Navigation and Positioning

Aids for Intracoronary Interventions

Navigations- und Positionierhilfen

für intrakoronare Interventionen

Von der Fakultät für Elektrotechnik und Informationstechnik
der Rheinisch-Westfälischen Technischen Hochschule Aachen
zur Erlangung des akademischen Grades einer
Doktorin der Ingenieurwissenschaften genehmigte Dissertation

vorgelegt von

Bárbara Martín Leung

Ingeniera Superior de Telecomunicación aus Madrid

Berichter: Universitätprofessor Dr.-Ing. Til Aach
Universitätprofessor Dr.-Ing. Georg Schmitz

Tag der münlichen Prüfung: 10. Juli 2009

To my mom, the bravest person I know.

Acknowledgements

This thesis is the result of the work carried out during my PhD studies at the Philips Research Laboratories in Aachen (Germany) and Philips Medical Systems in Best (The Netherlands) through several years. Many of the results described in this dissertation were possible because of the support of a lot of people and I would like to use this opportunity to thank them.

First, I would like to thank Prof. Dr.-Ing. Til Aach for giving me the opportunity to do my PhD thesis in cooperation with an industrial partner under his guidance, first from Lübeck University and later at RWTH Aachen University. I am also grateful to Prof. Dr.-Ing. Georg Schmizt for being the co-examiner in the jury.

I want to thank Falko Busse and Michael Overdick for their continued interest in my work and the support they provided. I am indebted to my project colleagues, Kai Eck and Jörg Bredno, for their guidance. Alexandra Groth and Sabine Mollus, you contributed with countless fruitful discussions.

I am grateful for the very interesting discussions and ideas that ensued therefrom to our colleagues at Philips Medical Systems, Peter van den Houten, Wiet Lie and Rens Schoones. And then there is a multitude of colleagues I would like to thank at Philips Research Aachen: every single member of the Imaging Systems group for making work a really enjoyable experience, Salvador Boleko, Mark Carpaij, Toni López and David Sánchez for showing me how to look at things through different eyes.

My very special thanks go to my gang in Aachen, who have supported me through thick and through thin, specially to Daniel Totev, also a PhD student at the time, for teaching me how to look on the bright side of life. And last, but not least, to JL-Baobei, for enduring countless rehearsals and paying attention during each one of them!

Zusammenfassung

Perkutane transluminale koronare Angioplastie ist die derzeit meistgenutzte Behandlungsmethode für koronare Herzkrankheit. Zur Lokalisierung der Gefäßverengungen werden präinterventionelle Röntgenbilder aufgenommen, in denen die Koronararterien mit einem radio-opaken Kontrastmittel gefüllt sind. Eins dieser Bilder wird auf einem Monitor angezeigt und dient als Navigationshilfe während der Intervention. Gleichzeitig werden interventionelle Bewegt-Röntgenbilder auf einem zweiten Monitor in Echtzeit angezeigt, wobei es dem behandelnden Arzt obliegt, die beiden Bilddaten auf den beiden verschiedenen Monitoren zu korrelieren und die jeweilige Position des Katheters zu bestimmen. Die Navigationsaufgabe wird durch die kontinuierliche Bewegung des Herzens im Brustkorb und die statischen Gefäßbäume der präinterventionellen Röntgenbilder erschwert.

Ziel dieser Arbeit ist die Entwicklung von bildgestützten unterstützenden Systemen zur Navigation und Positionierung von Herzkathetern. Dies wird erreicht durch Synchronisieren der präinterventionellen Cine-Angiographie mit der interventionellen Fluoroskopie, und Überlagern der so synchronisierten Angiographie mit der Fluoroskopie. Zur weiteren Verbesserung der Katheter-Positionierung kann Information aus intravaskulärem Ultraschall (IVUS) eingebracht werden. Die Synchronisation von Angiographie und Fluoroskopie basiert auf der Registrierung von Atmung und Herzzyklus. Erstere erfolgt durch multimodale Bildregistrierung auf der Basis von Mutual Information, letztere auf Basis des EKG-Signals. Atmungs- und Herzzyklus beider Bilddaten werden dann in ein zweidimensionales Diagram eingetragen, in dem benachbarte Punkte dann die Auswahl von Einzelbildern aus den beiden Modalitäten mit ähnlicher Herz- und Atemphase erlauben. Die IVUS-Bilddaten werden ebenfalls durch Kompensation von Atmung und Herzzyklus auf die Position der IVUS-Sonde in den Angiogrammen bezogen. Unvermeidbare Lokalisierungsfehler werden durch eine Feinregistrierung von Sonde und Gefäßen reduziert, darüber wird ein Ansatz zur Rekonstruktion des Pullback-Wegs der Sonde vorgestellt. Diese Algorithmen werden mittels eines Herz-Phantoms unter Röntgenbildgebung getestet, wobei sich ein Restfehler von ca. 2mm ergibt. Schließlich wird eine flexible Software-Architektur vorgestellt, die den Transfer der Ergebnisse in einen klinischen Echtzeit-Demonstrator ermöglicht.

Abstract

Percutaneous Transluminal Coronary Angioplasty is currently the preferred method for treatment of coronary artery disease in vivo. X-ray images of the coronary arteries filled with a radio-opaque contrast dye are acquired to localise the lesion. One of these images is displayed on a monitor and will serve as roadmap to navigate the interventional instruments under constant x-ray surveillance. The interventional images are displayed in real-time on a second monitor next to the roadmap monitor. It depends on the physician's ability to correlate the images on both monitors and estimate the actual position of the guidewire within the vessel tree in order to steer it further. The navigation task is hindered by the constant movement of the heart within the chest cavity and the static nature of the angiographic roadmap.

The main goal of this work is to present a novel image registration framework to enable navigation and positioning aids for cardiac catheterisation procedures to reduce interventional time. This task is accomplished by enhancing the visualisation of the procedure by means of overlaying onto each interventional frame a roadmap that exhibits the coronary vessels in a similar position and with a similar shape, and by fusing information acquired from Intravascular Ultrasound (IVUS). To this end, we characterise each acquired frame by a heart state vector, which main components are respiration state and heart contraction state . The latter is calculated from the ECG signal, and we propose and evaluate a multimodality similarity measure based on mutual information to measure the respiration state from the image content. With these two characteristic parameters, each image can be represented in a two-dimensional graph where nearby points represent images that show the heart in a similar state and position. Furthermore, methods to refine the registration of the vessels visible in the roadmap and the instrument visible in the interventional frames by respiration compensation and accurate device to vessel registration are proposed. A flexible software architecture is developed to ease transfer from research activities to a clinical setting and test the real-time performance of these algorithms.

To aid accurate positioning, we present a novel multimodality application to fuse IVUS and projection x-ray images. IVUS images are related to their acquisition locations in the diagnostic angiograms or to instrument positions during the intervention by means of respiration compensation and accurate device to vessel registration. A reconstruction algorithm of the

three-dimensional pullback path for projection x-ray sequences is also proposed. The two-dimensional registration algorithm for IVUS and angiography is tested in a heart phantom.

Contents

List of Acronyms

ACC	American College of Cardiology
AMI	Acute Myocardial Infarction
AVID	Angiography Versus Intravascular ultrasound Directed coronary stent placement
CABG	Coronary Artery Bypass Graft
CAD	Coronary Artery Disease
CHF	Congestive Heart Failure
CPU	Central Processing Unit
CRUISE	Can Routine Ultrasound Influence Stent Expansion
CT	Computed x-ray Tomography
CRPM	Cardiac and Respiratory Parametric Model
CRPP	Cardiac Respiratory Phase Plane
CSA	Cross-Sectional Area
CV	Coefficient of Variation
CVD	Cardiovascular Disease
DICOM	Digital Imaging and COmmunications in Medicine
DSA	Digital Subtraction Angiography
EBCT	Electron Beam Computed Tomography
ECG	ElectroCardioGram
EEM	External Elastic Membrane
GMI	Generalised Mutual Information
GUI	Graphical User Interface
JPEG	Joint Photographic Expert Group
IPC	InterProcess Communication
IVUS	IntraVascular UltraSound
LAD	Left Anterior Descending coronary artery
LAS	Lumen Area Stenosis
LC	Local Correlation
LCA	Left Coronary Artery
LCX	Left Circumflex artery
LDL	Low-Density Lipoprotein
LE	Lumen Eccentricity
MDD	Mean Distance between Diaphragm points
MDCT	Multi-Detector Spiral CT
MI	Mutual Information
MLD	Maximum Lumen Diameter
mld	Minimum Lumen Diameter
MoG	Mixture of Gaussians
MPM	Maximum Plaque plus Media thickness

mpm	Minimum Plaque plus Media thickness
MPS	Medical Positioning System
MRI	Magnetic Resonance Imaging
MSCT	Multi-Slice Spiral CT
MSD	Maximum Stent Diameter
msd	minimum stent diameter
NURD	Non-Uniform Rotational Distortion
OCT	Optical Computed Tomography
PB	Plaque Burden
pdf	Probability Density Function
PTCA	Percutaneous Transluminal Coronary Angioplasty
RCA	Right Coronary Artery
RMS	Root Mean Squared error
SM	Smart Mask
SNR	Signal to Noise Ratio
SS	Stent Symmetry
TULIP	Thrombocyte activity evaluation to determine the effect of Ultrasound guidance in Long Intracoronary stent Placement
US	UltraSound

List of Figures

Chapter 1

Introduction

The heart is a muscle that must work constantly and hence requires a constant supply of nutrients and oxygen, which are delivered by the blood flowing through the coronary arteries. *Atherosclerosis* is a disease in which an agglomeration of fatty substances or *plaque* builds up inside the vessel walls, thus narrowing the vessel lumen. Such a constriction of the lumen diameter is known as a *stenosis*. In *Coronary Artery Disease* (CAD), atherosclerosis affects the coronary arteries and results in a reduction of the blood supply to the heart muscle (*cardiac ischaemia*). Further complications might occur and lead to a clot, which completely cuts off the blood flow to the area of the myocardium supplied by the affected coronary artery, hereby causing an *acute myocardial infarction* (AMI) or heart attack, which can be fatal. CAD is the most common form of heart disease and the leading cause of death in USA and Europe.

Factors that are related to the incidence of CAD are: age, high serum cholesterol, hypertension, tobacco, alcohol consumption, diabetes mellitus, overweight and physical inactivity. The increasing tendency in the industrialised countries towards unhealthy lifestyles and the lengthened lifespan of the population will result in an increase in the incidence of chronic diseases, including CAD [8]. Table 1.1 shows the overall national bill for the top five diagnoses in USA hospitals ranked according to aggregate charges in the year 2005 [1]. We can observe that CAD holds the first place of the ranking and that other two heart diseases, *Acute Myocardial Infarction* (AMI) and *Congestive Heart Failure* (CHF), are among the three top-ranking diseases. The increase in CAD incidence throughout the years can be observed in figure 1.1, which displays the national bill for CAD, CHF and AMI from 1997 until 2005. A steady increase in the charges originating from these three cardiac diseases is noticeable.

Rank	Primary diagnosis	National Bill
1	CAD	$52.623.837.490
2	AMI	$35.114.820.856
3	CHF	$32.762.953.824
4	Liveborn	$31.289.090.319
5	Septicemia	$30.320.496.884
6	Osteoarthritis	$28.038.010.569

Table 1.1: The USA national bill for the top five diagnoses treated in USA hospitals in 2006 ranked according to national bill expenditure [1]. It can be observed that heart-related diseases occupy the first three positions.

.

CAD is usually diagnosed and treated performing a minimally invasive intracoronary intervention, Percutaneous Transluminal Coronary Angioplasty (PTCA), which employs a *Cathlab* system. A Cathlab system acquires x-ray images of the heart throughout the whole procedure, which are made available to the physician on monitors. The physician heavily relies on these images to diagnose and treat the lesion. This thesis presents a novel registration framework and algorithms to process the images acquired by the Cathlab system and augment the information made available to the physician in order to improve the intervention results and shorten the interventional time.

1.1 Motivation

In the past few decades, minimally invasive surgery has revolutionised CAD treatment, reducing the number of highly invasive surgery procedures, like *coronary artery bypass grafting* (CABG). Figure 1.2 shows the number of in-hospital deaths for both types of intervention from 1993 until 2002. The national bill for minimal-invasive procedures has tripled from the year 1997 to the year 2005, whereas the national bill for CABG has increased 21.39% (see table 1.2). We can therefore conclude that minimally invasive procedures are nowadays the preferred methods for CAD treatment.

In minimal-invasive procedures, the plaque location is accessed via *catheters* (long, thin, flexible tubes) which are inserted in peripheral arteries and guided through the vascular system into the affected coronaries. These catheters are

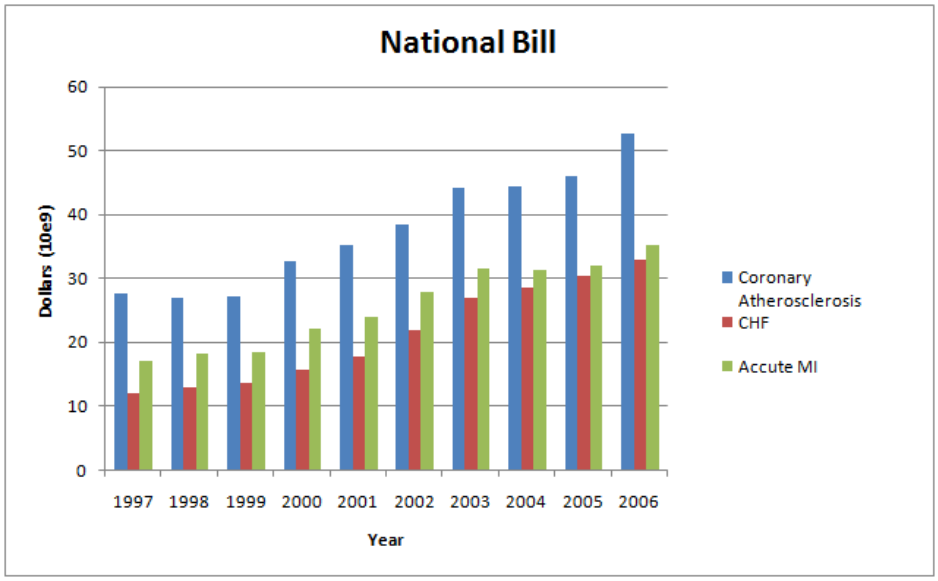

Figure 1.1: National bill for CAD, CHF and AMI between 1997 and 2006. This graphic represents data provided by [1].

used as guiding conduits for other instruments that are used to widen the lumen from within the diseased artery. The procedure is assisted by projection x-ray images of the heart acquired with a Cathlab system and the occasional injection of radio-opaque contrast agent into the coronary arteries.

The main tasks during minimal-invasive procedures are

1. the diagnosis of the plaque characteristics,
2. the navigation of the instruments through the coronary vessel tree to the lesion site, and
3. the correct positioning of these instruments within the vessel.

For diagnosis, navigation and treatment the physician relies on projection x-ray images acquired throughout the procedure. For diagnosis a number of images that display the coronary arteries filled with radio-opaque contrast agent are inspected to assess the lesion. Later on during navigation and instrument positioning images are being constantly acquired by the Cathlab system while the physician manipulates the interventional instruments. The duration and result of the intervention heavily relies on the skill of the physician to visually process the images presented on the monitors and manipulate the interventional instrument accordingly. There is an uncertainty

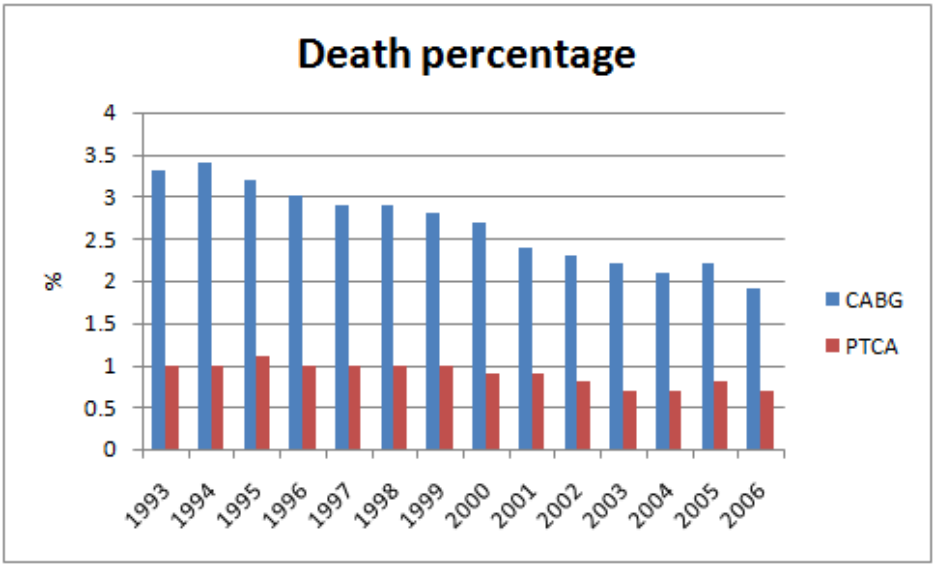

Figure 1.2: Number of in-hospital deaths in the USA from 1997 till 2005 for PTCA and CABG interventions. This graphic represents data provided by [1].

Year	Minimal-invasive	CABG
1997	$11.471.144.347	$18.415.324.002
1998	$12.095.222.688	$17.453.743.633
1999	$12.948.913.058	$16.915.400.837
2000	$16.624.789.630	$16.624.789.630
2001	$20.046.276.130	$20.949.896.887
2002	$23.600.884.492	$22.353.879.823
2003	$26.544.181.361	$24.429.744.205
2004	$31.825.054.479	$21.901.579.530
2005	$35.962.749.498	$22.246.051.703

Table 1.2: National bill in USA for minimal-invasive procedures and CABG [1].

in diagnosing, navigating and positioning the instrument caused by the constant movement of the heart inside the thoracic cavity, which may affect the intervention outcome as follows:

1. An inaccurate diagnosis leads to the wrong decisions as to treatment and instrument choice.
2. The uncertainty in the instrument navigation increases the intervention time and therefore, the x-ray dose received by patient and physician, especially for tortuous vessel-tree anatomy.
3. Inaccurate positioning of the instrument at the lesion leads to a suboptimal treatment result.

If would be desirable for the patient and the physician to

1. reduce intervention time, in order to reduce the exposure of the medical staff and the patient to x-ray radiation,
2. reduce the amount of contrast agent injected to the patient, in order to reduce the risk of an allergic reaction or renal deficiency in the patient, and
3. improve treatment, in order to reduce further complications or prevent the lesion from reappearing.

These goals can be achieved by providing the physician with the tools to better navigate and position the interventional instruments. The physician relies on the Cathlab images for such tasks. By processing the acquired images with the algorithms presented in this thesis it is possible to effectively achieve the goals listed above.

1.2 Statement of the problem

Currently, navigation and positioning during minimal-invasive procedures are supported by a single static x-ray image where the vessels, filled with contrast dye, are visible. This image or *roadmap* is displayed next to another monitor where live x-ray images of the procedure are shown, in which the instrument is visible. This roadmapping technique is known and used in non-cardiac intravascular interventions. When used in cardiac interventions the physician has to visually process the information displayed in one static image and live images to extract the information necessary for instrument navigation and

positioning. The ability to correlate the roadmap and the current live intervention images relies on the cardiologist's experience and ability to *estimate* the position of the instrument if it could be displayed on the roadmap monitor. This task is constantly impeded by the movement of the heart inside the thorax caused by breathing and heart contraction. The overlay of the static roadmap to the live images improves little the information provided to the physician to aid navigation because of the constantly changing state and position of the heart.

Enhancing the information offered to the physician to support diagnosis, help to navigate and position the instruments will reduce the necessary x-ray and contrast dye dose, improve the outcome of the procedure and increase patient throughput. In this thesis, a new image registration framework and novel methods that provide navigation and positioning support for intracoronary interventions is described. The presented methods fuse x-ray information acquired before and during the intervention, and ultrasound images acquired from within the vessels (*intravascular ultrasound*) in order to support diagnosis and aid navigation and accurate positioning of the interventional instrument.

1.3 Document structure

This thesis is organised as follows:

Chapter 2 describes the necessary background information about CAD and its diagnosis and treatment.

Chapter 3 discusses navigational and positioning support for minimal-invasive procedures and introduces the proposed image registration framework.

Chapter 4 presents the novel methods proposed to aid navigation and instrument positioning during PTCA interventions and analyses their performance on clinical data.

Chapter 5 explains further work on accurate vessel to device registration algorithms.

Chapter 6 presents a novel method for fusion of x-ray and intravascular ultrasound images.

Chapter 7 explains the real-time demonstrator built in order to test the proposed algorithms in a hospital setting.

Chapter 8 summarises the conclusions and discusses further work.

Chapter 2

Background

In this chapter we will present the necessary background information to understand current techniques for CAD treatment and diagnosis, and the rationale behind our proposed solution. We will give an outline of coronary artery anatomy, CAD symptoms and development, and the current diagnosis and treatment techniques. An overview of currently used imaging techniques and modalities will be presented, as well as an introduction to currently utilised Cathlab systems. Section 2.1 presents the anatomy of the heart and the coronary vessels. In section 2.2 the mechanisms for plaque formation are presented. Section 2.3 addresses the issue of CAD diagnosis and treatment, focusing on minimal-invasive procedures aided by x-ray imaging. Alternative imaging modalities are presented in section 2.4. In section 2.5 we summarise and draw conclusions as to how to aid navigation and positioning of instruments during intracoronary interventions.

2.1 Anatomy

The heart is supplied with blood by the coronary arteries, which branch off from the aorta at its commencement into several secondary vessels and capillaries (see figure 2.1). The two main coronary arteries are referred to as the *right* and *left coronary artery* (RCA and LCA). The LCA branches into the *left anterior descending* (LAD), that supplies blood to the front of the heart, and the *left circumflex* (LCX), that encircles the heart. The RCA divides into the *right posterior descending artery* and a large marginal branch.

The coronary arteries consist of three layers [10] (see figure 2.2):

1. The *intima layer* is the innermost layer. It is a lining of endothelium,

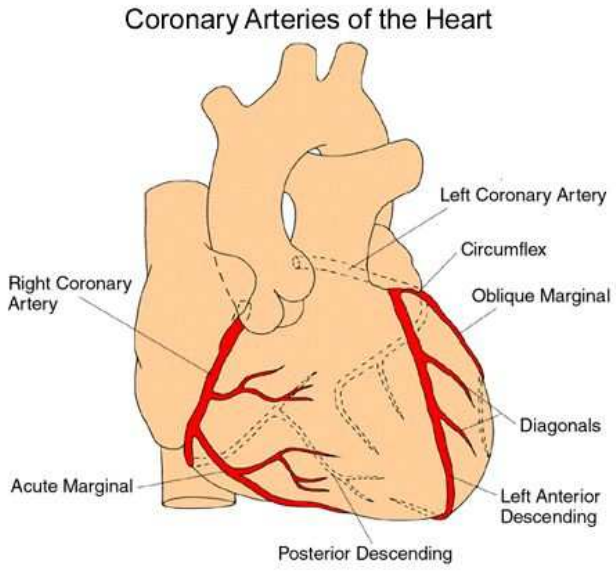

Figure 2.1: Anatomy of the heart and the coronary vessels. (Image courtesy of the Stanford Hospital [9]).

collagen and elastin that provides a smooth surface for the blood to flow over. It is separated from the next layer by the *internal elastic membrane*.

2. The *media layer* is composed of smooth muscle cells and elastic fibrous tissue that contract and relax to control the blood pressure and flow.

3. The *adventitia* is the outer layer and is composed of connective fibrous tissue. It is separated from the media by the *external elastic membrane* (EEM).

2.2 Plaque formation

Plaque is an agglomeration of fatty substances inside the vessel wall that results in the formation of a *stenosis*, a constriction of the lumen. In the presence of a stenosis the blood supply to the muscles is affected. When such a constriction is detected in a coronary artery, coronary artery disease is diagnosed. Recent research about plaque morphology and development has helped to improve intracoronary interventional devices to better adapt

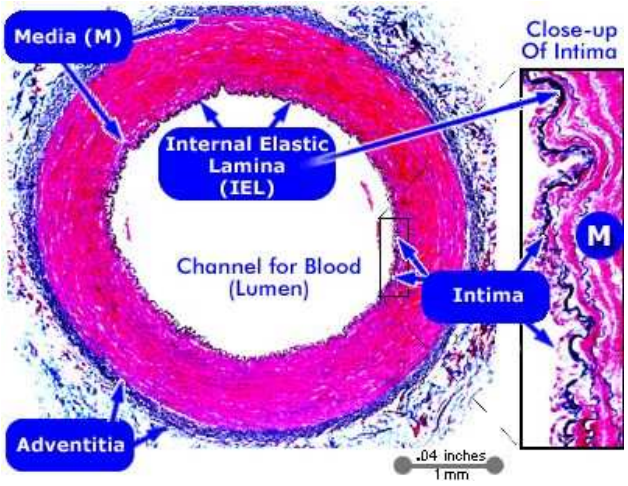

Figure 2.2: Coronary artery structure (image courtesy of TCT [2]): the intima, the media and the adventitia layer are visible in this image.

to the lesion characteristics.

The build-up of plaque inside a coronary artery wall is an inflammation response triggered by an excessive concentration in the blood of low-density lipoprotein (LDL) (a.k.a. bad cholesterol) [11][12][13]. LDL can pass in and out of the blood stream into the vessel wall, but at excessive concentrations it might remain inside the intima. The trapped lipids experience oxidation, which is interpreted as a danger sign by the cells from the vessel wall. Defensive cells pass from the blood into the *endangered* area and ingest the oxidated LDL, releasing healing molecules which favour the replication of smooth muscle cells in the media.

During the early stages of CAD, the affected vessel segment expands outwardly (*positive remodelling*) and does not hinder the blood flow through the artery despite plaque accumulation. The generated muscle cells eventually migrate from the media to the top of the intima, forming a fibrous cap on top of the original atherosclerotic zone. As the disease further advances, the plaque begins to grow inwardly (*negative remodelling*) until a lipid or necrotic core is formed under the fibrous cap. It is the process of negative remodelling that narrows the artery lumen and hinders the blood flow to the heart muscle. The lesion is then referred to as a *stenosis*. In figure 2.3a we can observe an histology sample where we can see how the artery lumen has been narrowed by the growth of a lipid core (white), which is covered by a

fibrous cap.

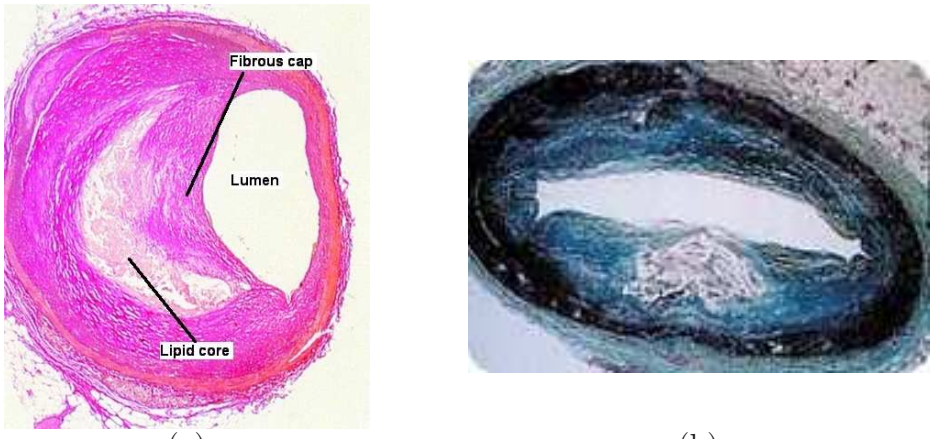

(a) (b)

Figure 2.3: Histology samples of plaque in a coronary artery. (a) The white area is the lipid core which is covered by the fibrous cap formed by smooth muscle cells. (Image taken from [14]). (b) Vulnerable plaque characterised by a thin fibrous cap. (Image taken from [15]).

The knowledge about plaque formation has helped to understand the process that causes an acute myocardial infarct (AMI). Until recently, it was believed that a heart attack was caused by a stenosis big enough to block completely the lumen of a coronary artery. It remained a mystery why 85% of heart attacks occurred without the presence of previous symptoms [11]. Now it is known that an AMI occurs when the thin fibrous cap covering the lipid core of a *vulnerable plaque* ruptures, thereby releasing the lipid core into the blood stream, which blocks the artery lumen downstream [11]. In figure 2.3b we can see an example of vulnerable plaque. On the lower part of the lumen we can observe a light lipid core covered by an almost non-discernible cap. If this cap were to break, the lipid core would be released into the blood stream probably causing an AMI. Imaging systems for detection of vulnerable plaque are an active field of research.

Histology studies have concluded that there is a direct relation between coronary calcium deposits and vulnerable plaque [16][17]. Coronary calcium is sometimes used as an indicator of the presence of vulnerable plaque although it may occur that no calcium in the coronary arteries appears despite the presence of vulnerable plaque. Although calcification is an indication atherosclerosis, the plaques causing acute coronary syndromes may have less calcium than those causing stable angina [12].

2.3 Minimally invasive techniques for diagnosis and treatment of CAD

Usually a patient suffering from CAD goes to the doctor complaining about chest pain (*angina pectoris*) and shortness of breath (*dyspnea*). Depending on the symptoms and the medical history of the patient a number of tests are performed. CAD can be diagnosed by indirect measurements on some physiological variables that depend on the heart performance. Nevertheless, only by acquiring images of the coronary arteries can CAD be accurately diagnosed and the location of the stenosis identified. A number of imaging modalities are available, like x-ray angiography [18], echocardiography [18], magnetic resonance imaging (MRI) [18][19][20][21][22][23] and computed x-ray tomography (CT) [18][24][25][26][27][28][29].

X-ray angiography is a minimally invasive procedure that remains the gold-standard for CAD diagnosis and ulterior navigational aid during treatment. Diagnosis relies on x-ray images or *angiograms* acquired after injecting the affected coronary artery with a radio-opaque contrast dye. The acquired images are inspected looking for a constriction in the visible lumen. If it is decided that the lesion should be treated with minimally invasive techniques, these angiograms are later utilised by the physician as guidance throughout the intervention.

The following subsections give an overview of x-ray imaging and cardiac x-ray imaging systems used for minimally invasive CAD interventions. Subsection 2.3.1 introduces some basic concepts about x-ray imaging. In subsection 2.3.2 we explain x-ray imaging for cardiac applications. In subsection 2.3.3 the catheterisation laboratory, where the minimal-invasive procedures are performed, is presented. In subsections 2.3.4 and 2.3.5 we discuss the characteristics of two different cardiac x-ray imaging modalities used for diagnosis and treatment of CAD and the corresponding diagnostic and treatment procedures.

2.3.1 X-ray imaging

In x-ray imaging an x-ray beam is generated in an *x-ray tube*, passes through the patient, and is partially absorbed and scattered in the process. The beam that leaves the patient is detected by an image receptor (image intensifier or flat-panel detector) and the received x-rays are transformed into an intensity

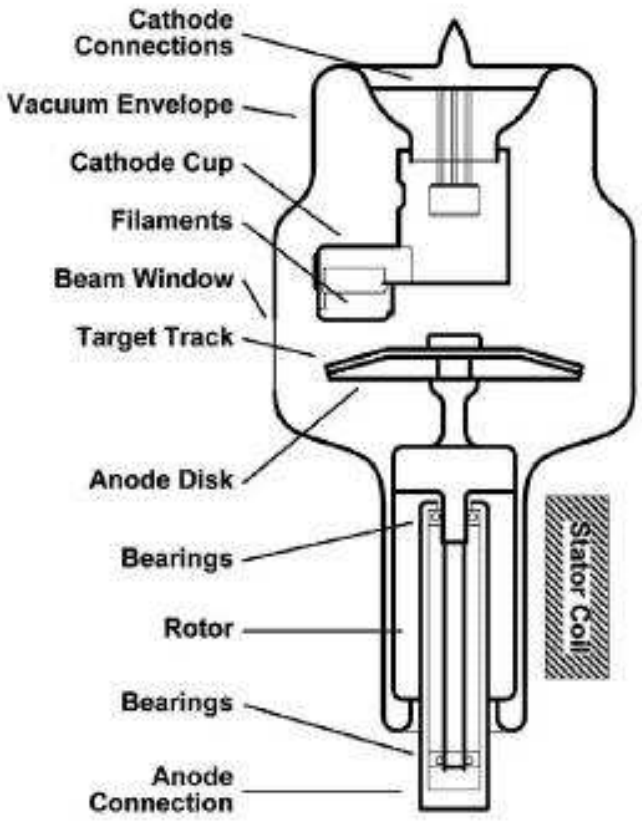

Figure 2.4: X-ray generator (image courtesy of the American College of Cardiology [3]).

image that is displayed on a monitor. A structure is visible in the resulting image if its x-ray absorption is different from that of the surrounding structures.

The x-ray tube has a cathode and a metallic anode (see figure 2.4). The cathode consists of a metallic filament that is electrically heated and thereby releases electrons. A voltage difference between cathode and anode accelerates the electrons. When these high energy electrons hit the anode, they interact with the anode's material and a portion of the electrons' energy is transformed into x-ray photons [3]. The cathode current determines the number of photons produced, whereas cathode-anode voltage determines the energy of the x-rays produced [3].

Dose, image contrast and noise are influenced by the current and voltage

settings of the x-ray tube [3]. The *dose* is a measure of the x-ray radiation absorbed by the patient's tissue. A higher voltage setting produces higher energy x-ray photons (*hard* x-ray) that are less attenuated (absorbed) by the passage through the patient, therefore decreasing the dose to the patient. However, if less photons are absorbed by the tissue, the contrast of the resulting image is degraded, as the difference in absorption between tissue types is reduced. The current of the cathode is related to the amount of x-ray photons created. If the current is increased, the noise present in the image decreases, and the dose increases. If a pulsed current is used to generate the x-rays, the overall absorbed dose depends on the pulse rate and the total duration of the imaging process. If a continuous sequence of images is acquired, the acquisition frame rate is usually a multiple of the pulse rate and therefore the higher the frame rate, the better the temporal resolution and the higher the absorbed dose.

Unfortunately, x-ray is harmful for the human body. Therefore, there must be a tradeoff between x-ray exposure dose at the entrance of the patient or *skin dose* and image quality. Imaging requirements as to contrast, temporal resolution and noise are different during the CAD diagnosis procedure and during CAD treatment as will be explained in the following subsections. Also the duration of the imaging procedure during diagnosis and intervention should be taken into account when choosing voltage, current and frame rate, in order to minimise the amount of dose a patient is exposed to and maximise the image quality. During diagnosis, a high image quality and low quantum noise are desired, whereas during the intervention, a reduction of x-ray dose absorbed by the patient is advisable due to its long duration.

It is an objective of current medical-image processing techniques applied to x-ray images, such as the ones presented in this thesis, to reduce the intervention time and the x-ray dose to the patient without degrading the intervention outcome. The solution presented in this thesis applies novel image processing algorithms to aid navigation and positioning by enhancing the information presented to the physician, thereby speeding up the procedure. The framework and algorithms presented in chapters 3, 4 and 5 have been tested in a clinical environment (see chapter 7) and received positive feedback from the physicians.

2.3.2 X-ray for cardiac imaging

Nowadays, x-ray is the preferred method to image the coronary arteries because of its high temporal and spatial resolution. During image acquisition the heart is constantly beating. The most significant effect of heartbeat is a deformation of the heart muscle and therefore the vessels on its surface, which further hinders the evaluation of the stenosis.

For diagnosis, short image sequences are acquired in order to locate the lesion and estimate its characteristics, most notably its dimensions. Typically during the diagnostic procedure, sequences are first acquired and are afterwards reproduced and closely studied in order to examine the contrasted arteries looking for lumen constriction sites or abnormal movement of the coronary arteries. Therefore high spatial resolution is necessary to accurately diagnose the severity of the lesion. An imaging modality with high temporal resolution is also desirable because the heart is constantly beating and moving because of patient breathing.

During intracoronary treatment of CAD high spatial resolution is required to accurately place the instruments at the lesion site and thus optimise treatment outcome. An imaging modality with very high temporal resolution is not needed because the heart is constantly beating and moving inside the patient's thorax due to respiration, and the acquired images are being examined live in order to navigate the instruments to the lesion site. Thus, a high-enough temporal resolution is sufficient during this stage of the procedure to allow the physician to examine the heart throughout the different stages of the heart cycle.

When a patient is imaged with x-ray, different tissue types absorb a different amount of incident radiation. The coronary arteries are composed of soft tissue and are not discernible from the surrounding soft tissue in x-ray images. During diagnosis of CAD using x-ray cardiac imaging a radio-opaque contrast dye is injected into the coronary arteries in order to visualise them in the acquired x-ray images. During minimally invasive intracoronary treatment, the radio-opaque or partially radio-opaque instruments that have been introduced into the coronary vessels are visible in the images. Contrast agent may also be injected during the treatment procedure in order to visualise the coronary arteries while navigating the instruments to the stenosis. Depending on the acquisition geometry, ribs, diaphragm and spine might be distinguishable. Acquisition geometry might be chosen so that the soft tissue below the diaphragm, which may obscure the view of the coronaries or the instruments,

is kept out of the field of view.

2.3.3 The catheterisation laboratory

The catheterisation laboratory or *Cathlab* is a special room in the hospital equipped with an x-ray imaging system were patients can be diagnosed and treated from CAD. Throughout this document, we will refer to the room as *catheterisation room* and to the system as *Cathlab system* or simply *Cathlab*. A Cathlab system is employed to acquire x-ray images of the heart for diagnosis and CAD treatment. The images or sequences acquired with the system are displayed on medical monitors next to the interventional table for the physician to examine.

A *monoplane Cathlab* is a system equipped with an x-ray source and a detector mounted on a C-shaped frame or *C-arm* for improved mobility and accessibility to the patient (see figure 2.5a). Monoplane Cathlabs provide 2D x-ray projection images of the heart. A *biplane Cathlab* system (figure 2.5b) consists of two C-arms that can acquire projection images simultaneously, so that 3D reconstruction of the coronary tree is theoretically possible. Biplane systems are not common in cardiac catheterisation rooms, because of their high cost [12], and their inadequacy to guide the treatment procedure. Although 3D imaging is theoretically possible with a biplane system, the constant movement of the heart makes it difficult to present to the physician a clear vessel tree representation with sufficient accuracy for guidewire navigation during cardiac procedures. Throughout this thesis we will refer to monoplane systems and present results based on monoplane acquisitions, although the framework and algorithms presented in this thesis are also applicable to biplane systems.

2.3.4 X-ray (cine)angiography for diagnosis

In this subsection the workflow for CAD diagnosis in the catheterisation room will be presented. At the beginning of the diagnostic procedure, a small incision is performed either at the right leg of the sedated patient to access the femoral artery, or at the arm to access the brachial artery or at the wrist to access the radial artery. A sheath is then inserted to provide access to the vascular system. Through this sheath, a guidewire is advanced up to the aortic arch. A diagnostic catheter is advanced over the guidewire until it reaches the entrance of the coronary artery to be treated. Once the catheter

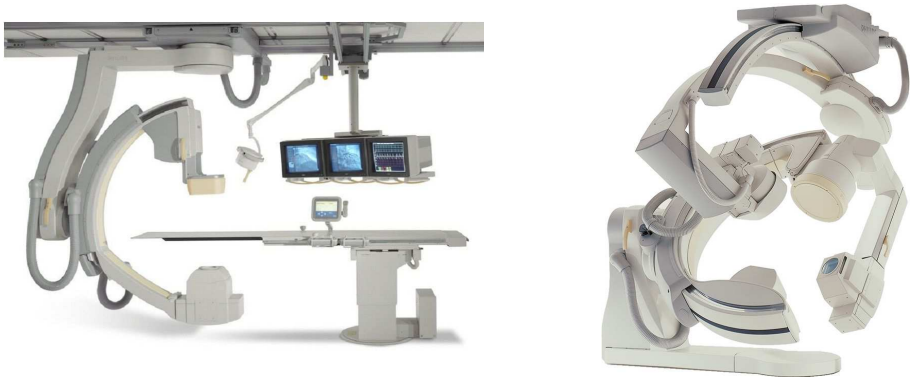

(a) Monoplane Cathlab system (b) Biplane Cathlab system

Figure 2.5: Cathlab systems: (a) a monoplane Cathlab system consists of an x-ray source and a detector mounted on a C-arm (image courtesy of Philips Medical Systems [30]); (b) a biplane Cathlab [31] system consists of two C-arms that acquire images simultaneously (image courtesy of Philips Medical Systems [31].)

is in place, a radio-opaque contrast dye is injected. Throughout this procedure x-ray images are continuously acquired with the Cathlab and this set of acquired angiograms $\{\mathcal{A}\}$ is stored. It might be decided to acquire sets of angiograms with different geometry settings of the C-arm.

When the radio-opaque contrast agent is injected into the coronary artery, the lumen filled with contrast agent appears visible in the acquired x-ray images (see figure 2.6) for a few heart beats, until the contrast agent is washed out by the blood. These sequences are usually acquired in breath-hold, to keep the diaphragm and the tissue below it away from the field of view and prevent the heart displacement caused by continuous respiration. The acquired images are examined for abnormal lumen constrictions which indicate the presence of a stenosis (figure 2.7). In order to diagnose very eccentric plaque or vessels with a difficult geometry when using a monoplane system, several x-ray diagnostic sequences with different geometry settings of the C-arm may be acquired. The selected treatment might be medication, bypass surgery or to treat the patient directly using catheter-based techniques.

Diagnostic frames or *angiograms* are meant to be visualised as a continuous sequence of frames or as individual frames for closer examination. The image quality should be high to perform and accurate diagnosis of the lesion, i.e. contrast and temporal resolutions should be high, and noise levels low.

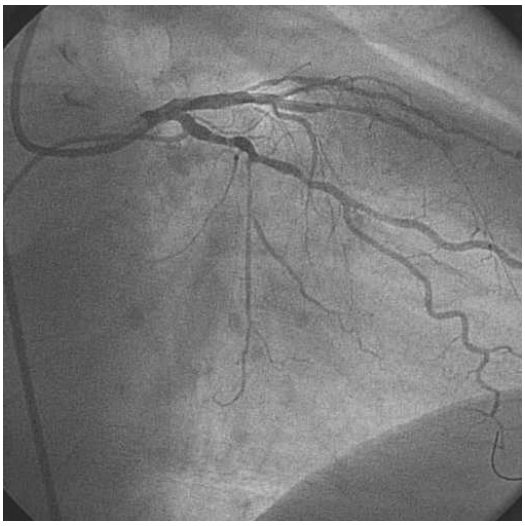

Figure 2.6: Diagnostic x-ray image. The coronary arteries filled with radio-opaque contrast dye are visible in x-ray. (Image courtesy of Catharina Hospital Eindhoven)

To produce hight-quality diagnostic images the Cathlab is set to *angiography* mode. In order to improve contrast, the voltage between the cathode and anode should be low to produce *soft* x-rays. The cathode current is set high, to reduce noise in the resulting images. In Philips Cathlab systems, the default voltage for diagnostic acquisitions is 80kV and changes dynamically depending on the absorption of the patient, exposure and contrast of the resulting images. The current is adjusted depending on the voltage. These settings (low cathode-anode voltage, high cathode current) cause a high dose per second absorbed by the patient. A high frame rate (high pulse rate) also increases the overall dose to the patient. The used frame rate should be low enough to reduce dose and high enough to provide a suitable temporal resolution when imaging the moving heart. The typical frame rate for adults used during x-ray angiography is 12.5 frame/s.

