registration phase. We also would have had a larger pool of data to draw from. 32 patients is far smaller data set than we would have liked; even the full 89 that were enrolled in the study would have been pushing the bounds of clinical significance; merely not having to exclude patients for bad gating would have brought the number to 42, with which we could have generalized more effectively, and our observations would have had more statistical power.

Because the DICOM and the VH data were collected from 2 different workstations in the cathlab for the generation of Volcano equipment used in this study, there were occasionally synchronization issues between the two sides. When this happened, it was sometimes possible for the DICOM frame (which we segmented) to vary slightly from the corresponding VH frame (upon which the VH analysis was performed). This creates unknown uncertainty with regard to the accuracy of the VH analysis; figure 6.1 provides an illustration.

Many other patients were excluded because we never received the complete data set. On the clinical side, data was lost or destroyed, and frequently we received mislabeled data. With persistent followup, we were able to address many of the issues with mislabeled data; however it was a recurring problem, and we were never able to close all the gaps in the raw data provided. Thus for a combination of gating and organizational issues, we only included 32 out of 89 patients.

Roughly half of the angiograms presented a significant challenge in the segmenta-

Figure 6.1: Slight border mismatch in the VH application due to lack of synchronization.



tion, because with the equipment the clinicians used initially, the catheter was insufficiently radio-opaque to be visible in the presence of contrast dye. The workaround method involved estimating the position of the catheter from another frame in the image at the same phase, where the catheter was visible instead of the vessel. This improves the quality of what is essentially a blind guess, but is far inferior to the segmentation that can be performed when both are clearly visible.

This was alleviated with the introduction of a vastly superior catheter, termed the *Miracle Wire*, which shows up in the x-ray even when the vessel is fully inundated with contrast dye. In the case of these subsequent angiograms, segmentation is trivial, and highly accurate. The automated segmentation flawlessly finds the vessel and catheter borders every time, and the resultant 3-D reconstructions are much more reliable.

The projections were frequently taken with an eye towards only clinical relevance. This is fine for diagnostic purposes, but if the vessel is severely foreshortened, the accuracy of the 3-D reconstruction, and our ability to estimate the absolute orientation, are severely limited. This hampers our ability to get accurate WSS estimates, and limits our ability to accurately register the two pullbacks for analysis.

Other angiograms were inadequate because the protocol was neglected. In several cases, only one projection was taken with a catheter at all. This significantly impacts the reliability of the 3-D reconstruction, and adds considerable uncertainty to any indices that depend upon the 3-D geometry.

6.3 Future lines of inquiry

6.3.1 Acquisition protocol

Studies have been done as to the best angiographic projections for branch visibility. If we were to do a subsequent study, attention to a small handful of details would dramatically improve the quality of the data, and, presumably, classifier accuracy:

1. Miracle-wire used for all angiograms, to ensure simultaneous visibility of catheter

and lumen.

2. Angio projections chosen carefully to limit foreshortening, and ensure clear branch visibility.
3. Two orthogonal projections taken with catheter and transducer in full view prior to commencing automated pullback for each angiogram (some straightening and tensioning is usually required to eliminate slack in the wire before the pullback begins).
4. Careful examination of ECG signal prior to beginning procedure, to ensure gating can be done consistently and reliably.

Patient data is difficult and expensive to acquire, thus it is always good if even marginal data can be used to increase the statistical significance of a study. But the 4 above quality control issues are significant in terms of guaranteeing the utility of data in a study with many steps and many sources of uncertainty.

6.3.2 3-D reconstruction

One ongoing ambiguity in the 3-D reconstruction that was left open-ended in the work of [39] concerns the absolute orientation of the vessel: the algorithm, as implemented, is constrained in its ability to handle *wrapping* of the angle values. This is currently fudged in the software with a button on the interface alternate the representation between $[-180,180]$ and $[0,360]$. Often, the software will come to an estimate that is only a degree or two off after the inversion button is pressed; in these cases, no wrapping has occurred and the orientation is the same in either representation. When the estimates differ significantly, however, the current approach is to generate a VRML scene representation of each version, and then view both juxtaposed against the angiograms.

The trajectory of the catheter is then examined in both, in comparison to where

it appears on the angiographic projections. Bifurcations, if marked in the IVUS, are also represented in the VRML scene as an additional set of fiducial points. The human user then must choose which appears to be a better representation of the true orientation of the vessel. If there is ambiguity, the RMS error as computed by the reconstruction software can also be used as a hint, but as it is prone to being skewed by a small number of misaligned frames, it cannot be trusted completely.

Horn[61] developed an alternate, quaternion-based approach to the geometric analysis used in the reconstruction which avoids this ambiguity. It may be worthwhile to investigate reimplementing portions of the reconstruction software to use this approach, as mathematical elimination of the ambiguity would at least introduce consistency into the process, and possibly also eliminate a certain amount of human error. Obviously, both of these approaches fall short if the catheter is invisible in the angiograms, as mentioned in section 6.2, thus the importance of *Miracle Wire* or something similar.

The work presented in section 3.2 also offers some potential for reducing uncertainties in the 3-D reconstruction. Our results in section 3.2.4 show that our *detwisted* versions are comparable in accuracy to the originals, they don't offer a compelling improvement. If the angiograms consistently featured a specific, identifiable subset of branches, which could be matched to specific IVUS branches based on clinical knowledge, this method may begin to emerge as useful, especially if the bifurcations could be segmented in an automated fashion, which would allow them to be matched to an atlas heuristically.