The resulting high dose to the patient is necessary to obtain good quality images of the lesion, and is tolerable because diagnostic imaging is usually short compared to the intervention, so that the overall exposure dose to the patient is kept within acceptable levels [3].

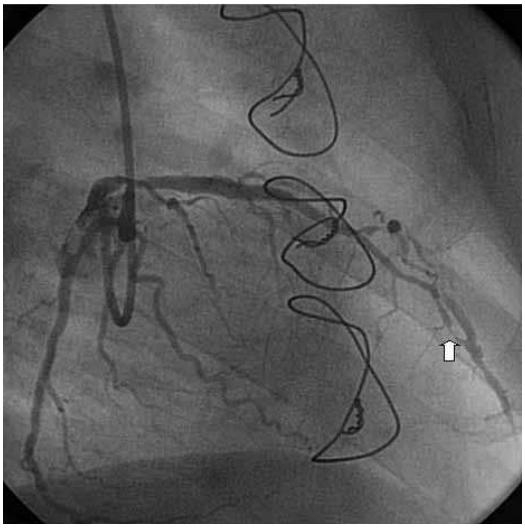

Figure 2.7: A stenosis (arrowed) is recognised in the x-ray diagnostic images as a constriction of the residual lumen. (Image courtesy of Catharina Hospital Eindhoven)

2.3.5 X-ray fluoroscopy for treatment

If it is decided to treat the lesion with minimally invasive techniques, then the patient remains on the Cathlab table. The physician proceeds to replace the diagnostic catheter by a guiding catheter which will be used to deliver a flexible guidewire and the interventional instruments to the lesion site. One of the x-ray diagnostic images is selected and displayed on one of the monitors to serve as a *roadmap* throughout the intervention, in order to aid the physician to navigate the instruments (see figure 2.8). During the treatment procedure a continuous stream of x-ray frames is provided on a second monitor. These images are used by the physician to monitor the progress of the guidewire, which is partially radio-opaque with a radio-opaque tip, and the position of the instruments, which have radio-opaque markers. Because of the difficulties in navigating and positioning the instruments at the stenosis, treatment procedures can be lengthy.

The first step is to navigate the interventional guidewire through the vessel-tree and position its tip past the stenosis. While navigating the guidewire to the lesion the physician can see, on one monitor, the selected static roadmap, and on the other monitor, the live interventional images. Sometimes a small amount of contrast dye is injected in order to correlate the current location with the local vessel-tree anatomy [32], or to determine whether

the vessel wall has been perforated. Once the guidewire is positioned past the stenosis, it is used as a rail to deliver the necessary instruments to the lesion site and proceed with the treatment. After the treatment is completed, the catheter and the instruments are withdrawn and the patient incision is closed.

Figure 2.8: Cathlab procedure (image courtesy of Philips Medical Systems [33]): on one monitor, the physician can see a static angiogram, with the coronary vessels filled with contrast agent. The live x-ray images acquired during the procedure are fed to the other monitor.

To reduce the amount of radiation absorbed by the patient and the physician, the Cathlab is used in *fluoroscopy* mode. The voltage settings in fluoroscopy are usually higher than in angiography [34] (with a starting value of approximately 95kV), and therefore harder x-rays are produced which are less absorbed by the human body. The frame rate used during fluoroscopy is also decreased, to lower the overall dose absorbed by the patient and the physician. The typical frame rate for adults used during x-ray fluoroscopy is 8.3 frames/s.

A higher voltage setting and a lower frame rate degrade the quality of the image, and increase the jerkiness of the sequence. Fortunately, the ability of the human eye to integrate between successive images (typical integration period is of about 0.2 ms [34]), decreases the perceived noise and diminishes the effects of the low frame rate. Therefore, during fluoroscopy, the Cathlab is set to provide the lowest input dose rate needed to generate a usable image.

2.3.6 Percutaneous Transluminal Coronary Angioplasty

Percutaneous Transluminal Coronary Angioplasty (PTCA) is a minimally invasive treatment that widens the lumen using an interventional catheter with a balloon on its tip. Once the balloon is positioned at the stenosis site, it is inflated in order to compress the plaque against the vessel wall (see figure 2.9). Unfortunately, about 35% of the patients develop *restenosis* 6 months after angioplasty [35]. Restenosis consists of a constriction of the lumen at the same site of a previously treated lesion. It is caused by the natural inflammation response triggered by the compression of the plaque against the vessel wall.

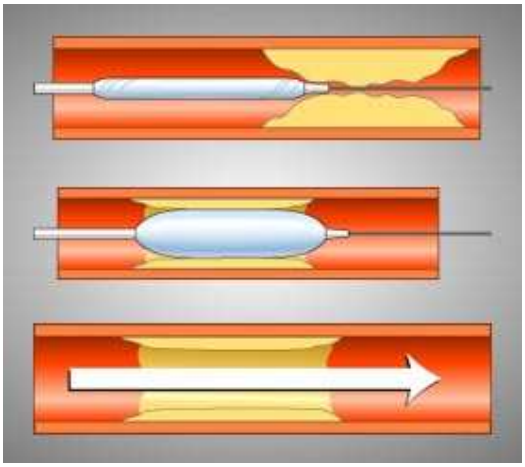

Figure 2.9: PTCA procedure: a balloon is advanced until it reaches the lesion site and is then inflated. (Image courtesy of Mid-Atlantic Surgical Associates [4]).

There are a number of instruments that can be used for CAD treatment that lower the restenosis ratio [18][35][36]. A *stent* is a metal mesh mounted on a balloon (figure 2.10a). Stents were introduced to reduce the restenosis rate of angioplasty. When the stent is navigated up to the lesion site, the balloon is inflated, the stent is deployed and its metal struts are (ideally) compressed against the vessel walls (figure 2.10b). The catheter, the balloon, and the guidewire are removed, leaving the stent bracing the vessel walls. *In-stent restenosis* ratios for stents are around 15 to 20% [11]. Important factors related to in-stent restenosis are accurate positioning, appropriate stent dimensions and complete stent deployment.

A novel approach to reduce restenosis is the use of *drug-eluting stents*[1], which are thinly coated with drugs that impede endothelial cell growth. Once the stent is in place, the drug is slowly administered directly on-site, lowering the restenosis rate from 20-35% to less than 10% [11]. Another approach to hinder cell proliferation at the lesion site after angioplasty or stent placement is *brachytherapy* [35][37][38] in which a low radioactive dose is delivered to the site using a special brachytherapy catheter, thus preventing tissue growth.

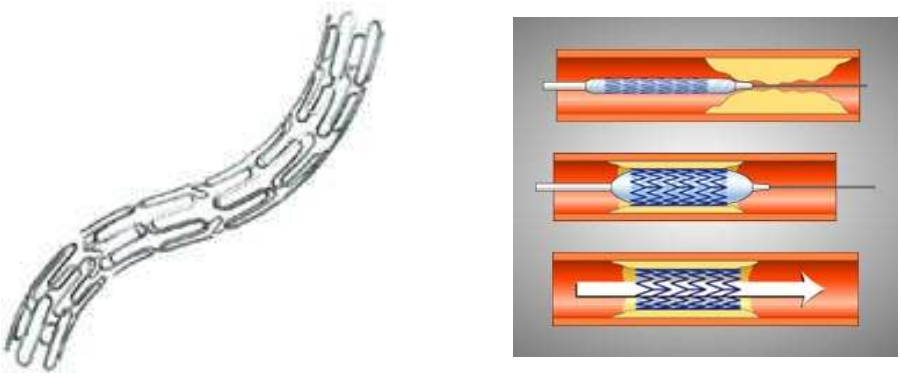

(a) Stent (b) Stenting procedure

Figure 2.10: (a) A stent is a metal mesh that will support the vessel walls. (Image courtesy of Hokkaido Ohno Hospital [39]). (b) A stent is mounted on a balloon and is deployed as the balloon is inflated.(Image courtesy of Mid-Atlantic Surgical Associates [4]).

An important requirement in the mentioned procedures is an accurate positioning of the instruments. The beneficial effect of drug-eluting stents or brachytherapy can be reached by delivering the drugs or radiation to the proper spot. Instrument positioning is currently guided by one static angiogram and live fluoroscopic frames which are displayed on two different monitors. The positioning of the guidewire and instrument as seen on the fluoroscopic frames are affected by patient respiration and heart contraction. It is therefore the task of the physician to visually process the images presented in the Cathlab monitors and correlate the position of the instrument in the fluoroscopic images with the static vessels visible in the angiogram. It is the objective of this thesis to present such image registration framework and processing algorithms to automatically relate instrument position and vessel

[1]For further information about drug-eluting stents, we refer the reader to http://www.tctmd.com

anatomy and thus improve the visual information presented to the physician.

2.4 Other imaging techniques

In this section we will briefly discuss non-invasive imaging techniques available for diagnosis of CAD which nevertheless have limitations in spatial, temporal or contrast resolution [29] when compared with x-ray. For CAD treatment x-ray fluoroscopy is still the chosen imaging technique. Moreover, because angiograms are previously acquired, these can be used during treatment to guide the procedure. Images acquired with other non-invasive imaging techniques serve exclusively as diagnostic tools. If the diagnosis is positive and it is decided to perform PTCA, the patient is then taken to the Cathlab room and angiograms are acquired in order to guide the subsequent treatment.

2.4.1 Computed tomography

Computed tomography imaging consist on the acquisition of a large series of two-dimensional x-ray images taken around a single axis of rotation to reconstruct an image of the traverse plane, orthogonal to the axis of rotation. When such reconstructions are taken from contiguous planes, a 3D reconstruction of the object is possible. Such 3D reconstructions are possible with spiral CT, in which the patient is slowly moved through the gantry while acquiring information. Voltage settings are set higher than those used for x-ray fluoroscopy such that the patient dose is lowered but image quality can be superior than x-ray because of the multiple acquisitions of images and their subsequent processing to yield one image.

Assuming that the size of the heart along the rotation axis is about 13cm and that an adult heart beats at 60 beat/min [29], the patient should be moved at a speed of 13 cm/s through the gantry to be able to image the heart in the time it takes to beat once. Because each slice would be acquired at a different time of the cardiac cycle, severe cardiac motion artifacts would appear on the reconstructed image.

CT is very sensitive at detecting high-density structures like calcium build-up inside the vessels [40], which is an indication of atherosclerosis, but it does not detect vulnerable plaque [12]. Studies have shown that there is

correlation between calcium score values and vulnerable plaque [12][29], although a low score in coronary calcium does not mean that vulnerable plaque is not present.

Multislice Spiral CT Multislice computed tomography (MSCT), also known as Multi-Detector Computed Tomography (MDCT), consists in multislice image acquisition of the body, with increased rotation time. Combined with ECG gating [29][41] temporal resolution is improved. The ECG gating technique has nevertheless a few shortcomings [29]:

- Good spatial resolution can be achieved only for regular and slow heart rates, but patients undergoing the examination typically have some type of arrythmia.
- It is of the utmost importance that the patient maintains a stable breathhold state for 20 to 30 seconds, in order not to displace the heart in the field of view. Failure to do this would result in respiration motion artifacts.
- The contrast to noise ratio of coronary arteries diminishes from proximal to distal end.

X-ray angiography provides good spatial and temporal resolution irrespective of patient heart rate or breath-holding capabilities. An irregular heart beat does not affect the quality of the images acquired. Furthermore, the angiogram sequence provides information as to how the catheter moves with the target, which can provide extra-information to the physician. Besides, if the patient is not able to maintain breath-hold, the only consequence is that the diaphragm and the tissue bellow it might obscure a few frames from the acquisition (depending on acquisition geometry). Angiograms are usually acquired with different C-arm geometries, with preference to the geometry settings that positions the detector parallel to the stenosis for optimum view of the contrasted lumen and length estimation.

MSCT can be used as a non-invasive method to measure coronary calcium score in patients, although it has been noted that reproducibility of calcium score values is still low when using MSCT [29]. Therefore an MSCT scan is not conclusive as to CAD and further examinations using angiography might be necessary.

Electron beam computed tomography Electron beam computed tomography (EBCT) is a form of computed tomography in which the x-ray source is not mechanically rotated. The electron beam current within the x-ray tube is electronically rotated around the object to be imaged therefore achieving faster image acquisition time than MSCT. Faster imaging times are beneficial for heart imaging. As MSCT, EBCT can detect calcium deposits in the vessels which are an indicator for CAD.

EBCT with ECG gating can be used for cardiac image acquisition for improved temporal resolution. But as MSCT, it performs best for regular heartbeats [29]. Another disadvantage of EBCT is that it is not possible to vary the scanning parameters to adapt the image quality for different patient typologies or coronary artery characteristics [29][42]. But probably the main drawback for this system is its very high cost which limits its availability to very few sites worldwide [29][43]. According to [42] EBCT scans are also not conclusive as to diagnosis of CAD and there seems to be a high variability as to the conclusions drawn from the scan results.

2.4.2 Magnetic Resonance Imaging

Magnetic Resonance Imaging (MRI) is different from the imaging techniques discussed so far because it does not use x-ray but magnetic fields for examination. A strong magnetic field is applied to alter the properties of the hydrogen nucleus (very abundant in water and fat). When the magnetic field is switched off the nuclei return to their original state radiating energy in the form of radiofrequency signals, which are detected and converted to gray scale values.

MRI can identify lipid-rich tissue like vulnerable plaque as described in section 2.2 although its capability for detecting coronary calcium is limited [12]. A major disadvantage of MRI for non-invasive characterization of coronary plaques is the length of time required to complete a full study, and that the contrast resolution is insufficient for characterization of CAD [12]. The injection of contrast agents can improve the contrast resolution of MRI images of coronary arteries and new contrast agents are being developed for applications such as functional/metabolic studies [12]. ECG gating is also necessary for MRI of cardiac arteries to avoid cardiac motion artefacts in the reconstructed images and therefore patient arrythmia causes motion artefacts in the resulting reconstructed images.

2.4.3 Ultrasound

UltraSound (US) is an imaging modality that uses the reflection of high-frequency pressure waves to construct an image of a body organ. An US probe emits pulses of high-frequency pressure waves and receives the returning echoes. When an US wave encounters an interface between different types of tissues, part of the wave's energy is reflected. This echo is received by the US probe, transformed into an electrical signal and processed to be displayed.

Intravascular Ultrasound *IntraVascular UltraSound* (IVUS) is an imaging modality that delivers 2D cross-sectional US tomographic images of the vessel from within the vessel (see figure 2.11). An IVUS transducer is introduced in the vascular system as any other interventional instrument and advanced into the coronary arteries until the lesion, where it emits pressure waves and detects the signal reflected from the vessel walls [18]. Calcified vessels, diffused plaque and vulnerable plaque cannot be easily detected in angiographic images [18], but appear clearly in IVUS images. If an IVUS device is inserted inside an artery and pulled back from the distal to the proximal end of the stenosis, it delivers cross-sectional views along the vessel, and a number of quantitative measurements of the lesion relevant for diagnosis can be performed (see appendix B). These include vessel volume, lesion extension, and plaque geometry. Clinical trials like CRUISE [44], AVID [45] and TULIP [46] support the utilisation of IVUS-guided stenting. To this effect we have found some patents of catheters that include an IVUS transducer in [47][48][49][50].

Table 2.1 summarises the main features of diagnostic x-ray and IVUS. The leftmost column lists a number of desirable imaging and navigation capabilities in current catheterisation interventions. The central and rightmost columns state whether angiography and IVUS, respectively, can provide said capabilities. Coronary angiography helps to diagnose and localise a stenosis and provides a suitable map of the vessel-tree topology that is afterwards utilised to navigate and position the instruments during treatment. Monoplane angiography provides a 2D projection of a 3D volume, which might hinder the visualisation of very eccentric and tortuous plaque. Furthermore, angiography does not provide information about the lumen area, plaque composition and morphology, the presence of calcification or the distribution of diffuse plaque, which are of diagnostic value.

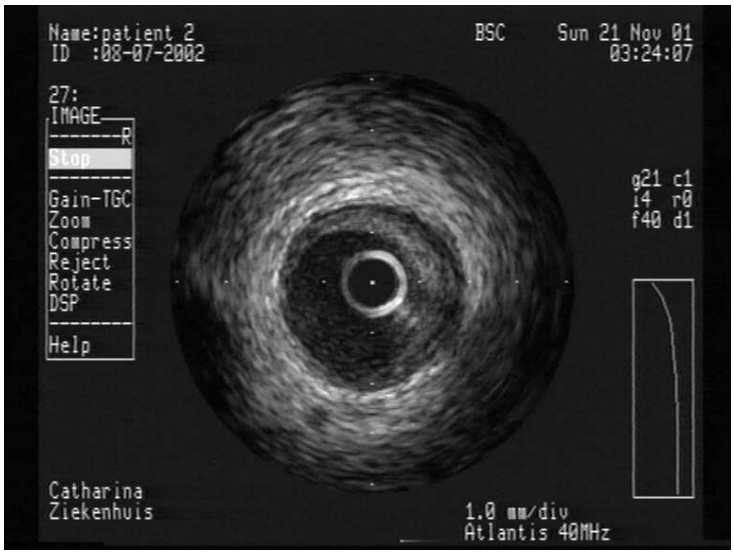

Figure 2.11: IVUS image as displayed on an IVUS system monitor. At the centre of the image is the catheter shadow. Different tissue types display a different echogenecity.

2.5 Summary

X-ray is the modality of choice for diagnosis and treatment of CAD *in vivo* because of its high spatial and temporal resolution. X-ray CAD procedures takes place using a Cathlab system which provides projection images of the heart. The Cathlab system is used in angiographic mode for acquiring diagnostic angiograms and in fluoroscopic mode for navigation and instrument positioning. It is known that angiograms provide an underestimated measure of lumen diameter [51]. IVUS provides information about lumen geometry, plaque morphology and can be used for vulnerable plaque assessment. Currently Cathlab and IVUS are independent systems.

Diagnostic angiography is performed during breath-hold in order to eliminate breathing motion from the sequence of acquired images, and to keep the diaphragm and underlaying tissue, which can obscure the view of coronaries, away from the field of view. A contrast dye is injected and a stenosis is recognised in the angiograms as a constriction of the visible lumen. The analysis of angiograms provides information about geometric characteristics of the stenosis: length, severity, position.

X-ray fluoroscopy is the only available imaging modality for intracoro-

X-ray vs. IVUS	Angiography	IVUS
Vessel topology	$\surd$	$\times$
3D capability	not accurate	needs x-ray
Lumen area	underestimated	$\surd$
Plaque morphology	$\times$	$\surd$
Diffuse plaque	difficult	$\surd$
Calcium	$\times$	$\surd$
Intervention outcome	inaccurate	$\surd$

Table 2.1: Comparison of IVUS and angiography. A tick ($\surd$) signifies the ability of supporting the desired capability, whereas a cross ($\times$) stands for the inability of the corresponding imaging modality of providing the desired capability.

nary treatment of CAD. During treatment, one angiogram is selected and displayed on one of the Cathlab monitors to serve as a roadmap during the intervention. The interventional fluoroscopic frames are displayed in real-time on a second monitor next to the roadmap (see figure 2.12). The current interventional technique relies solely on images acquired with the Cathlab system. It depends on the physician's ability to correlate in real-time the live fluoroscopic images, where the manipulated instrument is visible, to the static angiogram, where the coronary arteries filled with contrast dye are shown. For navigation and accurate instrument positioning it is common to inject some contrast dye to visualise the local vessel structure and relate the instrument's current position to the stenosis.

To maximise the effectiveness of minimally invasive intracoronary CAD treatment in the Cathlab, it is of utmost importance to properly position the instrument (stent, brachytherapy catheter or other) at the lesion site. The goal of the work presented in this thesis is twofold:

- to aid instrument navigation to minimise intervention time and therefore reduce the amount of x-ray dose received by patient and physician, and
- to aid accurate instrument positioning treatment to improve procedure outcome.

Navigation and accurate positioning is impeded by constant movement of the heart and patient respiration. This thesis proposes a new image registration framework and novel algorithms to compensate for such movement and present the physician with images in which an updated suitable angiography

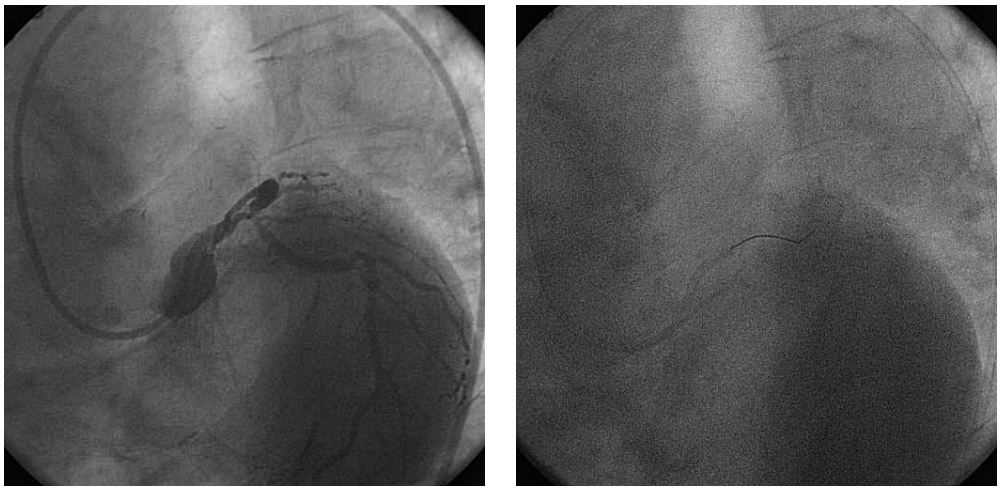

Figure 2.12: Typical Cathlab display: on the left-hand side, the angiography displaying the contrast-filled vessels; on the right-hand side, the live fluoroscopy frame, where the guidewire tip is visible. By looking at both, the physician estimates the position of the guidewire within the vessel tree.

is overlaid to the real-time fluoroscopy images, such that the lesion site can be accurately related to the instrument position and thus aid the navigation and positioning of said instrument.

IVUS acquisitions are currently used for lesion morphology and evaluation. IVUS can potentially improve instrument positioning during Cathlab interventions by providing images of the lesion site during instrument manipulation under x-ray fluoroscopy. A novel method is proposed to fuse angiography and fluoroscopy with IVUS images to aid diagnosis and accurate positioning of the instrument while minimally altering the current Cathlab workflow procedures.

Chapter 3

Diagnosis and treatment support

In the past chapter Cathlab systems and the workflow for intracoronary CAD treatment was presented. Coronary angiography is the preferred imaging modality for supporting minimally invasive interventions. During intracoronary treatment of CAD x-ray fluoroscopy is utilised. Currently the two principal causes for long interventional times are

- Misguidance of the instrument during navigation, which requires to pull back the instrument and attempt again to introduce it in the right coronary artery, and
- difficulty to accurately position the instrument.

Furthermore, during the intervention, additional contrast agent is usually required to clarify the actual position of the guidewire in overlapping or nearly parallel vessels as well as during the crossing of partial vessel occlusions and the positioning of interventional devices.

The main cause for the above mentioned issues is the fact that the angiographic roadmap the physician is looking at throughout the intervention is static, whereas in the fluoroscopic live feed the instrument is constantly moving because of manipulation, patient breathing and heart contraction. Consequently, a static angiogram provides a sub-optimal navigational and positioning aid to the physician. It is desirable to reduce intervention time in order to reduce the amount of the radiation the physician and the patient are exposed to, and to minimise the amount of contrast dye that is injected to the patient.

Currently most research and development effort lies on improving individual image quality by intensity- or feature-based image processing. In this thesis we present a novel approach that increases image quality by improving the diagnostic or interventional value of the displayed images based on anatomy and context information to fuse different sources of information (diagnostic angiogram, IVUS, interventional fluoroscopy). Our framework is based on a new algorithm that establishes whether two Cathlab images present the heart in the same state. This algorithm bases on ECG analysis and image processing, and is independent of the image modality of the images being compared. All the algorithms involved are real-time capable in order to be used during the intervention for navigation and positioning aid of guidewire and instruments. Section 3.1 comments current state of the art for navigation and positioning aids for Cathlab systems. Section 3.2 presents the new applications proposed in this thesis and section 3.3 gives an overview of the algorithms needed to realise the proposed features. These new algorithms will be presented in the following chapters.

3.1 State of the art

Alternative CAD diagnosis methods were presented in the past chapter. During treatment x-ray remains the imaging modality of choice and the Cathlab the preferred system for minimally invasive interventions. The current *Smart Mask* (SM) feature in Philips Cathlab systems (called *Fluoro Fade* in Siemens' systems) is devised for peripheral angiography and interventions. On one monitor, the live interventional images are displayed. On a second monitor, an angiogram or *roadmap* chosen by the physician is permanently overlaid to the real-time fluoroscopic frames. If applied during cardiac interventions, cardiac motion artifacts appear on the resulting overlay image because of the influence of breathing and heart contraction [52] (see figure 3.1). Therefore, the difference in position, orientation and shape of the heart between the overlaid roadmap and the live fluoroscopic frames prevents an estimation of the current guidewire or instrument position related to the vessel-tree or the stenosis location.

An IVUS probe can acquire cross-sectional images of the artery at a given location with no relation to vessel topology. If an IVUS probe is pulled back from the distal to the proximal end of a stenosis, IVUS images along the pullback path are acquired. Commercially available IVUS systems stack the IVUS images along a straight axis (figure 6.1) assuming that the catheter

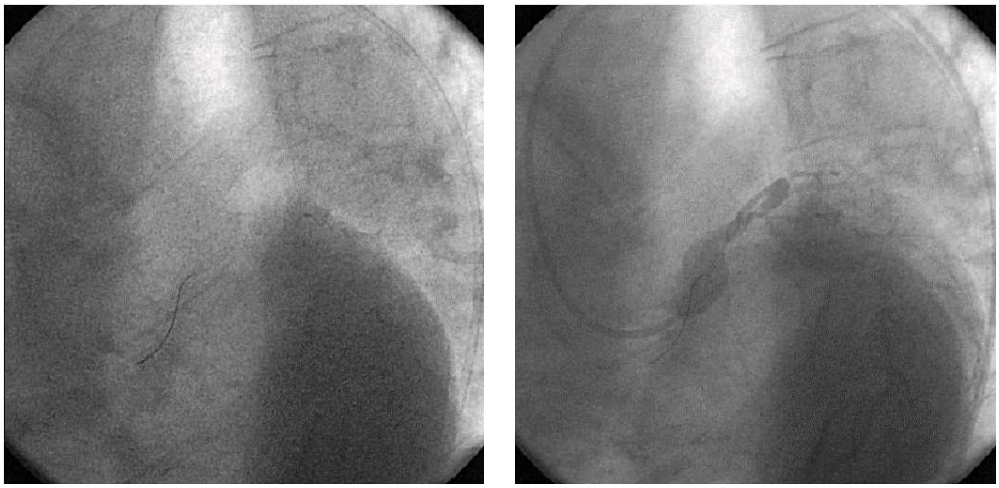

(a) (b)

Figure 3.1: Non-ideal overlay: (a) a fluoroscopic frame, where the guidewire tip is visible; (b) result of overlaying an angiographic mask to the frame in (a). In (b), the position of the instrument is unrelated to the vessel-tree.

follows a straight path within the vessel, which is far from reality.

There are a number of publications dealing with 3D vessel reconstruction from biplane angiography and IVUS for diagnostic purposes [53][54][55][56][57] [58][59] but there are currently no such systems available for clinical use. Furthermore, the proposed application utilises biplane Cathlab systems, which are not common in Cathlab interventions for their high costs [12]. None of the previous literature deals with x-ray and IVUS fusion for accurate instrument positioning.

3.2 Proposed clinical features for Cathlab systems

In the following two subsections the new Cathlab applications proposed in this thesis are presented. These consist of real-time anatomical registration of angiograms to live fluoroscopic images, and x-ray angiography/fluoroscopy to IVUS fusion. The novel algorithms applied to achieve the proposed new Cathlab features will be commented in the next section and will be presented

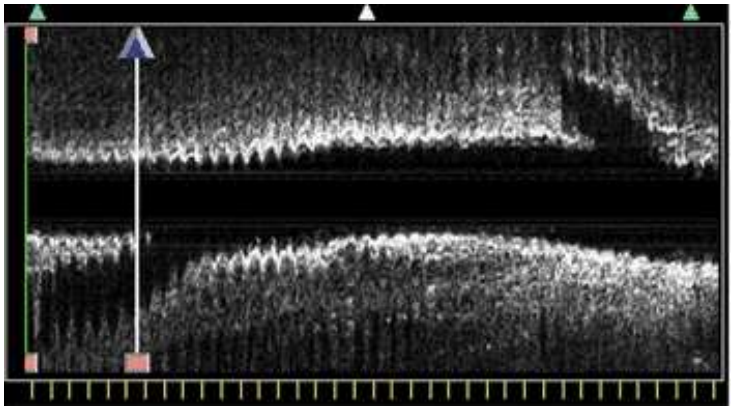

Figure 3.2: IVUS vessel reconstruction in JOMED's *In-Vision Gold Intravascular Ultrasound System.* [5].

in detail throughout this thesis.

3.2.1 Navigational and navigation support by angiographic mask overlay

As mentioned earlier, in present systems a monitor presents a static angiogram to the physician while another monitor displays live fluoroscopic frames. It is possible to overlay the static angiogram onto the fluoroscopic frames and present such overlay to the physician.

In this thesis a novel method is proposed that provides updated angiographic overlays for the interventional x-ray sequence to aid guidewire navigation and instrument positioning. Overlaying an angiogram to the fluoroscopy frame will result in an image that facilitates relating the instrument position, which is visible in the fluoroscopy frame, to the lesion site, which is visible in the angiogram. In figure 3.3 we can see an example of such an updated overlay. Figure 3.3a shows an interventional frame acquired during guidewire navigation. Figure 3.3b shows an overlay between an angiogram appropriately chosen and the fluoroscopy in figure 3.3a. We can see how the guidewire appears inside the overlaid contrasted coronary artery. Providing the physician with such images during the intervention can reduce interventional time, thus reducing the fluoroscopic x-ray dose, and the amount of injected contrast agent .

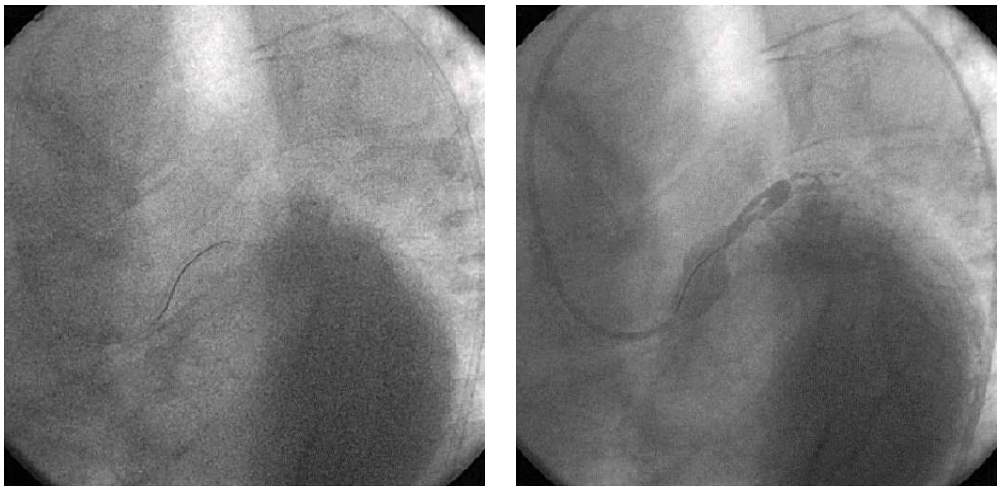

(a) (b)

Figure 3.3: Ideal overlay: (a) a fluoroscopic frame, where the guidewire tip is visible; (b) result after overlaying a carefully selected angiogram to the image in (a). In (b), the instrument appears inside a contrasted vessel, showing its position with respect to the coronary tree.

3.2.2 Diagnosis and positioning support through x-ray and IVUS fusion

Diagnosis can be based on the examination of angiograms that display the vessels filled with a radio-opaque contrast agent. These images offer information about the residual lumen at the site of the stenosis, but not about the plaque properties. On the other hand IVUS provides additional information about plaque structure and composition, more accurate information about lumen diameter and stenosis length (see section 2.4). Current IVUS systems are used solely for diagnosis and independently of Cathlab systems. Previous literature proposes to fuse biplane angiography with IVUS [56][57][60][61][62][63][64][65][66][67][68].

In this thesis it is proposed to integrate IVUS imaging into the procedure by fusing IVUS frames acquired during an IVUS probe pullback along the affected coronary artery with:

- diagnostic angiograms, for vessel reconstruction to aid diagnosis by providing information as to plaque structure and stenosis geometry, and

- interventional fluoroscopy, to aid accurate positioning by providing IVUS images from the current instrument location, provided it has been navigated to the correct coronary artery.

An integration of IVUS in the Cathlab can help the physician to reach a more accurate diagnosis, place more precisely the instrument or check the outcome of the intervention. This synergy has the ability of supporting the intervention process with reduced radiation dose and improved outcome.

3.3 Thesis contribution

In the previous sections we have exposed the proposed new applications for aiding navigation and accurate positioning of instruments during PTCA interventions and integrating IVUS into Cathlab procedures. The novel algorithms developed to achieve the desired results are briefly introduced in the following paragraphs and will be explained in the following chapters.

The objective of this thesis is to present algorithms for anatomical-based image fusion that do not base on image features but on the state of the imaged anatomical structures. A similarity in state results in valid image fusion, which provides the physician with valuable information.

For cardiac imaging, we propose to characterise each image by the heart state, which is influenced by:

- heart contraction state, which affects the shape of the heart;
- respiration state, which influences the position of the heart within the thorax;
- table panning, which causes an overall translation of the imaged objects;
- patient movement, which induces an affine transformation of the imaged objects;
- imaging geometry, which affects the perspective from which the objects are imaged.

The heart state is thus defined as a *state vector* $\vec{\xi}$ with the above mentioned components, which affect the shape and position of the heart within the thoracic cavity and therefore in the acquired images. Images acquired

with different imaging modalities that have the same $\vec{\xi}$ are suited for overlay and anatomical multi-modality registration. Specifically, an angiogram overlaid onto a fluoroscopic frame with the same $\vec{\xi}$ would display the contrasted vessels at a position and contraction state that coincides with the instrument's.

Furthermore, due to the discrete nature of image acquisition and the constant movement of the imaged body structures it occurs that two images with the same $\vec{\xi}$ are not always available for registration nor overlay. To improve image display and enhance the information presented to the physician, methods to compensate for a difference in state vector $\Delta\vec{\xi}$ are needed. These include:

- Heart contraction compensation: In cardiac imaging, a difference in heart contraction translates into a difference in shape of the coronary arteries that can be modelled by a complex affine transformation. In projection imaging, it is not possible to compute the parameters to calculate such transformation.

- Respiration compensation: Respiration induces a translation of the heart, which can be modelled as a planar shift in projection imaging [69].

- Table panning compensation: Table panning causes a planar shift in the imaging plane. Compensation strategies are landmark identification and registration. In cardiac imaging such landmarks might be wire stitches, the spine, ribs or other structures known to be static.

- Patient movement compensation: Involuntary patient and organ movement is difficult to accurately model. It can be assumed that in projection images it results in a planar shift that can be aggregated into table panning.

- Imaging geometry compensation: In 3D imaging, where a 3D model of the imaged structure can be reconstructed, a change in geometry settings like rotation or angulation of the imaging plane can be compensated for by 3D registration techniques. In 2D projection imaging, a change in imaging geometry would cause too big a change in image content to be compensated for.

As mentioned in section 3.2, some applications of the proposed technique are:

- Overlay of angiograms onto interventional fluoroscopic frames for guidewire navigation and accurate instrument positioning. Such a method yields images in which the contrasted vessels acquired during a previous angiographic acquisition are presented overlaid onto the live interventional frames.

- Overlay of angiograms and IVUS pullback frames for enhanced diagnosis. Such a method enables combined reconstruction of vessel lumen and plaque characteristics with diagnostic angiograms.

- Overlay of interventional fluoroscopic frames and IVUS pullback frames for accurate positioning and outcome result evaluation. Such a method shows IVUS data for a position of interest in a vessel, which is the actual position of the interventional device in each incoming fluoroscopic frame.

In this thesis we present a novel image registration framework that enables improved diagnosis and instrument navigation and positioning for PTCA interventions, thus reducing the radiation dose and intervention time. Methods to calculate $\vec{\xi}$ and compensate for $\Delta\vec{\xi}$ in the context of PTCA interventions using Cathlab projection images are presented. It is assumed that the images to be registered are acquired with the same imaging geometry. Novel algorithms to characterise each image by heart contraction and respiration state will be presented in section 3.3 and explained in chapter 4. Solutions for calculating the heart contraction state are presented in [20], [70], [71] and [72]. Based on previous work we propose to characterise each image by a *heart phase* which is calculated from the ECG signal available in Cathlab systems. To characterise the respiration state we propose to calculate a *respiration phase*. Previous work relies on external respiration measurements [70] and biplane diaphragm tracking [71][72]. Neither meets our requirements: projection image-based and independent of image content. A novel imaging processing method based on multi-modality registration is presented in chapter 4. Heart contraction compensation in projection angiography and fluoroscopy by affine transformation modelling cannot be performed. Instead we choose to define a reasonable range for heart contraction matching as will be explained in the next section. Methods for respiration compensation by respiration modelling and novel anatomical vessel to guidewire/instrument registration are explained in chapter 5. All these methods are able to run in real-time so that the search for an angiogram with the same state vector as the current live fluoroscopic frame can be found.

3.4 Heart state

As stated above, a method to establish whether the heart in two images is in the same state has been developed in this thesis. By *heart state* it is understood the position and shape of the heart within the image characterised by the feature vector $\vec{\xi}$. The n features identified to affect heart state are heart contraction, breathing, table motion, patient motion and C-arm geometry, which define an n-dimensional space. If an angiogram characterised by heart state $\vec{\xi}_a$ is overlaid onto a fluoroscopy frame characterised by the same heart state vector we would get a result similar to figure 3.3b.

Heart contraction induces a complex affine transformation of the coronary arteries, as can be seen in figure 3.4. Figures 3.4a and 3.4b display the difference image between two angiographic frames, $I_{\vec{\xi}_i}(x, y)$ and $I_{\vec{\xi}_j}(x, y)$, characterised by state vectors $\vec{\xi}_i$ and $\vec{\xi}_j$ which differ on the heart contraction state. In the difference image, we can observe how the background approximately cancels because all other components of the state vectors are roughly equal and the main difference is caused by the heart contraction, which affects the shape of the coronary arteries.

Respiration involves a cyclic motion of the diaphragm, the lungs and the ribs. The contraction and relaxation of the diaphragm shortens or lengthens the chest cavity with a corresponding change in volume. Elevation and depression of the ribs increase or decrease the diameter of the lungs, and the lung volume changes as they fill with air. Figures 3.5a and 3.5b display the difference image between two frames which state vectors differ in the breathing state. It can be noticed that a difference in respiration state causes a global translation of the coronary tree. This observation corroborates the conclusions of Wang *et al.* [69]. Figure 3.5a and 3.5b also suggest that if the respiration state in the images were the same, the position and the shape of the vessels would coincide. A simple experiment is to calculate in said figures the global translation that minimises the mean squared error of the gray levels. In figure 3.5a the mean squared error between the gray levels of both images, excluding the borders is $mse_{orig} = 2.1e - 3$ whereas a global translation in the vertical direction of the subtracted image expressed by the vector $\vec{t} = 10\vec{j}$, where $\vec{j}$ expresses the vertical direction, reduces the MSE to $mse_t = 0.4e - 3$. In this example, both images have similar image content, that is, both display vessels filled with radio-opaque contrast agent. For the example in figure 3.5b, the calculated MSE before applying any global translation is $mse_{orig} = 3.9e-3$ which is reduced to $mse_t = 1.3e-3$ after a global

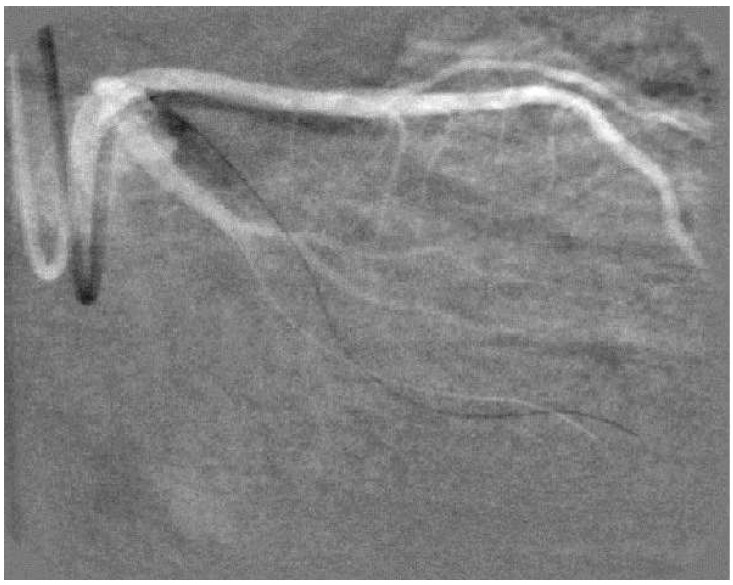

(a)

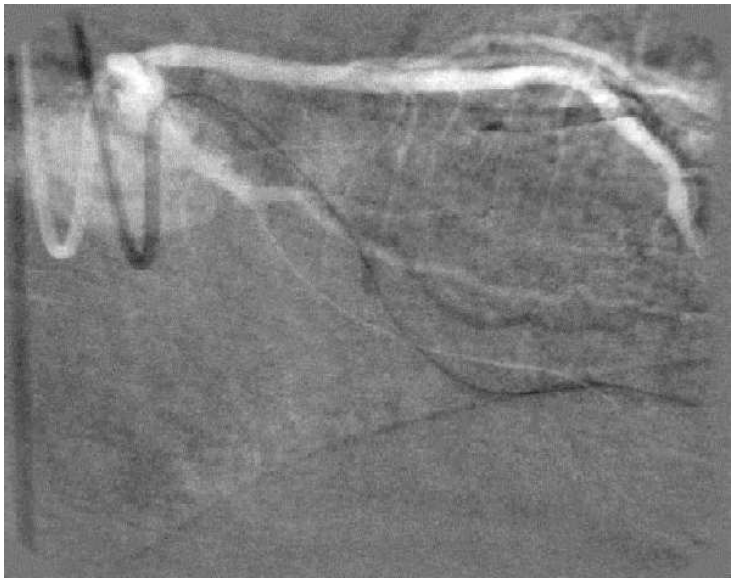

(b)

Figure 3.4: Influence of the heart contraction on the coronary vessels. (a) and (b) show the difference image between two images that have a similar respiration state but different heart phase. In (a) the two images have exactly the same respiration phase and we can see how the background disappears. In (b) the diaphragm in the two images do not match exactly, but are very close.

translation of $\vec{t} = -22\vec{i} + 12\vec{j}$, where $\vec{i}$ expresses the horizontal direction.