6.3.3 TCFA analysis

The approaches to TCFA labeling discussed in section 4.2.1.4 offer an intriguing tool set for identifying at-risk lesions in CAD patients; however, they are limited in the sense that they require the formulation of increasingly complex and difficult to implement rules and offer no guarantee of improvement in classification performance.

Assuming a human observer can identify TCFA lesions, at least based upon the VH data (it's generally considered impossible based upon the B-mode alone), this is an ideal application for machine learning/object recognition approaches.

Rather than prepare a training set in advance, which is potentially labor intensive, it makes more sense to use the automated, rule-based classification as a starting point. A human arbiter would be presented with a succession of randomly chosen VH overlay images, chosen 50/50 between slices for which the rule-based TCFA labeling was each of *true* and *false*. The user would choose between *yes*, *no*, and *unsure*, where the third category excludes the image in question from the training set. When a stopping criterion had been reached (either sufficient images selected, or if the training is performed on-line, the training error from newly inserted images falls below a predetermined threshold), the classifier would be regarded as trained.

A classifier would presumably operate on a log-polar version of the VH data in order to guarantee rotation invariance, and then employ an approach based upon small local features to identify specific plaque textures common to images labeled as TCFA. The best known approach of this form is SIFT (Scale Invariant Feature Transform) [62]; other promising approaches include Hotta's work with *visual words* based on Gabor filters [63] and Zhang's hybrid SVM/KNN approach [64].

CHAPTER 7 CONCLUSION

Out of 89 patients from which baseline and followup data were collected, 32 were analyzed to completion. These were used to train and test a series of classifiers, in an effort to characterize the development and progression of atherosclerotic coronary artery disease on a one-year timescale. Initially, we had hoped to be able to reliably identify 5 output classes (− − / − /0/ + / + +). However, due to the distribution of the data, and the small size of the data set, we found we achieved much better classification results when restricting the classes to 3, by merging − − with − and ++ with +.

Results are promising. They are considerably better than chance. There is reason for optimism that with more data, and better quality data, better, more reliable predictions would be a possibility.

Out of 5 aims laid out in section 1.3, the results can be summarized as follows:

1. Compensation for rotational artifacts of the catheter (section 3.2) was achieved in a way that was difficult to either evaluate conclusively or utilize with the data available in this study, given the dependence upon bifurcations visible in both the angiograms, and the IVUS. Nonetheless, the method developed does appear to have normalized the absolute orientation estimate (figures 3.12,3.13). It shows promise for future studies that take it into account at acquisition time.
2. The point-to-point registration method was a complete success. Figure 3.1b demonstrates that the method achieves correct results from the point of view of visual inspection; the results in table 3.1 and figures 3.2 and 3.3 show that the results are quantitatively convincing as well.
3. The data from all patients made available to this study was thoroughly examined; the 32 found to be suitable were fully analyzed and used for the training

and testing of the classifier. This aim can be considered to be a complete success.

4. Both through the action of the decision tree algorithms used for classification (section 5.3.1), and through the SFFS feature selection procedure used to refine the set of inputs to said classifier (section 5.2.1), a suitable subset of features was identified for the training and testing of a classification system. This aim can be considered a success.
5. The classifier was limited in its ability to detect large changes, due to the heavy bias towards zero change in the input data. The data also seemed to present a bias towards reduced disease, and the classification results, in particular, performed better in samples where the disease regressed. This may well owe to the fact that the patients enrolled in this study were on statins and other aggressive cholesterol reducing drugs, which if the results are to be believed, work rather well. This aim can also be considered to be a success.

A larger study, with more patients, and perhaps more stringent controls on the acquisition protocol should be planned, to solidify the statistical significance of the work in this study before any jumps to clinical applications are made. However, a considerable amount of work was done in this study to the end of advancing the state of the art of image analysis and model construction based upon IVUS and x-ray angiography, which should lower the bar for a subsequent study, and moreover open the door to the inclusion of OCT data as an additional modality, which has particular relevance to the inner elastic lamina and the disease state closest to the interior of the vessel.

Regardless of the limitations, we feel this project makes a significant contribution to the body of work in the field of cardiac image analysis. The identification and modeling of bifurcations, as well as the work on mesh merging and automation of

the WSS analysis is wholly novel and unlike other approaches in the literature. Virtually all vascular modeling software available presumes an image volume that can be segmented into a vessel tree with level set methods; this is completely impossible with IVUS, thus a method for joining disjoint structured grids in space had to be developed.

The landmark based registration method works reliably, and is immune to variations in the border produced either by inter-observer error in manual tracing, by variations in an automated segmentation, or by physiological changes between image acquisitions. It adapts robustly to rotational movement within the images, and resamples in a way that yields correct, reliable results.

Moreover, the collation of the data from these 32 patients into a unified system for the analysis of the multi-modality image sets, alongside demographic and biomarker data, was a considerable undertaking. Through the construction of the classification pipeline, we have a) built a nontrivial body of software that can be easily used for additional patient data in subsequent studies, and, more importantly, b) demonstrated the feasibility and utility of using decision-tree based machine learning techniques to predict physiological changes in the assessment of coronary artery disease.