A change in acquisition geometry between angiography and fluoroscopy alters the perspective point from which the heart is imaged and therefore changes significantly the image content. Once a C-arm geometry has been chosen for optimum viewing of the stenosis (usually the imaging geometry is chosen such that the stenosis is parallel to the imaging plane), table panning to bring the lesion to the centre of the field of view may occur. Patients are normally conscious during the intervention and may move. This movement can be simplified as being within the table plane and can be combined with the table panning if present.

If for two x-ray acquisitions the system geometry is the same, there was not any table or patient movement then the state vector is reduced to a two dimensional vector with the following components: heart contraction state and heart displacement caused by breathing. We propose to represent each image by its heart and respiration phase as a point in a 2D graph, the *phase-space*. The search for a suitable angiogram to be fused with a fluoroscopic frame reduces to the 2D search of the closest point in the phase-space according to a distance measure. This search method is further explained in the next subsection.

3.4.1 Phase-space graph

Let's assume the same acquisition geometry for the diagnostic and interventional image acquisitions and no table nor patient movement. If these conditions are met $\vec{\xi}$ defines a two-dimensional space or *phase-space* [73][74] where each point is defined by a *heart* and *respiration phase* coordinates, δ and β respectively. Any two images acquired at similar respiration and heart contraction states present the heart in a similar position within the image and with the same contraction state [70][73][71]. When a diagnostic angiogram and an interventional frame acquired at similar respiration states and with the heart in the same contraction state are overlaid, the instrument can be related to the coronary vessel structure. Figure 3.6 shows a difference image between a fluoroscopic image and an angiographic image. The vessels filled with contrast agent are visible in white, and a dark guidewire is also visible. We can see how in this difference image the guidewire shows inside a corresponding coronary artery.

Thus, each acquired frame $I_{\delta,\beta}(x, y)$ can be represented by its heart and

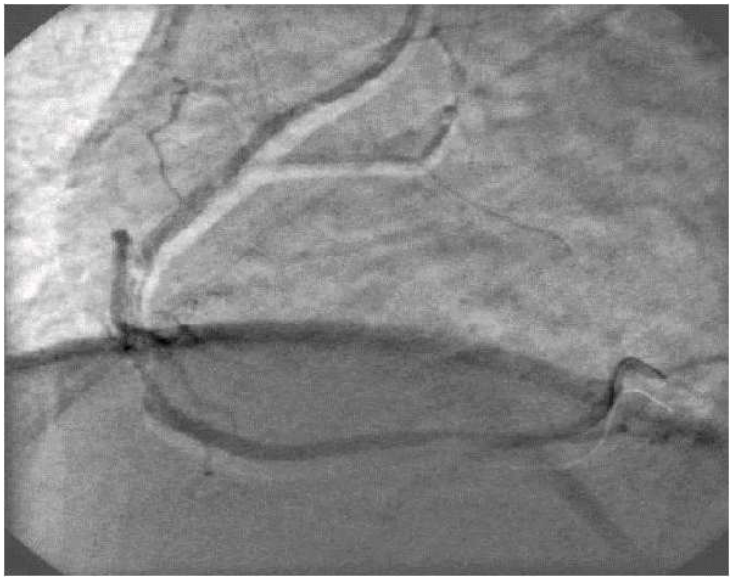

(a)

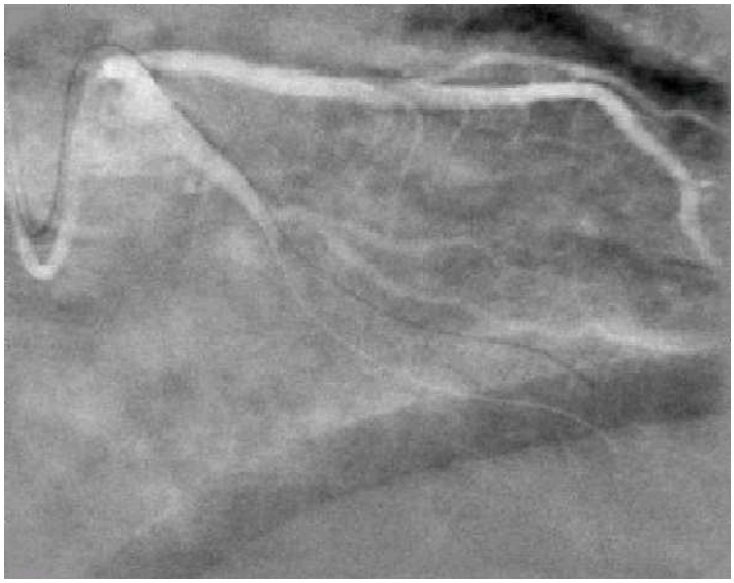

(b)

Figure 3.5: Influence of respiration on the coronary vessels. (a) and (b) show the difference image between two frames that have a similar heart phase but different respiration phase. We can appreciate how respiration induces a translation of the whole coronary tree.

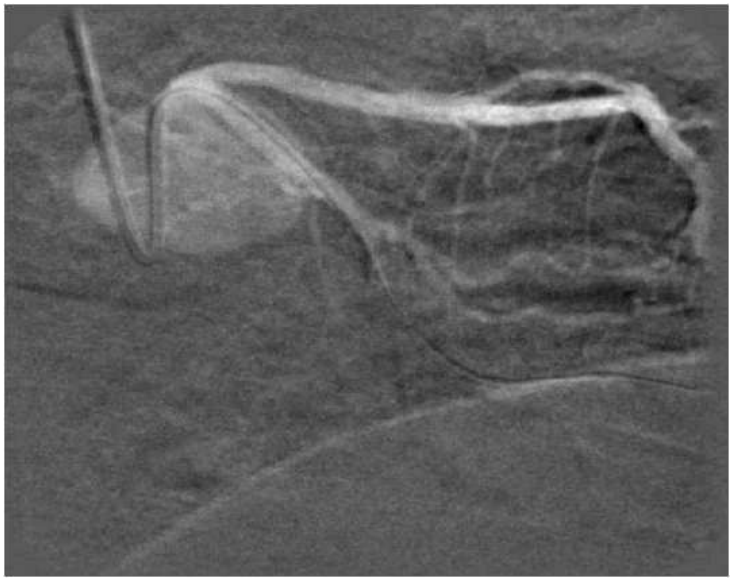

Figure 3.6: Difference image between two images with a similar heart and respiration state.

respiration phase as a point in a two-dimensional graph or phase-space [73][74]. In figure 3.7 we show an example of such a phase space, where the heart phase is calculated from the ECG signal (see section 4.1) and the respiration phase by measuring the similarity between each frame to a reference, as will be explained in section 4.5.

Because breathing and heart contraction are periodic activities, the possible values of heart and respiration phase are limited to the interval $[-1, 1]$. In a phase-space graph, nearby points represent images which show the heart in a similar state. In figure 3.7 we have connected consecutive frames by solid to highlight the periodicity of breathing and heart contraction. We can observe how the heart phase axis is swept from left to right with a periodicity of about 10 frames in this concrete example, which corresponds to a heart rate of 75 beats per minute. The lines also show the respiration phase evolution within one heart contraction period. We can notice that in this particular example, the breathing period is approximately double the heart contraction period.

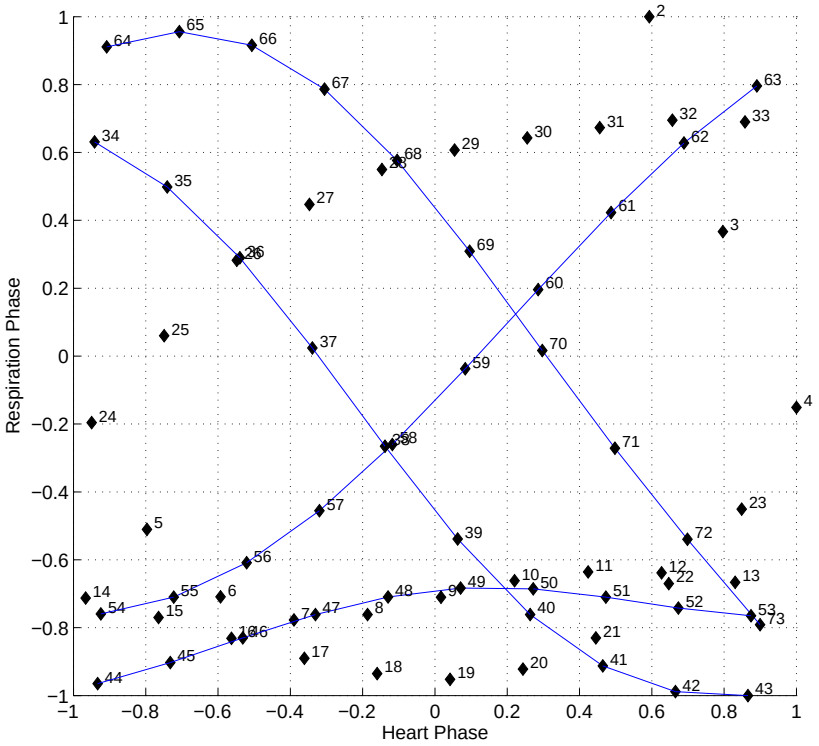

Figure 3.7: Phase-Space example. The horizontal axis represents the cyclic ECG phase (normalised by π) and the vertical axis the cyclic respiration phase (normalised by π respectively). Each point represents an image, and the number next to it is the frame number within the sequence. The lines connect consecutive frames. For the shake of clarity, not all the frames have been connected.

3.4.2 Diagnostic roadmap search

For each frame acquired during diagnosis or intervention and assuming that all other components of $\vec{\xi}$ are equal, our algorithm calculates a heart phase δ and a respiration phase β in real time so that each frame can be represented as a point on the phase-space.

The sampling of the phase-space for typical angiographic sequences is quite sparse. Therefore, it is uncommon that a perfect matching angiogram $A_{\delta_a,\beta_a}(x,y)$ will be available for each interventional frame $F_{\delta_f,\beta_f}(x,y)$. To overcome this problem, we propose in the following chapters different strategies:

1. Not to look for equality between $\vec{\xi_a}$ and $\vec{\xi_f}$, but define a wider search space based on anatomical characteristics of the heart;

2. Provide a stroboscopic update of the overlaid angiogram with subsequent fading;

3. Compensate for differences in respiration phase $\Delta\beta$;

4. Perform guidewire to vessel registration based on anatomical features.

In this subsection we will discuss the first two points. Respiration compensation and guidewire to vessel anatomical registration will be presented in chapter 5.

A matching roadmap is defined as the angiogram that presents the coronary vessels in a state as similar as possible to the state in a given fluoroscopic frame. This similarity is measured in terms of respiration and heart phase. In order to search for a matching angiographic roadmap, we search in the phase-space for the closest point, with the distance defined as:

$$D = \sqrt{\alpha_h(\delta_a - \delta_f)^2 + \alpha_r(\beta_a - \beta_f)^2} \tag{3.1}$$

where

1. D is the *distance* between points images represented on the phase-space,

2. δ is the heart phase of an angiogram δ_a or fluoroscopic frame δ_f,

3. β is the respiration phase of an angiogram β_a or fluoroscopic frame β_f,

4. α_r is a weighting factor for the Cartesian distance in respiration phase, and

5. α_h is a weighting factor for the Cartesian distance in heart phase.

α_r and α_h in equation 3.1 reflect the fact that the distance along the respiration phase axis and the distance along the heart phase axis have different clinical relevance for the visualisation of overlay images. Two frames with very similar respiration phases but somewhat different heart phases will exhibit a different shape of the coronary arteries. If the difference in heart phase is big, the quality of the overlay will be greatly affected, since the affine transformation from one heart phase to the other involves rotation and deformation and cannot be resolved visually by the physician. On the other hand, frames with similar heart phase but different respiration phases

$\{(A_{\delta_a,\beta_a}(x,y), F_{\delta_f,\beta_f}(x,y)), \delta_a \simeq \delta_f\}$ will present the same heart contraction but a different translation of the heart within the thoracic cavity. It was our observation and used as hypothesis that a similarity in heart phase was of more importance for an enhanced perception of the overlay, since the physician can visually compensate for the respiration phase by a translation of one frame with respect to the other. A difference in respiration can also be compensated for by computation [73] as will be explained in section 5.1. The maximum allowed distance is fixed *a priori*, to prevent the overlay of angiograms that present the heart in a very different state from the interventional frame. If no suitable angiogram is found in this reduced search space, no roadmap overlay is available for the current interventional frame.

Qualitative subjective tests were conducted to establish whether a stroboscopic update of the overlaid roadmap or a fade-out effect should be used in the interventional room. When using stroboscopic update, the overlaid roadmap is changed only when a new matching mask is available. This yields a very jerky effect to the displayed overlay sequence. Therefore, the fade-out effect was chosen. Whenever a suitable mask is found, it is overlaid onto the current fluoroscopic frame. As time elapses, the overlaid mask is faded until a new matching mask can be provided. This overlay procedure enhances the perception of the physician during the intervention.

3.5 Related work

Solutions have been presented in previous literature to compute δ and β [70][71][72]. Neither of the presented methods for respiration phase measurement is suitable for our requirements. These methods either lack real-time capability or are not suited for Cathlab interventions because they require additional machinery that might not be available or degrades the interventional images.

Close *et al.* presented in [75] a stabilised display for cardiac x-ray to enhance the visibility of a selected feature. By manually or semiautomatically selecting an area of interest in an angiography sequence, and augmenting the frame rate of the display, they obtain a stabilised display where the feature of interest is always presented on the middle of the monitor. This method enhances visualisation of lesions but does not address the problem of aiding navigation and positioning during the intervention. It also requires the presence of an operator in the control room to perform the ROI selection after

the sequence has been acquired. Therefore, this method is not suited for real-time applications.

Ro *et al.* proposed in [70] an algorithm to avoid the effects of cardiac and respiratory motion in coronary subtraction angiography by synthesising angiographic masks. They introduce the concept of characterising each image in an acquired sequence with two variables representing the cardiac and the respiratory state. Contraction and respiration are periodic processes and each image can be assigned a phase representing its position inside these periods. The heart phase is calculated from the ECG and an impedance respiratory monitoring instrument measures the breathing state. With these two variables they construct a two-dimensional *scatter graph*, where each point represents an acquired image. They propose a functional decomposition of motion to compute a subtraction image $S_a(\delta_a, \beta_a)$ (where δ and β represent heart and respiration phases respectively), from the acquired masks $A_b(\delta_b, \beta_a)$, $A_c(\delta_a, \beta_c)$ and $A_d(\delta_b, \beta_c)$ and a live image $I_{live}(\delta_a, \beta_a)$. Using an external device to measure respiration phase is not suited for minimally invasive procedures because the external device would appear in the field of view on the fluoroscopic images thereby blocking the view of certain structures in the field of view. Furthermore, Ro's approach is restricted to a 2-dimensional feature vector to represent the heart-state, whereas during an intervention there are several other factors that influence the heart state and should be taken into account, like table panning, patient movement, etc. The image registration framework presented in this thesis allows for the definition of the a n-dimensional feature vector $\vec{\xi}$ to represent the heart state. Moreover, our objective is not to synthesise angiographic masks to overlay to live fluoroscopic frames during instrument positioning, but to manipulate the acquired angiograms in order to better display and match each acquired fluoroscopic image.

Shechter analysed in [71][72] the respiratory motion of the heart in order to develop strategies for motion correction in MR imaging. The heart state is characterised like in [70] by a heart and a respiration phase, which are represented graphically in a *cardiac respiratory phase plane* (CRPP). The heart phase is computed from the ECG, and the respiratory phase is calculated by tracking a profile through the diaphragm in biplane coronary angiographies. Each point in the CRPP represents a previously acquired and semiautomatically segmented 3D coronary tree. The author proposes a *cardiac and respiratory parametric model* (CRPM) that parameterises the deformation field and decomposes it into separate cardiac and respiratory components. This method of respiration phase calculation is not suited for our purposes

since it is not real-time capable and during a catheterisation intervention it is uncommon to have a biplane Cathlab system.

Neither the methods presented by Ro *et al.* in [70] nor Shechter *et al.* in [71][72] is suited for our purposes. A new method of respiration phase calculation based solely on the acquired images is proposed in this thesis.

3.6 Conclusion

This thesis proposes anatomical-based image fusion applied to catheterisation procedures to support diagnosis and aid navigation and positioning of interventional instruments. The imaged structure, the heart, is characterised by a heart state vector. Heart state is relevant to estimate the position and shape of the heart within the thoracic cavity and therefore to decide whether a diagnostic angiogram can be overlaid to a fluoroscopic frame in order to render an overlay image that is useful for instrument navigation and positioning. It has been shown that heart phase and respiration phase are the major factors to take into account. Each frame from an x-ray sequence is characterised by its heart and respiration phase $I_{\delta,\beta}(x, y)$, and a two-dimensional phase-space is constructed where each frame is represented by a point with coordinates (δ, β). A distance measure is defined to evaluate whether any two frames, and in particular an angiogram $A_{\delta_a,\beta_a}(x, y)$ and a fluoroscopic frame $F_{\delta_f,\beta_f}(x, y)$ present the heart in a similar state. Methods are required that allow heart and respiration phase calculation and compensation in real-time with clinically relevant frame rate [76].

Our approach has in common with Ro *et al.* [70] and Shechter *et al.* [71][72] that we construct a two-dimensional graph where each point represents an image with the heart in a certain contraction and respiration phase. Our goal is not to generate artificial masks like in [70] or to model the heart motion like in [71][72], but to register x-ray images acquired at different points in time during the intervention and with different x-ray dose in order to aid navigation and accurate positioning. Furthermore, our framework bases on constructing a feature space based on anatomical features of the heart, and can be applied to other imaging modalities.

Chapter 4

Heart and respiration phase

In the previous chapter the concept of a two-dimensional *phase-space* as a representation of all images in an acquired x-ray sequence in terms of the heart and respiration phase was introduced. In order to find a suitable angiographic mask for each interventional frame, we characterise each image by a heart and respiration phase. Previous literature [20][70][71][72] presents methods for heart phase calculation of angiographic acquisitions based on ECG analysis and respiration motion detection based on external measurements [70] or 3D acquisitions [71][72].

Previously presented heart phase algorithms do not deal with real-time heart phase calculation. We here present a heart phase definition based on ECG analysis that can be calculated off-line for an angiographic sequence as well as in real-time for each incoming fluoroscopic frame. Previous work on respiration phase depends on external measurement devices that might interfere with the imaged area, or depend on the presence of specific image features. We present a novel method for respiration phase calculation from the image content [74]. We hypothesise and prove that the field of view contains various anatomical structures imaged in projection that are affected by respiration and heart beat. The main goal is to identify and compare regions of interest in two images where region similarity represents respiration similarity, with as little influence from heart beat and contrast injection as possible. Unlike other approaches we propose a respiration phase measure based solely on the image content and that does not depend on the presence of the diaphragm in the image nor relies on the presence of additional measure systems. The proposed measure is also capable of comparing images with different image content, i.e. angiograms where the vessel tree is visible and interventional frames where only the interventional device can be discerned.

One angiogram acquired during end-inspiration is chosen as the respiration reference. All other acquired images are compared to the respiration reference angiogram and the similarity is used as a measure of respiration phase. When comparing diagnostic and interventional x-ray frames to a reference, we have to consider that they are acquired with different x-ray doses (angiography and fluoroscopy). This difference in x-ray dose translates into an non-linear relation between the gray-levels of the same structure in the images for each modality. Therefore x-ray fluoroscopy and angiography can be regarded as different imaging modalities and the image content comparison is related to multimodality image registration. *Mutual information* (MI) was first proposed as a measure of multimodality image alignment in [77]. We here propose mutual information in selected regions of interest as a similarity measure between different angiographic and fluoroscopic frames and a fixed reference in order to compute a respiration phase [74]. We expect the result of our similarity measure to be periodic and to have frequency components related to breathing and heart contraction. The influence of the heart movement is removed by frequency analysis by a Hamming-based FIR filter and the remaining signal is used as respiration phase [74].

In section 4.1 we explain the calculation of the heart phase based on the ECG signal. In section 4.2 we comment previous work on respiration detection. In sections 4.3 and 4.4 we discuss our assumptions and the feasibility of the calculation of a respiration phase measure from the image content. In section 4.5 we propose and evaluate the performance of MI for respiration phase calculation. In section 4.6 we compare the proposed MI respiration phase measure with other similarity measures used for multimodality registration [78]. Practical implementation issues for clinical applications will be commented in chapter 7.

4.1 Heart Phase

During contraction, the heart experiences a series of complex elastic deformations that affect its shape and orientation. The contraction state of the heart affects the configuration of its walls and, therefore, the shape of the coronary vessels. Heart-phase calculation from the image content is not feasible because the coronary arteries are not visible during fluoroscopy. *ECG-gating* is a common approach to avoid movement artefacts caused by heart contraction

during cardiac imaging [56][70][71][79][80][81]. Image acquisition is performed only at a certain heart phase, which is usually end-diastole. Contrary to ECG-gating techniques, image acquisition in the Cathlab is continuous and for phase-space population and search we need to assign to each angiographic and fluoroscopic frame a heart phase. Since it is standard practice in PTCA interventions to record the ECG signal of the patient during the intervention, it is natural that this signal is used and a heart phase derived from it. A number of electrodes positioned on the patient's chest and limbs detect the changes in electrical potential caused by the excitation of the heart muscle. An ECG signal is represented in figure 4.1a, where a number of landmarks can be identified within each period. These landmarks are called the P, Q, R, S and T waves as marked on figure 4.1b. The P wave indicates atrial contraction, whereas the QRS complex signals ventricular contraction and the T wave indicates ventricular relaxation. Variations in size and duration of these waves are an indication of abnormal heart rhythm.

We assume that similar contraction states correspond to a similar normalised distance from the R peaks that mark the boundaries of a contraction period. An R peak is detected as the maximum within a running square window with length of a configured minimum ECG period. We calculate the heart phase δ as:

$$\delta(t) = 2\pi \frac{t - t_{j-1}}{t_j - t_{j-1}} \tag{4.1}$$

where

- t represents discrete time,
- t_{j-1} is the time at which the previous R peak occurred,
- t_j is the time at which the next R peak will occur.

Figure 4.2 shows the ECG from one of the DICOM files from our database. Even though the ECGs from our database have been stored sometimes with as little as three bits precision instead of the two bytes available for each sample, the R peaks are still clearly distinguishable and therefore the heart phase can be calculated as defined in equation 4.1. During the intervention, $t_j - t_{j-1}$ for each frame is estimated as a moving average of the pulse rate unless a new R peak occurred during the current frame acquisition, in which case, t_j can be directly known and the estimated pulse rate is updated. To assign a heart phase to each frame, we take into account that the ECG and the images are acquired with different sampling rates (in figure 4.3 the

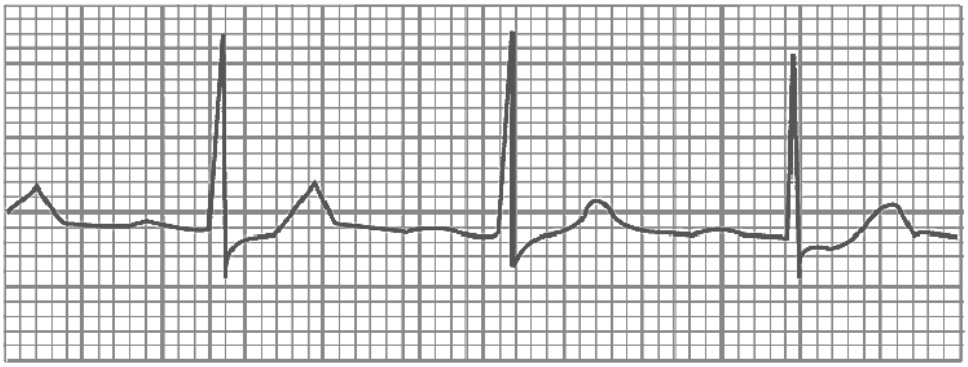

(a) ECG signal

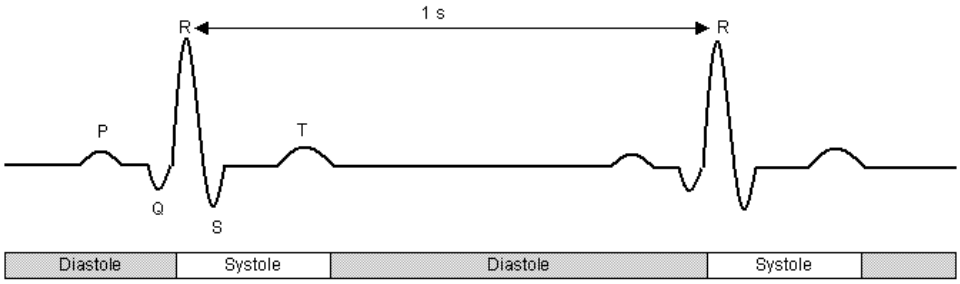

(b) Theoretic ECG segment.

Figure 4.1: ECG signal. (a) The signal amplitude is related to the electric potentials that excite the heart muscle (image courtesy of Stanford Hospital [82]). (b) The P, Q, R, S and T waves are marked on this model of one typical ECG period (image courtesy of S. Mollus [83]).

dots represent the sampling instants from the image sequence). For an ECG sampling frequency of f_{ecg} and an image frame rate acquisition of f_{img} we get $n_{\text{ecg}} = f_{\text{ecg}}/f_{\text{img}}$ ECG samples for each acquired frame. All n_{ecg} samples are utilised for R peak detection, whereas equation 4.1 is calculated at $\{t = \text{round}(n_{\text{ecg}}k - n_{\text{ecg}}/2), \quad k = 1, 2, \dots\}$, where k is the frame number.

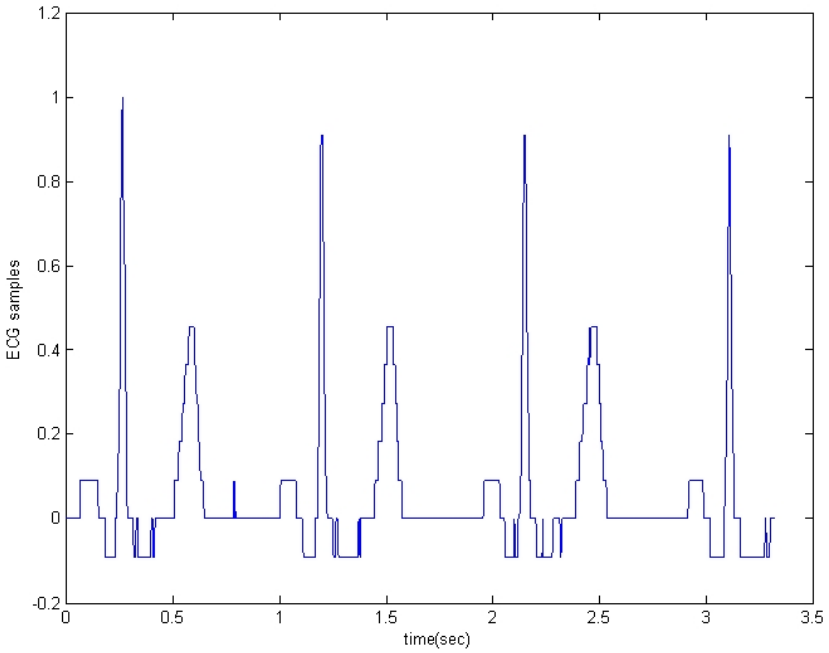

Figure 4.2: ECG from one of the DICOM files in our database. Despite the low bit resolution, the R peak is still clearly distinguishable.

Figure 4.4 shows an ECG from an angiographic sequence of our database and its corresponding heart phase calculated according to equation 4.1. The dotted line in the lower figure indicates the heart phase value calculated for each frame in the sequence.

4.1.1 Alternative ECG phase

Figure 4.1b represents the theoretical shape of an ECG signal. By looking at it we might appreciate within each ECG segment a rough symmetry around the centre of the R-R segment. It is known that during contraction and relaxation, the heart follows a hysteresis cycle, i.e. the shape of the heart evolves

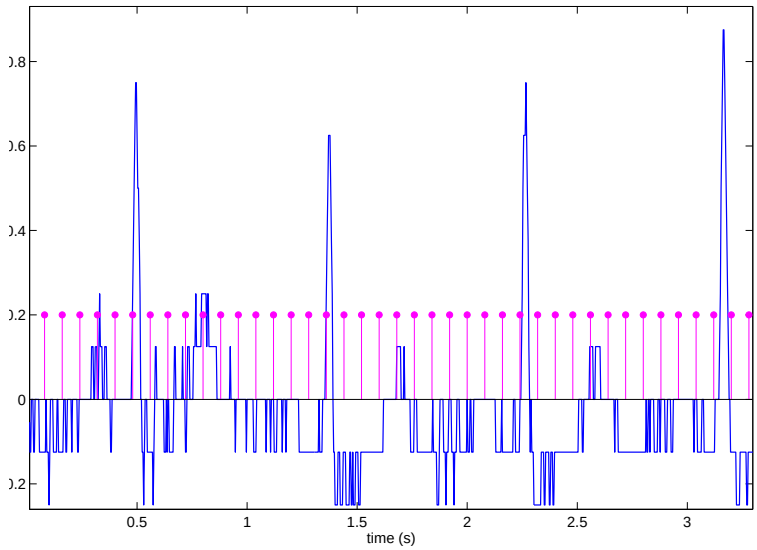

Figure 4.3: ECG from one of the DICOM files in our database. The dots represent the time instants when an image was acquired.

differently during contraction and relaxation. We conducted some experiments to establish whether the shape of the heart is substantially different during contraction or relaxation, or whether angiograms with a heart phase different from that of the current interventional frame might still be overlaid to aid navigation. Because of the depolarisation of the muscle cells and the time necessary to empty the heart chambers, there is a delay of around $0.6s$ between the occurrence of an R peak in the ECG signal and its visible effect on the heart [83]. When calculating an alternative heart phase that exploits the (imperfect) symmetry around the centre of each $R-R$ segment, this effect must be taken into account. Thus the phase-space cannot be *folded* directly around the centre of the heart phase axis, but it is necessary to introduce this offset when synchronising the ECG to the frame acquisition. In this case, the heart phase was defined as

$$\varphi(t) = 2\pi \left\| \frac{(t-t_0)-n_c}{(t_{j-1}-t_j)/2} \right\|, \qquad n_c = \frac{t_{j-1}+t_j}{2} \tag{4.2}$$

where

- t represents discrete time,
- t_0 is the offset caused by a delay of $0.6s$, measured in samples,

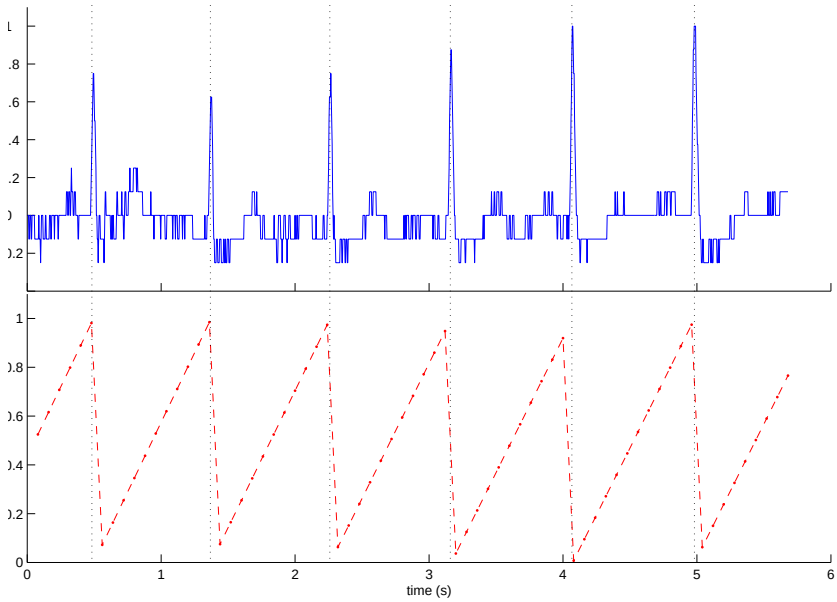

Figure 4.4: The upper figure shows an ECG from our database, while the lower figure shows the resulting heart phase calculated according to equation 4.1

- t_{j-1} is the position of the previous R peak, and
- t_j is the position of the next R peak.

Figure 4.5 shows an ECG from an angiographic sequence of our database (the same as in figure 4.4) and its corresponding heart phase calculated according to equation 4.2. The solid dots in the lower figure indicate the alternative heart phase value calculated for each frame in the sequence.

In figure 4.6, an example of registration with the new definition of heart phase in equation 4.2 is illustrated. An expert was presented several examples and was asked to evaluate how unnatural the result looked when overlaying not a matching mask but one with a heart phase symmetric around the centre of the heart phase period. The overlay looks good to unexperienced eyes, but to an expert it looks artificial. Nevertheless, we were advised that in the case of not having a roadmap with a corresponding heart phase available, this newly defined heart phase could be used in case of a possible lack of matching angiographic overlay images. We can conclude that although the heart deformation undergoes a hysteresis cycle, the difference in deformation

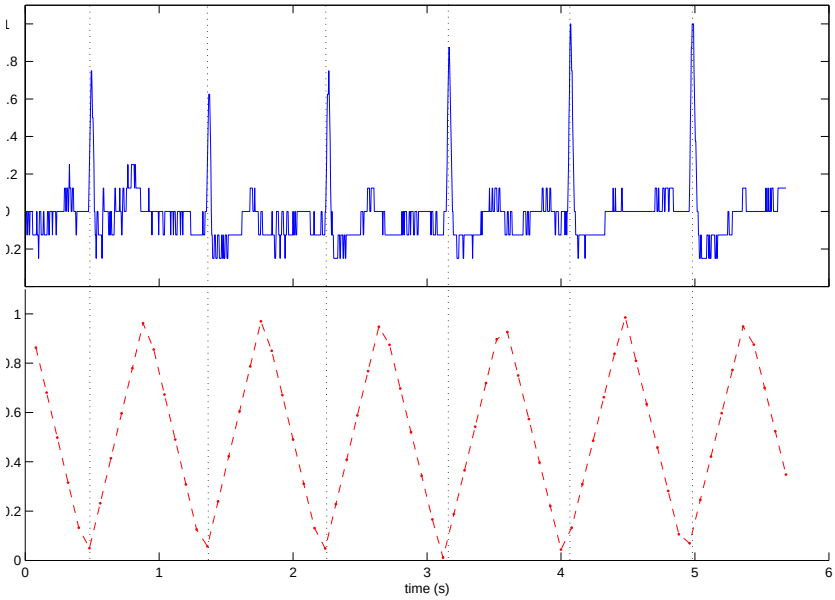

Figure 4.5: The upper figure shows an ECG from our database, while the lower figure shows the resulting heart phase calculated according to equation 4.2

during contraction and relaxation is low enough to allow an overlay of masks with the same heart phase as defined in equation 4.2 if a matching mask was not found using equation 4.1.

4.2 Related work in respiration phase detection

The most common approach for avoiding respiration motion artefacts in cardiac imaging is to avoid breathing, i.e. to perform the acquisition completely under breath-hold [84]. Unfortunately, most patients (here, a patient is by definition a person with a heart disease) are unable to keep a breath-hold for 10 seconds or more. Even if the patient manages to maintain the breath-hold, there still exist involuntary motions (digestive and cardiac motion) [24][52] that affect the quality of the breath-hold and translate into motion artefacts in the acquired images or reconstructed volumes. When using mul-

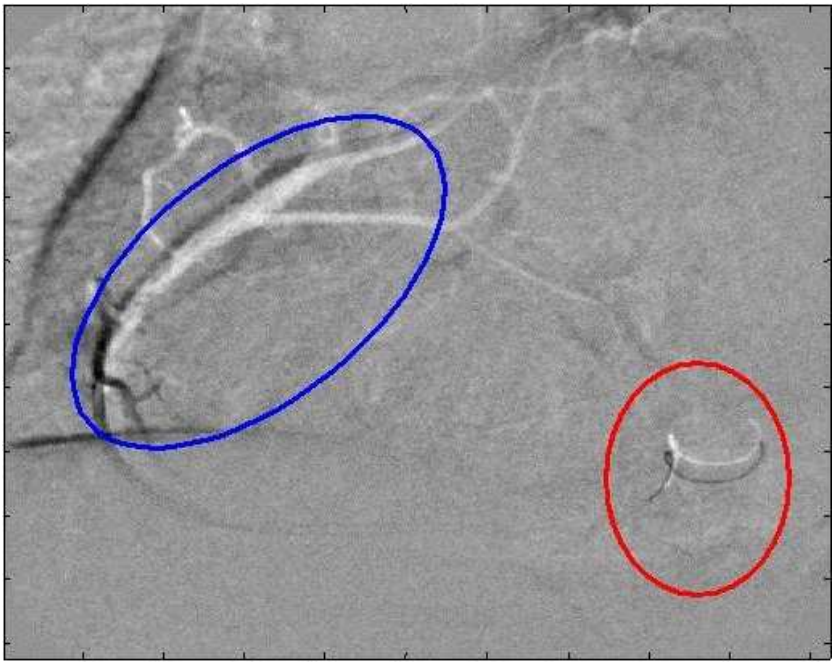

Figure 4.6: Difference image of two angiograms with the same heart phase as defined in equation 4.2 and same respiration phase. The circled areas indicate the main differences in the anatomy of the coronary vessels.

tiple breath-holds during long acquisitions, registration artefacts appear [71]. Our method benefits from dense phase-space population, i.e. acquiring angiograms for different heart and respiration states for ulterior overlay with live fluoroscopic frames. For our application it has been proposed to allow and encourage free breathing during diagnostic angiography.

Methods for respiration detection can be classified in three categories: direct measurements, indirect measurements and image-derived respiration detection methods. Direct measurements, like those provided by nasal thermocouplers and spirometers, interfere with normal respiration [85] and will not be commented here. Indirect respiration measurements (see subsection 4.2.1), like thorax volume or transthoracic impedance, measure the effect of respiration on some body variable. The third type of respiration detection methods measure the respiration from the acquired image frames (see subsection 4.2.2).

4.2.1 Indirect respiration measurement

Respiration can be measured with electrodes placed on the patient's thorax [70], which measure the change in thorax impedance as the lung volume changes with breathing. However, these electrodes will appear in the angiographic frames and obstruct the field of view. Proper placement of the electrodes to measure the respiratory phase often impedes the acquisition of a useful ECG [70], because different locations are required. In [85][86][87][88], methods to derive the respiration from the ECG are analysed. These methods do not supply a calibrated respiration signal, but are nevertheless useful for detection of apneas and estimation of the respiration rate.

Another method for measuring respiration uses a rubber bellows placed around the thorax or abdomen. The bellows is attached to a pressure transducer that records body movement associated with respiration. These methods measure movement related to respiration effort and lack the accuracy needed for our purpose. We also aim at interfering as little as possible with the current Cathlab setting and intervention, and therefore, do not favor extra machinery in the catheterisation room.

In [80], a method for static 3D biplane x-ray roadmap reconstruction is tested on a heart phantom. The respiration phase is measured by a magnetic sensor attached to the front panel of the heart phantom. Based on this respiration phase, a respiration motion is modelled as an affine translation similar to [20]. [80] does not solve the problem of extracting a respiration phase from clinical data and the method of attaching a magnetic sensor anywhere on the patient is impractical for our purposes.

4.2.2 Image derived respiration detection

If the diaphragm is present in the image, it can be used as a landmark and a respiration phase could be computed from its position in each frame (angiographic and fluoroscopic) relative to a reference position. Shechter [71][72] uses an image profile through the diaphragm in biplane x-ray frames to calculate the respiration phase, which is defined as the displacement of the diaphragm-lung interface relative to end-expiration. Condurache *et al.* model the diaphragm border as a circle [89] and detect and track the diaphragm based on edge detection and Hough transform in 2D projection x-ray images. In our application, the diaphragm is not always present in the images, since it appears as a dark structure and deteriorates the contrast between the coronary arteries and the background. Therefore, the acquisition geometry may

be chosen such that the diaphragm does not enter the field of view.

In [90] the respiration motion is indirectly measured from the position of an IVUS probe. The catheter path is imaged continuously with biplane angiography during a constant speed pullback. The x, y, and z components of the motion of an IVUS catheter are modelled with three sinusoids with frequencies for the curvature, the cardiac motion and the respiratory motion. The model is controlled by ten parameters: three frequencies, three phases, three amplitudes and a constant offset. These parameters are estimated using a trust region reflective Newton algorithm. The catheter path can be constructed taking into account only the curvature term for each spatial component. In this approach the respiration motion is eliminated by tracking continuously the manually segmented catheter position with biplane angiography. In PTCA this method is not regularly feasible because a constant surveillance of a catheter (or an instrument) that moves at a constant speed is not carried out, neither during diagnostic angiography nor during interventional fluoroscopy, unless IVUS imaging is performed. Even in this case, an angiography follow-up during pullback is unfeasible because of the high x-ray dose delivered to the patient and the physician.

In MR, navigator techniques are well established [19][20][21][22][23]. A pencil-beam navigator is placed on the dome of the right hemidiaphragm, and delivers the diaphragm position during scanning. Only frames corresponding to a certain respiration phase window are accepted. An approach similar to that of the navigator echo used in MR is also unpractical on angiograms, because a constant pencil of x-ray radiation through the patient diaphragm would mean an increase of radiation on the patient and of scattered radiation for the physician.

In [91] Nakai *et al.* propose an inter-image subtraction method of visual light images for respiration monitoring during sleep. This method performs image substraction on a selected ROI and analyses the power spectrum of the resulting image histogram. This method has been designed for same-modality images and in our framework we want to measure respiration state from sequences acquired with different image modalities, like x-ray angiography and x-ray fluoroscopy.

4.3 Assumptions for respiration phase detection

Given a sequence in which the diaphragm is visible (see figure 4.7), suppose we take a one-pixel wide column that at some point crosses the diaphragm frontier. If we observe the evolution of the pixels' gray values in time, we will see a periodicity in the intensity values of the pixels located around the diaphragm border (see figure 4.8). This confirms the assumption that respiration is a periodic activity. It also indicates that the diaphragm is a good respiration marker, since it moves according to breathing. A possibility to calculate the respiration phase would be to detect the diaphragm in the fluoroscopic and angiographic images, as in [71], [89] or the MR navigator echo techniques, and use its position as an indicator of the respiration phase. Unfortunately, the diaphragm is not always present during the acquisition and another method for extracting a respiration phase from the image content must be considered.

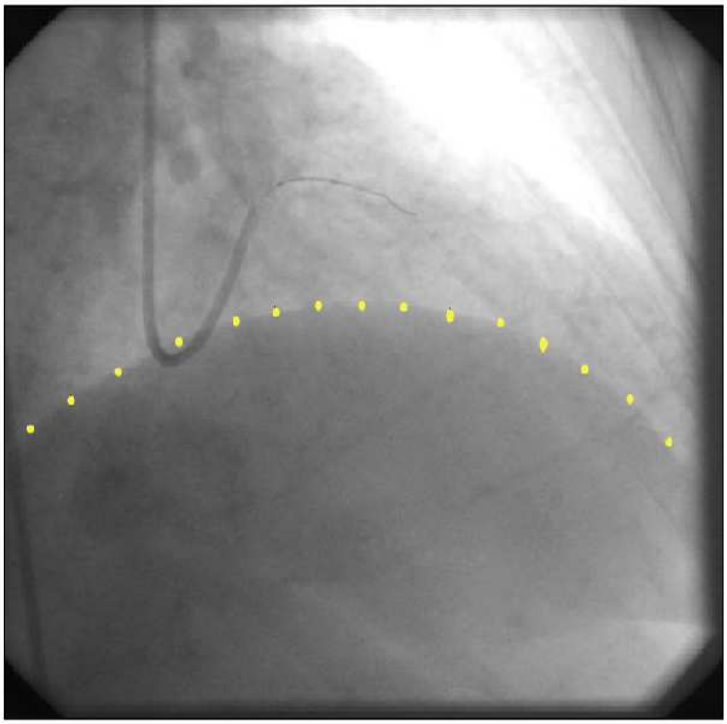

Figure 4.7: Frame from an interventional sequence in which the diaphragm is visible. The diaphragm has been manually segmented and is indicated by yellow dots.

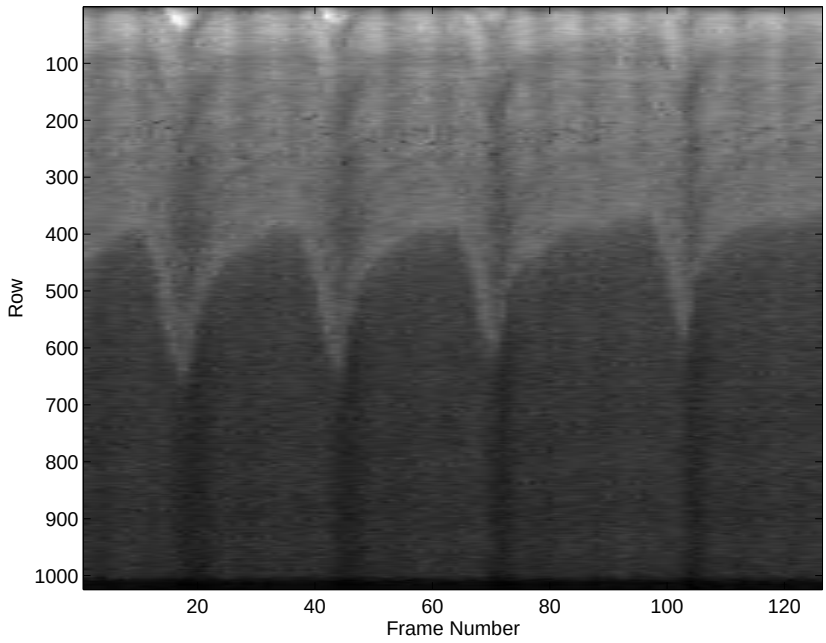

Figure 4.8: Evolution of a vertical line of pixels that crosses the diaphragm frontier in the sequence depicted in figure 4.7. The periodic movement of the diaphragm border can be seen between rows 400 and 600, where a periodic pattern in the gray level of the pixels appears.

In order to find another way of measuring the respiration phase for each frame, it is important to determine what other image content is influenced by breathing. The tissue surrounding the heart is mainly lung tissue and, consequently, is also affected by respiration. The approach taken to detect the breathing state assumes that [74]:

- Breathing affects directly the gray-level content of the image. A respiration phase can therefore be computed directly by comparing the image gray-level to that of a reference, without extracting any specific feature from the images.

- The difference in image content between two frames is caused by a different heart phase and a different breathing state (see section 3.4 in the previous chapter). We assume that it is possible to separate the respiration from the heart influence, both being periodic but with different frequencies.

If it is feasible to extract a respiration phase measuring the change effected by breathing on the tissue around the heart (see next section), it is no

longer required to segment the diaphragm or any other landmark to measure respiration induced movements. This adds to the robustness of our method and makes it applicable to all acquisitions regardless of geometry settings.

In order to estimate a respiration phase for each image, we need to define a zero-phase reference where $\beta = 0$. Assuming respiration phase over time $\beta(t)$ is a periodic function of t, the reference should show extremal respiration phase for a respiration-phase measure to be representative of respiration state. Therefore a frame at begin- or end-inspiration would be suitable as respiration phase reference. For the remainder of this chapter, a frame at end inspiration was manually chosen for each sequence on which the presented respiration phase measures were performed. In chapter 7 an algorithm to automatically chose a reference frame is presented.

4.4 Effect of breathing on the gray-level content of acquired frames

To prove that our first assumption (that a respiration phase can be measured from the image content without segmenting any specific landmark) is valid, we calculated the effect of breathing on different regions of interest of fluoroscopic sequences. This was done by measuring the similarity between each area on the frames with those of a reference image from the sequence and observing the resulting similarity signal. The reference image has respiration phase 0 and is chosen from the acquired fluoroscopies such that it is acquired at end inspiration.

To observe the effect of respiration on the soft tissue visible in the images, we chose a similarity measure previously proposed for the registration of angiographic images in Digital Subtraction Angiography (DSA) [92][93][94]. The similarity between two images is measured based on the histogram of the difference image. When the images are registered, we expect the histogram of the subtraction image to present pronounced peaks close to 0 corresponding to registered areas. If the images are misregistered, we obtain a dispersed histogram. Buzug *et al.* propose in [92][93][94] a similarity measure that quantifies how dispersed the histogram of the difference image is by calculating its energy, which is defined as

$$E = \sum_{k=0}^{N-1} p(k)^2 \tag{4.3}$$

where N is the number of histogram bins and $p(k)$ is the fraction of pixels with gray value k. The square function is convex and super-additive, which translates in the absolute minimum of the energy being reached when $p(k) = 1/N$ and increasing energy as peaks appear in the histogram [93]. A quick proof of this property can be seen by considering the gray levels of the difference image to be uniformly distributed, so that $p(k) = 1/N$ for all gray values k. In this case, the energy is $E = 1/N$. If now we consider that a peak appears in the histogram such that $p(k) = 1/N$ for $N - 2$ bins, $p(k) = 0$ for one bin and $p(k) = 2/N$ for another bin, the resulting energy is

$$E' = \frac{N-2}{N^2} + \frac{2^2}{N^2} > E = \frac{1}{N} \tag{4.4}$$

This simple proof can be generalised [93] to prove that the histogram energy (or the similarity of the images) increases when peaks appear in the histogram (i.e. the images are properly registered).

The histogram energy measure in equation 4.3 is appropriate to register images acquired with the same imaging modality even at different dose levels, because an offset between the gray values of the images to be compared leads to a shift of the histogram peaks [93]. We used a fluoroscopic sequence for this experiment, to eliminate the need of a multimodality measure and avoid the influence of the flowing contrast agent present in angiography.

The respiration phase was calculated for different rectangular regions of interest (ROI) with the same area (see figure 4.9), to prevent computationally expensive calculations that are not suitable for real-time applications, and avoid non-relevant parts of the images. β is better determined from areas where the image content depends mostly on respiration movement and less on heart beat or contrast agent presence. But the field of view is usually centred at the coronary arteries and interventional devices, so that noise related to heart beat will be present in the measure. For each ROI in each frame in the sequence, we subtracted the same ROI in the reference frame, and calculated equation 4.3 using 256 bins for the histogram. The ROIs where manually placed around the catheter stem, in an area centred on some coronary artery, at the diaphragm border, inside the diaphragm area and around some lung tissue (see figure 4.9). Previous to the calculation of the respiration phase according to equation 4.3, each area was smoothed with a median filter to remove quantum noise. Additionally, the resulting respiration phase was normalised to the interval $[0, 1]$ and low-pass filtered with an order 3 windowed linear-phase FIR filter.

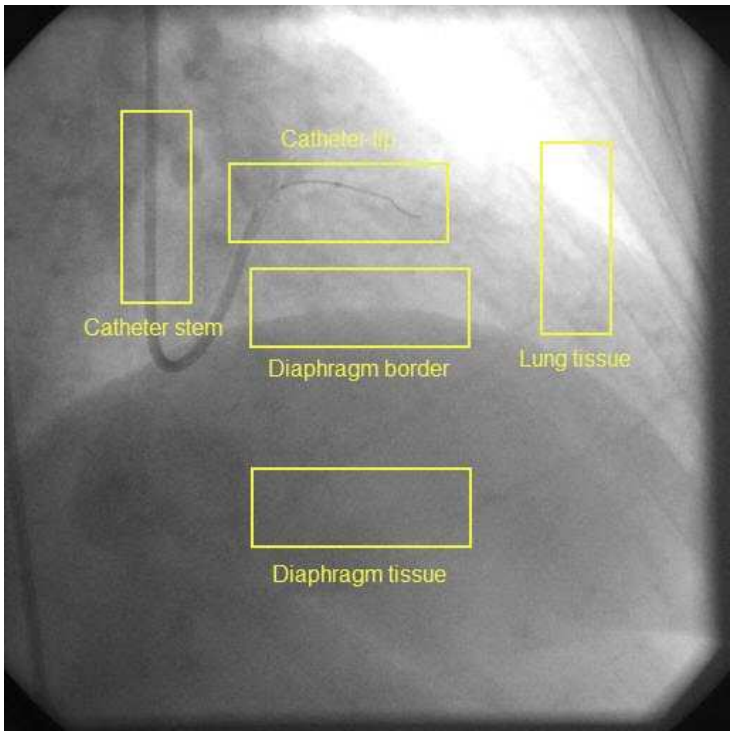

Figure 4.9: Areas selected to prove that a respiration phase can be derived from different contents present in the image.

The results for each area appear in figures 4.10 to 4.14. The reference image at end inspiration and the corresponding ROI, (a), are depicted above the calculated histogram energy, (dotted line in (b)), superimposed to the result from low-pass filtering the histogram energy (solid line in (b)) with a cutoff frequency of 1Hz in order to eliminate the interference of the heart beat (a normal heartbeat is between 60 and 100 beats per minute). We can observe a low-frequency periodic component in the calculated histogram energies, which reflects the changes in similarity to the reference image due to breathing, whereas the higher frequencies are caused by noise and the high-frequency changes that heart contraction induces in the surrounding soft tissue.

In figure 4.10 and 4.11 we see the results of the similarity measure for the diaphragm border and the diaphragm tissue, respectively. The influence of respiration on the image gray levels can be seen on the low frequency component of the similarity signal, and the high-frequency heart-contraction influence is almost completely removed by low-pass filtering with a cut-off

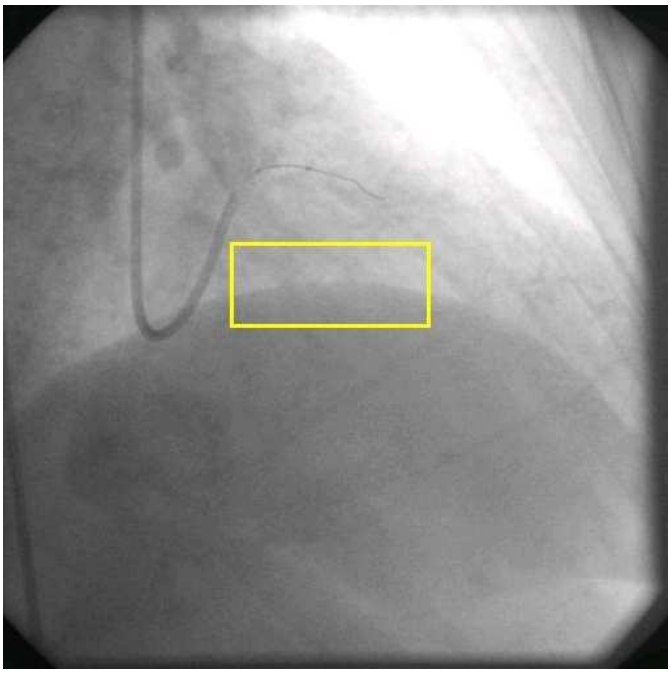

(a) ROI around the diaphragm border.

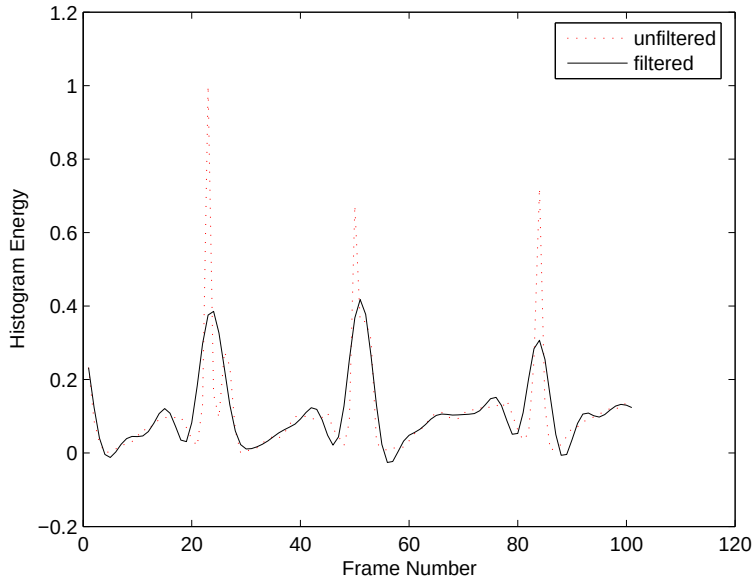

(b) Histogram energy.

Figure 4.10: Respiration phase calculated with the histogram energy criteria for an area in the diaphragm border, (a). The calculated histogram energy in (b) is represented with a dotted line, while the low-pass filtered result is depicted with a continuous line. The resulting respiration phase is dominated by a low-frequency component.

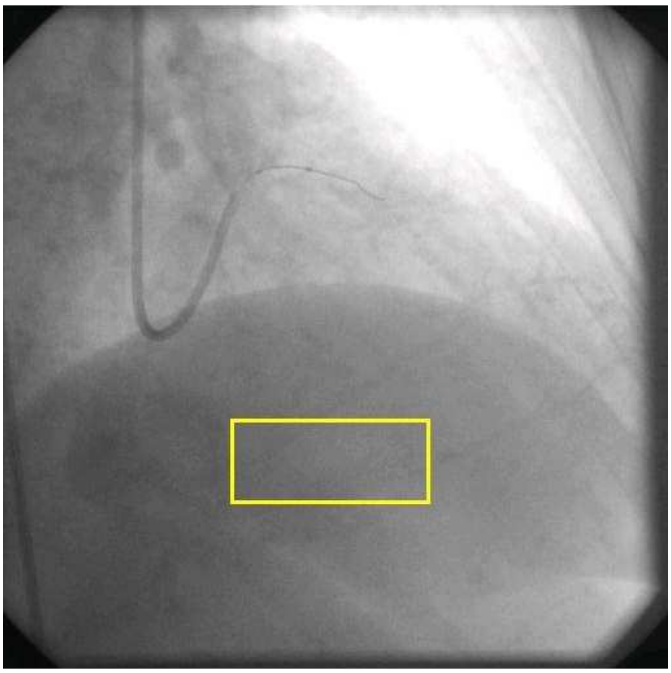

(a) ROI composed of diaphragm tissue.

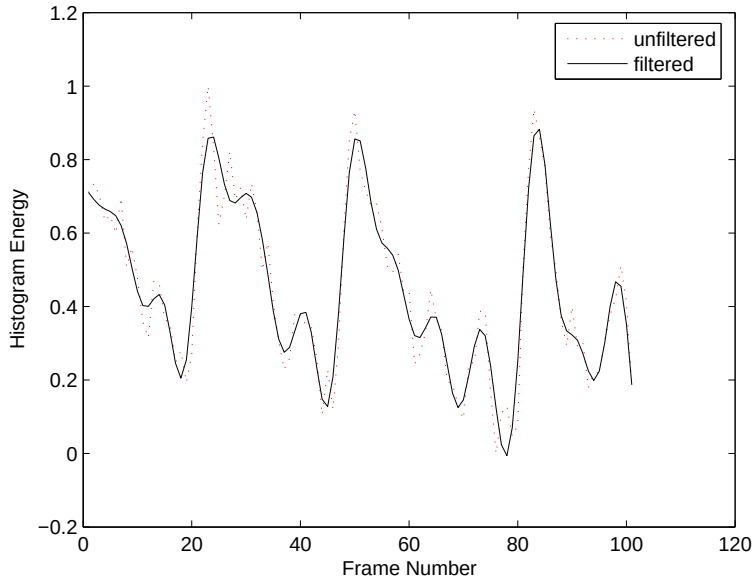

(b) Histogram energy.

Figure 4.11: Respiration phase calculated with the histogram energy criteria for an area inside the diaphragm, (a). The calculated histogram energy in (b) is represented with a dotted line, while the low-pass filtered result is depicted with a continuous line. The resulting respiration phase is dominated by a low-frequency component.

frequency of 1Hz. As expected, this result confirms that the diaphragm is a good respiration indicator. In figures 4.12 and 4.13 we have selected two ROIs around the catheter tip (inside a coronary artery) and around the catheter stem, respectively. The resulting histogram energy curves (dotted lines in figures 4.12b and 4.13b) have more high-frequency components than in figures 4.10 and 4.11 because of the higher influence of the heart contraction in the content of the ROIs. The resulting respiration phase for an area containing lung tissue was also investigated and is illustrated in figure 4.14. Figure 4.14b presents a distinctive low-frequency component related to breathing.

As expected, the closer to the diaphragm, the smaller is the influence of heart contraction and the more dominant is the breathing on the similarity signal. But our results also indicate that it is not absolutely necessary to centre our measurements around the diaphragm border (as in [23][71] or the navigator echo technique used in MR [19][20][21][22][23]) in order to detect the periodic movement induced by breathing, because other parts of the image are also affected by it. We can therefore conclude that by comparing the image content with that of a reference, a valid respiration phase can be calculated without placing any respiration detector on the patient's chest or segmenting any specific landmark in the images. This result facilitates the task of finding a suitable area in the image to perform respiration state detection.

The histogram energy similarity measure used in this experiment has a major drawback. This measure depends on the histogram of the difference image, which depends directly on the intensity value of the pixels. This implies that frames with similar breathing states from an angiographic and a fluoroscopic sequence, in which the radiation dose and the presence of contrast agent is different, would be assigned different values of respiration phases when compared to the same reference image. This is a major drawback for our present application, in which we need to compare frames from diagnostic angiography and interventional fluoroscopic sequences in terms of respiration phase. The problem is related to that of multimodality image registration, where 2D or 3D images acquired with different and complementary imaging modalities are registered. The geometric transformation applied to one image or volume that maximises its similarity to other images or volumes is calculated [95][96][97][98]. Extensive literature exists on multimodality registration of images using local correlation [95][99][100], mutual information [96][101][102][103][104][105], and other entropy-based similarity measures [78][104][106]. Mutual information is discussed and evaluated as a respiration phase measure [74] in the next section. In section 4.6, we compare

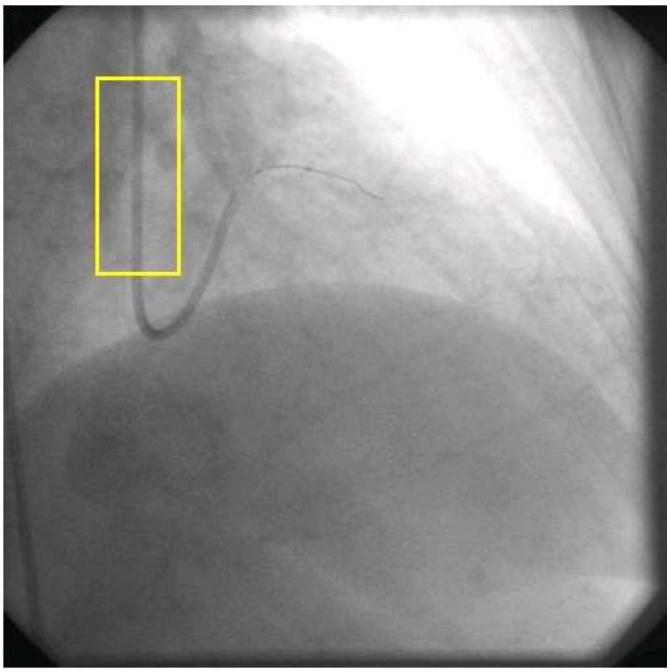

(a) ROI around the catheter stem.

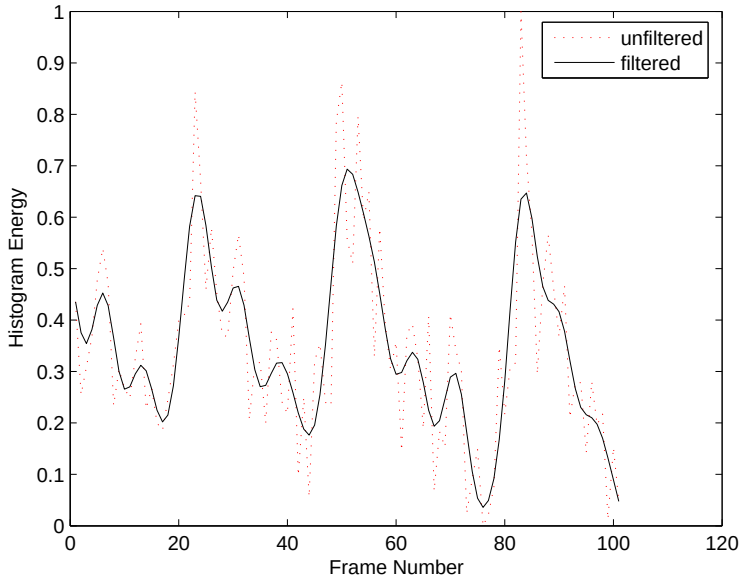

(b) Histogram energy.

Figure 4.12: Respiration phase, (b), calculated with the histogram energy criteria for an area centred around the catheter stem, (a).

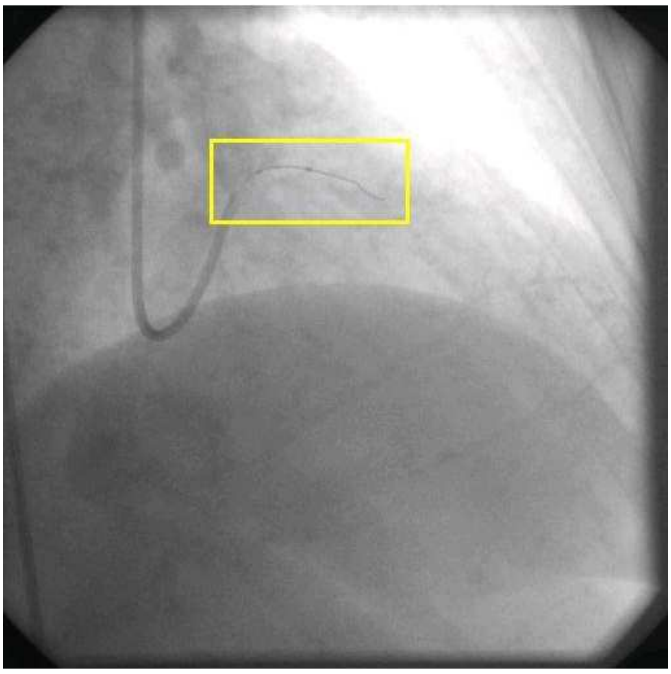

(a) ROI around a coronary artery.

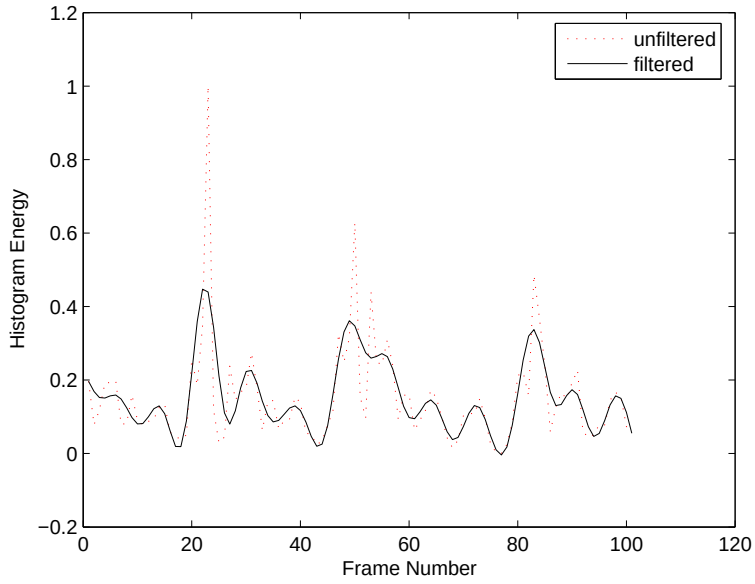

(b) Histogram energy.

Figure 4.13: Respiration phase, (b), calculated with the histogram energy criteria for an area around the coronary arteries, (a).

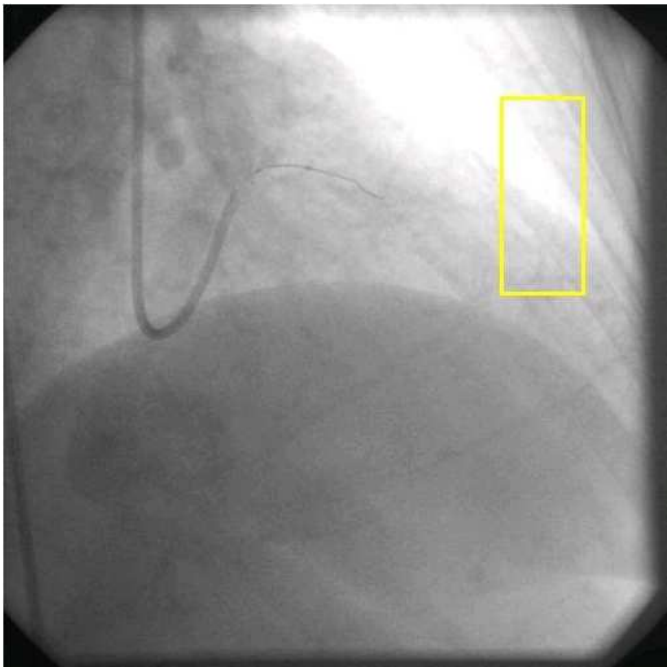

(a) ROI composed of lung tissue.

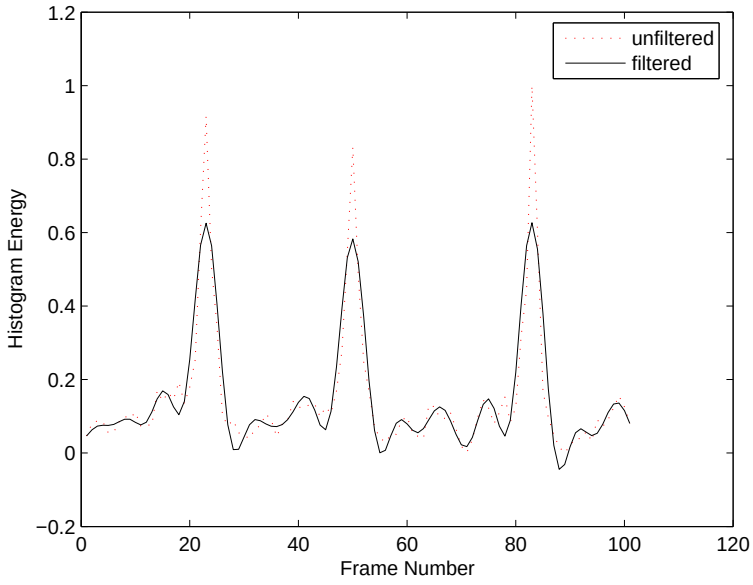

(b) Histogram energy.

Figure 4.14: Respiration phase, (b), calculated with the histogram energy criteria for an area containing lung tissue, (a).

MI with local correlation [95][99][100] and the measures proposed in [78] as respiration phase measures.

4.5 Mutual Information

MI is a measure used in information theory that measures the statistical dependance between two stochastic variables. MI as a similarity measure between pairs of images for multimodality registration was introduced by Paul Viola in his PhD thesis [77] and has been used in several multimodality registration applications [96][102][103][104][105][107]. MI does not assume a linear relation between the intensity levels of the images, but that the co-occurrence of the most probable intensity values is maximised when the images are registered. In subsection 4.5.1 we define the necessary information theory concepts that lead to MI. In subsection 4.5.2 we evaluate MI as a respiration phase measure. In section 4.6 we compare MI to other multimodality measures proposed in the literature.

4.5.1 Mutual information as a similarity measure

The Shannon entropy of a random variable X, with a discrete alphabet $\chi = \{x_1, x_2 \dots x_k\}$, and a marginal probability distribution $p_X(x)$, is defined as [108]

$$H(X) = -\sum_{x \in \chi} p_X(X = x) \log[p_X(X = x)] \tag{4.5}$$

where $p_X(X = x)$ is the probability of X taking the value x. The entropy can be interpreted as the mean of $I(p_X) = -\log[p_X(X = x)]$, where $I(p_X)$ is regarded as the amount of *information* related to variable X [108].

The entropy is a measurement of the uncertainty of a random variable. If the random variable has a uniform distribution, $p_X(X = x) = 1/N \quad \forall x \in \chi$, then the entropy reaches its upper limit, because the uncertainty as to the value of X is maximum. This is a consequence of the sub-additivity property of the entropy, i.e.

$$-(p_i + p_j) log(p_i + p_j) \leq -\big(p_i log(p_i) + p_j log(p_j)\big) \tag{4.6}$$

where $p_i = p_X(X = x_i)$ and $p_j = p_X(X = x_j)$, $i \neq j$.

Given images B and T, the MI between them is defined as [77]

$$MI(B;T) = \sum_{x,y} P_{XY}(x,y) \log_2 \frac{P_{XY}(x,y)}{P_X(x)P_Y(y)} \tag{4.7}$$

where

- X and Y are two random variables taking values in $\mathcal{X}$ and $\mathcal{Y}$, which are the set of intensity values of images B and T respectively,
- $P_{XY}(x,y)$ is the joint probability distribution of B and T,
- $P_X(x)$ is the probability distribution of B, and
- $P_Y(y)$ is the probability distribution of T.

The probability distributions are estimated as the normalised histograms or joint histogram of the images.

MI can also be expressed in terms of entropy as

$$MI(B;T) = H(B) - H(B|T) = \tag{4.8}$$

$$H(B) + H(T) - H(B,T) \tag{4.9}$$

According to equation 4.8, MI can be interpreted as the decrease in the uncertainty of stochastic variable B, $H(B)$, when the conditional entropy between B and T, $H(B|T)$, is known. To understand the usefulness of MI for image alignment, let's consider equation 4.9. If images B and T are aligned, the joint histogram (and therefore $P_{XY}(x,y)$) has defined peaks for the pixel combinations of the registered areas. As a consequence, the term $H(B,T)$ in equation 4.8 decreases (remember that the entropy reaches its maximum if all variables have the same probability) and the MI increases. If the images are not aligned, the joint histogram is more dispersed and the MI decreases.

4.5.2 Evaluation of MI as respiration phase measure

To evaluate the performance of MI as a respiration phase measure, we compared the resulting respiration phase for different ROIs of the image with a bronze standard [74]. We manually segmented the diaphragm in each frame of two fluoroscopic and two angiographic sequences acquired with different geometry settings, and selected for each a reference frame at end inspiration.

In the case of the fluoroscopic sequences, the reference frame was taken from an angiographic sequence acquired with the same system geometry. We calculated in each case the mean distance between diaphragm points (MDD) in each frame of the sequences and their respective reference frame (see figure 4.15). This distance variation was adopted as a respiration phase bronze standard. We plotted the normalised respiration phase derived from MI for different areas (see figure 4.16) against the MDD and fitted a cubic curve to the resulting plot for reference. If the ROI selected to measure the respiration phase accurately represents the movement of the diaphragm border, we expect the MDD *vs.* MI plot to result in a straight line. A deviation from a straight line implies that the tissue contained in the corresponding ROI does not follow the movement of the diaphragm border present in the image. We use the correlation coefficient between the respiration phase calculated with MI and the normalised MDD to evaluate the accuracy of MI as a respiration phase measure in each area. The absolute value of the correlation coefficient between the normalised MI and normalised MDD should be close to 1. In the following figures (figures 4.17 to figure 4.19), we show the results for an exemplary sequence. We summarise the results for all four test sequences at the end of this section.

In figure 4.17a the respiration phase for a ROI that crosses the diaphragm border is shown. When this respiration phase is plotted against the MDD, the resulting fitted curve (figure 4.17b) is monotonically decreasing and approximately linear. The quasi-linearity of the fitted curve (the correlation coefficient is -0.9686) indicates that the MDD is indeed an appropriate measure for the quality of the MI as respiration phase, and that, as expected, the MI respiration phase for an area crossing the diaphragm border reflects the translation experienced by the diaphragm border as expressed by the MDD.

Figure 4.18 displays the results obtained for an area that contains diaphragm tissue but does not cross the diaphragm border at any time. The fitted curve in figure 4.18b is monotonically decreasing but its linearity is degraded (the correlation coefficient is -0.8751) as the difference in respiration phase increases. This can be explained by the hysteretic nature of respiration [52] and the fact that respiration does not induce a simple global translation like suggested by [69], but a more complicated affine motion [20]. The hysteresis cycle can be noticed in the interval $30 < MDD < 100$ of figure 4.18b, where the points located above the fitted curve correspond to inspiration, and the points located below the fitted curve correspond to expiration. When measuring the respiration phase in any area that does not cross the diaphragm border, we expect the effects of the affine motion caused by

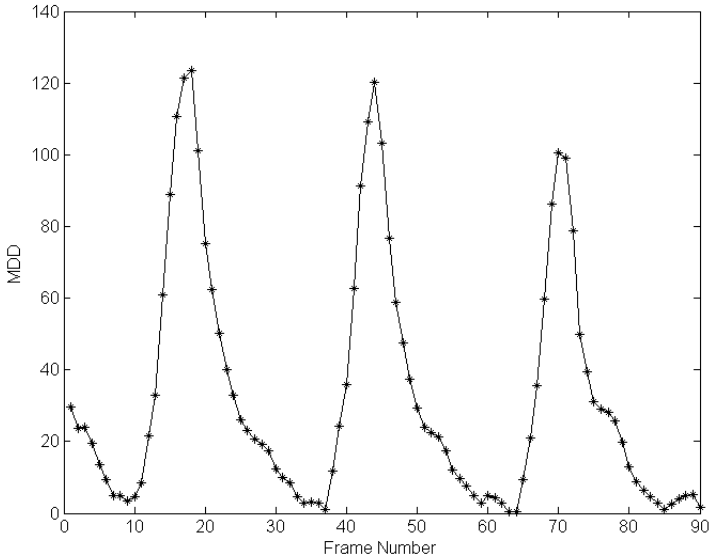

Figure 4.15: Bronze standard: mean distance between diaphragm points (MDD). After manually segmenting the diaphragm in all frames of the sequences and in the reference image, the respiration phase was calculated as the mean distance between segmented diaphragm points.

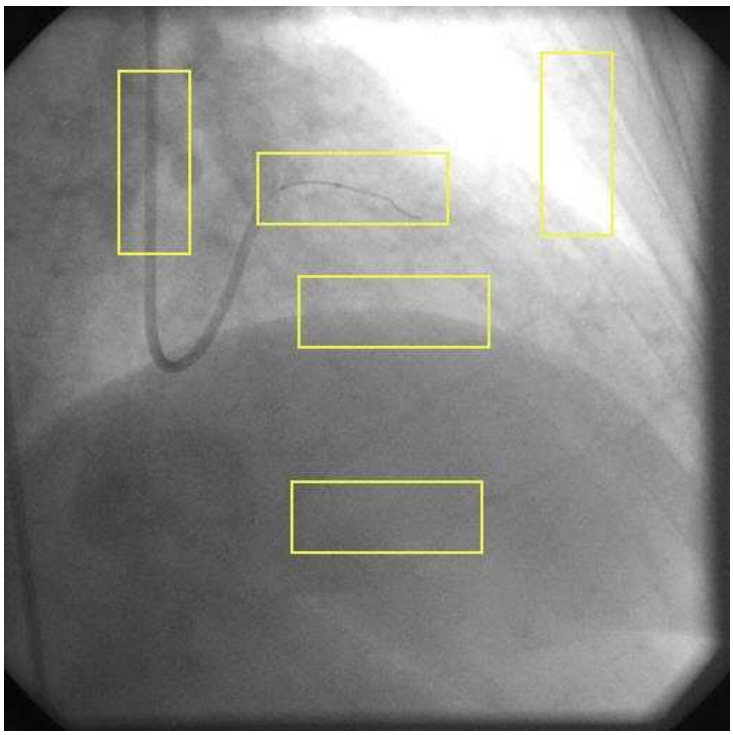

Figure 4.16: Areas for which the MI was evaluated as a respiration phase measure.

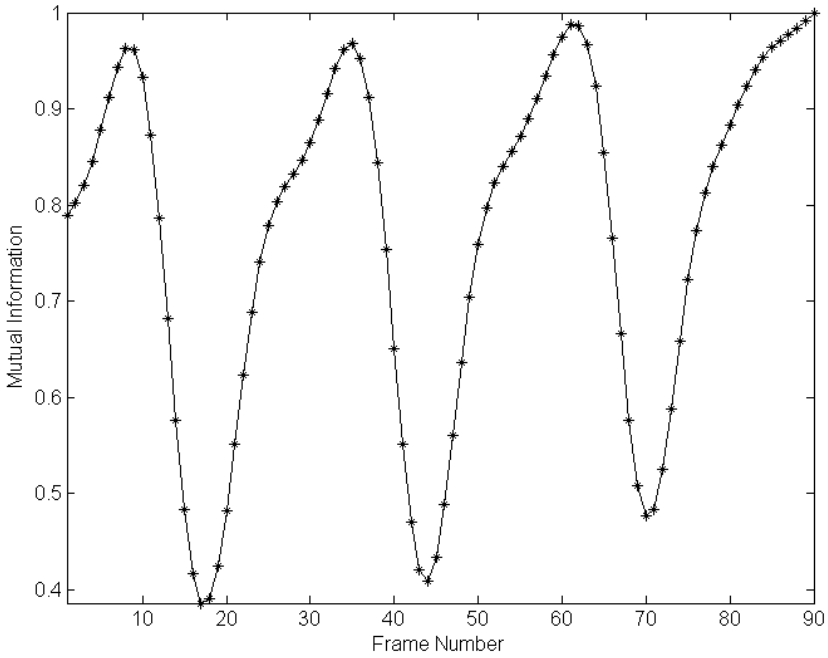

(a) MI-based respiration phase.

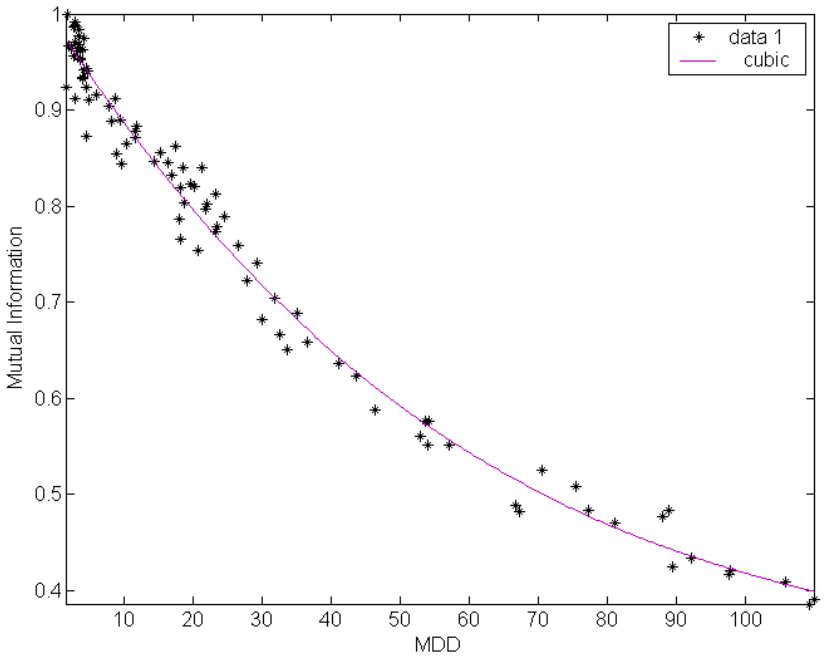

(b) MI-based respiration phase *versus* MDD

Figure 4.17: Evaluation of the MI-based respiration phase at the diaphragm border. (a) Respiration phase based on MI. (b) Calculated respiration phase plotted against the MDD.

breathing on the heart and the hysteretic nature of respiration, to be present.

As we focus on regions farther away from the diaphragm and nearer to the heart, the high-frequency influence caused by heartbeat increases and a phase offset appears on the similarity measure. This offset is explained by the physiology of breathing: the diaphragm and the lung tissue are affected before the rest of the tissue because of their active role in breathing. In figure 4.19a the MDD and the resulting MI-based respiration phase for an area around the catheter stem are plotted. A phase offset of approximately 1 frame ($\approx 0.12s$) is present. When comparing the MI with the MDD in this area, this offset must be taken into account. The resulting plot is shown in figure 4.19b where the hysteretic nature of breathing can still be appreciated, but the points are more dispersed than in figure 4.18b. The correlation coefficient improves from -0.8163 to -0.8412 when including this offset in the calculation.

In figure 4.20 the effect of respiration in an area around the coronary arteries can be seen. The low frequency effect caused by breathing can be appreciated, but we can also observe an interfering low frequency signal superimposed to respiration as can be seen in figure 4.20a. We hypothesise that this signal is caused by some low-frequency contraction movement of the heart. The effect of this interference can be clearly seen in figure 4.20b for $MDD < 40$. In this particular example, the correlation coefficient between the normalised MDD and MI respiration phases is -0.2032.

In figure 4.21a the respiration phase calculated for a ROI placed around the lung tissue area and the MDD are shown. The MI measure shows a phase offset of approximately 2 frames ($\approx 0.24s$) ahead of MDD. As we can see on figure 4.21, the resulting signal is weak, because of overexposure of the area. When plotting the MI respiration phase *versus* the MDD we observe a linear monotonical decrease (the correlation coefficient is -0.8373), but the measured values are very dispersed with respect to the fitted curve, which compared to the MDD would have a correlation coefficient of -0.9645.