REFERENCES

- [1] J. H. Reiber, C. J. Kooijman, C. J. Slager, J. J. Gerbrands, J. C. Schuurbiers, A. Den Boer, W. Wijns, P. W. Serruys, and P. G. Hugenholtz, "Coronary artery dimensions from cineangiograms methodology and validation of a computer-assisted analysis procedure," IEEE Trans Med Imaging, vol. 3, no. 3, pp. 131–141, 1984.
- [2] J. D. Klingensmith, P. Schoenhagen, A. Tajaddini, S. S. Halliburton, E. M. Tuzcu, S. E. Nissen, and D. G. Vince, "Automated three-dimensional assessment of coronary artery anatomy with intravascular ultrasound scanning," Am. Heart J., vol. 145, no. 5, pp. 795–805, May 2003.
- [3] M. Olszewski, "Assessment of coronary atherosclerosis using intravascular ultrasound and multidetector computed tomography," Ph.D. dissertation, University of Iowa, 2005.
- [4] S. Rathore, M. Terashima, H. Matsuo, Y. Kinoshita, M. Kimura, E. Tsuchikane, K. Nasu, M. Ehara, Y. Asakura, O. Katoh, and T. Suzuki, "Association of coronary plaque composition and arterial remodelling: A optical coherence tomography study," Atherosclerosis, Oct 2011.
- [5] A. Wahle, G. P. M. Prause, S. C. DeJong, and M. Sonka, "Geometrically correct 3-D reconstruction of intravascular ultrasound images by fusion with biplane angiography — methods and validation," IEEE Transactions on Image Processing, vol. 18, no. 8, pp. 686–699, August 1999.
- [6] A. Wahle, S. C. Mitchell, S. D. Ramaswamy, K. B. Chandran, and M. Sonka, "Virtual angioscopy in human coronary arteries with visualization of computational hemodynamics," in Medical Imaging 2001: Physiology and Function from Multidimensional Images, C. Chen and A. V. Clough, Eds., vol. 4321. Bellingham WA: SPIE Proceedings, 2001, pp. 32–43.
- [7] —, "Four-dimensional coronary morphology and computational hemodynamics," in Medical Imaging 2001: Image Processing, M. Sonka and K. M. Hanson, Eds., vol. 4322. Bellingham WA: SPIE Proceedings, 2001, pp. 743–754.
- [8] A. Wahle, J. J. Lopez, M. E. Olszewski, S. C. Vigmostad, K. B. Chandran, J. D. Rossen, and M. Sonka, "Plaque development, vessel curvature, and wall shear stress in coronary arteries assessed by X-ray angiography and intravascular ultrasound," Medical Image Analysis, vol. 10, no. 4, pp. 615–631, Aug. 2006.

- [9] A. Nair, J. D. Klingensmith, and D. G. Vince, "Real-time plaque characterization and visualization with spectral analysis of intravascular ultrasound data," Stud Health Technol Inform, vol. 113, pp. 300–320, 2005.
- [10] M. Sonka, C. R. McKay, J. J. Lopez, A. Wahle, K. B. Chandran, T. L. Burns, and Y. G. Lai, "Revision of NIH research grant 1R01HL63373-01: 3-D and 4-D coronary hemodynamics and local atherosclerosis," The University of Iowa, Jun. 1999.
- [11] M. Olszewski, "Integrated system for multi-modality image data fusion and quantitative analysis," Master's thesis, University of Iowa, 2002.
- [12] G. P. M. Prause, S. C. DeJong, C. R. McKay, and M. Sonka, "Semiautomated segmentation and 3-D reconstruction of coronary trees: biplane angiography and intravascular ultrasound data fusion," vol. 2709. SPIE Proceedings, 1996, pp. 82–92.
- [13] K. Li, X. Wu, D. Z. Chen, and M. Sonka, "Optimal surface segmentation in volumetric images—a graph theoretic approach," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 28, no. 1, pp. 119–134, January 2006.
- [14] M. E. Olszewski, A. Wahle, S. C. Vigmostad, and M. Sonka, "Multidimensional segmentation of coronary intravascular ultrasound images using knowledge-based methods," in Medical Imaging 2005: Image Processing, J. M. Fitzpatrick and J. M. Reinhardt, Eds., vol. 5747. Bellingham WA: SPIE Proceedings, 2005, pp. 496–504.
- [15] R. Medina, A. Wahle, M. E. Olszewski, and M. Sonka, "Volumetric quantification of coronary arteries reconstructed by fusion between intravascular ultrasound and biplane angiography," in Proc. 2002 IEEE International Symposium on Biomedical Imaging. Piscataway NJ: IEEE Press, 2002, pp. 891–894.
- [16] —, "Three methods for accurate quantification of plaque volume in coronary arteries," International Journal of Cardiovascular Imaging, vol. 19, no. 4, pp. 301–311, Aug. 2003.
- [17] —, "Curvature and torsion estimation for coronary artery motion analysis," in Medical Imaging 2004: Physiology, Function, and Structure from Medical Images, A. A. Amini and A. Manduca, Eds., vol. 5369. Bellingham WA: SPIE Proceedings, 2004, pp. 504–515.

- [18] R. W. Downe, A. Wahle, T. Kovárník, H. Skalická, J. J. Lopez, J. Horák, and M. Sonka, "Segmentation of intravascular ultrasound images using graph search and a novel cost function," in The 2nd International Workshop on Computer Vision for Intravascular and Intracardiac Imaging (MICCAI 2008), G. Únal, I. A. Kakadiaris, N. Navab, and M. Sonka, Eds. Istanbul: Sabanci University, 2008, pp. 71–79.
- [19] A. Wahle, J. J. Lopez, M. E. Olszewski, S. C. Vigmostad, K. C. Braddy, T. M. H. Brennan, S. W. Bokhari, J. G. Bennett, E. M. Holper, J. D. Rossen, K. B. Chandran, and M. Sonka, "Inverse relationship between local wall shear stress and plaque thickness in coronary arteries is retained by compensatory enlargement in early atherosclerosis," American College of Cardiology, 54th Scientific Sessions, Orlando FL, Journal of the ACC (Supplement), vol. 45, no. 3/A, p. 416A, Feb. 2005, abstract 1126-138.
- [20] —, "Relationship between plaque development and local hemodynamics in coronary arteries," in Medical Imaging 2005: Physiology, Function, and Structure from Medical Images, A. A. Amini and A. Manduca, Eds., vol. 5746. Bellingham WA: SPIE Proceedings, 2005, pp. 223–232.
- [21] A. Wahle, J. J. Lopez, M. E. Olszewski, S. C. Vigmostad, K. B. Chandran, J. D. Rossen, and M. Sonka, "Analysis of the interdependencies among plaque development, vessel curvature, and wall shear stress in coronary arteries," in Functional Imaging and Modeling of the Heart (FIMH '05), ser. Lecture Notes in Computer Science, A. F. Frangi, P. I. Radeva, A. Santos, and M. Hernandez, Eds., vol. 3504. Berlin/New York: Springer, 2005, pp. 12–22.
- [22] C. M. Gibson, L. Diaz, K. Kandarpa, F. M. Sacks, R. C. Pasternak, T. Sandor, C. L. Feldman, and P. H. Stone, "Relation of vessel wall shear stress to atherosclerosis progression in human coronary arteries," Arteriosclerosis and Thrombosis, vol. 13, no. 2, pp. 310–315, Feb. 1993.
- [23] J. V. Soulis, T. M. Farmakis, G. D. Giannoglou, and G. E. Louridas, "Wall shear stress in normal left coronary artery tree," J Biomech, vol. 39, pp. 742–749, 2006.
- [24] H. Samady, P. Eshtehardi, M. C. McDaniel, J. Suo, S. S. Dhawan, C. Maynard, L. H. Timmins, A. A. Quyyumi, and D. P. Giddens, "Coronary artery wall shear stress is associated with progression and transformation of atherosclerotic plaque and arterial remodeling in patients with coronary artery disease," Circulation, vol. 124, pp. 779–788, Aug 2011.

- [25] R. W. Downe, A. Wahle, J. Garvin, T. Kovárník, J. Horák, J. J. Lopez, and M. Sonka, "Identification and 3-D modeling of coronary branches in intravascular ultrasound," in The 3rd International Workshop on Computer Vision for Intravascular and Intracardiac Imaging (MICCAI 2011), G. Ünal, I. A. Kakadiaris, N. Navab, and M. Sonka, Eds. Istanbul: Sabanci University, 2011.
- [26] T. Kovarnik, G. S. Mintz, H. Skalicka, A. Kral, J. Horak, R. Skulec, J. Uhrova, P. Martasek, R. W. Downe, A. Wahle, M. Sonka, V. Mrazek, M. Aschermann, and A. Linhart, "Virtual histology evaluation of atherosclerosis regression during atorvastatin and ezetimibe administration: HEAVEN study," Circ. J., vol. 76, pp. 176–183, 2012.
- [27] C. Gatta, S. Balocco, F. Ciompi, R. Hemetsberger, O. Rodriguez Leor, and P. Radeva, "Real-time gating of IVUS sequences based on motion blur analysis: method and quantitative validation," Med Image Comput Comput Assist Interv, vol. 13, pp. 59–67, 2010.
- [28] C. Gatta, O. Pujol, O. Rodriguez Leor, J. Mauri Ferre, and P. Radeva, "Fast rigid registration of vascular structures in IVUS sequences," IEEE Trans Inf Technol Biomed, vol. 13, pp. 1006–1011, Nov 2009.
- [29] B. Christoudias, C. Georgescu and P. Meer, "Synergism in low-level vision," in 16th International Conference on Pattern Recognition, vol. 4, August 2002, pp. 150–155.
- [30] K. Hollig, Finite Element Methods with B-Splines. Society for Industrial and Applied Mathematics, 2003.
- [31] D. Comanicu and P. Meer, "Mean shift: A robust approach toward feature space analysis," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 24, no. 5, pp. 603–619, May 2002.
- [32] P. Meer and B. Georgescu, "Edge detection with embedded confidence," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 23, no. 12, pp. 1351–1365, December 2001.
- [33] "CGAL, Computational Geometry Algorithms Library," <http://www.cgal.org>.
- [34] W. Nef, Beiträge zur Theorie der Polyeder: mit Anwendungen in der Computergraphik, ser. Beiträge zur Mathematik, Informatik und Nachrichtentechnik. Lang, 1978. [Online]. Available: <http://books.google.com/books?id=3Lm0AAAAIAAJ>