In all four sequences evaluated, the ROI containing the diaphragm border resulted in the absolute value of the correlation coefficients closest to 1. The deviation from an absolute value of the correlation coefficient of 1 was evaluated calculating the RMS for each ROI. The results are given in table 4.1. The better performance of the ROI located around the coronaries compared to those around the lung or the catheter can be explained because despite the increased influence of heart contraction, the ROI is also located

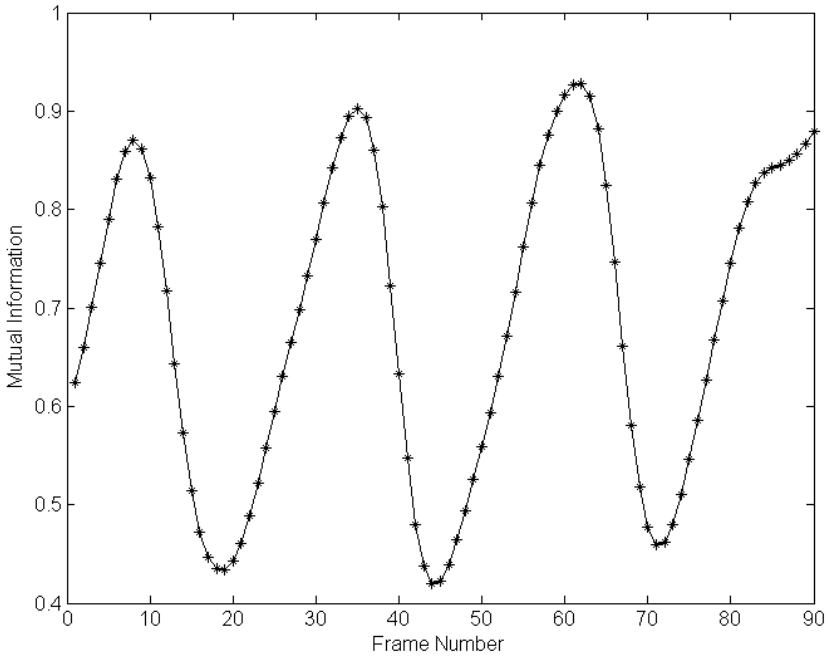

(a) MI-based respiration phase.

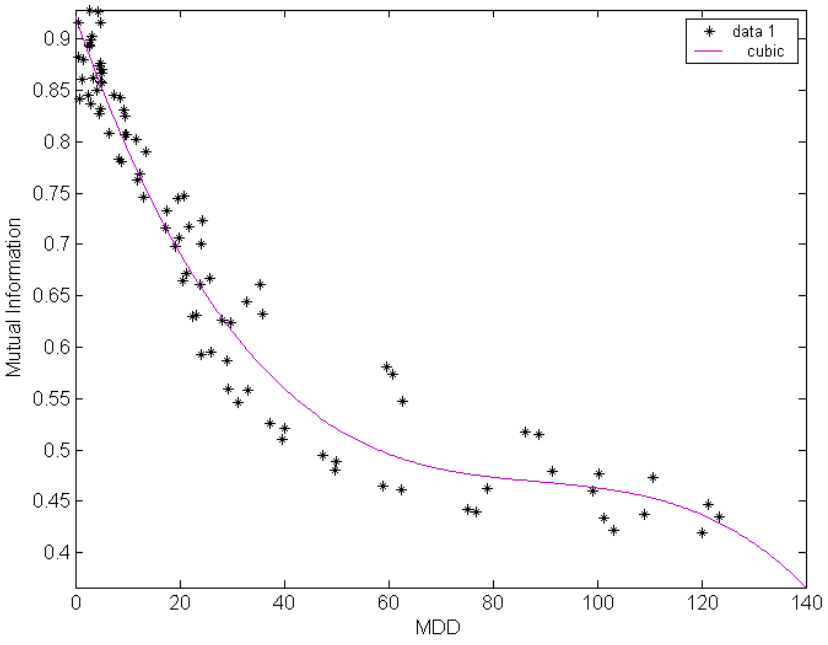

(b) MI-based respiration phase *versus* MDD

Figure 4.18: Evaluation of the MI-based respiration phase for an area that contains diaphragm tissue. (a) Respiration phase based on MI. (b) Respiration phase plotted *versus* the MDD.

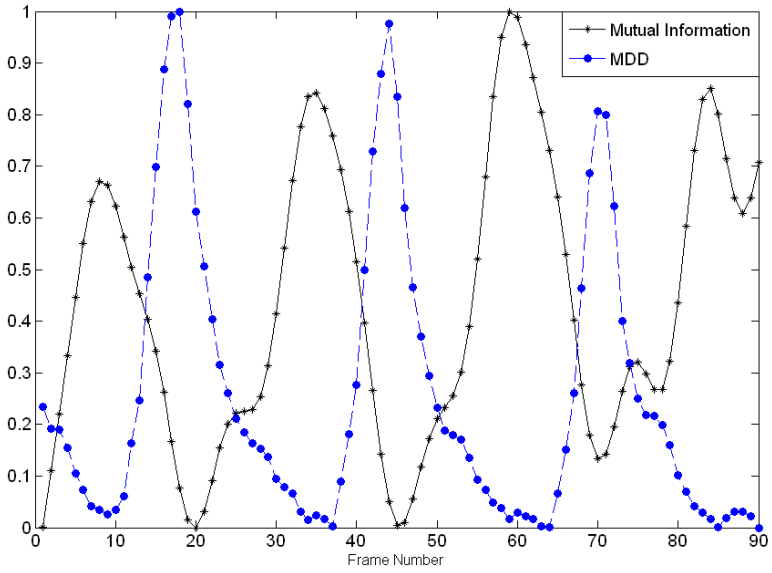

(a) MI-based respiration phase.

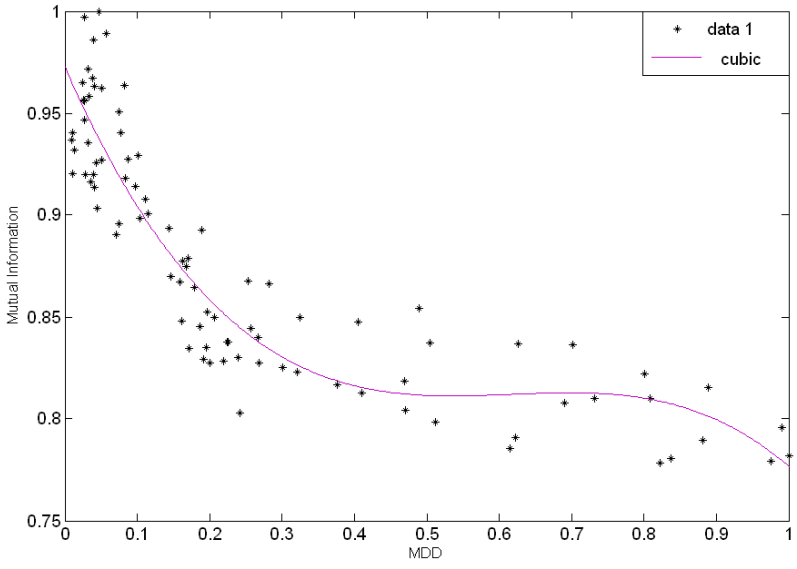

(b) MI-based respiration phase *versus* MDD

Figure 4.19: Evaluation of the respiration phase for an area centred around the catheter stem. (a) Respiration phase based on MI. (b) Respiration phase plotted *versus* the MDD.

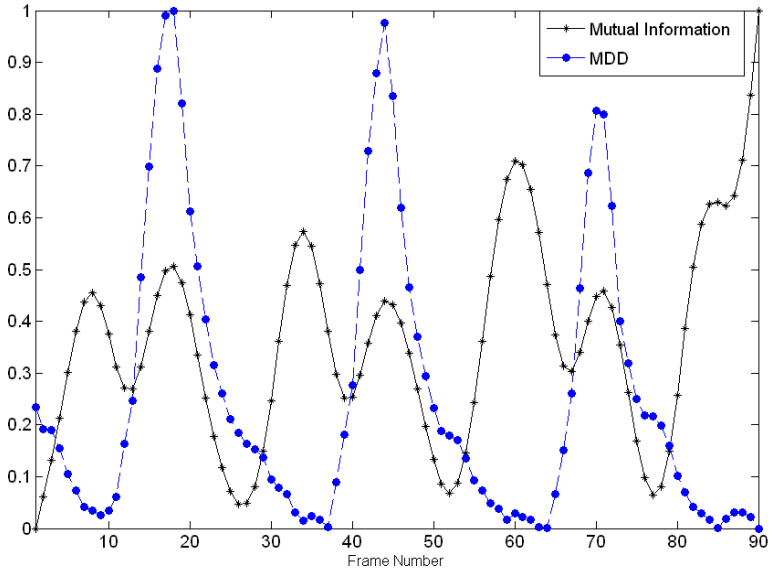

(a) MI-based respiration phase.

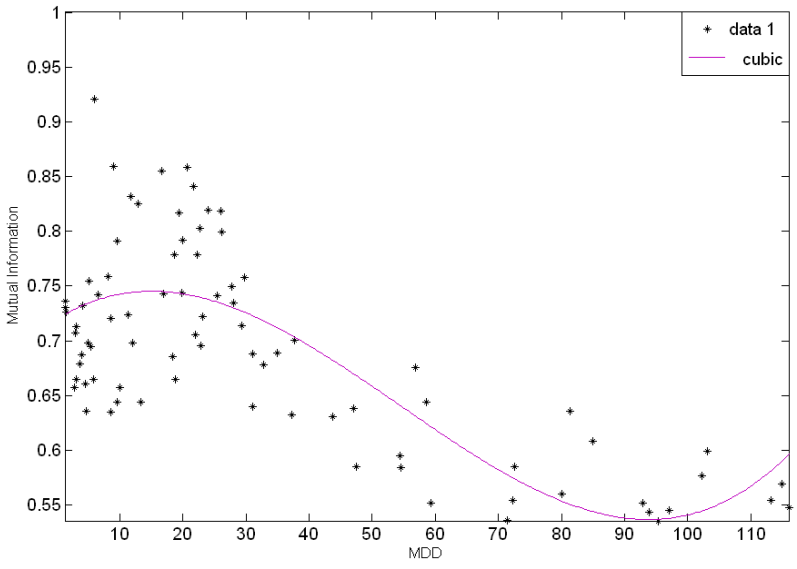

(b) MI-based respiration phase *versus* MDD

Figure 4.20: Evaluation of the MI-based respiration phase for an area centred at the coronaries. (a) Respiration phase based on MI. (b) Respiration phase plotted *versus* the MDD.

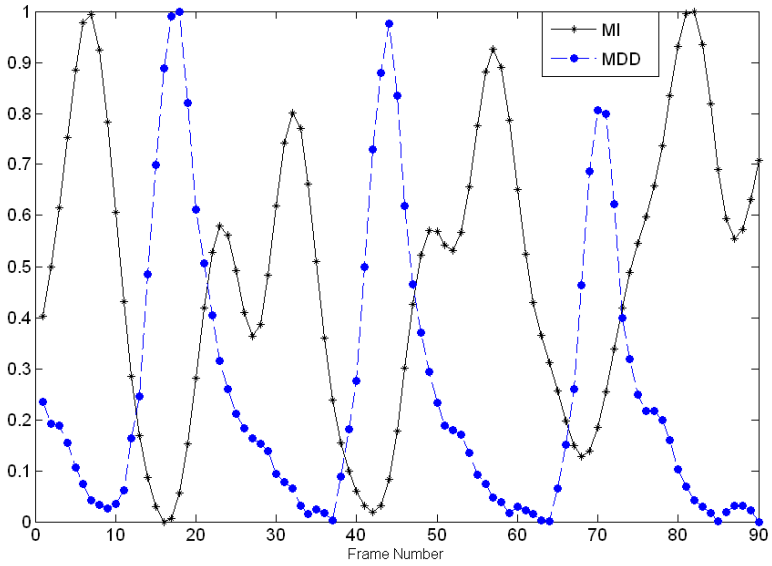

(a) MI-based respiration phase.

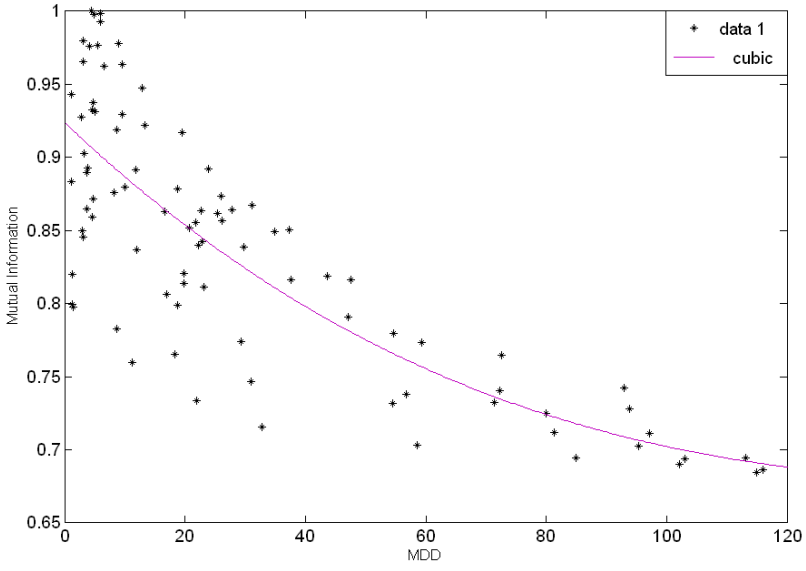

(b) MI-based respiration phase *versus* MDD

Figure 4.21: Evaluation of the MI-based respiration phase for an area located at the lung. (a) Respiration phase based on MI. (b) Respiration phase plotted *versus* the MDD.

at an object that is moving according to respiration, and therefore, the respiration phase measured has a very close relation to the actual movement of the heart caused by breathing. Furthermore, the lung area suffers from x-ray overexposure during the intervention which greatly affects the image content in said region and discourages its selection for respiration phase estimation based on image gray-values.

ROI	RMS
Diaphragm frontier	0.1054
Diaphragm tissue	0.3010
Catheter	0.3766
Coronaries	0.2750
Lung	0.4914

Table 4.1: Considering that in the ideal case, the absolute value of the correlation coefficient between the MDD and the MI should be 1, this table shows the RMS of the correlation coefficients for all examples in each ROI considered for evaluation.

In this section we have studied the performance of MI for respiration phase detection and evaluated which area of the image results in the best quality respiration measure. Our results show that the diaphragm border is a good ROI for respiration phase detection but that other areas of the image perform also well when the diaphragm is not imaged. In the following section, we compare MI with other similarity measures.

4.6 Comparison with other multimodality similarity measures

In order to study the performance of MI as a respiration phase measure, we compare it in the following sections with five other multimodality similarity measures that have been proposed in the literature as an alternative to mutual information for multimodality image registration.

In [109][110] Weese *et al.* propose *local correlation* as a multimodality measure based on the observation that the correlation coefficient is not an appropriate similarity measure for multimodality image registration because of the non-linear relation of intensity levels in the images. The images (or ROI considered in each image) are divided in small neighbourhoods where they assume a linear relation between intensity levels, and local correlation

is calculated as the sum of all individual correlation coefficients between the neighbourhoods considered in the two images. In [78] Bardera *et al.* propose two measures for multimodality image registration based on information theory. These measures quantify the dissimilarity between two random variables, whereas mutual information measures the dependence between the two variables.

We compared the measures proposed by Weese *et al.* [109][110] and Bardera *et al.* [78] with MI (subsections 4.6.1 and 4.6.2 respectively) and evaluated their performance for respiration phase calculation in an area around the diaphragm border by comparing them with the MDD (mean distance in pixels between diaphragm points) bronze standard for our four exemplary sequences. After evaluating the performance of these alternative respiration phase measures, we conclude that mutual information is the most accurate when measuring the respiration phase from angiographic and fluoroscopic x-ray images.

4.6.1 Local Correlation

A very common approach to find the similarity between two linearly-related signals, x and y is to use the cross-correlation or correlation coefficient, which is defined as

$$r^2 = \frac{[\sum_{i=1}^{N-1}(x_i - \bar{x})(y_i - \bar{y})]^2}{[\sum_{i=1}^{N-1}(x_i - \bar{x})]^2[\sum_{i=1}^{N-1}(y_i - \bar{y})]^2} \tag{4.10}$$

r gives the quality of a least square fitting between x and y. It is close to 1 if $y \approx a \cdot x + b$ (a and b constants) and is 0 if the signals are not linearly correlated.

In multimodality imaging the global cross-correlation is not a good measure of the similarity because the intensity values of the images are not linearly correlated. Local Correlation (LC) was proposed in [109][110] as a gray-scale based multimodality similarity measure for three-dimensional image registration. LC assumes a linear relation of image intensities in small neighbourhoods $n(j)$ of the frames to be registered [95]. Given images B and T which take intensity values in $\mathcal{X}$ and $\mathcal{Y}$ respectively, the cross-correlation coefficient between the gray values x_i of the voxels (or pixels) inside a small neighbourhood $n(j)$ in image B and the corresponding gray values in the

target image T is calculated and accumulated [95] to result in

$$LC(X,Y) = \frac{1}{N}\sum_{j}\frac{[\sum_{i\in n(j)}(x_i-\bar{x}_j)(y_i-\bar{y}_j)]^2}{[\sum_{i\in n(j)}(x_i-\bar{x}_j)]^2[\sum_{i\in n(j)}(y_i-\bar{y}_j)]^2} \tag{4.11}$$

where

- $n(j)$ is a neighbourhood around voxel (or pixel) j in the image B,
- x_i and y_i are the gray values of voxels (or pixels) i in $n(j)$ for images B and T, and
- $\bar{x}_j$ and $\bar{y_j}$ are the average gray values of the voxels (or pixels) in the neighbourhood $n(j)$ for images B and T.

If all the voxels (or pixels) in $n(j)$ have the same value, the contribution to LC is defined to be 0.

Note that the term inside $\sum_j$ corresponds to the squared correlation coefficient of the voxels (or pixels) inside the neighbourhood $n(j)$. This coefficient is affected by the intensity values present in the image and therefore a few large differences of intensity values in the neighbourhood (like those caused by edges) can affect the similarity measure. The contribution of aligned edges is greater than the contribution of aligned homogeneous neighbourhoods, and therefore LC should be maximum when the structures in the images are aligned [95].

4.6.1.1 Comparison of MI and LC

In figure 4.22 we can see the results from calculating the respiration phase using LC for the example sequence used throughout this chapter, with a neighbourhood of 7×7 pixels and with a ROI that contains the diaphragm border (see figure 4.11a). Unlike figure 4.17b, the slope of the curve does not change linearly. The hysteresis cycle of the breathing process can be noticed in the interval $30 < MDD < 50$. The correlation coefficient between the LC-based respiration phase and the MDD is -0.9070, which indicates a worse performance than the MI in the same area, which gave a correlation coefficient of -0.9686. In our application, the interventional images do not present strong edges nor defined objects in the image that should be registered, so we do not expect LC to perform as well as in the applications

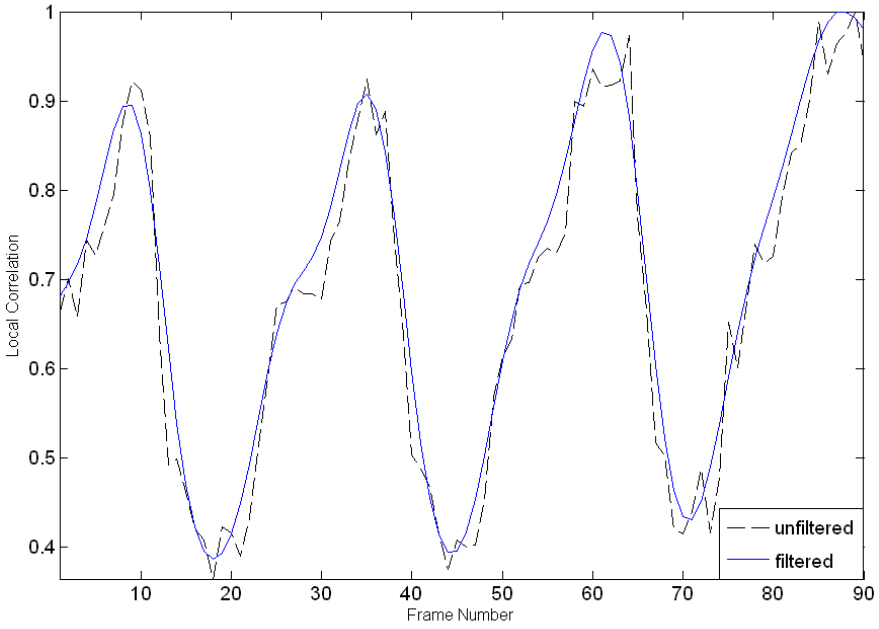

(a)Local Correlation

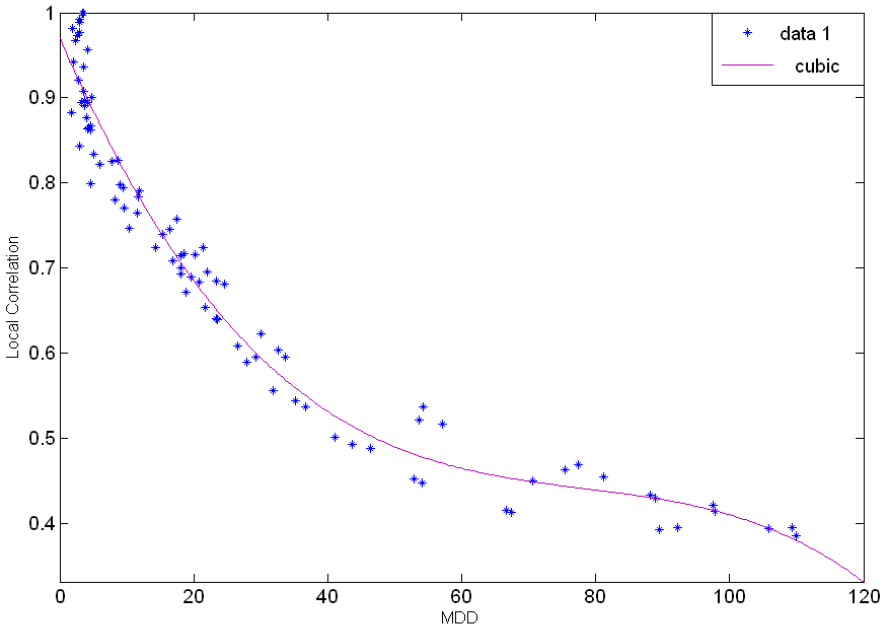

(b) LC *versus* MDD

Figure 4.22: LC evaluation for a region at the diaphragm border. (a) depicts the resulting respiration phase, and (b) compares this respiration phase against the MDD respiration phase.

presented in [95] (registration of CT-MR and MR images).

The performance of LC was compared to that of MI for two fluoroscopic sequences and two angiographic sequences for a manually selected ROI containing the diaphragm border (the area of the ROI was the same in all cases). In all four sequences, MI performed better. Considering that our bronze standard is the MDD, and that a good similarity measure should give a correlation coefficient of absolute value close to 1 when compared with the MDD in a ROI containing the diaphragm border, the MI gave for our test sequences a RMS of 0.1054, whereas LC resulted in a RMS of 0.1353. As expected, these results confirm that MI performs better than LC in the absence of defined edges in the image.

In figure 4.23 the influence of the neighbourhood size on the accuracy of the measured respiration phase is illustrated for all four test sequences. Figure 4.23 displays the correlation coefficient between the normalised LC respiration phase and the normalised MDD respiration phase *versus* the neighbourhood size. We can see that, the bigger the neighbourhood, the closer to -1 the value of the correlation coefficient, and therefore, the more linear the curve is. This can be explained by the good behaviour of LC in aligning neighbourhoods that contain edges. For comparing MI and LC, we placed the ROI at the diaphragm frontier and therefore two distinct areas with differentiated gray values contribute to the LC. This explains the improved behaviour as the neighbourhood size increases. Nevertheless, although performing the respiration phase calculation with a ROI favorable to LC, the results do not outperform those for MI as seen in section 4.5.2.

4.6.2 Information-theory-based similarity measures

Mutual information is a measure of the dependence between two stochastic variables based on Shannon's definition of entropy, which is a measure of the average uncertainty of a variable. *Jensen's divergence* is a measure of the divergence between two probability distributions and has been proposed for multimodality image registration in combination with alternative entropy definitions in [78][106][111][112].

The *Rényi entropy* [106][113] generalises Shannon's entropy definition by generalising the definition of the mean of the amount of information $I(p_X)$. The *Tsallis entropy* [114] is also a generalisation of Shannon's entropy used in physics, that describes non-extensive physical systems (i.e. systems in which

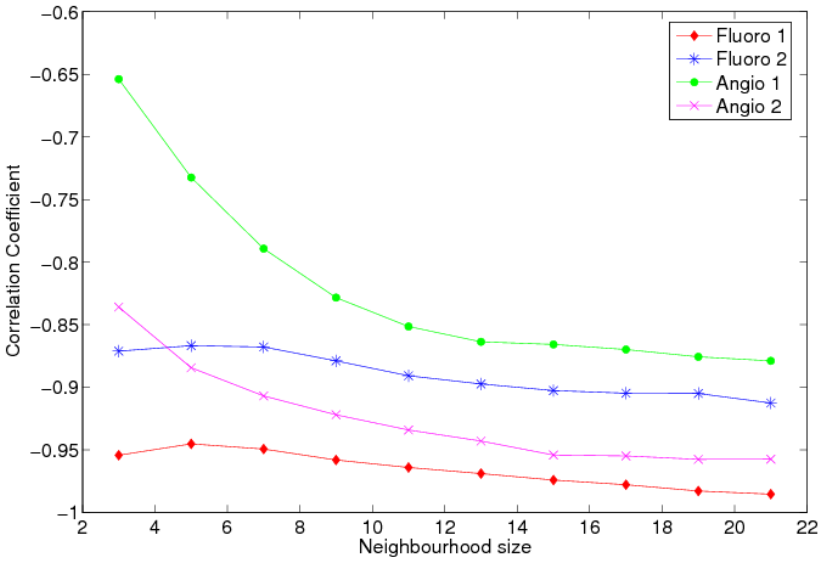

Figure 4.23: Correlation coefficients of the LC respiration for different neighbourhood sizes.

independence between random variables does not translate to the additivity of entropy). In subsections 4.6.2.1 and 4.6.2.2 we first introduce the definitions of Rényi and Tsallis entropies. In subsection 4.6.2.3 we introduce Jensen's divergence measure which has been proposed as a similarity measure for image registration [115][106].

Bardera *et al.* suggest in [78] four similarity measures for multimodality image registration based on Jensen's divergence employing Rényi and Tsallis entropies. These similarity measures are explained in 4.6.2.4. In subsection 4.6.2.5 these multimodality similarity measures have been evaluated and compared to MI as respiration phase estimation measures.

4.6.2.1 Rényi entropy

Shannon's entropy (see equation 4.5) $H(X) = -\sum_{x\in\chi} p_X(X = x)\log[p_X(X = x)]$, where $\chi = \{x_1, x_2 \dots x_k\}$, can be interpreted as the mean of the amount of information $I(p_X) = -\log[p_X(X = x)]$ on a stochastic variable. The Rényi entropy is a generalisation of Shannon's entropy in that it considers a more general definition of the mean function. The generalised or Rényi entropy is

defined as [78][106]

$$H_R^\alpha(X) = \frac{1}{1-\alpha} \log \sum_{x \in \chi} [p_X(X = x)]^\alpha \tag{4.12}$$

where

- $0 < \alpha < 1$ (for $\alpha > 1$, H_R^α is neither concave nor convex), and
- $\{p_X(X = x),\ x \in \chi\}$ is the probability distribution of the stochastic variable X.

The following information theory procedure that leads to the definition of the Rényi entropy as a generalised entropy measure (of which Shannon's entropy is just a specific case) has been taken from [113]. Given a discrete marginal probability distribution $\{p_X(X = x),\ x \in \chi\}$, this analytical procedure derives the Rényi entropy by assuming that:

1. an information measure $\mathcal{I}$ should be additive for two independent events,
2. $\mathcal{I}$ should only depend on p_X, and
3. if different amounts of information occur with different probabilities, the total amount of information is the average of the individual information weighted by their probabilities.

The first two conditions lead to

$$\mathcal{I}(p_i, p_j) = \mathcal{I}(p_i) + \mathcal{I}(p_j) \tag{4.13}$$

where $p_i = p_X(X = x_i)$ and $p_j = p_X(X = x_j)$ are the probabilities of two independent events $X = x_i$ and $X = x_j$ occurring. This is Cauchy's functional equation and has a unique class of solutions: $\mathcal{I}(p_X) = -\kappa \log(p_X)$. The constant κ is fixed by setting appropriate boundary conditions. Given two independent events, the worst case occurs when both are equally probable $p_i = p_j = 1/2$, in which case we need 1 bit of information to transmit the result of the experiment through a communication channel. By setting $\mathcal{I}(1/2) = 1$, we obtain $\kappa = 1$ and thus $\mathcal{I}(p_X) \equiv I(p_X)$.

To express mathematically the third assumption, a general mean is employed. Let f be a real function with an inverse f^{-1}. The generalised mean value of $\{y_1, y_2, \ldots, y_k\}$ is defined as

$$f^{-1}\left(\sum_{k=1}^{n} p_k f(y_k) \right) \tag{4.14}$$

Thus defining the mean and taking $y_k = \mathcal{I}(p_k)$ where $p_k = p_X(X = x_k)$, leads to the following form for the mean amount of transmitted information:

$$\mathcal{H}_R(p_X) = f^{-1}\left[\sum_{x\in\chi} p_X(X=x) f\Big(-\log[p_X(X=x)]\Big)\right] \tag{4.15}$$

It can be shown [113] that the function f can only be linear or exponential to comply with the first assumption on additivity. If f is linear, $f(y) = ay$, by substituting in equation 4.15 we obtain Shannon's information measure or Shannon entropy:

$$\mathcal{H}_R(p_X) = -\sum_{x\in\chi} p_X \log(p_X) \equiv H \tag{4.16}$$

If the function f is exponential, $f(y) = C \cdot (2^{(1-\alpha)y} - 1)$ (C is a constant), it leads to

$$\mathcal{H}_R(p_X) = \frac{1}{1-\alpha}\log\left(\sum_{x\in\chi} p_X^{\alpha}\right) \equiv H_R^{\alpha}(X) \tag{4.17}$$

The parameter α controls the weight given to the probability distribution and when $\alpha \to 1$, H_R^{α} reduces to Shannon's entropy. Thus, the Rényi information measure or entropy is a generalisation of the Shannon entropy. For $\alpha < 1$, Rényi entropy is concave as a consequence of $log_2(x)$ and x^{α} being concave for $\alpha < 1$.

4.6.2.2 Tsallis entropy

The Tsallis entropy was introduced by Tsallis to describe non-extensive physical systems [112], i.e. systems for which independence does not translate into the sum of entropies. Therefore the first assumption made when deriving the Rényi entropy in the previous subsection does not hold, and is substituted by a pseudo additive axiom

$$\mathcal{I}(p_i, p_j) = \mathcal{I}(p_i) + \mathcal{I}(p_j) + (1-\alpha)\mathcal{I}(p_i)\mathcal{I}(p_j) \tag{4.18}$$

for two independent events. The parameter α is a measure of the degree of non-extensivity.

The Tsallis-Havrda-Charvát entropy is then defined as [78][116]

$$H_T^{\alpha}(X) = \frac{1}{\alpha-1}\left(1 - \sum_{x\in\chi} p_X^{\alpha}\right) \tag{4.19}$$

where

- $\alpha > 0$ and $\alpha \neq 1$, and
- $\{p_X(X = x),\ x \in \chi\}$ is the probability distribution of the stochastic variable X.

When $\alpha \to 1$, $H_R^\alpha(X) \to H(X)$ [78], the Shannon entropy.

4.6.2.3 Jensen's divergence

In the previous sections mutual information and entropy have been defined. These are measures of the dependance of two statistical distributions. The Jensen divergence is a measure of the divergence of two probability distributions.

Let $\{\boldsymbol{p_1}, \ldots, \boldsymbol{p_n}\}$ be n probability distributions. The generalised Jensen-Shannon divergence is a measure of the divergence between probability distributions and is defined as

$$JS(\boldsymbol{p_1, \ldots, p_n}) = H\Big(\sum_{i=1}^{n} \omega_i \boldsymbol{p_i}\Big) - \sum_{i=1}^{n} \omega_i H(\boldsymbol{p_i}) \tag{4.20}$$

where $\{\omega_i \ / \ \omega_i > 0, \quad \sum_i \omega_i = 1\}$ is a set of weights. $JS(\boldsymbol{p_1, \ldots, p_n})$ is non-negative, convex, symmetric, and equals 0 if and only if $\boldsymbol{p_1} = \cdots = \boldsymbol{p_n}$ [117]. The first term in equation 4.20 calculates the entropy of a weighted mean of all probability distributions, whereas the second term in the equation calculates a weighted sum of the entropy of each individual distribution.

If instead of the Shannon entropy in equation 4.20 we employ the Rényi entropy (equation 4.12), we obtain the Jensen-Rényi divergence measure [106][115]:

$$JR_\alpha^\omega(\boldsymbol{p_1, \ldots, p_n}) = H_R^\alpha\Big(\sum_{i=1}^{n} \omega_i \boldsymbol{p_i}\Big) - \sum_{i=1}^{n} \omega_i H_R^\alpha(\boldsymbol{p_i}) \tag{4.21}$$

Similarly, Tsallis entropy (equation 4.19) can be substituted in equation 4.20 to obtain

$$JT_\alpha^\omega(\boldsymbol{p_1, \ldots, p_n}) = H_T^\alpha\Big(\sum_{i=1}^{n} \omega_i \boldsymbol{p_i}\Big) - \sum_{i=1}^{n} \omega_i H_T^\alpha(\boldsymbol{p_i}) \tag{4.22}$$

Equations 4.21 and 4.22 are divergence measures between a finite number of weighted probability distributions [106]. JS, JR_α^ω and JT_α^ω as disparity

measures reach their maximum value when $\boldsymbol{p_1, \ldots, p_n}$ are degenerate distributions, i.e. $\boldsymbol{p_i} = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise [106].

If $\mathcal{X}$ and $\mathcal{Y}$ are the set of intensity values of images B and T and X and Y are two random variables taking values in $\mathcal{X}$ and $\mathcal{Y}$ respectively, JR_α^ω and JT_α^ω can be used for image registration if $\boldsymbol{p_i} = (p_{ij})_{1 \leq j \leq n}$, where p_{ij} is defined as [106][78]

$$p_{ij} = P(Y = y_j | X = x_i),\ j = 1, 2, \ldots, n \tag{4.23}$$

By thus defining the probability distribution, the similarity measure is dependent on the joint histogram and the Jensen divergence as defined in equations 4.20, 4.21 and 4.22 becomes the *Generalised Mutual Information* as proposed by Bardera *et al.* [78] and explained in the next subsection. The Jensen divergence measure presents the advantage of being a weighted measure of the probability distributions (or image histograms) and thus allowing to control the measurement sensitivity to the joint histogram [115][118].

4.6.2.4 Generalised Mutual Information

Bardera *et al.* propose in [78] normalised similarity measures between two images, B and T, based on equations 4.21 and 4.22, where the weights are the *a priori* probabilities, $w_i = P(X = x_i)$, and $\boldsymbol{p_i}$ are the rows of the conditional probability matrix $p_{ij} = P(Y = y_j \mid X = x_i)$ calculated from the histogram of both images. Thus, the *Generalised Mutual Information* (GMI) between images B and T is defined as

$$GMI_h^\alpha(B;T) = H_h^\alpha\Big(\sum_{i=1}^{n} w_i p_{ij}\Big) - \sum_{i=1}^{n} w_i H_h^\alpha(p_{ij}) \tag{4.24}$$

where

- $w_i = P(X = x_i)$,
- $p_{ij} = P(Y = y_j | X = x_i)$ are the rows of the probability conditional matrix, and
- h denotes either the Rényi ($h = R$) or the Tsallis ($h = T$) entropy.

When images B and T are exactly matched, $\boldsymbol{p_i} = (\delta_{ij})_{1 \leq j \leq n}$ is a degenerate distribution. The parameter α controls the sensitivity of the similarity measure, and the weights ω_i can be adapted to each particular image registration problem. If $\alpha = 1$ and $\omega_i = P(X = x_i)$, the Jensen-Rényi is the generalised

Jensen-Shannon divergence (see equation 4.20)[106].

In [78] two ways of normalising GMI in equation 4.24 are proposed. One normalisation is performed by the marginal distributions. The resulting measure is the *generalised correlation* and is defined as

$$nGMI_h^\alpha = \max\left\{ \frac{GMI_h^\alpha(B;T)}{H_h^\alpha(T)}, \frac{GMI_h^\alpha(B;T)}{H_h^\alpha(B)} \right\} \tag{4.25}$$

where

- B and T denote the images to be registered,
- $\alpha > 0$ and $\alpha \neq 1$, and
- h denotes either the Rényi or the Tsallis entropy.

When $\alpha \to 1$, then $nGMI_{R,T}^1 \to \frac{I(B;T)}{min(H(B),H(T))}$.

Another possibility is to normalise the GMI by the joint entropy, $H(B,T)$. This normalisation is given by

$$NGMI_h^\alpha = \frac{\max\{GMI_h^\alpha(B;T), GMI_h^\alpha(T;B)\}}{H_h^\alpha(B,T)} \tag{4.26}$$

When $\alpha \to 1$, then $NGMI \to \frac{I(B;T)}{H(B,T)}$.

4.6.2.5 Evaluation

The performance of nGMI and NGMI to measure the respiration phase was evaluated for the Rényi and Tsallis entropies for two diagnostic angiographic sequences and two interventional fluoroscopy sequences. The value of α was chosen based on the results in [78], who suggest $\alpha = 0.6$ for the Rényi entropy, and $\alpha = 1.5$ for the Tsallis entropy. In figure 4.24a the resulting respiration phase for $nGMI_R^{0.6}$ is shown. The resulting fitted curve, figure 4.24b, is very similar to that in figure 4.17b but the correlation coefficient of $nGMI_R^{0.6}$ has a lower absolute value, 0.9313. For the respiration phase calculated as $nGMI_T^{1.5}$ in figure 4.25a, the absolute value of the correlation coefficient of the points in figure 4.25b is 0.8637.

To establish the optimum value of the parameter α, $nGMI_R^\alpha$ for the interval $\alpha \in [0,1]$ and $nGMI_T^\alpha$ for the interval $\alpha \in [0,2]$ was calculated. The correlation coefficients are plotted for $nGMI_R^\alpha$ and $nGMI_T^\alpha$ in figures 4.26

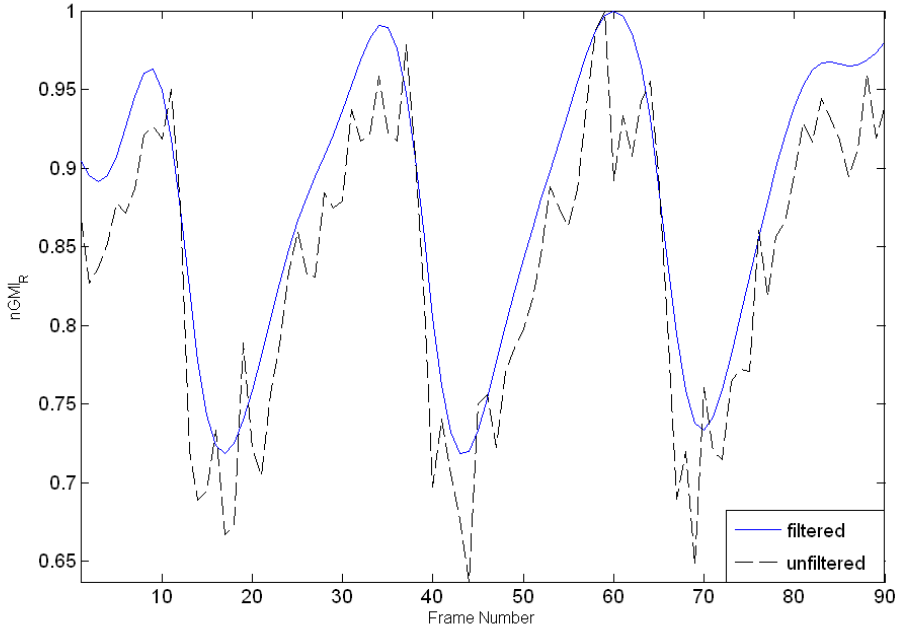

(a) nGMI Rényi.

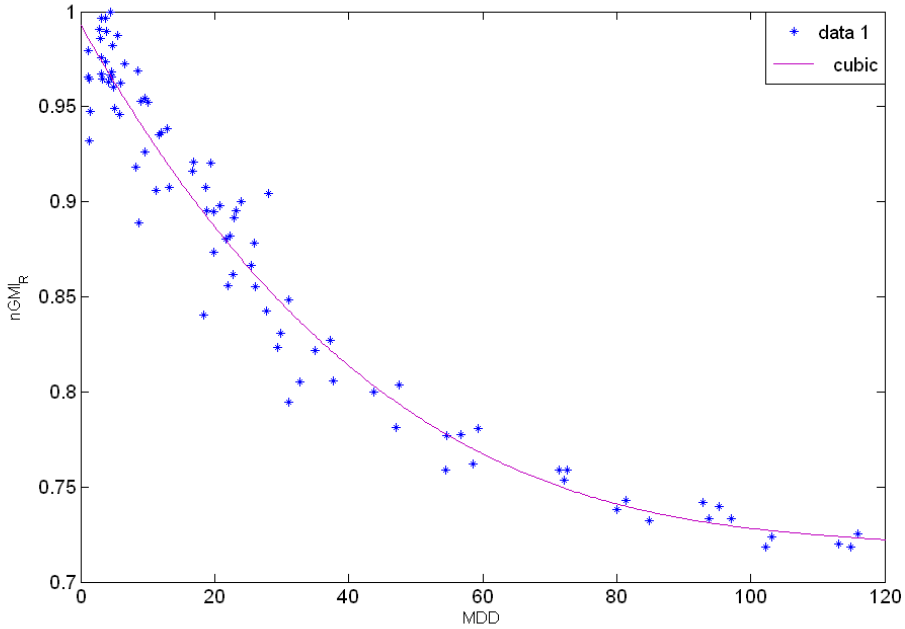

(b) nGMI Rényi *versus* MDD.

Figure 4.24: (a) Respiration phase calculated for $nGMI_R^{0.6}$ for an area on the diaphragm border. (b) The respiration phase in (a) is compared to the MDD. The points above the fitted cubic line do not correspond only to inspiration nor expiration, therefore a hysteresis cycle is not responsible for the scatter in the points.

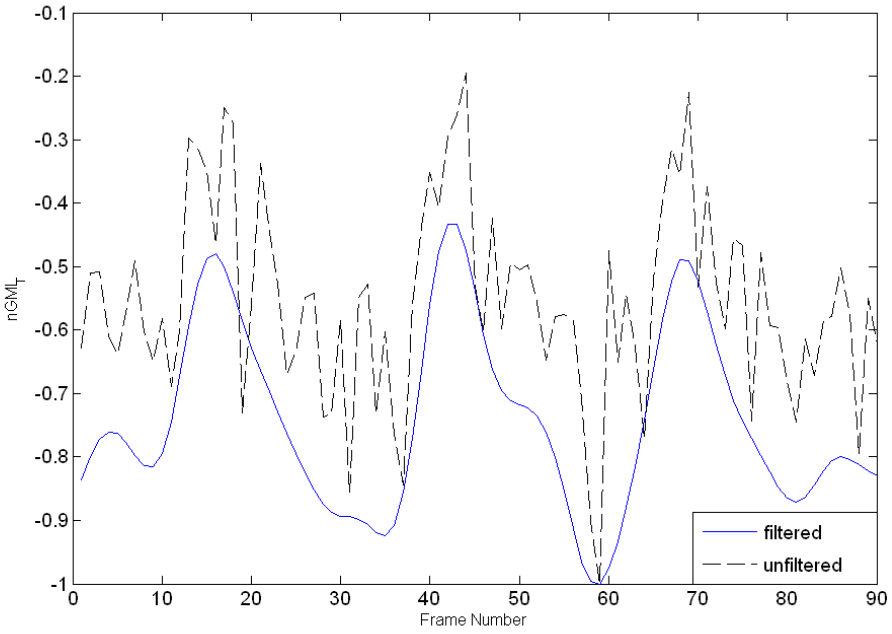

(a) nGMI Tsallis.