- [35] L. Rineau and M. Yvinec, "A generic software design for delaunay refinement meshing," Computational Geometry Theory and Applications, vol. 38, pp. 100–110, 2007.
- [36] P. Hachenberger and L. Kettner, "Boolean operations on 3-D selective nef complexes: Optimized implementation and experiments," in ACM Symposium on Solid and Physical Modeling (SPM), Boston, MA, June 2005.
- [37] M. Audette, R. Brooks, R. Funnell, G. Strauss, and T. Arbel, "Piecewise affine initialized spline-based patient-specific registration of a high-resolution ear model for surgical guidance," in MICCAI Workshop on Image Guidance and Computer Assistance for Soft-Tissue Interventions, 01 2008.
- [38] B. W. Silverman, Density estimation: for statistics and data analysis, Chapman and Hall, Eds., London, 1986.
- [39] A. Wahle, G. P. M. Prause, C. von Birgelen, R. Erbel, and M. Sonka, "Automated calculation of the axial orientation of intravascular ultrasound images by fusion with biplane angiography," in Medical Imaging 1999: Image Processing, K. M. Hanson, Ed., vol. 3661. Bellingham WA: SPIE Proceedings, 1999, pp. 1094–1104.
- [40] S. A. Sheynin and A. V. Tuzikov, "Formulae for polytope volume and surface moments," in Proc. International Conference on Image Processing (ICIP 2001), vol. 3. Piscataway NJ: IEEE Press, 2001, pp. 720–723.
- [41] T. Thim, M. K. Hagensen, D. Wallace-Bradley, J. F. Granada, G. L. Kaluza, L. Drouet, W. P. Paaske, H. E. B?tker, and E. Falk, "Unreliable assessment of necrotic core by virtual histology intravascular ultrasound in porcine coronary artery disease," Circ Cardiovasc Imaging, vol. 3, pp. 384–391, Jul 2010.
- [42] G. W. Stone and G. S. Mintz, " ," Circ Cardiovasc Imaging, vol. 3, p. e4; author reply e5, Sep 2010.
- [43] L. Antiga, M. Piccinelli, L. Botti, B. Ene-Iordache, A. Remuzzi, and D. A. Steinman, "An image-based modeling framework for patient-specific computational hemodynamics," Med Biol Eng Comput, vol. 46, no. 11, pp. 1097–1112, Nov 2008.
- [44] P. R. Moreno, "The high-risk thin-cap fibroatheroma: a new kid on the block," Circ Cardiovasc Interv, vol. 2, pp. 500–502, Dec 2009.

- [45] P. R. Moreno, V. H. Bernardi, J. Lopez-Cuellar, A. M. Murcia, I. F. Palacios, H. K. Gold, R. Mehran, S. K. Sharma, Y. Nemerson, V. Fuster, and J. T. Fallon, "Macrophages, smooth muscle cells, and tissue factor in unstable angina. Implications for cell-mediated thrombogenicity in acute coronary syndromes," Circulation, vol. 94, pp. 3090–3097, Dec 1996.
- [46] F. D. Kolodgie, A. P. Burke, A. Farb, H. K. Gold, J. Yuan, J. Narula, A. V. Finn, and R. Virmani, "The thin-cap fibroatheroma: a type of vulnerable plaque: the major precursor lesion to acute coronary syndromes," Curr. Opin. Cardiol., vol. 16, pp. 285–292, Sep 2001.
- [47] S. P. Marso, A. D. Frutkin, S. K. Mehta, J. A. House, J. R. McCrary, V. Klauss, A. Lerman, M. B. Leon, A. Nair, P. Margolis, R. Erbel, K. Nasu, F. Schiele, and J. Margolis, "Intravascular ultrasound measures of coronary atherosclerosis are associated with the Framingham risk score: an analysis from a global IVUS registry," EuroIntervention, vol. 5, pp. 212–218, Jun 2009.
- [48] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: An update," 2009.
- [49] S. Daskalaki, I. Kopanas, and N. Avouris, "Evaluation of classifiers for an uneven class distribution problem," Applied Artificial Intelligence, vol. 20, pp. 1–37, 2006.
- [50] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: Synthetic minority over-sampling technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357, 2002.
- [51] P. Pudil, F. Ferri, J. Novovicova, and J. Kittler, "Floating search methods for feature selection with nonmonotonic criterion functions," in Pattern Recognition, 1994. Vol. 2 - Conference B: Computer Vision Image Processing., Proceedings of the 12th IAPR International. Conference on, vol. 2, oct 1994, pp. 279 –283 vol.2.
- [52] L. C. Molina, L. Belanche, and A. Nebot, "Feature selection algorithms: A survey and experimental evaluation," in Proceedings of the 2002 IEEE International Conference on Data Mining, ser. ICDM '02. Washington, DC, USA: IEEE Computer Society, 2002, pp. 306–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=844380.844722>
- [53] M. Gütlein, E. Frank, M. Hall, and A. Karwath, "Large-scale attribute selection using wrappers," 2006.
- [54] R. Quinlan, C4.5: Programs for Machine Learning. San Mateo, CA: Morgan Kaufmann Publishers, 1993.

- [55] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [56] Y. Freund and R. E. Schapire, "Experiments with a new boosting algorithm," in Thirteenth International Conference on Machine Learning. San Francisco: Morgan Kaufmann, 1996, pp. 148–156.
- [57] J. J. Rodriguez, L. I. Kuncheva, and C. J. Alonso, "Rotation forest: A new classifier ensemble method," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 28, no. 10, pp. 1619–1630, 2006. [Online]. Available: <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2006.211>
- [58] T. Landgrebe and R. P. W. Duin, "A simplified extension of the area under the roc to the multiclass domain," in 17th annual Symposium of the Pattern Recognition Association of South Africa, November 2006.
- [59] R. M. Everson and J. E. Fieldsend, "Multi-class roc analysis from a multi-objective optimisation perspective," Pattern Recogn. Lett., vol. 27, no. 8, pp. 918–927, Jun. 2006. [Online]. Available: <http://dx.doi.org/10.1016/j.patrec.2005.10.016>
- [60] J. E. Fieldsend and R. M. Everson, "Visualisation of multi-class roc surfaces," 2005.
- [61] B. K. P. Horn, "Closed-form solution of absolute orientation using unit quaternions," Journal of the Optical Society of America A, vol. 4, no. 4, pp. 629–642, 1987.
- [62] D. Lowe, "Object recognition from local scale-invariant features," in Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on, vol. 2, 1999, pp. 1150 –1157 vol.2.
- [63] K. Hotta, "Object categorization based on kernel principal component analysis of visual words," in Proceedings of the 2008 IEEE Workshop on Applications of Computer Vision. Washington, DC, USA: IEEE Computer Society, 2008, pp. 1–8.
- [64] H. Zhang, A. C. Berg, M. Maire, and J. Malik, "SVM-KNN: Discriminative nearest neighbor classification for visual category recognition," in Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 2, ser. CVPR '06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 2126–2136.