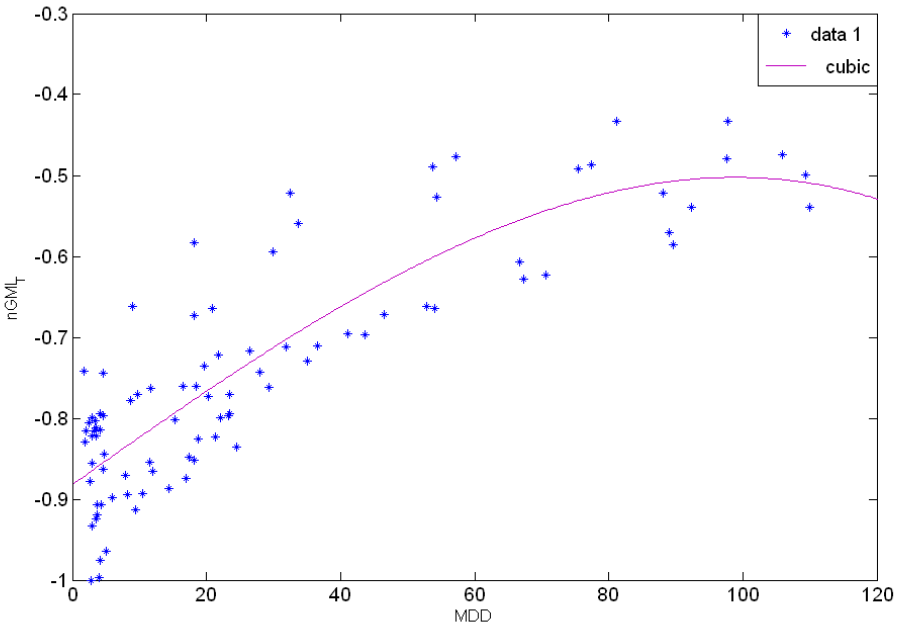

(b) nGMI Tsallis *versus* MDD.

Figure 4.25: (a) Respiration phase calculated for $nGMI_T^{1.5}$ for an area on the diaphragm border. (b) The respiration phase in (a) is compared to the MDD. The points above the fitted cubic line do not correspond only to inspiration nor expiration, therefore a hysteresis cycle is not responsible for the scatter in the points.

and 4.27 respectively. From these figures we conclude that the optimum performance for $nGMI_R^\alpha$ and $nGMI_T^\alpha$ is reached when $\alpha = 1$, in which case, $nGMI_{R,T}^1 = I(B;T)/\min(H(B), H(T))$. This suggests that weighting the probability distribution by α is unnecessary and does not result in a better similarity measure for respiration phase calculation than MI. In table 4.2 we show the correlation coefficients for all four sequences if the respiration phase is calculated in a ROI containing the diaphragm frontier using MI and $nGMI_{R,T}^1$. We conclude that MI is better suited for measuring the respiration phase from a sequence.

Sequence	$nGMI_{R,T}^1$	MI
Sequence 1	0.9545	0.9686
Sequence 2	0.9367	0.9539
Sequence 3	0.9209	0.9095
Sequence 4	0.5902	0.9365

Table 4.2: Correlation coefficients for MI and $nGMI_{R,T}^1$

In figures 4.28 and 4.29 respectively the results for $NGMI_R^{0.6}$ and $NGMI_T^{1.5}$ are shown. For $NGMI_R^{0.6}$ the correlation coefficient of the resulting respiration phase, figure 4.28a, compared to MDD (figure 4.28b) is -0.9241. The results when the Tsallis entropy and $\alpha = 1.5$ are used, $NGMI_T^{1.5}$, are shown in figures 4.29a and 4.29b. The correlation coefficient in figure 4.29b is -0.8676.

The optimum value of the parameter α for $NGMI_R^\alpha$ and $NGMI_T^\alpha$ was explored. The correlation coefficients are plotted for $NGMI_R^\alpha$ and $NGMI_T^\alpha$ in figures 4.30 and 4.31 respectively. In both cases, the highest absolute correlation coefficient value is obtained for $\alpha = 1$, in which case $NGMI = I(B;T)/H(B,T)$. Table 4.3 shows a comparison between the correlation coefficients between the MI and $NGMI_{R,T}^1$ similarity measures and the MDD for all four sequences. Except in one case, the MI gives better results for calculating the respiration phase.

4.7 Conclusions

Current state of the art does not address respiration motion artefacts during image acquisition, so that current instrument navigation and positioning techniques rely on the display of static angiograms. It has been proposed

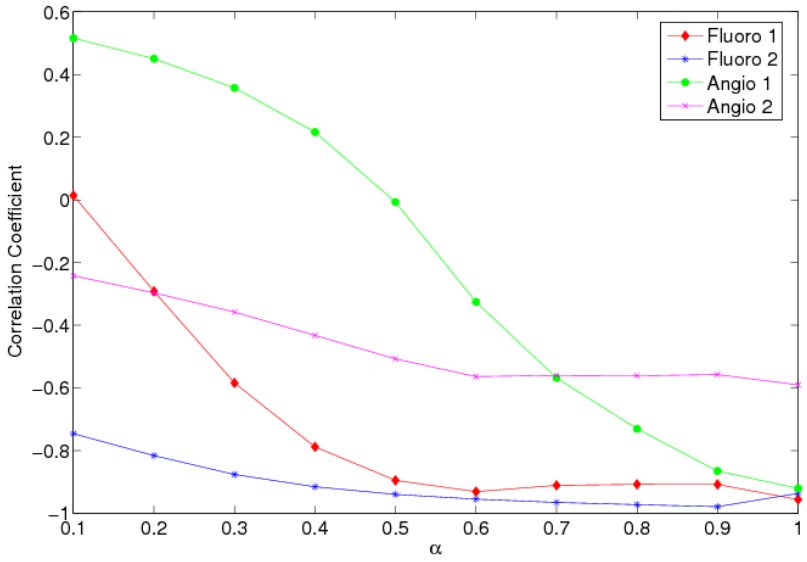

Figure 4.26: Optimum α for the Rényi generalised correlation $nGMI_R^\alpha$ for the four test sequences.

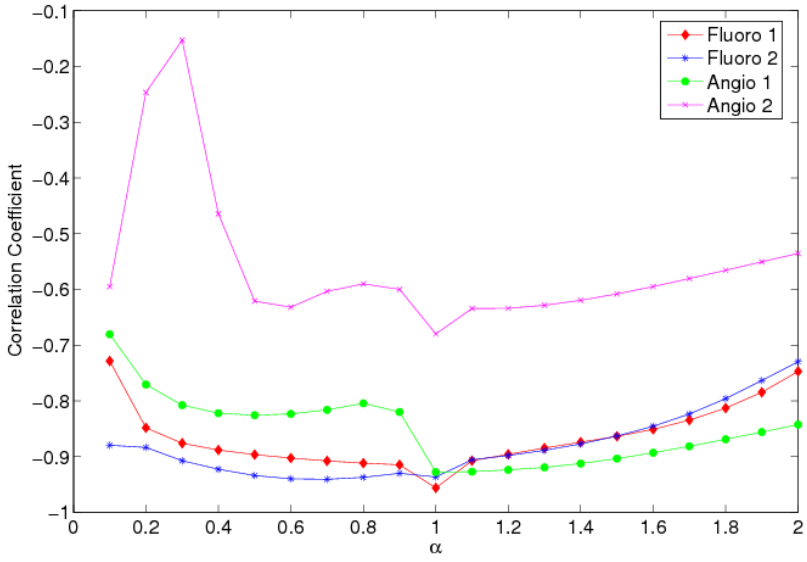

Figure 4.27: Optimum value of the parameter α for the Tsallis generalised correlation $nGMI_T^\alpha$ for four test sequences.

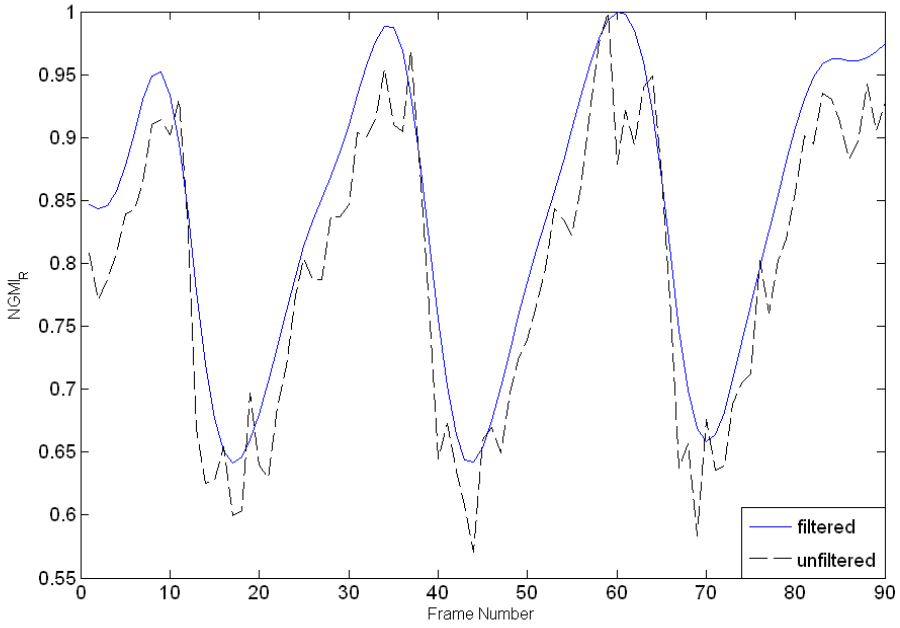

(a) NGMI Rényi

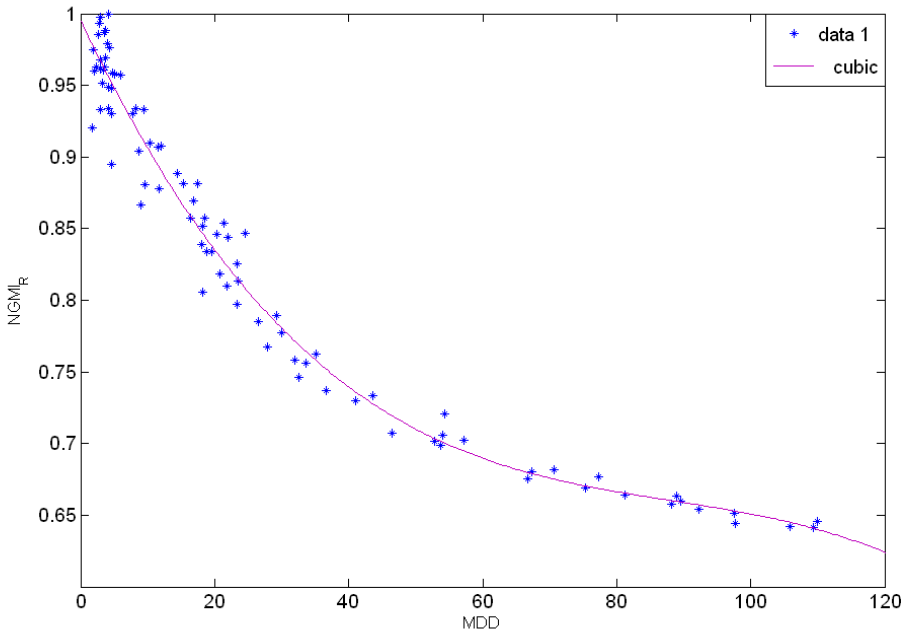

(b) NGMI Rényi *versus* MDD

Figure 4.28: (a) Respiration phase calculated for $NGMI_R^{0.6}$ for an area on the diaphragm border. (b) The respiration phase in (a) is compared to the MDD.

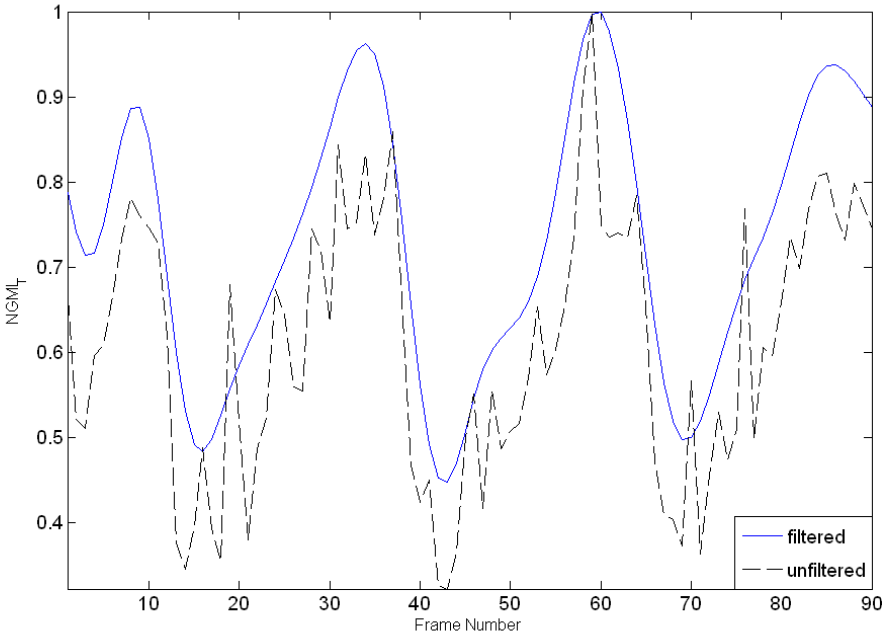

(a) NGMI Tsallis.

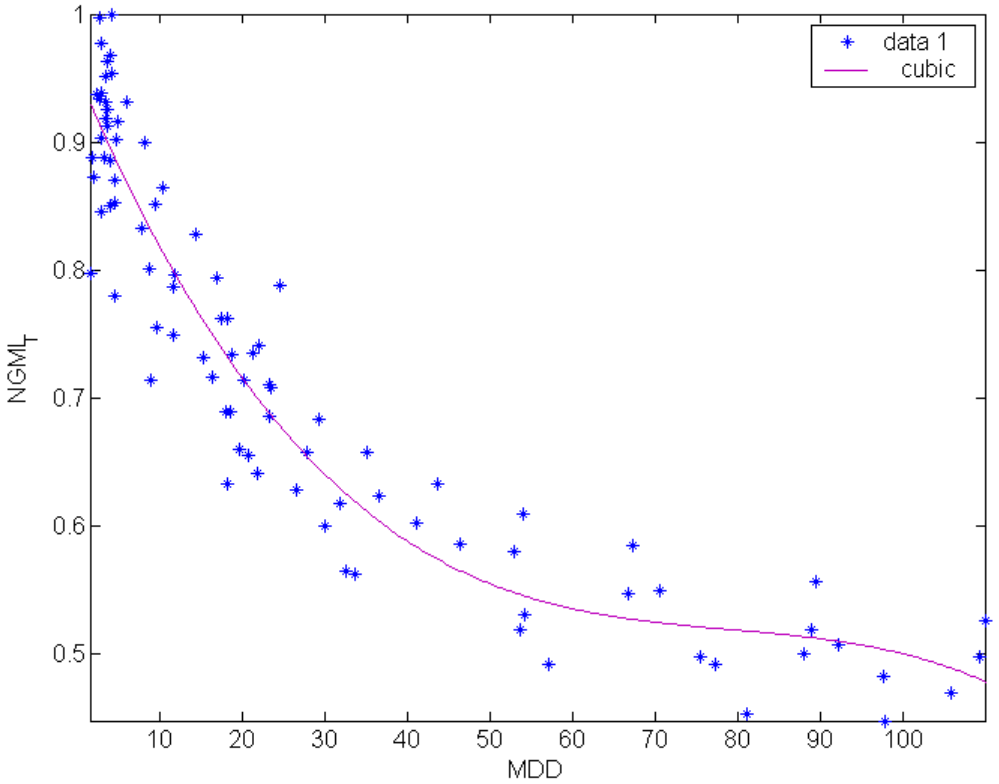

(d) NGMI Tsallis *versus* MDD.

Figure 4.29: (a) Respiration phase calculated for $NGMI_T^{1.5}$ for an area on the diaphragm border. (b) The respiration phase in (a) is compared to the MDD.

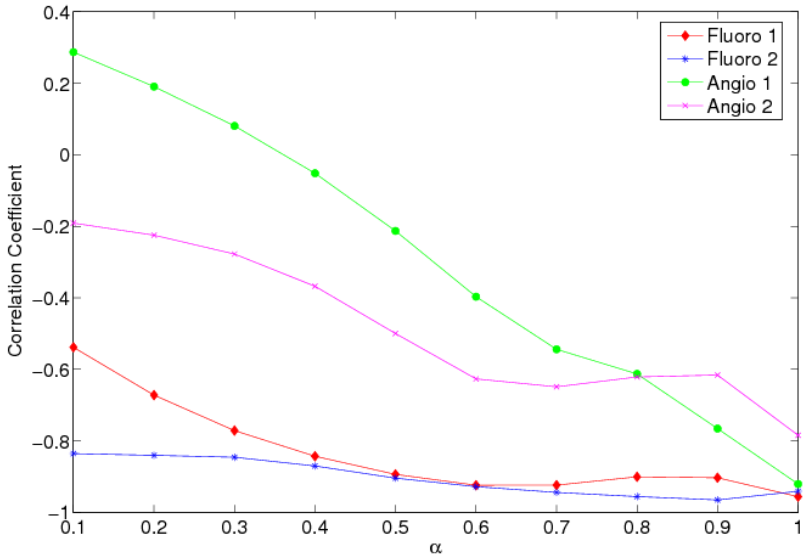

Figure 4.30: Optimum value of the parameter α for $NGMI$ calculated with the Rényi entropy for four test sequences.

Sequence	$NGMI^1_{R,T}$	MI
Sequence 1	0.9561	0.9686
Sequence 2	0.9413	0.9539
Sequence 3	0.9207	0.9095
Sequence 4	0.7843	0.9365

Table 4.3: Correlation coefficients for MI and $NGMI^1_{R,T}$

to characterise each image in an angiographic or fluoroscopic acquisition by its heart and respiration phase. A 2-dimensional phase-space graph is used to locate the most similar angiogram for each fluoroscopic frame for overlay and thus aid navigation and instrument positioning in PTCA interventions. Novel algorithms for calculating an ECG-derive heart phase and an image content derived respiration phase have been proposed in this chapter.

An algorithm to derive a heart phase from the ECG signal is described in section 4.1. Our method detects R peaks in the ECG signal and calculates the heart phase as the normalised distance to the previous R peak. The possibility of an alternative definition of the heart phase, to double the sampling of the phase-space along the heart phase axis, was explored in section 4.1.1.

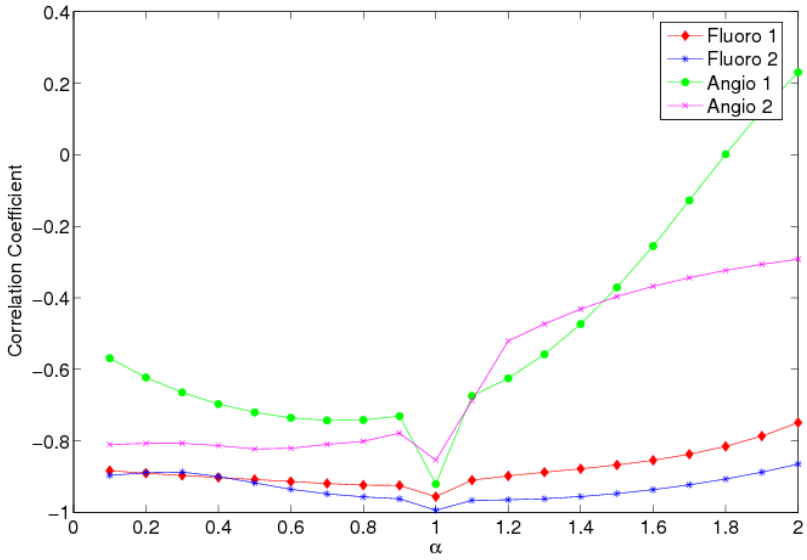

Figure 4.31: Optimum value of the parameter α for $NGMI$ calculated with the Tsallis entropy for four test sequences.

It was concluded that using the proposed alternative heart phase definition would render artificial-looking angiographic masks to experienced users but might still be useful when lacking a matching angiogram to be registered with the current fluoroscopic frame.

A novel approach to calculate the respiration phase from the image content was presented in section 4.5. This method measures the similarity between each frame in an angiography or fluoroscopy sequence and a reference using the mutual information measure introduced in [77]. The similarity measure is affected by breathing and heart contraction, which is removed by frequency analysis with a low-pass FIR filter with cutoff frequency 1Hz. Also the reference frame must be located in the extremes of the respiration cycle in order to avoid an artificial doubling of the respiration phase curve frequency.

The performance of MI in different areas of the image was compared with a bronze standard in subsection 4.5.2. It was established that although the diaphragm is the best area to measure the respiration phase, other areas of the image are also affected by breathing and the proposed similarity measure is able to measure the breathing state of the patient. By comparing the ef-

fects of breathing in different areas of the image, it can be concluded that, given a certain heart phase, breathing mainly causes an affine transformation on the heart as suggested by [20].

Mutual information for respiration phase measurement was compared to other multimodality measures in section 4.6: local correlation [95][109][110], normalised generalised correlation [78] and normalised generalised MI [78]. MI proved to have better performance than LC, $nGMI_h^\alpha$ and $NGMI_h^\alpha$ for the specified application.

Chapter 5

Accurate device to vessel registration

In the previous chapter, a method to select a matching angiogram from a collection of previously acquired diagnostic frames $\{\mathcal{A}\}$, for the overlay of said angiogram onto the current interventional frame to aid navigation was proposed. This method assumes that given a certain imaging geometry, and in the absence of table or patient movement, each image can be characterised by a heart and a respiration phase. Images with similar heart and respiration phases present the heart in the same state, and can therefore be overlaid to assist navigation. The search for a proper roadmap is aided by a 2D representation or phase-space, in which nearby points correspond to similar frames in terms of heart and respiration phase. During the intervention, each incoming frame is characterised in real-time by its heart and respiration phase and represented as a point in the phase-space. Then, the closest point corresponding to a previously acquired angiogram is identified. If found, the corresponding roadmap is overlaid.

Because the phase-space is sparsely sampled, an angiographic roadmap similar to each interventional frame might not be available. As already mentioned in section 3.4, images with the same heart phase but different respiration phase present the heart at a different position within the thorax. It has been proposed to model the heart movement caused by respiration as a rigid body translation [69][73]. Eck *et al.* propose in [73] a method to estimate the heart-position function from a finite number of frames with the same heart phase but different respiration phases. A translation function parameterised by the respiration phase is obtained, so that a rigid-body transformation of the roadmaps can be calculated for respiration phase compensation. This method will be summarised in section 5.1.

Nevertheless, due to involuntary movements of the patient, internal movement of the organs and soft tissue, and small discrepancies in the heart or respiration phase, the registration between an artery, visible in the angiogram, and the instrument, visible in the interventional frame, when an overlay is displayed is not always precise. An accurate registration would be useful to improve instrument positioning, aid detecting vessel wall perforations, or identifying the vessel that currently contains the instrument when more candidates are nearby. Such adjustment of the overlay can be achieved by further registering the vessel and the instrument. The proposed method requires vessel segmentation, guidewire and device marker segmentation and a registration algorithm, which will be presented in this chapter.

We propose a novel vessel filter, *directional stamping*, in section 5.2 for segmenting dark elongated structures like vessels or guidewires. A method to segment small moving dark device markers on interventional frames is presented in section 5.3. Section 5.4 defines a new multiplicative distance transform that applied to the result of a vessel probability filter like directional stamping provides a potential field around the detected vessels that attracts the segmented device and registers the vessel and the device. In section 5.5 we present the results of the vessel, and guidewire and device segmentation filters and the guidewire and device to vessel registration method.

5.1 Respiration compensation

Typically, the number of samples along the heart and respiration phase axes (δ and β respectively) in the phase-space is not high, i.e. diagnostic frames in all possible respiration and heart phase combinations are not available. Therefore, it might occur that we do not find a good match between the current fluoroscopic frame characterised by $\vec{\xi_f}(\delta_f, \beta_f)$ and any previously acquired angiographic roadmap. For general navigation purposes, this is not critical. But when tortuous vessel geometry is present, or two or more vessels run parallel to each other, an imperfect angiographic overlay may convey the wrong information to the physician. Two frames with the same heart phase but different respiration phase present the heart at a different position (see figure 3.5). Respiration induced movement is periodic and can be modelled as a translation [69][73].

Eck *et al.* propose in [73] a method for respiration compensation. Given

a fluoroscopic frame characterised by $\vec{\xi}(\delta_f, \beta_f)$, an angiogram with the closest state vector $\vec{\xi}(\delta_a, \beta_a)$ is selected and the residual difference in respiration phase $\Delta\beta = \beta_f - \beta_a$ is then compensated for. Their method assumes that breathing induces a rigid-body translation of the heart. By comparing pairs of diagnostic angiographies with similar heart phases but different respiration phases, they calculate the parameters of a rigid-body transformation to compensate for the difference in respiration phase for each pair. Then, they propose a piecewise iterative approximation to reconstruct the heart-position function from the measured individual parameters, so that angiographic images of arbitrary respiration statuses can be synthesised by calculating a translation vector.

5.1.1 Respiration-induced translations

Let $A_{\delta,\beta_1}(x, y)$ and $A_{\delta,\beta_2}(x, y)$ be two angiograms with the same heart phase δ but different respiration phases, β_1 and β_2. x and y denote the pixel coordinates in the images. If the motion induced by breathing can be modelled by a rigid-body translation, then there exist m_x and m_y such that

$$A_{\delta,\beta_1}(x, y) = A_{\delta,\beta_2}(x + m_x, y + m_y) \tag{5.1}$$

where $\vec{m} = (m_x, m_y)$ is a translation vector that depends on the respiration phase difference $\Delta\beta = \beta_2 - \beta_1$.

A number of translation vectors $\vec{m}$ are calculated from pairs of diagnostic angiograms with sufficient contrast agent [119] and similar heart phases as the translation that maximises the cross-correlation between salient regions. Because the assumed transformation is a translation, only one salient region is necessary. More regions should be used if the transformation between two different respiration phases were modelled as an affine motion.

These salient regions are typically vessel bifurcations or bends and are identified with the following algorithm:

1. The angiograms are pre-processed to enhance the elongated vessel structures using the method described in [120].

2. The images are Radon-transformed and local maxima, that correspond to straight lines, are identified.

3. The intersections of the identified lines are calculated in the spatial domain and weighted by the line strength and the intersection angle.

90 ° intersections are considered good intersections, whereas 0 ° or 180 ° intersections are considered not suitable.

4. The three highest-ranking intersection candidates are tested in the spatial domain by calculating the two-dimensional auto-correlation of the region surrounding each candidate intersection. The region with the most prominent auto-correlation peak is chosen as salient region to calculate the translation $\vec{m}$ that maximises the cross-correlation between salient features.

Let $\vec{r}$ be the position of the heart in frame $I_{\delta,\beta}(x, y)$, where the heart position $\vec{r}$ is a symmetric and monotonic function of the respiration phase β, $\vec{r} = \vec{r}(\beta)$. Then the calculated translation vector $\vec{m} = (m_x, m_y)$ for the chosen salient region is related to the heart position by

$$\vec{m} = \vec{r}_1 - \vec{r}_2 \tag{5.2}$$

where $\vec{r}_1$ and $\vec{r}_2$ are the position of the heart in two angiograms with the same heart phase δ acquired at two different time instants, $A_{\delta,\beta_1}(x, y)$ and $A_{\delta,\beta_2}(x, y)$. Because we are assuming a rigid-body translation, it is irrelevant for our application how the heart position vector is defined. Whether defined as the position of the salient region chosen for calculating $\vec{m}$ or the centre of mass of the heart, assuming respiration causes a rigid-body displacement, any of these definitions of $\vec{r}$ would render the same translation result.

5.1.2 Reconstruction of the heart-position function

Because each translation vector can be expressed as in equation 5.2, they can be regarded as an approximation of the slope of $\vec{r}$ between two different instants t_1 and t_2. Eck *et al.* method reconstructs $\vec{r}$ from a number of translation vectors $\{\vec{m}_k, \quad k = 1, 2, \dots\}$ by iteratively estimating the slope of $\vec{r}$, $\tilde{r}$. Function $\vec{r}$ is decomposed into its independent x and y components and the following reconstruction method is applied to each component, r_x and r_y, separately.

The initial assumption is $\tilde{r}_x = 0$ and $\tilde{r}_y = 0$. In the following we omit the subscript x or y for simplicity of notation so that $\tilde{r}$ refers to either $\tilde{r}_x$ or $\tilde{r}_y$ and m refers to m_x or m_y. For each iteration n, the corresponding component of the measured translation vector $\vec{m} = (m_x, m_y)$ between a pair

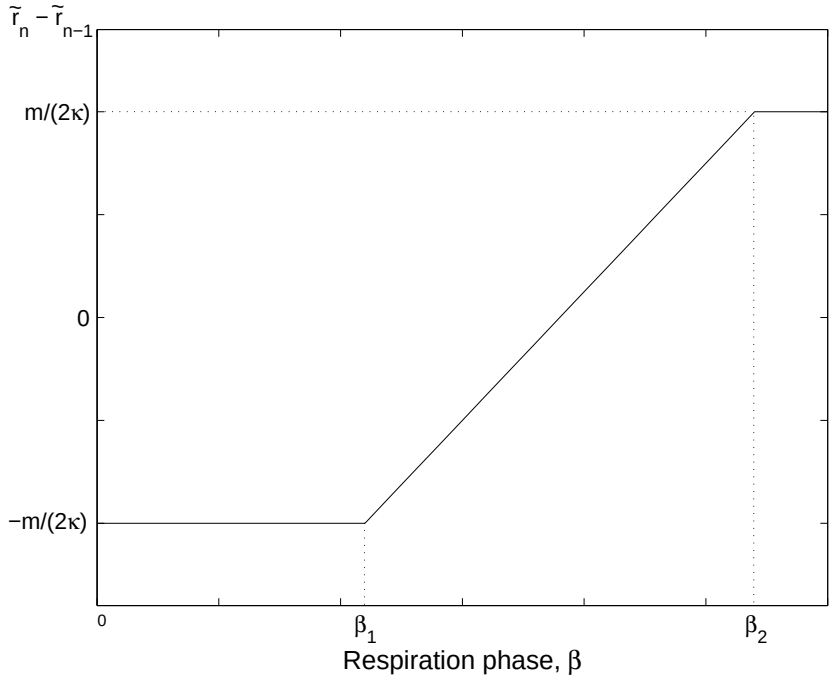

Figure 5.1: Reconstruction of the heart-position function. Representation of $\tilde{r}_n - \tilde{r}_{n-1}$ in equation 5.3.

of angiograms in respiration phases β_1 and β_2 is iteratively included to refine the assumption on $\tilde{r}$, $\tilde{r}_n$, following the equation

$$\tilde{r}_n = \begin{cases} \tilde{r}_{n-1} - m/(2 \cdot \kappa) & 0 \leq \beta \leq \beta_1 \\ \tilde{r}_{n-1} + \frac{m}{\kappa}\left(\frac{\beta-\beta_1}{\beta_2-\beta_1} - \frac{1}{2}\right) & \beta_1 < \beta \leq \beta_2 \\ \tilde{r}_{n-1} + m/(2 \cdot \kappa) & \beta_2 < \beta \leq 2\pi \end{cases} \tag{5.3}$$

where κ is a damping factor that depends on the number of samples available and the noise level of the input data pairs. For each component, the difference in position m is evenly distributed in the interval $[\beta_1, \beta_2]$ and the function $\tilde{r}_n$ is updated outside that interval by $\pm m/(2\kappa)$. This iterative piecewise linear approximation is illustrated in figure 5.1.

In figure 5.2a we can see the result of approximating the y component of a position function (solid line), with the shape of the raising part of a cosine function using equation 5.3. The simulation data contains a set of 30 randomly placed respiration states, $\{\beta_i,\ i = 1 \ldots 30\}$ and $r_y(\beta_i)$. Figure 5.2b shows pairs of m_y for the respiration function in figure 5.2a where the

lines connect pairs of samples with different respiration states and same heart phase. The dotted line in figure 5.2a is the result of the approximation after incorporating one m_y, and the dashed line is the result after incorporating twenty values of m_y.

The accuracy of $\tilde{r}_n$ depends strongly on the accuracy of the estimation of the arguments β_1 and β_2 and on the number of data pairs that are available for reconstruction. Choosing a damping factor κ extremely small leads to unstable results and non-converging behavior, choosing κ too big leads to systematic underestimation of the necessary changes that lead from the starting function hypothesis to the underlying function.

5.1.3 Respiration compensation results

Eck *et al.* tested their method of iterative reconstruction of $\vec{r}$ on a diagnostic sequence containing 32 angiograms with sufficient contrast dye (the amount of contrast agent was evaluated manually). For all image pairs whose heart phase differed less than $\pi/15$, the translation vector $\vec{m}$ was calculated by maximizing the cross-correlation of one salient region.

With their method, Eck *et al.* report a matching accuracy of 3.3 pixels [73], which is still not enough for accurate instrument to vessel registration. Possible reasons for this are that:

- they assume that breathing induces a rigid body translation on the heart, whereas it causes an affine transformation as established by [20] and [71] and our results in section 4.5.2, and
- they do not consider an hysteresis cycle during breathing, which contradicts the results in [52], our bronze standard measurement in figure 4.15 and our results from figures 4.18 and 4.19 in section 4.5.2.

This method is nevertheless useful to circumvent the problem of sparse sampling of the phase-space along the respiration-phase axis and as a starting point for further fine-tuning of instrument to vessel registration. To provide accurate device to vessel matching another method based on local image features is needed. A method we propose in [32][121] is explained in section 5.4.

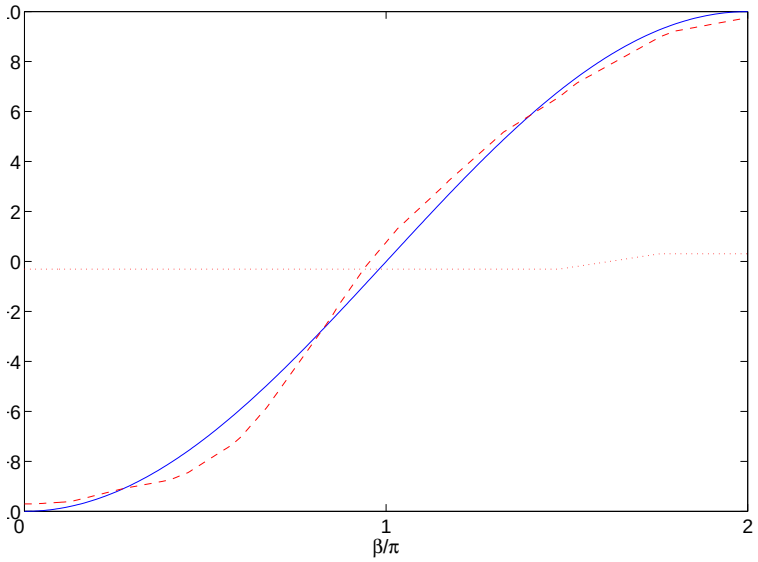

(a)

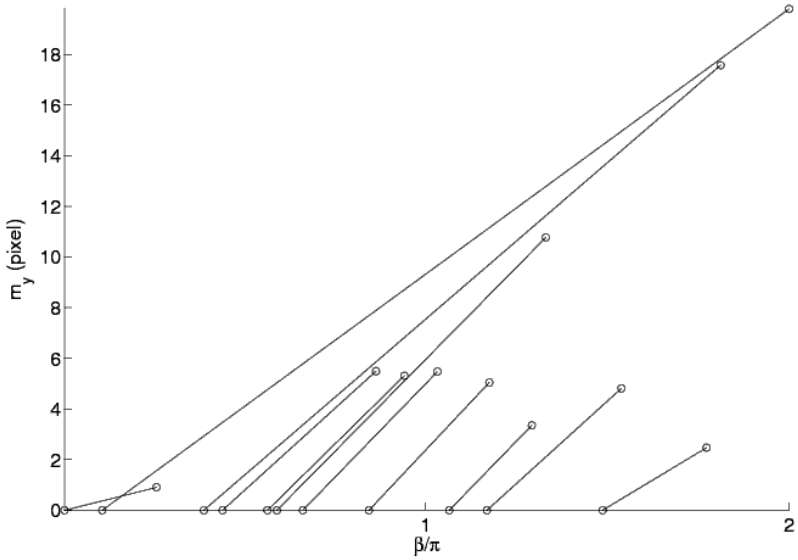

(b)

Figure 5.2: (a) Reconstruction of a simulated position function from pairs of sampled respiration states and displacements. (b) Pairs of m_y for the respiration function shown above (solid line), where the the lines connect two samples in different respiration states.

5.2 Vessel and guidewire segmentation

An overlay of a diagnostic roadmap $A_{\delta,\beta}(x, y)$ over an interventional frame $F_{\delta,\beta+\Delta\beta}(x, y)$ might not result in an image where the instrument can be seen inside the corresponding vessel, even if the respiration phase is very similar ($\Delta\beta \to 0$) or has been compensated for. This can be explained by a slight movement of the patient or the patient's organs, or by inaccuracies in the respiration measurement or compensation algorithm.

Our goal is to obtain a device to vessel registration algorithm robust to segmentation errors that works in frame rate and bases on extracted image features, taking into account that the instrument, which is visible in the interventional frames, must lie within one of the vessels visible in the diagnostic images. We propose in [32] and [121] a novel method that detects vessel-like structures in an angiographic roadmap and applies a distance transform to allow for fast registration of the segmented instrument and the transformed vessel tree. Because the instrument is already located nearby the corresponding vessel in the images, fast convergence times of the algorithm are possible and make it real-time capable [32][76].

This section is organised as follows: subsection 5.2.1 presents a vessel detection filter that provides for each pixel in an image the probability of it belonging to a vessel and in subsection 5.2.2 we apply a fast version of the previous algorithm to binary segment the guidewire in interventional frames.

5.2.1 Directional stamping

In this section a method to segment vessels filled with contrast dye is presented. *A priori* knowledge regarding the appearance of the coronary arteries in the diagnostic angiograms is available: an artery is a dark, oriented and elongated structure with spatial coherence. *Directional stamping* [32][121] is a coherence-enhancing diffusion vessel detection filter that performs a multiscale extraction of elongated dark image features by combining the pixel value and the local orientation at the pixel. The algorithm outputs high values for pixels that likely belong to an artery or any other elongated dark structure, whereas background pixels result in low filter outputs.

Given an image $I_{\delta,\beta}(x, y)$ and a certain scale $\sigma \in \{\sigma_i,\, i = 1, 2, \dots\}$, directional stamping reconstructs in the vicinity of each pixel, the dominant dark oriented structure (if present) along its main orientation. In the result-

ing image, all dark elongated structures of a certain scale σ_i are enhanced whereas all the other image content is suppressed. The final result of the filter is a sum of the reconstructed structures at all scales and is interpreted for each pixel as the probability of the pixel belonging to a dark elongated structure. The scales σ_i can be selected depending on the particular application and thus extract vessels or guidewires that are anatomically relevant for the interventional treatment. A block diagram of the directional stamping filter is provided in figure 5.3 and will be explained in the following paragraphs.

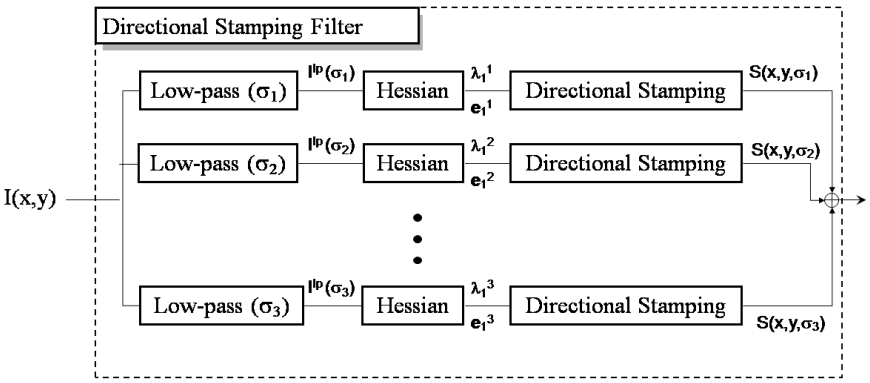

Figure 5.3: Directional stamping filter diagram.

For each scale σ_i, the image $I_{\delta,\beta}(x,y)$ is low-pass filtered with a Gaussian filter of variance σ_i^2 resulting in a low-pass image $I^{lp}_{\sigma_i} = I_{\delta,\beta}(x,y) * G_{\sigma_i}$. For each pixel of this low-pass image, the Hessian matrix is calculated

$$H_{\sigma_i}(x,y) = \begin{pmatrix} \frac{\partial^2 I^{lp}_{\sigma_i}}{\partial x^2} & \frac{\partial^2 I^{lp}_{\sigma_i}}{\partial x \partial y} \\ \\ \frac{\partial^2 I^{lp}_{\sigma_i}}{\partial x \partial y} & \frac{\partial^2 I^{lp}_{\sigma_i}}{\partial y^2} \end{pmatrix} = \begin{pmatrix} I^i_{xx} & I^i_{xy} \\ \\ I^i_{xy} & I^i_{xx} \end{pmatrix} \tag{5.4}$$

The corresponding eigenvalues of the Hessian matrix are given by

$$\lambda^i_{1,2} = \frac{1}{2}\left(I^i_{xx} + I^i_{yy} \pm \sqrt{(I^i_{xx} - I^i_{yy})^2 + 4(I^i_{xy})^2} \right) \tag{5.5}$$

The eigenvalues of the Hessian matrix are an indication of the elongated-ness of the structure in the vicinity of the considered pixel. If the pixel belongs to

a vessel, which is an elongated dark structure, the eigenvalues of the Hessian matrix for the vicinity of this pixel will reflect this fact. The first eigenvalue $\lambda_1^i(x,y)$ and eigenvector $\vec{e}_1^i(x,y)$ indicate the strength and direction of the elongated structure respectively. Only pixels that belong to a dark structure on a brighter background, i.e. $\lambda_1^i(x,y) > 0$, are considered [32][122]. In figure 5.4 we see an example of calculation of λ_1^i. In figure 5.4a we can see the original image, where elongated dark structures as vessels filled with radio-opaque contrast agent and metalwire stitches are present. Figure 5.4b shows the result from calculating the first eigenvalue of the Hessian matrix λ_1^i. Some of these values are negative. In figure 5.4c we can see the result from setting to zero the negative values in figure 5.4b.