APPENDIX SCHEMA DIAGRAM

The figures that follow depict the entity relationships present in the relational database used to organize this project. Because of the number of tables and the complexity of the full graph, it has been broken into 4 subgraphs, in order to make the text readable and the relationships visually navigable. Certain tables appear on multiple pages, and thus examining the relevant pages side-by-side may make it easier to follow the relationships.

The *IVUS* is present in both figure A.1 and figure A.2. Figure A.1 illustrates how the *IVUS* table fits into the root portion of the hierarchy, where high level objects such as *patients* are stored; figure A.2 shows how the various data items that depend upon the *IVUS* table relate to one another. Likewise with the angiogram table in figures A.1 and A.3.

Figure A.4 shows how the *angio* and *IVUS* entities reconnect for the 3-D reconstruction and the portion of the data modeling that depends upon the reconstruction.

Dotted lines are used to map foreign key relationships. These are labeled with text that contains both the source and destination key. All primary keys in this schema are integers. In most cases, they are bound to monotonically increasing sequences (this is evident when = *nextval* appears in the same line). In a small handful of cases, the primary key is the foreign key. This is the case in the *WSS* table, where there can be exactly 1 *WSS* per fusion. This is exploited in the *WSS* computation pipeline, in that the status of all *WSS* analyses can be queried with a simple join of the *WSS* and fusion tables.

The *finalsegmentations* table is used to allow someone manually tracing an *IVUS* data set to mark 1 particular version of the borders as definitive. This is important because with the tracing process being time consuming, people often save incrementally. This simplifies the process of selecting a segmentation to use in the downstream analysis.

Tracerassignments is used for personnel management – it permits the delegation of each IVUS data set to a particular expert tracer, so that they don't step on each others' toes, and so that there is accountability in the system.

Figure A.1: *Root* portion of the database schema hierarchy.