For each pixel (x,y), a square neighbourhood of size $6\sigma_i \times 6\sigma_i$ is extracted, normalised to mean 0 and variance 1, weighted by the strength of the elongated structure $W(x,y) = \lambda_1^i(x,y)$ and windowed with a Gaussian to avoid border artefacts. The resulting *stamp* $S(x,y,\sigma_i)$ is a weighted model of the *vessel-ness* of the pixel and its neighbourhood. Directional stamping can be combined with another vessel detection filter, $P_{out}(x,y)$ by weighting each stamp not by $W(x,y) = \lambda_1^i$, but by the output of said filter at each pixel $W(x,y) = P_{out}(x,y)$. Figure 5.5 shows an example of stamp calculation. In figure 5.5a we see the original image and a squared area that contains a vessel and has been chosen to illustrate stamp calculation. In figure 5.5b we show that squared area in detail. After normalisation and weighting we obtain the stamp as shown in figure 5.5c.

The output of the directional stamping filter is the cumulative addition of the stamp to a *reconstruction buffer* $R(\sigma_i)$ along a line centred at pixel (x,y), perpendicular to $\vec{e}_1^i$ with length $k^{St}\sigma_i$, where $k^{St} \in \mathbb{N}$. The resulting $R(\sigma_i)$ for the image in figure 5.4a for three different scales can be seen in figures 5.6a to 5.6c. It can be observed how the scale influences the minimum vessel diameter that can be detected and the amount of noisy false positives.

Figure 5.7 shows in detail the calculated cumulative buffer $R(\sigma_2)$ for the scale in figure 5.6b. We have pointed out different artifacts that occur in the result, and which might be present in other scales σ_i depending on said scale. Directional stamping enhances dark elongated structures, like metal wire stitches. Object borders might also appear in the resulting images, like rib boundaries which can be seen in the results for σ_1 and σ_2 in figure 5.6. The diaphragm border can be seen more clearly in the result for σ_3 in figure 5.6, as well as elongated lung tissue borders.

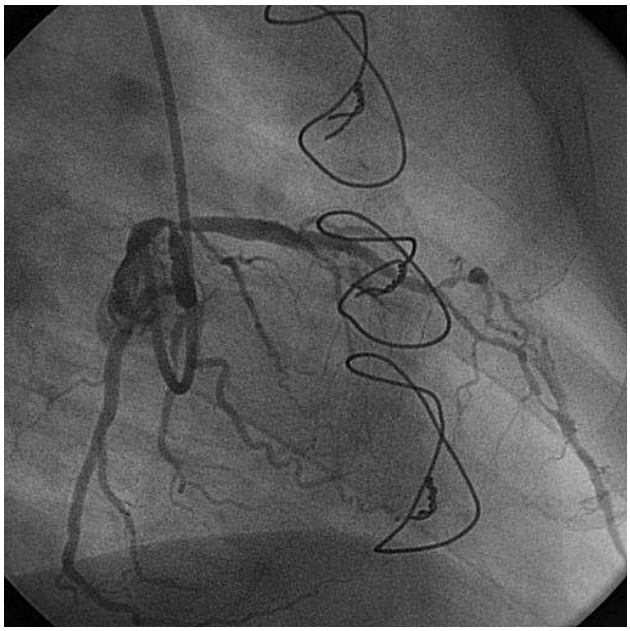

(a) Original fluoroscopy.

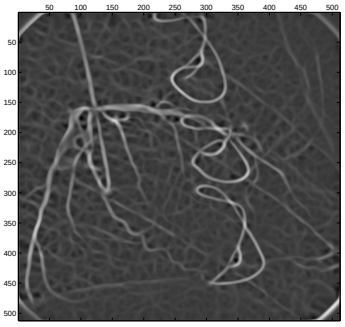

(b) λ_1

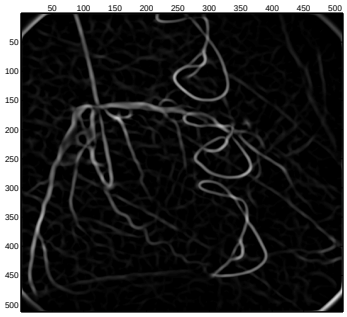

(c) $\lambda_1 > 0$

Figure 5.4: (a) Fluoroscopy frame chosen to exemplify the methods presented in this section. (b) Values of the first eigenvalue λ_1^i for $\sigma_i = 2.5$. For dark elongated structures on a brighter background $\lambda_1^i > 0$. (c) Positive values of λ_1^i.

For each scale, the values of the reconstruction buffer are normalised to the interval $[0, 1]$ and added to yield a vessel probability image $V(x, y) = \sum_{i=1}^{n}(1 - R(\sigma_i))$. The filter output corresponds to the values in $V(x, y)$, which are clipped to 1 if $V(x, y) > 1$. Three examples of the output of the filter can be seen in figure 5.8, where dark elongated structures have high pixel values.

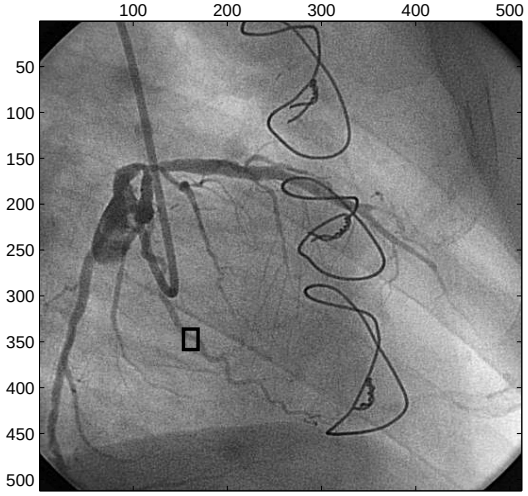

(a) Angiogram with vessels filled with contrast agent.

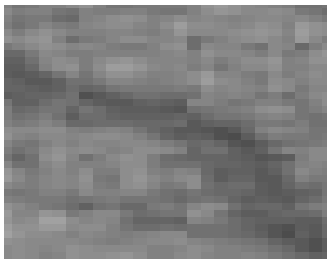

(b) Squared area

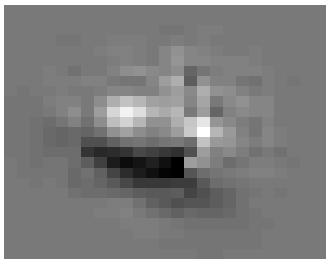

(c) Stamp

Figure 5.5: (a) An angiogram with contrasted vessels. The squared area contains a vessel and has been chosen to illustrate stamp calculation in (b) and (c). (b) Squared area in detail. (c) Stamp after normalisation to mean 0 and variance 1 and Gauss low-pass filtering of (b).

To avoid false positives caused by elongated structures of the background (like ribs, stitches or vertebrae), it is advisable to apply directional stamping on an estimated background. The background is estimated for each pixel as the 75-percentile brightest value of the pixel in the diagnostic sequence. The detected false positive image can be afterwards subtracted from the vessel probability image, in order to reduce the effect of false positives on the final detection result.

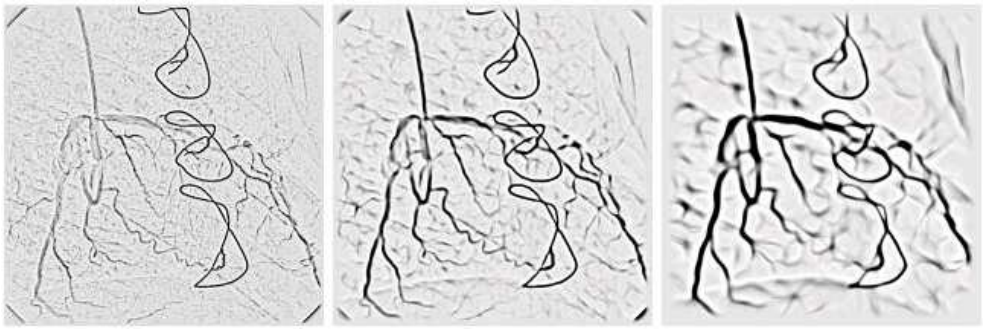

Figure 5.6: $R(\sigma_i)$ after applying directional stamping on figure 5.4a for three different scales, σ_1, σ_2 and σ_3 respectively.

5.2.2 Guidewire segmentation

In order to accurately register an interventional guidewire to a vessel it is necessary to first segment the vessel and the guidewire. We propose in the present subsection to use directional stamping as a guidewire segmentation algorithm. A guidewire is a dark elongated structure with a fixed size visible in live interventional frames. We propose to modify directional stamping so it can segment the guidewire in frame rate.

In [122] Baert *et al.* propose a 2D multi-scale guidewire tracking algorithm for improved guidewire visualisation and position estimation. Their method is based on B-spline optimisation, where line-like structures in an image are enhanced by analysing the first eigenvalue of the Hessian matrix (equation 5.5) and directional information. They propose an optional preprocessing step with coherence enhancing diffusion to reduce noise. They report that their method detects the guidewire correctly in 96% of the tested frames. They argue that not detecting the guidewire in one frame is not critical due to the high frame rate of the monitors ($12.5fr/s$). The approach in [122] was extended to biplane fluoroscopy in [123] by detecting the guidewire in each biplane projection independently, and corresponding guidewire points in both projections using the epipolar constraint. They report detection errors of up to $1.5mm$ after correction of the distortions of the biplane system. Koning *et al.* in [124] use manual segmentation to localise the instruments in fluoroscopic frames. In our application, the guidewire *must* give high segmentation results in order to be able to match it with a vessel and therefore the approach presented in [122] is not suitable. User interaction like in [124] is not to be considered for real-time applications like ours and the method in [123] is computationally very costly for real-time applications [32].

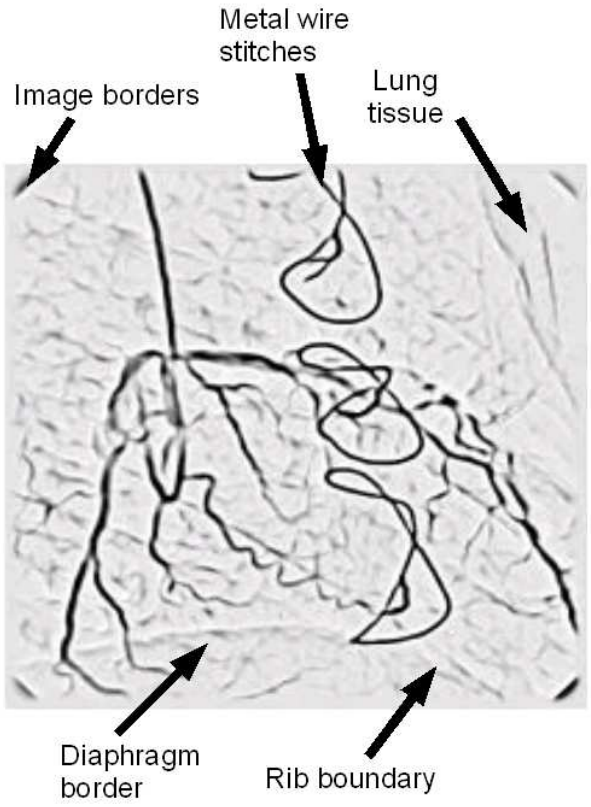

Figure 5.7: Cumulative buffer for scale σ_2 in figure 5.6.

Directional stamping, as explained in subsection 5.2.1, is a multi-scale vessel detection filter that provides, for each pixel, the probability of it belonging to a a dark elongated structure. This method analyses the principal orientation in the vicinity of a pixel at a certain scale and derives the probability of the pixel belonging to a dark oriented structure. In interventional frames, the guidewire is a semi-opaque elongated structure and should, therefore, give high detection results when applying directional stamping over an appropriate scale. In the present section the performance of the directional stamping filter for guidewire extraction is evaluated [32]. Again, to avoid interference of background structures on the guidewire detection, a background correction is required. For each current interventional frame, a sliding maxi-

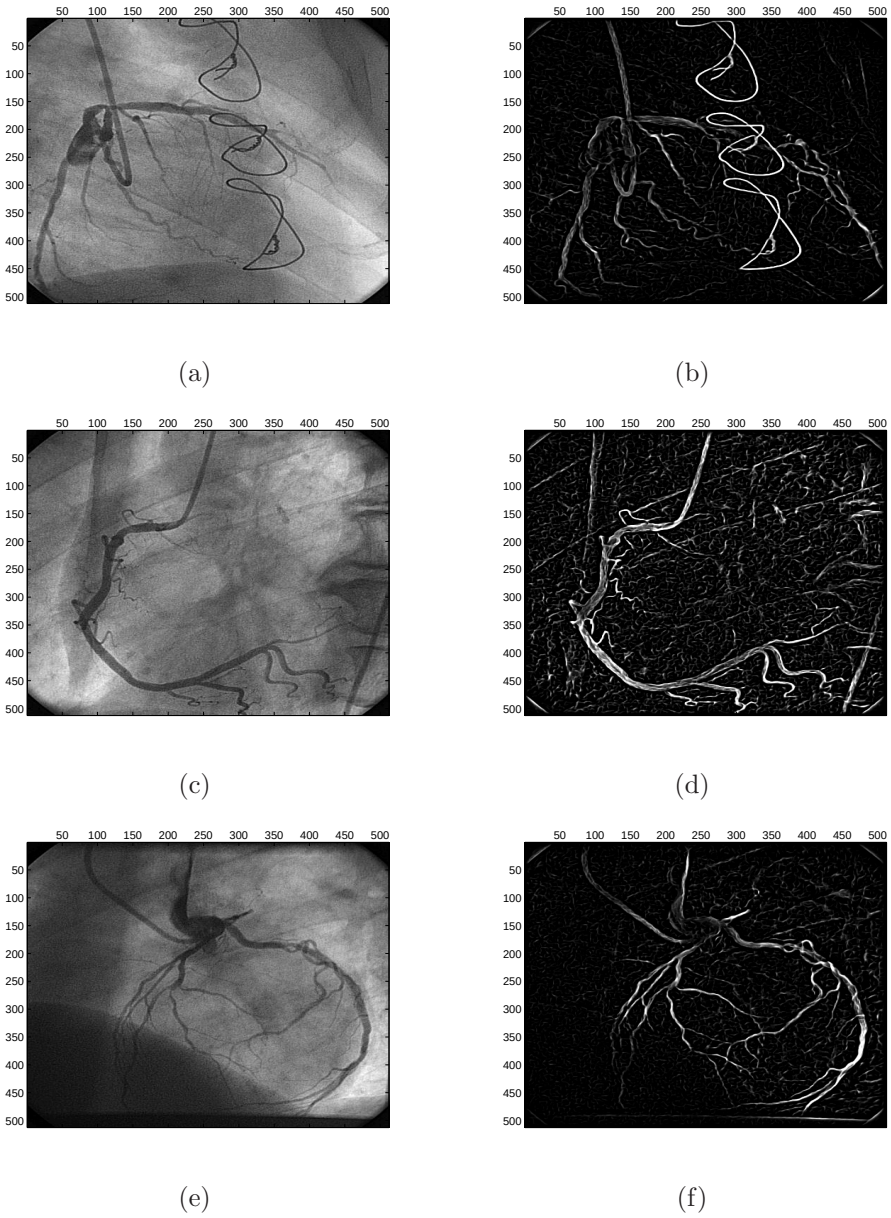

Figure 5.8: Three examples of angiograms (left column) and the resulting vessel probability filter based on directional stamping (right column). False positives are clearly visible.

mum image of a few previous images is calculated. The directional stamping filter is applied to this background estimation, and the output is subtracted from the filter results for each interventional frame.

Guidewire segmentation must run in frame rate because any added latency in the processing of the images will disturb the hand-eye coordination of the physician. To allow fast guidewire extraction, the derivatives in the Hessian (equation 5.4) are simplified and calculated with discrete filters

$$\frac{\partial}{\partial x} \rightarrow \frac{1}{2}(-1 \quad 0 \quad 1), \text{ and } \quad \frac{\partial}{\partial y} \rightarrow \frac{1}{2}\begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} \tag{5.6}$$

To further speed up the segmentation, the local orientation for each pixel is discretised to four orientations (up/down, right/left, up-right/down-left, and up-left/down-right). The computation time is further reduced by tracking the guidewire in a square ROI centred at the segmented guidewire's centre of gravity in the previous frame. The results of applying this fast version of directional stamping on figure 5.2.2a can be seen in figure 5.2.2b. Subsequently, thresholded segmentation of the most probable guidewire pixels is performed. *A priori* estimation of the number of pixels that should belong to the used guidewire given the imaging geometry was set manually and is used to adapt the threshold value (see figure 5.2.2c).

5.3 Radio-opaque device marker segmentation

Guidewire tip detection is used for more accurate angiogram to fluoroscopic frame registration in order to aid navigation to the lesion site. In the past section a method for segmenting the guidewire tip based on directional stamping [32][121] has been presented. Instruments used in PTCA procedures like stents or balloons are equipped with two small radio-opaque markers that allow to track their position in the acquired fluoroscopic x-ray images (see the examples in figure 5.10a and 5.10b). IVUS probes are clearly visible as small dark spots in the image, as can be seen in figures 5.10c and 5.10d. These markers provide a visual landmark that can be used to register the location of the device with the contrasted vessels in a corresponding angiogram. Instrument marker segmentation is used to aid accurate instrument positioning for improved intervention outcome.

There is little literature on device marker segmentation in catheterisation procedures. In [124] segmentation of the device markers is performed manually. In [125] a method to segment aortic graft markers in CT images is proposed. Their method requires the user to segment manually in each

(a) Fluoroscopic frame.

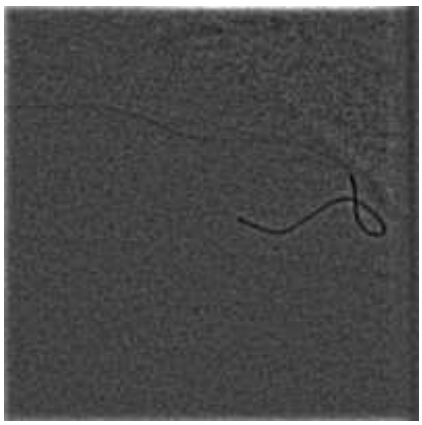

(b) Directional stamping output

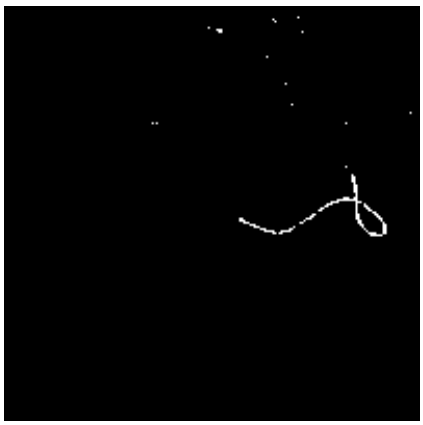

(c) Result after thresholding

Figure 5.9: (a) Original fluoroscopic image of a guidewire. (b) Output of the directional stamping filter for guidewire detection. (c) Result of thresholding using *a priori* knowledge about the size of the guidewire.

image the first and last graft marker for initialisation and uses prior knowledge as to the marker shape. Weide *et al.* report in [126] an intravascular magnetic marker detection based on the spatial intensity distribution. The proposed method localises the markers after subtraction of a reference image and coarse vessel centerline segmentation in MR images. During a catheterisation procedure, we lack the information about the vessel tree structure

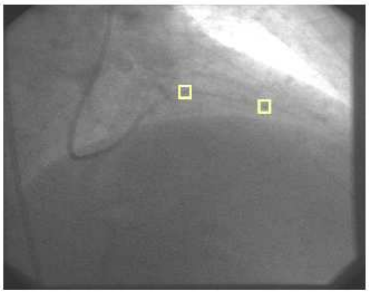

(a) Stent markers

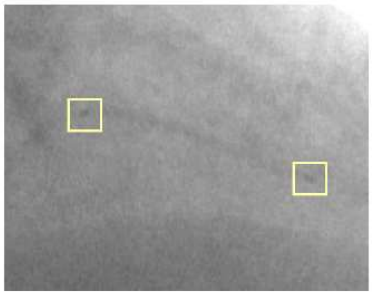

(b) Stent markers detail

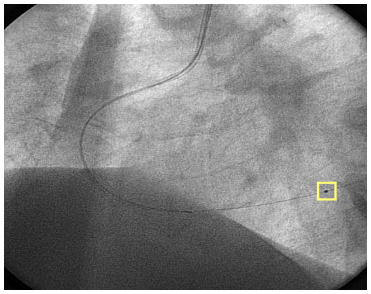

(c) IVUS probe

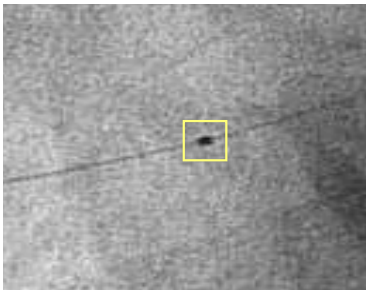

(d) IVUS probe detail

Figure 5.10: (a) Interventional frame where the stent markers (squared) can be seen. (b) Zoom from stent markers, which are radio-opaque but not really circular. (c) Interventional frame where an IVUS probe (squared) is visible. (d) Zoom from the IVUS probe.

and also a reference background image as in [126], because of the constant movement of the heart and surrounding tissue inside the thorax.

Spatial filtering with predefined marker templates highly depends on the similarity in size and shape between the template used and the marker. For the considered application, live acquisition of fluoroscopic frames, manual segmentation of the device marker for ulterior use as template is impractical. Furthermore, template matching applied to every pixel in the image would result in a very computationally intensive method.

We present a novel approach that combines temporal and shape features for radio-opaque device marker segmentation. The algorithm first performs temporal filtering that computes for each pixel, the probability of showing a marker based on temporal dynamics. Spatial filtering is applied subsequently

to the results of temporal filtering to yield a detection result that thus combines temporal and spatial features for detection.

In subsection 5.3.1 we explain the proposed method. In subsection 5.3.2 we present the theoretical background for our candidate pixel detection based on the temporal dynamics of a pixel's gray values. Subsection 5.3.3 explains the method proposed for radio-opaque marker detection in interventional frames that combines temporal features of the sequence of images and template matching on candidate pixels. We present our validation results and compare our method with template matching alone in subsection 5.5.3.

5.3.1 Method overview

A way of regarding each frame on a PTCA interventional sequence is as a composite of radio-opaque objects moving on the foreground with a constantly changing background. The background consists of soft tissue like the diaphragm or the lung, bone tissue like the ribcage or the spine, and anything else that is in the image and is not the guidewire or the interventional instrument. The dark objects we are interested in, *target objects*, are any interventional instruments (stent, balloon, IVUS probe, ...) that are being manipulated and are equipped with radio-opaque markers. Target objects are contained within a coronary artery, and therefore moving because of heart contraction and breathing, and also being pushed forward or pulled backward by the physician. The background is constantly changing due to respiration and heart beat, and therefore, standard background subtraction techniques would introduce movement artifacts in the resulting background-corrected images.

We propose a method to detect dark device markers from a moving background based on temporal statistical modelling of the pixels' intensity values and posterior spatial filtering. Target objects move due to heart beat, breathing and instrument manipulation by the physician. Our method first selects candidate target object pixels based on pixels' gray values temporal features $g(t)$. We hypothesise that the chest tissue visible in fluoroscopic frames does not present strongly discontinuous image intensity levels changes over space and time. A target object is recognised as a region where pixels momentarily take a low intensity value that is not plausibly explicable with the local observation of gray levels in previous frames, and it is therefore explicable by the x-ray attenuation of the instrument marker that is momentarily present in this region. In consequence, the detection of a non-continuous

drop in intensity can be attributed to a small attenuating marker, i.e. a radio-opaque marker. We can see an example of a pixel's gray level values over time $g(t)$ in figure 5.11. We can see in figure 5.11a how a pixel that shows background tissue (solid line) during several consecutive frames takes values around a mean, with low deviation from said mean. The same applies for a pixel that shows diaphragm tissue (dotted line), where the mean has a lower value. The dashed line shows $g(t)$ for a pixel that shows the diaphragm border. The variance for each of these curves is $\mathrm{var}(g_{\mathrm{backgr}}(t)) = 0.0016$, $\mathrm{var}(g_{\mathrm{diaph}}(t)) = 0.0008$ and $\mathrm{var}(g_{\mathrm{border}}(t)) = 0.0022$ respectively. On figure 5.11b we can see $g(t)$ for a pixel that occasionally shows the IVUS probe in figure 5.10c. Notice the sudden drop in intensity for frames 7 and 20. For this curve $\mathrm{var}(g_{\mathrm{probe}}(t)) = 0.0165$. Figure 5.12 illustrates the spatio-temporal distribution of the stent markers in figure 5.10a over forty frames, which approximately corresponds to 5s of time. In figure 5.12a we can see the first frame from the acquisition, where stent markers are visible. The white stars in 5.12b indicate the position of said stent markers in the subsequent forty frames.

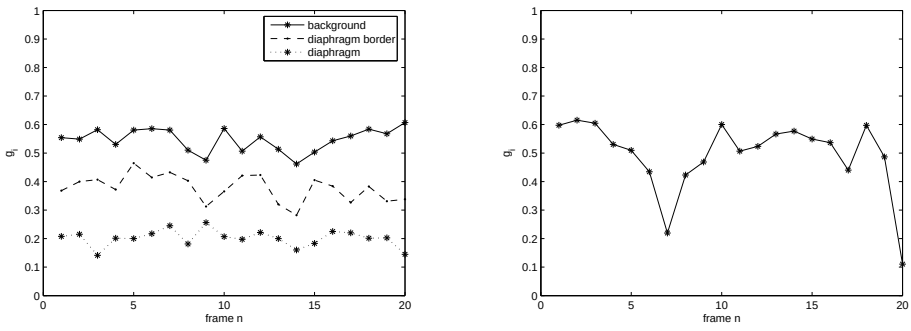

(a) Background and diaphragm pixel (b) IVUS probe pixel

Figure 5.11: (a) Gray values of a pixel showing background tissue (solid line), diaphragm border (dashed line) and diaphragm tissue (dotted line) in 20 consecutive frames. (b) Gray values of a pixel showing an IVUS probe in 2 frames out of 20 consecutive frames.

We observe that a pixel's intensity values are affected by quantum noise, and image content that varies in intensity and position with time. A pixel in the image might contain diaphragm tissue in one frame, other soft tissue in another frame, and a radio-opaque object in yet another frame. Because the objects we are trying to segment are small and move rapidly, there are not any pixels which permanently contain the target objects throughout the se-

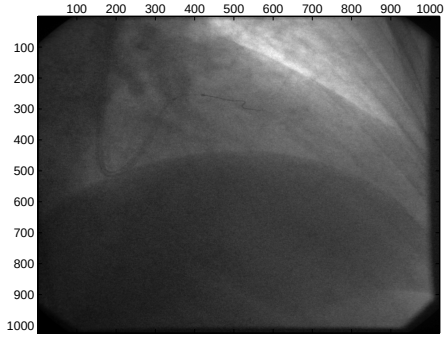

(a) Stent markers

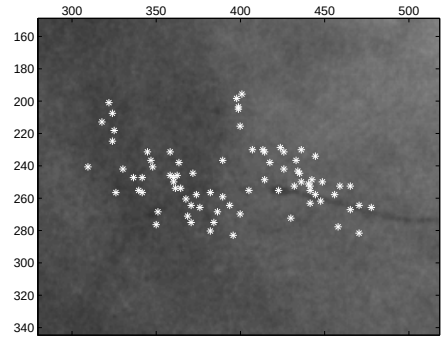

(b) Positions in time

Figure 5.12: (a) First image from a fluoroscopic sequence where stent markers are visible. (b) Zoom of an area around the stent markers. The white starts indicate the positions of said markers at different times.

quence. The gray values of pixels that have an approximately constant value throughout the sequence, are only affected by noise and will have low variance as can be seen in the background and diaphragm tissue gray pixel values in figure 5.11a. Pixels that belong to different objects at different points in time will have a higher variance like the pixels for the diaphragm border in figure 5.11a and an IVUS prove in figure 5.11b. Therefore, the variance of $g(t)$ for each pixel is used as a coarse estimator to select *candidate pixels* and reduce computational effort. We compute the variance of all the pixels in the image based on a sample of their W most recent values and threshold with the 75-percentile before proceeding to segment the device markers.

We assume that the values of $g(t)$ for each pixel on the image are the real-

isations of a stochastic process with an unknown probability density function (*pdf*). We propose to estimate the *pdf* of a pixel's gray values based on the observation of a number of W past values by non-parametric kernel density estimation, more specifically, Parzen window estimation using a Gaussian kernel (see section 5.3.2). This method does not make any *a priori* assumption as to the *pdf* of each pixel's gray values, but estimates the *pdf* from W past observations. For each new gray value of the pixel g_i, we use the pixel's estimated *pdf* to forecast the probability of the pixel $I(x, y)$ having a value lower than g_i, $P[I(x, y) < g_i]$. If the value is dark, and the probability for this intensity $I(x, y)$ is very low when no additional attenuation, like the one caused by a radio-opaque body, is assumed, then we assume that this pixel at this point in time is a candidate to contain the target object. Further spatial image processing based on *a priori* knowledge of the object is applied to finally segment the target.

5.3.2 Probability density estimation

In our scenario, there are different processes that influence the intensity value of the pixels: the moving background, moving target objects and quantum noise. Because each pixel might be affected by one or all of this processes depending on its location within the image and the imaging geometry, the probability distribution of their intensity values cannot be modelled as a single gaussian distribution. A mixture of Gaussians model (MoG) requires the *a priori* knowledge of the number of Gaussian processes generating the pixel intensity values, their parameters (mean and variance) and their weights.

Therefore, Parzen non-parametric *pdf* estimation was employed. In *Parzen window density estimation*, a kernel (usually Gaussian) is centred on each data point and added to a running sum. Parzen window density estimation makes no *a priori* assumptions about the underlying *pdf* to be estimated nor its number of nodes, and is therefore appropriate to estimate the *pdf* of poorly understood processes.

5.3.2.1 Parzen window density estimation

Parzen window density estimation is a non-parametric method that estimates the density of an unknown distribution given a finite number of samples without making any assumption of the underlying *pdf*. This method is mainly used when the uncertainty of the process generating the observed variable is high and it cannot be assumed that the *pdf* has a definite shape. The

unknown *pdf* is estimated as a weighted sum of window functions centred at each sample. This method has been widely used for classification problems [127][128][129][130][131].

Let R^q be a q-dimensional space ($q \geq 1$) and $\{\mathbf{x_i}, i = 1 \dots n\}$ be n q-dimensional training samples from which we wish to estimate a *pdf* $f(\mathbf{x})$. We define the unit hypercube centred at the origin as

$$\delta(\mathbf{u}) = \begin{cases} 1 & \|\mathbf{u_j}\| \leq 1/2 \, \forall 1 \leqslant j \leqslant q \\ 0 & otherwise \end{cases} \tag{5.7}$$

Then $\delta(\mathbf{x} - \mathbf{x_i}/h_n) = 1$ if $\mathbf{x}$ falls inside the hypercube $R_n \subset R^q$ centred at $\mathbf{x_i}$ with side h_n and volume $V_n = h_n^q$. The probability that a given sample $\mathbf{x}$ falls within the region R_n is

$$P[\mathbf{x} \in R_n] = \int_{\mathbf{x} \in R_n} f(\mathbf{x}) d\mathbf{x} \tag{5.8}$$

where $f(\mathbf{x})$ is the *pdf* of $\mathbf{x}$.

If $f(\mathbf{x})$ is continuous in the region R_n and R_n is such that $f(\mathbf{x})$ does not vary significantly within it, we can approximate the integral in 5.8 by

$$P[\mathbf{x} \in R_n] = \int_{\mathbf{x} \in R_n} f(\mathbf{x}) d\mathbf{x} \approx \int_{\mathbf{x} \in R_n} \hat{f}(\mathbf{x}) d\mathbf{x} \approx \hat{f}(\mathbf{x}) V_n \tag{5.9}$$

where

- $\hat{f}(\mathbf{x})$ is the estimated *pdf*, and
- V_n is the volume enclosed by R_n surrounding $\mathbf{x}$.

If we have a sample of size n, k_n points will fall inside the hypercube R_n and k_n/n is a good estimate of the probability of $\mathbf{x} \in R_n$. Therefore, substitution in equation 5.9 yields

$$\hat{f}(\mathbf{x}).V_n \approx k_n/n \tag{5.10}$$

where the number of samples contained in R_n is

$$k_n = \sum_{i=1}^{n} \delta\Big(\frac{\mathbf{x} - \mathbf{x_i}}{h_n}\Big) \tag{5.11}$$

Therefore, the estimation of the *pdf* can be written as

$$\hat{f}(\mathbf{x}) = \frac{k_n}{nV_n} = \frac{1}{nh_n^q} \sum_{i=1}^{n} \delta\Big(\frac{\mathbf{x} - \mathbf{x_i}}{h_n}\Big) \tag{5.12}$$

Equation 5.12 gives an estimation of the *pdf* as a sum of the window function $\delta(\mathbf{x})$ centred at each sample.

A common choice is not use as window function $\delta(\cdot)$ as defined in equation 5.7 but a Gaussian window function [128][129][131]

$$\phi(\mathbf{x}, h) = \frac{1}{(2\pi)^{q/2} h^q \|\Sigma\|^{1/2}} \exp\Big\{ -\frac{\mathbf{x}^T \Sigma^{-1} \mathbf{x}}{2h} \Big\} \tag{5.13}$$

where Σ is the covariance matrix of a q-dimensional vector of random variables $\mathbf{x}$.

The estimated density function $\hat{f}(\mathbf{x})$ is calculated as a sum of kernel functions centred on a set of random samples. The choice of the smoothing parameter or window width h determines the smoothness of the estimated probability density. Too small an h would lead to a spiky or noisy estimate, with each spike being a narrow window function centred at each training sample. If h is very large, each training sample contributes towards the density estimation at every point $\mathbf{x}$ and the result is an oversmoothed estimate of $f(\mathbf{x})$. For a Gaussian distribution, the optimum value of h that minimises the expected error between the estimated density and the actual density is given by [132]

$$h_{\text{opt}} = \Big(\frac{4}{q+2} \Big)^{\frac{1}{q+4}} \sigma n^{-\frac{1}{q+4}} \tag{5.14}$$

where σ is the standard deviation of the samples.

In our application we estimate the *pdf* of a random variable representing past observations of a pixel's gray values of selected pixels. We use a 1-dimensional Gaussian window

$$\phi(\mathbf{x}, h) = \frac{1}{(2\pi)^{1/2} h} \exp\Big\{ -\frac{\mathbf{x}^2}{2h} \Big\} \tag{5.15}$$

with width h_{opt} as in equation 5.14, since we expect the distribution of gray values to be unimodal.

In figure 5.13 we can see two examples of calculated *pdf*s. A sequence in which an instrument with radio-opaque markers are visible was selected. We selected manually two pixels that belonged to the background, or to one the device markers at certain instants.

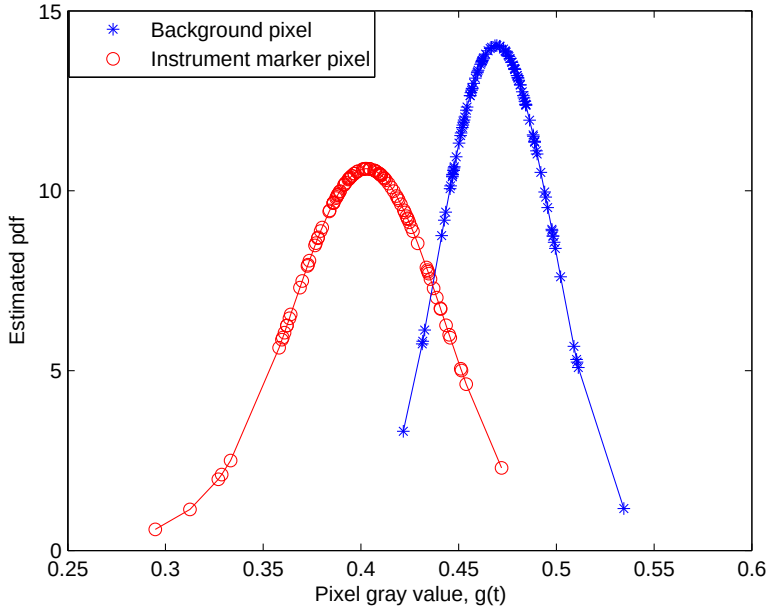

Figure 5.13: Parzen gray level $g(t)$ continuous *pdf* estimation results on a background pixel and a pixel that contains a stent-marker at some point in time from image5.12. Notice the greater variance of the distribution for the marker pixel.

5.3.3 Marker detection

In order to reduce computation, we first select a set of candidate target pixels. As mentioned in subsection 5.3.1 we hypothesise that the gray value of target pixels does vary between consecutive frames and during a small number of frames has a value lower than the surrounding background. By observing for each pixel a number W of past gray values and subtracting a background calculated as the mean of the Wpast values, we filter out the pixels with a low variance, i.e. a relatively constant value.

The *pdf* for each candidate pixel is estimated as explained in the previous subsection. Each candidate pixel is treated as a distinctive statistical process and modelled using Parzen-window density estimation calculated from the most recent W samples. For each candidate pixel $I(x, y)$ in frame n, the probability $P[I_n(x, y) < g]$, is calculated based on the estimated *pdf*. If the *pdf* is $f(g)$, then the probability of pixel $I_n(x, y)$ taking gray values in a region $g \in \mathcal{R}$ is

$$P[g \in \mathcal{R}] = \int_{g \in \mathcal{R}} f(g)dg \tag{5.16}$$

Our target objects are characterised by dark gray values and high mobility, which causes that $P[I_n(x, y) < g]$ should have a value close to 0 if the pixel at frame n indeed belongs to one of the target objects.

In figure 5.14a we see a frame in a sequence that contains an IVUS probe. A detail of the area around the IVUS probe we aim to segment can be seen in figure 5.14b. After selecting candidate pixels based on their variance as explained before we estimate the *pdf* based on a number W of past samples and calculate $P[I_n(x, y) < g]$. In figure 5.15a we can see a 2-dimensional plot of $P[I_n(x, y) < g]$ for the selected target pixels from figure 5.14a. The resulting image represents low probability values for pixels that are darker than their surrounding background and whose gray value changes rapidly along the sequence. We also observe in figure 5.15a that not only the pixels that belong to the device marker are considered, but the candidate pixels also correspond to pixels that contain guidewire pixels, diaphragm border pixels or other moving dark structures. By discriminating pixels based on the high variance of their observed values in W previous frames we are excluding pixels that belong to uniform areas of the background, but not pixels affected by movement of anatomical structures. In figure 5.15 we can see with more detail the calculated probability $P[I_n(x, y) < g]$ at the location of the IVUS probe in figure 5.15b.

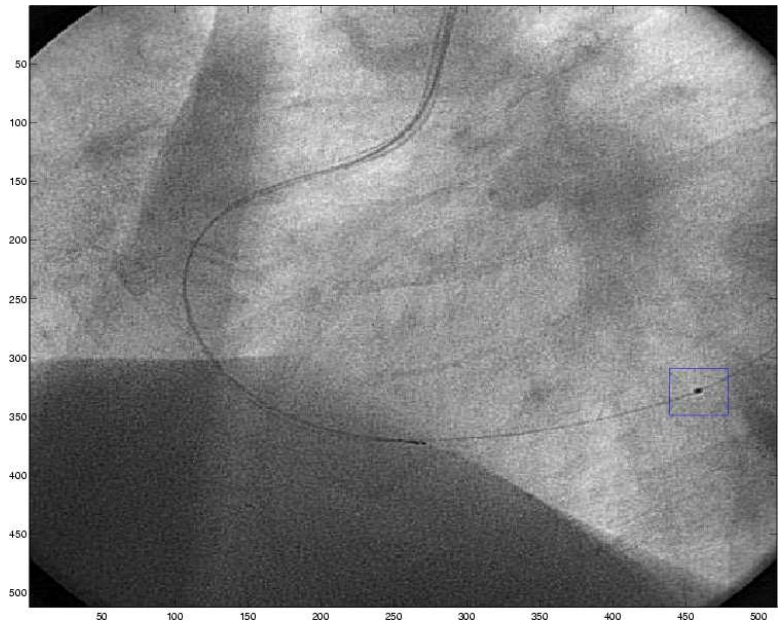

(a)

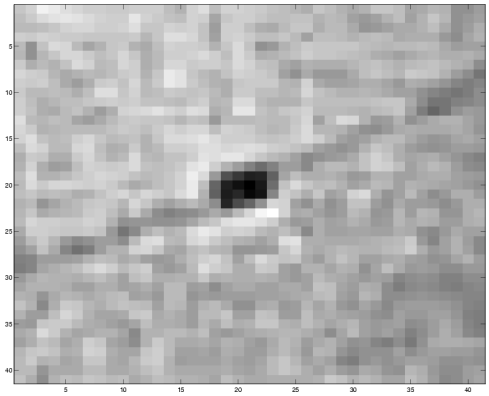

(b)

Figure 5.14: (a) Frame i in an interventional sequence. An area around the actual location of the target object is marked by a square. (b) Detail of the frame in (a) for the marked area.

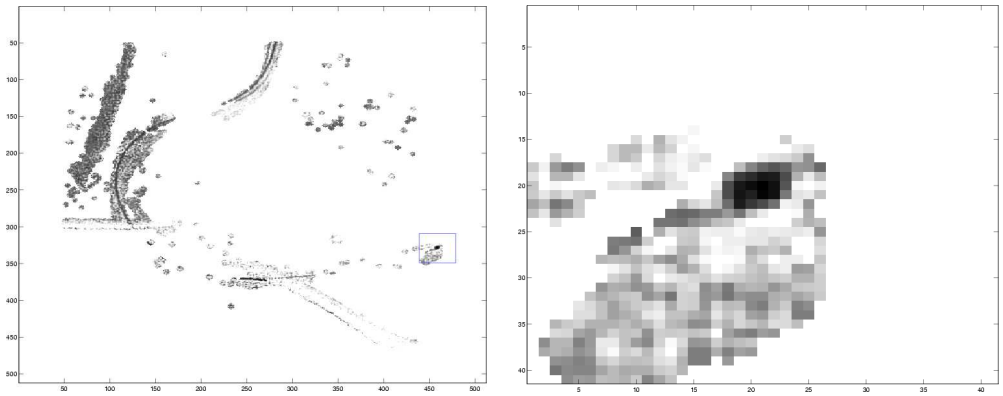

(a) (b)

Figure 5.15: (a) Result of calculating for each candidate pixel in frame i in figure 5.14(a) the probability $P[I(x,y) < g]$. An area around the actual location of the target object is marked by a square. (b) Detail of the resulting probability values in the area squared in (a).

In order to further enhance the results of computing $P[I(x,y) < g]$ for the pixels belonging to the target object, we apply *a priori* knowledge. The instrument markers have dark intensity values and a circular profile. We therefore first weigh $1 - P[I(x,y) < g]$ by the inverse of the original gray-value frame, $1 - I(x,y)$, where $I(x,y)$ is normalised to 1. This procedure further enhances the results for the target pixels, which correspond to high probability of being a marker pixel in a certain point in time. Then, we apply a matched 2D filter with a Gaussian profile and a radius matched to the expected target marker size in the image sequence. The radius of the Gaussian profile was selected manually based on the size of the device markers in our test images. We expect the target objects to result in the highest output values of the matched filter.