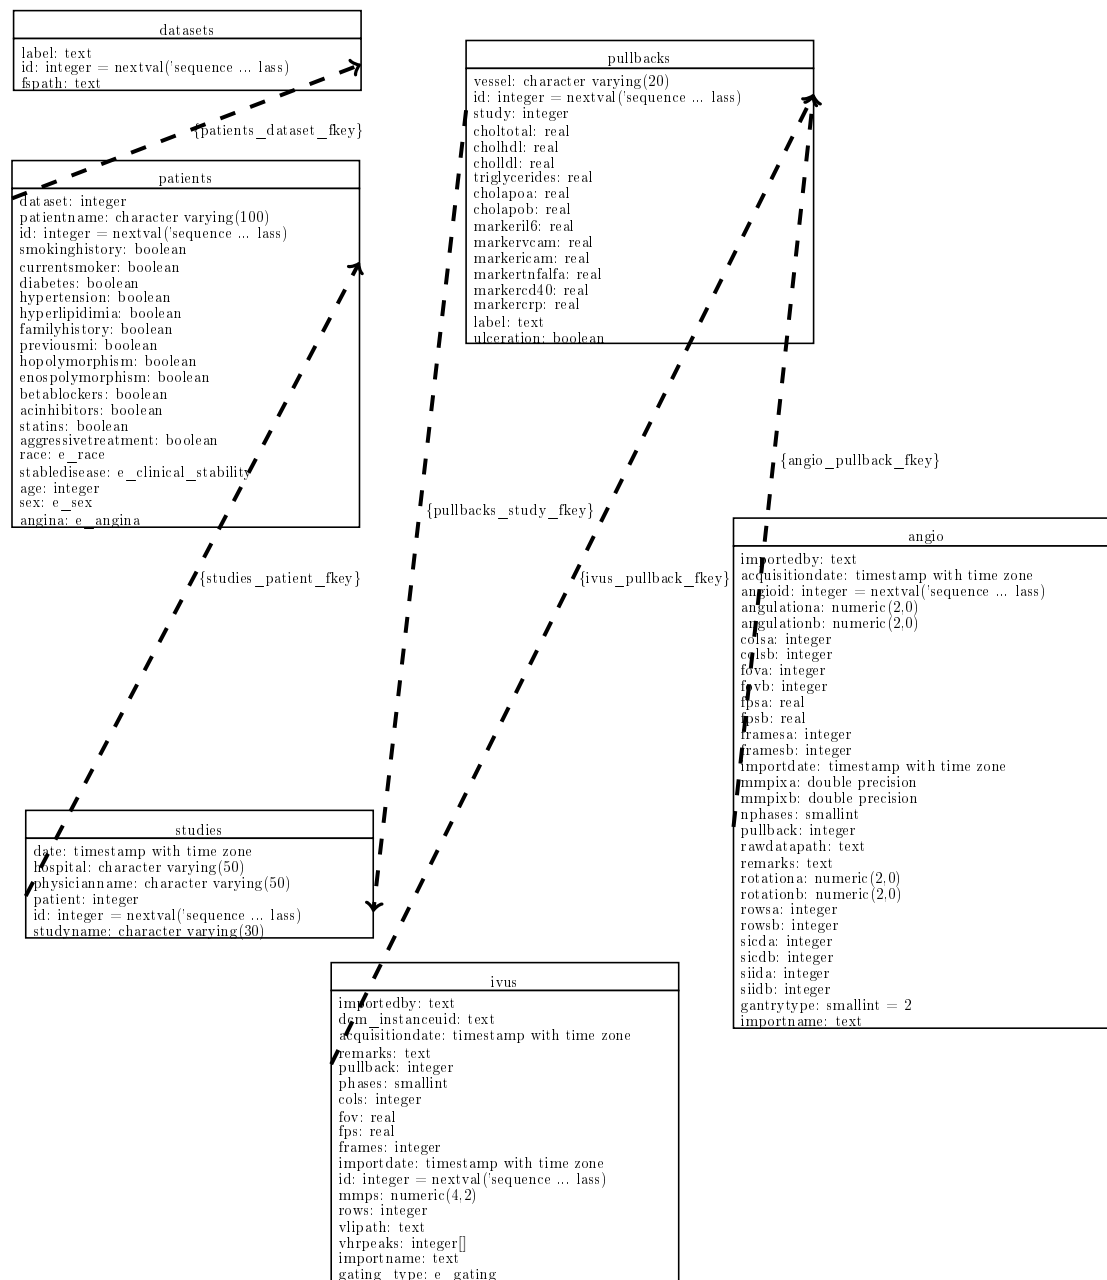


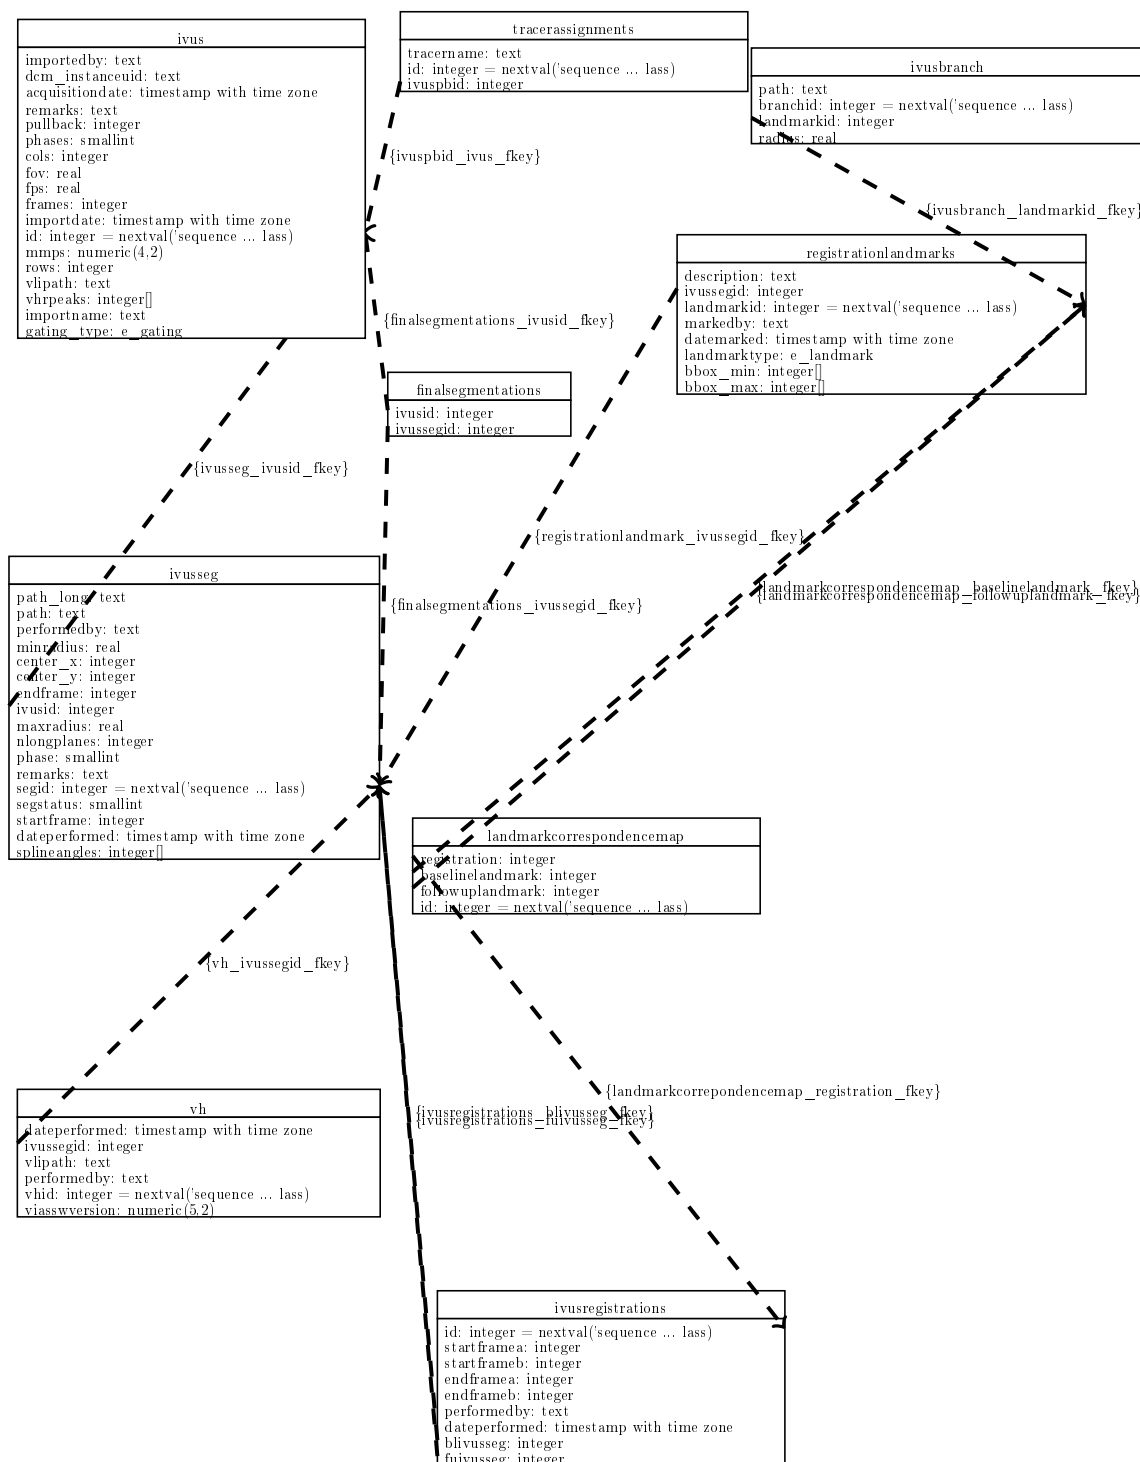
Figure A.2: *IVUS* portion of the database schema hierarchy.

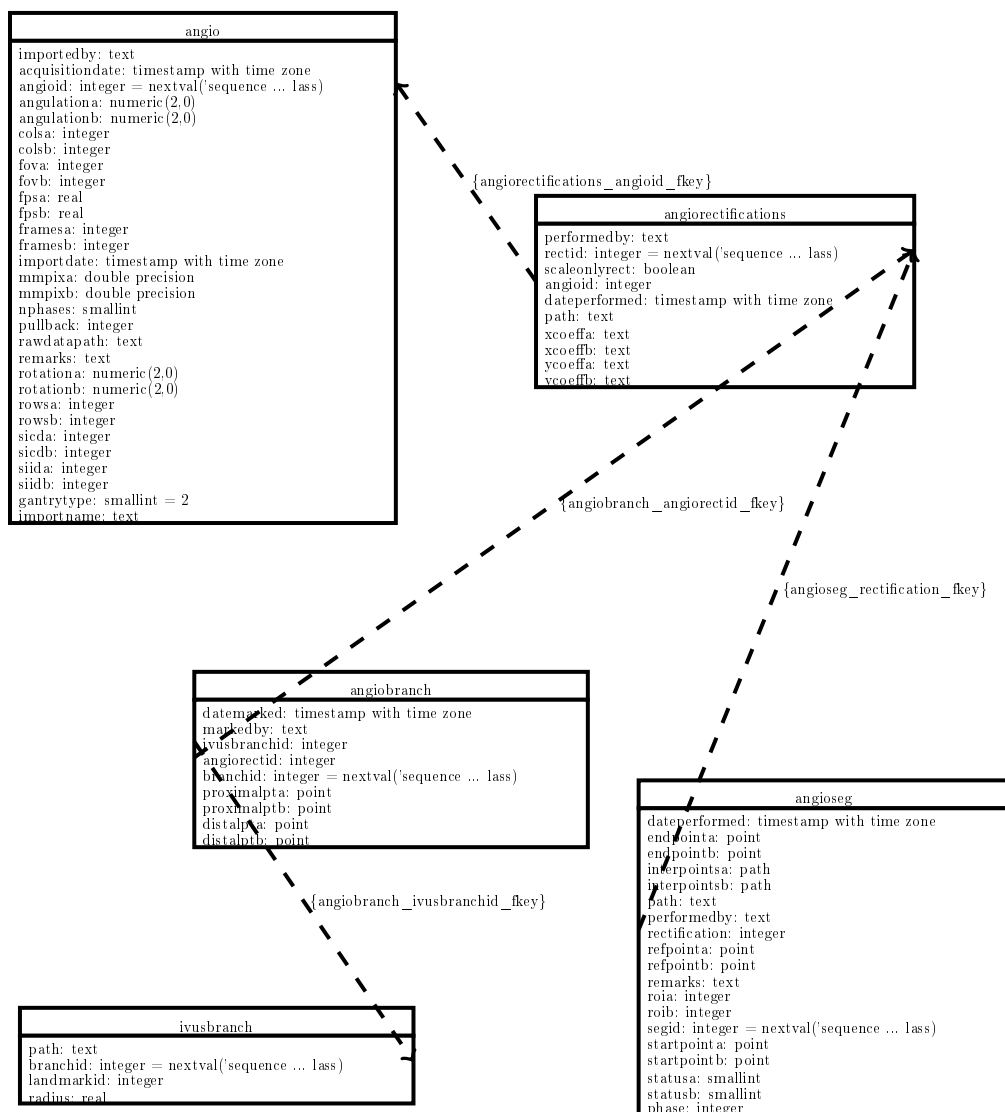
Figure A.3: *Angiogram* portion of the database schema hierarchy.

Figure A.4: *Fusion* portion of the database schema hierarchy.