This algorithm needs initialisation time in order to collect enough samples to estimate a meaningful *pdf*. During long interventions, old samples should be erased from memory to avoid a biasing of the estimated distributions.

5.4 Anatomical device to vessel registration

The objective of accurate device to vessel registration is to maximise the overlap of the segmented instrument or guidewire to the nearest segmented

vessel when overlaying an angiogram onto a fluoroscopic frame during the intervention to aid guiding and accurate positioning. The first step is to segment the objects of interest: vessels, guidewire or instrument markers. Methods have been proposed in sections 5.4 and 5.3. The next step would be to effectively optimise the overlap of the segmented vessels with the segmented device in one angiogram $A_{\vec{\xi_a}}(x, y)$ and a fluoroscopic frame $F_{\vec{\xi_f}}(x, y)$.

Vessels, guidewire and instrument markers are elongated or round structures with a narrow *influence radius*. By this we understand that an overlap algorithm that directly uses the result of a vessel filter as an attraction field to draw the segmented device marker or guidewire towards the vessel will encounter the problem that the distance from the vessel over which the attraction field is effective is relatively small. If the device and the vessel in both images run parallel and close to each other with no overlap, a direct optimisation algorithm is confronted with a difficult search space to optimise the overlap in real-time. Even if there is a partial overlap and the guidewire needs to be shifted along its main elongation direction so that the local curvature of guidewire and vessel fully overlap, this is also invisible if the overlap of the result of the vessel filter and the segmented guidewire are used directly for optimisation.

We want to provide a convex search space such that facilitates robust and fast convergence to the optimum overlap solution in real-time. Even small, dimly contrasted arteries, which are a usual target in catheterisation procedures, should have an attraction basin with the same influence radius as bigger, more contrasted vessels. Therefore, the influence radius of each vessel should be independent of the values of the vessel filter output.

The proposed solution computes an attraction field from the result of a vessel detection filter like the one proposed in section 5.4. To produce this attraction field, a distance transform is applied to the output of the directional stamping filter to create a potential field that is used to match the segmented device and the corresponding vessel [32][121].

Traditional distance transforms are usually applied on binary images. The output of vessel filters like directional stamping is a gray value image where each pixel represents the probability of the pixel belonging to a dark elongated structure. To obtain a binary image we would have to threshold this output, and we would probably leave out dimly contrasted vessels from the thresholded result. This is undesirable because those dimly contrasted vessel are usually the target in PTCA interventions.

We propose in [32] a multiplicative distance transform with an hyperbolic kernel. A rotationally symmetric template $D(r)$ of size $(2R+1) \times (2R+1)$ is defined as a section of a hyperbola with slope $m < 0$ at $r \to 0$ (see figure 5.16)

$$D(r) = \frac{R}{1 - mr} \tag{5.17}$$

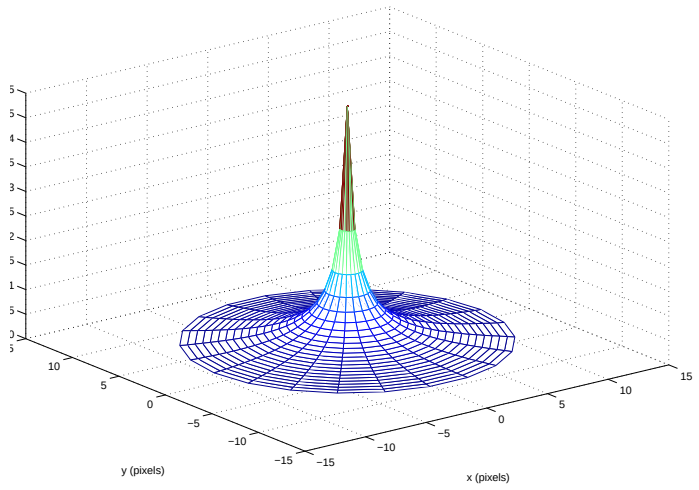

Figure 5.16: Hyperbolic template for the distance transform.

The template is individually multiplied with the vessel filter output and the result is input to a maximum determination of the distance transform. This results in a transformed distance map, where the valleys have a depth proportional to the vessel filter result but a width independent of the vessel filter intensity so that smaller vessels have an influence radius similar to that of bigger, more contrasted vessels. In figure 5.17a we can see the result of applying the proposed distance transform to the output of the directional stamping filter in figure 5.8b. Figure 5.17b shows a detail of an area in figure 5.17a (marked by a rectangle) but displayed in 3D where the deep elongated (and tortuous) valley corresponds to the potential field of a vessel visible in 5.17a.

Once the distance transform is computed and the instrument is segmented, fast overlap maximisation is performed [32]. For real-time device-to-vessel registration, the possible transformations are restricted to pixel-wise planar translations resulting in a shift $(\Delta x, \Delta y)$ in eight possible directions.

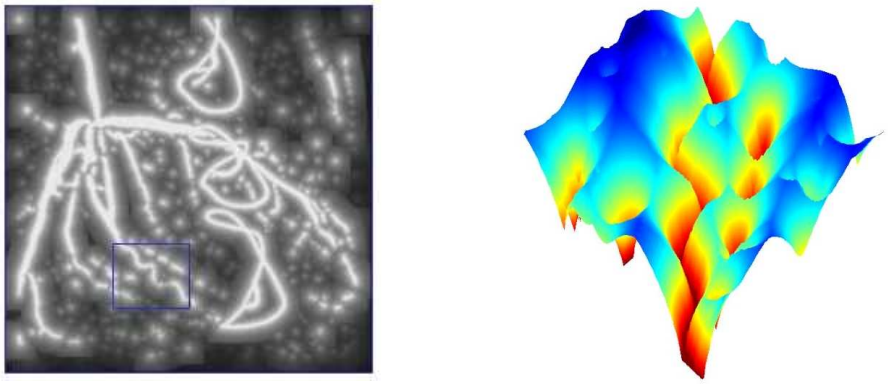

(a) Distance transform (b) Distance transform detail

Figure 5.17: (a) Distance transform applied to the results of the directional stamping filter in figure 5.8b. (b) 3D detail of the distance transform map from (a).

The overlap function $O(\Delta x, \Delta y)$ is defined as the sum of values in the possibly shifted distance map $\mathcal{D}(x, y)$ for all pixels that belong to the segmented device $DV(x, y)$

$$O(\Delta x, \Delta y) = \sum_{(x,y), DV(x,y) \neq 0} \mathcal{D}(x + \Delta x, y + \Delta y) \tag{5.18}$$

The optimisation process searches for $(\Delta x_{\max}, \Delta y_{\max})$ such that the overlap function is maximised. A multi-scale uphill algorithm is proposed for real-time capable optimisation of the overlap. The optimisation process is initialised with the translation necessary to register the previous interventional frame and angiographic roadmap, $\vec{m}_0 = (\Delta x_0, \Delta y_0)$ at scale σ_0 (the scale is related to the diameter of the detected vessels). For each iteration i and scale, $\{\sigma_i = \frac{\sigma_0}{2i}, i = 1, \ldots, n\}$, the overlap in the horizontal, vertical and diagonal directions with a shift of σ_i pixels is calculated and $(\Delta x_i, \Delta y_i)$ is updated. Optimisation ends when $O(\Delta x_i, \Delta y_i)$ is higher than the eight competing entries for $\sigma_i = 1$.

For the matching to converge rapidly to the desired solution, the instrument must be already located near a vessel. This would require for each fluoroscopic image $F_{\vec{\xi_f}}(x, y)$ to select an angiogram $A_{\vec{\xi_a}}(x, y)$ characterised by a similar state vector $\vec{\xi_a} \approx \vec{\xi_f}$ and then compensate for the differences $\Delta\vec{\xi}$. As in chapter 3 the phase-space is sparsely sampled and a sub-optimal choice would be to select an angiogram with the heart in the same state but

different respiration phase $\Delta\beta \neq 0$. Respiration compensation based on prior knowledge (section 5.1) may be applied to further speed-up the convergence of the matching algorithm. Subsequent device to vessel registration using the methods explained in this chapter would optimise the overlap on the treated vessel, visible in an angiogram, to the manipulated device, visible in the live fluoroscopic images, thus improving the overall overlayed image presented to the physician during the intervention.

5.5 Experiments and results

The vessel filter, guidewire and marker detection, and the matching were subject to extensive tests for a representative set of clinical data. For the presentation of the results we give median and quartile-differences $\xi_{0.5} \pm (\xi_{0.75} - \xi_{0.25})$.

5.5.1 Vessel detection

One angiogram $A(x, y)$ was selected from each of ten diagnostic acquisitions taken from coronary interventions from six patients. Angiograms are 512^2 pixels with 1024 gray levels and all visible contrasted vessels were manually segmented. This bronze standard represents the *ground truth* for filter validation and is used to determine a robust subset of true positives and negatives. The subset of true negative pixels consists of the complement of the binary ground truth morphologically eroded by 5 pixels. The subset of true positives consists of the skeleton (vessel centerlines) of the ground truth, further dilated by one pixel. With μ^{TP}, σ^{TP}, μ^{TN}, and σ^{TN} describing mean and standard deviation of the filter results for true positive (TP) and true negative (TN) pixels, the quality measure Q^V of a vessel detection result is defined as the absolute mean of the true positives, normalised to the distribution of true negatives,

$$Q^V = \left| \frac{\mu^{TP} - \mu^{TN}}{\sigma^{TN}} \right| \tag{5.19}$$

For a good segmentation, the true positives should have $\mu^{TP} = 1$, and the vessel filter should result in 0 for true negatives, in which case, $Q^V \to \infty$. Q^V decreases with high filter results for true negative pixels and low filter results for true positive pixels.

In the experiment, directional stamping was compared to the first eigenvalue only and the vessel resemblance filter (VR) described in [120], which assigns to each point in the image a vessel-ness value based on the Hessian matrix (equation 5.4) $H_\sigma(x, y)$, where σ is the considered scale. Let $\{\lambda_k^i, k = 1, 2\}$ be the eigenvalues of $H_{\sigma_i}(x, y)$ for scale σ_i (see equation 5.5). λ_2^i is the eigenvalue with the lowest absolute value, and is associated with the direction of minimum intensity variation and represents the strength of this variation. λ_1^i is associated with the eigenvector orthogonal to this direction. Therefore, vessel-like structures should be characterised by $|\lambda_2^i| \ll |\lambda_1^i|$. Because we are considering dark vessels on a bright background, only pixels with $|\lambda_1^i|$ are taken into account. Furthermore, background pixels are located in areas of low contrast variations and therefore have low values of λ_k^i. The norm of the Hessian, defined as $\sqrt{(\lambda_1^i)^2 + (\lambda_2^i)^2}$, is used to quantify this.

The output of vessel resemblance filter for scale σ is

$$V_\sigma(x, y) = \begin{cases} 0 & \text{if } \lambda_1^i < 0 \\ exp\left(\frac{-R_\sigma^2}{2\beta_1^2}\right)\left[1 - exp\frac{-S_\sigma^2}{2\beta_2^2}\right] & \text{otherwise} \end{cases} \tag{5.20}$$

where

$$\begin{aligned} R_\sigma &= \frac{|\lambda_2^i|}{|\lambda_1^i|} \\ S_\sigma &= \sqrt{(\lambda_1^i)^2 + (\lambda_2^i)^2} \end{aligned} \tag{5.21}$$

and $\beta_1 > 0$ and $\beta_2 > 0$.

R_σ is a measure of $|\lambda_1^i| \gg |\lambda_2^i|$, i.e. the elongated-ness of the area around the pixel, and S represents the overall curvature at the considered pixel. Both operands in the multiplication in equation 5.20 are in the interval $[0, 1]$ and are maximum for dark elongated structures, and therefore $V_\sigma(x, y) \in [0, 1]$ and maximum when a vessel is present. Parameter β_1 controls the ability of distinguishing between line- and blob-like structures (sensitivity to R_σ), whereas β_2 controls the sensitivity to low contrast background structures and noise (sensitivity to S_σ).

The vessel resemblance $V_\sigma(x, y)$ was tested with $\beta_1 = \{0.2, 0.4, \ldots 0.8\}$ and $\beta_2 = \{2, 4, \ldots 8\}$. Directional stamping was tested with a fixed $k^{St} = 1$. To extract the stamp for each scale, we tested as weighting functions $W(x, y) = \lambda_1(x, y)$ and $W(x, y) = V_\sigma(x, y)$. The scale for all three filters was varied within $\sigma = \{1, 1.5, 2, 2.5, 3, 4, 6, 8\}$. Table 5.5.1 shows the parameters that were used for evaluation of directional stamping and the vessel filter in

section 5.5. For each filter or filter combination, the best parameterisation for each frame was individually selected.

λ_1		
Parameters	Best	Q^V
$\sigma \in \{1, 1.5, 2, 2.5, 3, 4, 6, 8\}$		3.46 ± 1.02
V_σ		
Parameters	Best	Q^V
$\beta_1 \in \{0.2, 0.4 \dots 0.8\}$ $\beta_2 \in \{2, 4 \dots 8\}$ $\sigma \in \{1, 1.5, 2, 2.5, 3, 4, 6, 8\}$	$\beta_1 = 0.6$, $\beta_2 = 8$	7.45 ± 3.28
Directional stamping		
Parameters	$W(x, y)$	Q^V
$k^{St} = 1$ $\sigma \in \{1, 1.5, 2, 2.5, 3, 4, 6, 8\}$	$\lambda_1(x, y)$	6.40 ± 2.86
$k^{St} = 1$ $\sigma \in \{1, 1.5, 2, 2.5, 3, 4, 6, 8\}$	$V_\sigma(x, y)$	10.33 ± 3.30

Table 5.1: Vessel filter results. Directional stamping was compared to the vessel resemblance filter in [120] and the first eigenvalue of the Hessian matrix. The analised parameters are shown on the table.

We can see in table 5.5.1 that for the vessel resemblance $V_\sigma(x, y)$, $\beta_1 = 0.6$ and $\beta_2 = 8$ gave the best results on all scales with a final $Q^V = 7.45 \pm 3.28$. Directional stamping with $W(x, y) = \lambda_1(x, y)$ and $W(x, y) = V_\sigma(x, y)$ resulted in $Q^V = 6.40 \pm 2.86$ and $Q^V = 10.33 \pm 3.13$, respectively. These results suggest that the directional stamping method presented earlier in section 5.4 combined with the vessel resemblance filter is the most suitable vessel filter for our purposes.

5.5.2 Guidewire segmentation

To test the guidewire extraction method presented in subsection 5.2.2, 25 frames were extracted from 8 interventional sequences of 5 patients. A tracking ROI of size 200×200 pixels around the guidewire was manually selected. A manual bronze standard segmentation of the guidewire was used as ground truth.

The threshold was varied and the resulting binary image was compared to the manual ground truth. The quality Q^{GW} of a guidewire extraction algorithm is set to $Q^{GW} = \max_T (p^{TP} - p^{FP})$, i.e. the maximal difference of the probability of true positive p^{TP} and false positive p^{FP} pixels in the image thresholded with T, where the threshold is related to the number of pixels the segmented guidewire should have for the given imaging conditions. $Q^{GW} = 1$ indicates the possibility for error-free thresholding whereas $Q^{GW} = 0$ shows that a filter result is completely random with regard to the considered object.

All filters run on the original images and on downsampled frames with factors 1/3, 1/5, and 1/7 to test speed-up capabilities. The pixels in the downsampled images are set to the minimal pixel value in a square region of size 3×3, 5×5, or 7×7. The filter results are nearest-neighbour upsampled for validation. Directional stamping based on the simplified Hessian (equation 5.6) was tested and compared to the vessel resemblance filter with parameters $\beta_1 = \{0.2, 0.4, \ldots 0.8\}$ and $\beta_2 = \{2, 4, \ldots 8\}$.

Filter	Parameter	Q^V
λ_1		0.68 ± 0.09
Vessel resemblance	down 1/5, $\beta_1 = 0.8$ and $\beta_2 = 9$	0.70 ± 0.10
Directional stamping	$W(x, y) = \lambda_1$, down 1/5	0.96 ± 0.05

Table 5.2: Guidewire segmentation results. Fast directional stamping was compared to VR and the first eigenvalue of the Hessian matrix.

As we can see on table 5.2, directional stamping with the simplified Hessian gave a maximal $Q^{GW} = 0.96 \pm 0.05$ for a downsampling by 1/5. The vessel resemblance filter, applied to guidewire frames, gave a quality of $Q^{GW} = 0.70 \pm 0.10$ for $\beta_1 = 0.8, \beta_2 = 9$, and downsampling by 1/5.

5.5.3 Device marker segmentation

To validate the presented device marker segmentation method, six interventional sequences, containing stents, a balloon, and IVUS probes were used. A frame from three of these sequences can be seen in figures 5.18a (containing an IVUS probe), 5.18b (containing a stent) and 5.18c (containing a balloon). Detailed areas around the target objects of each sequence can be found in figures 5.18d, 5.18e and 5.18f respectively. We manually segmented the centre of the device markers and defined a square neighbourhood of size 3 pixels around each point for round markers and a 3×5 for elongated markers.

The pixels within this area were considered true positives. Our method was compared to template matching using circular templates of size the manually estimated marker size on a set of candidate pixels selected according to their variance in time.

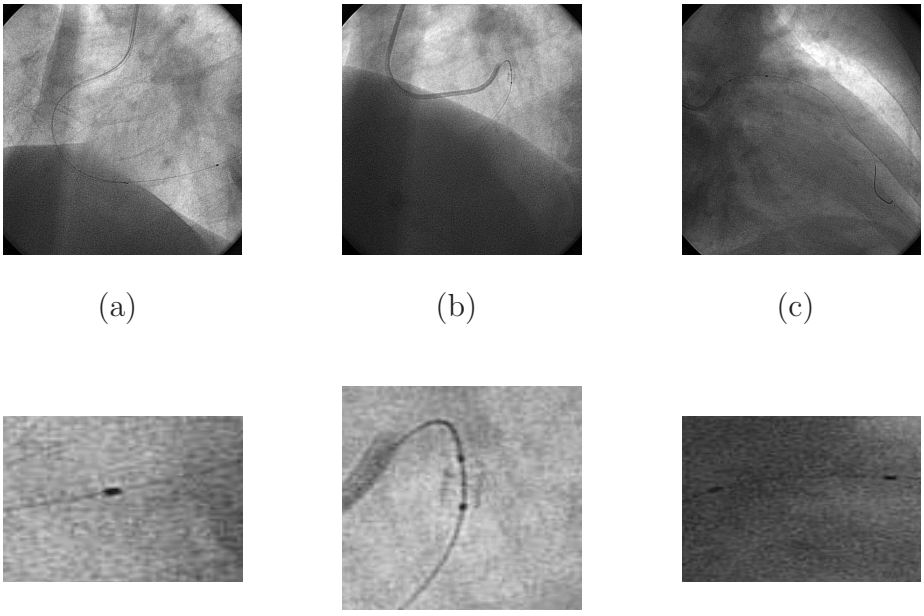

(a) (b) (c)

(d) (d) (f)

Figure 5.18: (a), (b) and (c) : Frames from the test sequences where the target objects are visible. (d), (e) and (f): Details from each frame in (a), (b) and (c) respectively, of an area around the target objects.

The quality Q^{MK} of a device marker segmentation algorithm is set to $Q^{MK} = \max_W (p^{TP} - p^{FP})$, i.e. the maximal difference of the probability of true positive and false positive pixels in the image for a window of past frames of length W. Running windows of $W = 5, 7, 10, 15$ past frames were tested.

To test speed-up capabilities, the frames were resized to sizes 1/3, 1/5 and 1/5 of the original, by characterising each pixel in the resized image by its standard deviation divided by its mean of each 3×3, 5×5 and 7×7 block, normalised to one over the entire resulting frame. The marker segmentation results are nearest-neighbour up-sampled for validation. We also performed marker tracking by using a search area centered on the previous detection

results.

The edges of the image were excluded from the analysis to avoid artifacts created by the shutters and image borders. For each frame, candidate pixels were selected according to their variance applying a 75-percentile threshold. The *pdf* of each candidate pixel was calculated based on the last W gray values as explained in subsection 5.3.2 using Parzen window estimation and the object was detected as explained in subsection 5.3.3. The coordinates of the marker centre were found by locating the maximum values of the resulting image after matched filtering with a Gaussian profile, and then compared with the set of true positives. If two device markers were present, then the two highest values with a distance greater than 10 pixels (or according to scale) between them were selected as the centres of the device markers. Table 5.3 shows the results from our marker segmentation algorithm.

Method	Scale	Q^{MK}
Template matching		0.11 ± 0.17
Marker segmentation (one marker)	1/5 (W = 10)	0.89 ± 0.14
Marker segmentation (two marker)	1/5 (W = 7)	0.79 ± 0.54

Table 5.3: Instrument marker segmentation results.

If more than one device marker was present (like in a stent or a balloon), at first only one of the markers was considered. In this case, the quality measure was $Q^{MK} = 0.89 \pm 0.14$. For the sequences with devices with two markers, we also tested the simultaneous detection of both markers in each frame, which resulted in a quality measure of $Q^{MK} = 0.79 \pm 0.54$.

Parameters that influence the success rate of the presented algorithm are:

- Window length W,
- device marker contrast,
- synthetic template size σ.

The window length W introduces a delay before the marker detection can start and therefore it is desirable to keep it as low as possible. Furthermore, for high W the probability of the device marker being at a certain position might increase and therefore one of the main assumptions for the presented method, i.e. that the target object is located at a certain position rarely, would be rendered invalid. The synthetic template σ affects the success rate

of marker detection. If the size and shape of the template deviate substantially from the shape and size of the device marker, the detection results worsen.

We also compared our method with template matching with a synthetic template. Selecting manually a template for the markers from the image is not feasible since we aim at automatic segmentation of instrument markers in interventional images. A circular template with a gaussian profile was used. The standard deviation σ was adjusted manually for each sequence to obtain the best possible detection results. The quality meassure is 0.11 ± 0.17. The detection result were specially bad in the sequence depicted in figure 5.18c, because the template used did not match exactly the shape of the markers. Template matching depends on an appropriate selection of the size and shape of the template. In our approach, the combination of temporal features (*pdf* estimation) and shape features (template matching) allows for higher success rates than when using template matching alone.

5.5.4 Device to vessel registration

Tests were performed to quantify the ability of the distance transform to register the device to the vessels visible in an angiogram, and to measure the convergence radius of the matching algorithm. To perform this test independently of the performance of vessel and guidewire and device marker segmentation algorithms, synthesised filter results were created for 25 manually segmented guidewire images and their respective manually segmented vessel centerlines. From these centerlines, synthetic filter results for a vessel of 2 pixels radius were created for a base intensity of 100 gray levels and additive Gaussian noise with σ^V. The synthesised guidewire extraction result was created by thresholding the binary image resulting from a manual segmentation after adding Gaussian noise σ^{GW}.

The synthetic detection quality was varied over a plausible range seen in the respective filter validations. The test regarding the vessel detection quality used a fixed $\sigma^{GW} = 0.4$ ($Q^{GW} \approx 0.8$) and $\sigma^V \in \{1, 1.25, \ldots 3\}$ (Q^V ranging from 17 down to 5.6). The test regarding the guidewire detection quality used a fixed $\sigma^V = 1.75$ ($Q^V \approx 10$) and $\sigma_{GW} \in \{0.1, 0.2, \ldots 0.9\}$ (Q^{GW} ranging from 0.99 down to 0.44). The distance transform of the synthesised vessel filter used $R = 20$ pixel and $m = -2$. From a known manual optimal position $\Delta x = 0$, $\Delta y = 0$ of guidewire and vessel, initial offsets $\Delta x_0 = \{-20, 18, \ldots 20\}$ and $\Delta y_0 = \{-20, 18, \ldots 20\}$ were used as initial misregister, a register was accepted for resulting $|\Delta x| < 3$ and $|\Delta y| < 3$.

Figure 5.19a shows the probability p of a successful registration over the initial misalignment $(\Delta x_0, \Delta y_0)$. This probability was calculated for a $\sigma^{GW} = 0.4$ ($Q^{GW} = 0.8$, slightly lower than the result for directional stamping for guidewire detection) and $\sigma^V = 1.75$ ($Q^V = 10$). As expected, p decreases with increasing distance to the target, i. e., a correct match is less probable if the initialisation is far from the solution.

In figures 5.19b and 5.19c, p is plotted for wide ranges of the initial radial distance $r = \sqrt{\Delta x^2 + \Delta y^2}$ for varying qualities of the vessel detection and the guidewire detection filters, respectively. We can see how the best results (high probability) occur when the distance between the guidewire and the vessel is low and the quality of the output of the vessel and guidewire detection filters is high (or inversely, when the noise has a low σ). Figures 5.19b and 5.19c also show that the impact of the vessel filter quality is higher than the impact of the guidewire detection quality and that the rate of failed matches raises dramatically when the vessel detection quality decreases.

5.6 Conclusion

Heart and respiration phase characterisation of the images acquired during diagnostic angiography results in a sparsely sampled phase-space representation. Differences in the state vector $\vec{\xi}$ between an angiogram and a fluoroscopic frame result in suboptimal overlay results for navigation and positioning support. In this chapter, methods to handle a difference in the state vector $\Delta\vec{\xi}$ have been presented. These include

- compensation of respiration phase differences $\Delta\beta$ and determination of motion as a function of β [73], and
- registration and fusion based on image content to compensate for residual changes.

Respiration compensation (see section 5.1), proposed by Eck *et al.* [73] consists of the estimation of a heart position function parameterised by the respiration phase in order to compensate for respiration phase differences between two frames. Although the respiration correction performed with this function does not achieve sub-pixel accuracy, it provides a good starting point for device to vessel registration algorithms.

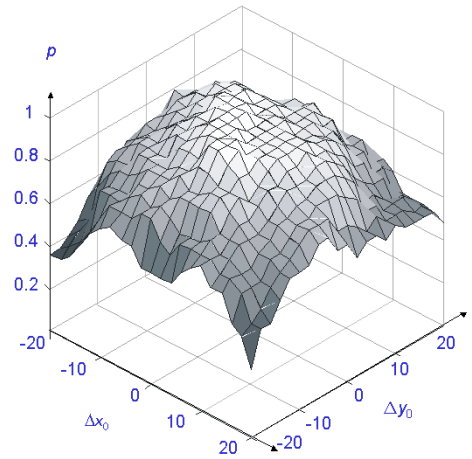

(a)

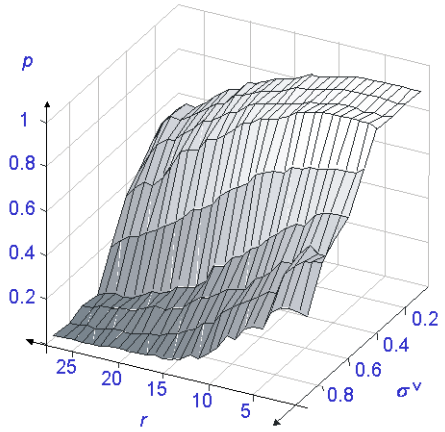

(b)

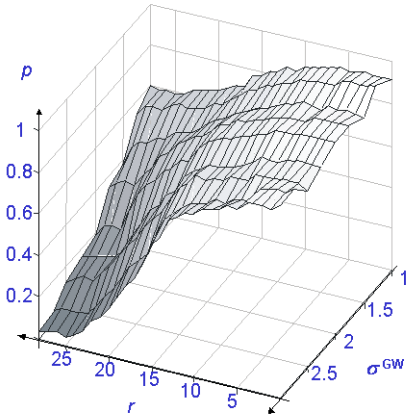

(c)

Figure 5.19: (a) Probability of a successful registration as a function of the initial misalignment $(\Delta x_0, \Delta y_0)$. (b) Dependence of the probability from the initial radial distance and the vessel detection quality. (c) Dependence of p from the initial radial distance and the guidewire detection quality.

We propose a device to vessel registration algorithm capable of running in frame rate (see section 5.4). Previous processing steps require vessel segmentation and guidewire or device marker segmentation. A novel vessel detection filter, *directional stamping*, is presented. This method is a multi-scale vessel detection filter that outputs for each pixel the probability of this pixel belonging to a vessel. It is shown that directional stamping combined with a vessel resemblance filter based on the Hessian matrix outperforms using directional stamping only. The proposed method is compared with other detection filters and it is shown that directional stamping outperforms current state of the art.

In section 5.2.2 a fast version of directional stamping is used to automatically segment the guidewire, which is also a dark elongated structure. The guidewire segmentation is applied only in a ROI centred around the guidewire centre of gravity. Unlike previous state of the art, the proposed method results in high values for guidewire detection and is capable of running in frame rate for real-time guidewire detection.

In section 5.3 a new method based on temporal dynamics (*pdf* estimation from past gray values) and shape features (template matching) is presented to automatically segment radio-opaque device markers in x-ray interventional frames. We showed that spatial filtering is highly dependant on the correct selection of the shape and size of the template. Our method outperforms standard template matching and does not rely on user interaction.

This method needs N frames to be initialised and be able to calculate an estimation of the pixels *pdf*. This initialisation requirement leads the method to false detections until a representative sample of pixel values has been stored. The real-time capabilities of the algorithm after initialisation have not been tested. Possible improvements would be to further reduce the computation by performing the calculations only around a square ROI centred at the tracked feature.

We presented in section 5.4 a method to perform device to vessel registration. This method bases on the previously presented algorithms for vessel, guidewire and instrument marker segmentation. A gray-level multiplicative distance transform is presented that provides a convex search space for robust and fast convergence to the optimum overlap solution.

Chapter 6

X-ray and IVUS fusion (XIVUS)

As mentioned in chapter 3, angiography yields a good overview of the artery tree, but does not provide information about the plaque type or its exact geometry. IVUS exploration can be performed relatively easy [12][133] and provides information about vulnerable plaque, calcium and plaque morphology [12]. By fusing x-ray and IVUS, the ability to provide accurate cross section information about the local (IVUS) and global (x-ray) shape of the vessel is combined. Previous work addresses 3D vessel reconstruction for diagnosis by biplane angiograms and IVUS fusion. In this chapter we present XIVUS (X-ray and IVUS) [134][135], an approach to fuse IVUS and x-ray *images* for enhanced diagnosis and positioning during PTCA interventions.

In the previous chapters we have presented the necessary framework and algorithms to characterise each acquired image by a state vector $\vec{\xi}$ that characterises the heart state (chapter 3). Our main concern has been to provide an estimation for heart and respiration phases δ and β (section 3.6). Assuming the other components of $\vec{\xi}$, such as table position or patient position, do not vary during and between acquisitions, δ and β are enough to characterise the heart state and represent each acquired frame in a two-dimensional phase-space, where nearby points represent images in which the heart has a similar state. When overlaying an angiogram and a fluoroscopic frame with similar heart state, $\vec{\xi}_a \approx \vec{\xi}_f$, navigation and instrument positioning are effectively aided. For accurate overlay a respiration compensation method was presented in [73] and the past chapter was devoted to novel instrument and guidewire to vessel registration algorithms.

Arising from the presented algorithms for angiogram to fluoroscopic over-

lay based on estimation of $\vec{\xi}$ we propose in this chapter to apply said algorithms to fuse angiograms and interventional fluoroscopic frames with fluoroscopic frames acquired during IVUS probe pullback, hereafter referred to as *IVUS pullback frames* or *pullback frames*. XIVUS can improve diagnosis, by reconstructing the lesion and segmenting its most notable features; and aid precise treatment, by displaying the IVUS frame that corresponds to the current instrument position, which would prove useful in interventions like brachytherapy or during positioning of drug-eluting stents.

The proposed method requires fluoroscopy acquisition at the beginning, end, and optionally during a constant speed pullback of an IVUS probe. We propose to aid diagnosis by overlaying angiograms to IVUS pullback frames, relating the IVUS probe positions on the fluoroscopic frames to specific locations along the vessel the probe is being pulled back through, and therefore the acquired relating IVUS frames to said positions. Likewise, by overlaying interventional fluoroscopic frames with IVUS pullback frames or angiograms from the previous step we can relate the current instrument position with IVUS frames.

IVUS pullback acquisition requires one additional fluoroscopic acquisition and therefore, additional x-ray dose for the patient and the physician, without any validated clinical benefit available. Therefore the algorithm was tested and validated on a heart phantom featuring independent heart and breathing motion, and equipped with an ECG signal generator.

In section 6.1 we give a short introduction to IVUS benefits as a complement to x-ray angiography. In section 6.2 we comment on previously proposed methods for IVUS and angiography fusion. The interventional workflow and the advantages of combining IVUS with angiographic and fluoroscopic images are discussed in section 6.3. In sections 6.4 and 6.5 new applications and algorithms for XIVUS as a diagnostic tool and accurate positioning aid, respectively, are presented. The proposed algorithms are tested on a heart phantom and the results presented in section 6.6. Section 6.7 discusses the 3D capabilities of XIVUS. We present a three-dimensional active surface lumen segmentation method in section 6.8 and we conclude in section 6.9.

6.1 Introduction

Table 6.1 summarises the main features of diagnostic x-ray and IVUS. The leftmost column lists a number of desirable imaging and navigation capabilities in catheterisation interventions. The central and rightmost columns state whether angiography and IVUS, respectively, can provide said capabilities. A tick ($\surd$) signifies the ability of supporting the wished feature, whereas a cross ($\times$) stands for the inability of the corresponding imaging modality of providing the desired capability. Coronary angiography helps to diagnose and localise stenoses and provides a suitable map of the vessel-tree topology that is afterwards utilised to navigate and position the instruments during treatment. Monoplane angiography provides a 2D projection of a 3D volume, which might hinder the visualisation of very eccentric and tortuous plaque. Furthermore, angiography does not provide information about the lumen area, plaque composition and morphology, the presence of calcification or the distribution of diffuse plaque, which are of diagnostic value.

X-ray vs. IVUS	Angiography	IVUS
Vessel topology	$\surd$	$\times$
3D capability	not accurate	needs x-ray
Lumen area	underestimated	$\surd$
Plaque morphology	$\times$	$\surd$
Diffuse plaque	difficult	$\surd$
Calcium	$\times$	$\surd$
Intervention outcome	inaccurate	$\surd$

Table 6.1: Comparison of IVUS and angiography.

An IVUS probe can acquire cross-sectional images of an artery at a given location with no relation to vessel topology. If an IVUS probe is pulled back from the distal to the proximal end of a stenosis, IVUS images along the pullback path are acquired. Some commercially available IVUS systems stack the IVUS images along a straight axis (see figure 6.1) assuming that the catheter follows a straight path within the vessel, which is far from reality.

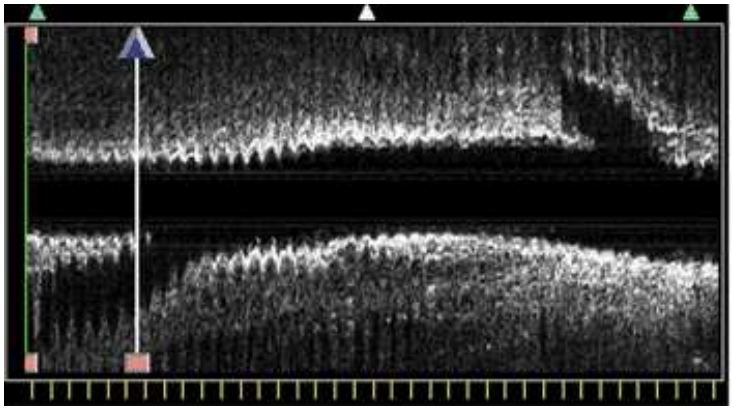

Figure 6.1: IVUS vessel reconstruction in JOMED's *In-Vision Gold Intravascular Ultrasound System.* [5].

6.2 Previous work

Current methods for fusion of angiography and IVUS [56][57][60][61][62][63][64][65][66][67][68] aim at providing an accurate 3D reconstruction of the imaged vessel for stenosis characterisation. These algorithms rely on a constant speed pullback, such that the IVUS images are acquired at equally spaced intervals. Different approaches have been proposed to calculate the pullback path:

1. methods which execute a constant surveillance of the pullback with fluoroscopy [60][61] or biplane angiography [62];
2. methods that perform acquisitions at the beginning and the end of the pullback and calculate the vessel centerline in a pair of biplane [63] or stereo angiograms [64];
3. methods that perform acquisitions at the beginning and the end of the pullback and segment the IVUS catheter from biplane angiograms [56][57][65][66][67];
4. methods that perform acquisitions at the beginning and the end of the pullback and calculate a path of minimum bending energy in the arteries segmented from two biplane angiographies [68].

Once the pullback path is established, the IVUS frames are placed at the segmented probe positions (methods of type 1) or equally spaced along the calculated pullback path (methods of type 2, 3 and 4).

Another consideration about the geometry of the pullback path must be taken into account to properly reconstruct the imaged artery. The pullback path can be modelled as a 3D curve with non-zero curvature and torsion. The curvature is a measure of how fast the curve turns, and the torsion describes how a vector tangent at each point along the curve *cork-screws* about the curve. The torsion of a curve is 0 if and only if the curve sits inside a fixed plane [65], which is in general not the case for the pullback path. Assuming the catheter is a torsion-free generalised cylinder, the effect of it being pulled along a non-zero torsion curve is that the catheter tip rotates whenever it enters a new plane [65]. This effect causes a relative rotation between consecutive IVUS images along the pullback path. By defining at each point a reference frame that rotates and changes orientation according to the curve's torsion and curvature, the relative rotation between IVUS frames can be calculated [56][65][66][67][57]. Such a reference frame is the Frenet-Serret frame [136] (see section 6.7.1), that defines at each position along the curve a positively oriented orthonormal frame composed of a tangent, a normal and a binormal vector that depend on the curve's curvature and torsion.

Once the pullback path is calculated and the relative orientation of the IVUS images known, the absolute orientation of the reconstructed vessel segment is still undetermined, because the first IVUS frame is positioned at the beginning of the pullback path on a plane perpendicular to the tangent at that point, but with a random rotation of this plane around the tangent. Subsequently, only relative rotations between IVUS frames are calculated with the Frenet-Serret frame. It has been proposed to estimate the absolute orientation by either backprojecting the vessel segment onto the acquired biplane angiograms [62] or calculating the minimum displacement energy [56][66][67][57].

Previously proposed methods to fuse IVUS and angiography result in a 3D model of the imaged artery segment to aid diagnosis. These methods rely on biplane systems [62][56][66][67][57], which are expensive and are not widespread for cardiac interventions [133][137]. Furthermore, these methods are ECG-gated and require breath hold during the angiographic and IVUS pullback acquisition. We propose a method to fuse projection x-ray with IVUS images for diagnosis and accurate positioning from images acquired during free respiration of the patient. To the author's knowledge, no algorithm has been proposed that employs previously acquired IVUS images to aid instrument positioning or creates fusion images in the way described below in this chapter.

In this chapter methods for IVUS and x-ray fusion of projection images for diagnosis *and* accurate instrument positioning are presented.

6.3 XIVUS workflow

One of the goals of XIVUS is to provide maximum functionality interfering as little as possible with the current Cathlab workflow. The need for acquiring IVUS images interferes with the current standard interventional workflow (see section 2.3.4). In order to apply our method we need angiography, pullback fluoroscopy and interventional fluoroscopy acquisitions with the same geometry settings and their associated ECG signal. Our intended workflow proposes the following steps:

1. acquire diagnostic angiography sequences of the coronary arteries to identify and locate the candidate lesions;
2. select an angiographic sequence in which the lesion is most clearly depicted;
3. select the imaging geometry of said angiographic sequence as the imaging geometry for the rest of the procedure;
4. place the guidewire past the lesion under fluoroscopic survey;
5. insert an IVUS probe until the distal end of the stenosis, and execute automated pullback with synchronous fluoroscopic image acquisition;
6. perform fluoroscopy-guided intervention, supported by angiographic and IVUS data.

In step 5, IVUS frames that will be later fused with the angiographic and fluoroscopic frames are acquired. This extra step lasts a variable amount of time depending on the length of the lesion. If we consider an average stenosis length of 12.4mm [138], then the pullback will last $24.8s$ at a pullback speed of 0.5mm/s. As will be explained later, continuous fluoroscopy surveillance during the pullback as well as fluoroscopic acquisition at the beginning and the end of the pullback for a whole heart cycle are contemplated in the proposed algorithm. By using an automated constant speed pullback device, the position of the IVUS probe along the vessel can be derived from the time function of the pullback [139] and thereby minimising the amount of extra x-ray radiation employed.

In the existing literature about IVUS and angiography fusion [56][60][62] [65][66][67][57][68][140] the workflow differs from the one suggested here after step 4. In the most common approach [56][65][66][67][57] a sheathed IVUS catheter is used to provide a stable pullback path. After locating the lesion and placing the IVUS probe at its distal end, a biplane angiographic acquisition with diluted contrast dye is performed and the IVUS catheter is segmented from the acquired images. The sheath provides a stable pullback path and remains in place while the IVUS probe is pulled back. Sheath and vessel reconstruction are ECG-gated and all acquisitions assume breath-hold so that respiration motion is avoided. This extra angiographic acquisition is performed to obtain the pullback path because the IVUS catheter deforms the vessel and therefore, an angiography taken prior to the introduction of the IVUS catheter would yield inaccurate geometric information for accurate pullback path reconstruction [65].

It could be debated whether we should take a similar approach and perform an extra angiography acquisition once the IVUS catheter is in place. If our goal were to perform a very exact three-dimensional reconstruction of the vessel for diagnosis and stenosis measurements, we would indeed need an accurate pullback path estimation. But the method we propose does not aim at an accurate 3D reconstruction of the imaged vessel (the 3D possibilities of our method are discussed on section 6.7) but at the registration of each IVUS image to a position in an angiogram for subsequent fused visualisation. It should also be considered that the stenosis length can already be measured without any 3D reconstruction by multiplying the time needed to image the lesion by the constant pullback speed. Furthermore, we aim at registering the instrument position during an intervention with a previous position of the IVUS probe. In either case, the deformation caused by the IVUS catheter and the interventional device will be different. Therefore, even after estimating a very accurate pullback path, the deformation of the vessel due to the interventional device would introduce errors in the registration. We consider that the benefit of very accurate reconstruction, compared to that of accurate-enough reconstruction is not sufficient to justify additional x-ray and contrast agent dose for the patient and the physician and the interference in the established Cathlab workflow.

6.4 XIVUS as diagnostic aid

The objective of fusing pre-interventional projection x-ray angiograms with IVUS is to provide the physician with a diagnostic tool that relates a match-

ing IVUS image to each point along a diseased artery. Previous approaches aim at a 3D vessel reconstruction in breath-hold. XIVUS as diagnostic aid registers, for each heart phase, pullback images with an angiogram acquired with the same system geometry while the patient is allowed to breath freely.

Each acquired pullback fluoroscopy frame is characterised by its heart and respiration phase, and the position of the IVUS probe obtained with the methods presented in sections 4.1, 4.5 and 5.3 respectively. The heart phase is quantised in N values. For each quantised heart phase, the acquired pullback fluoroscopy images are registered to an angiogram with the same heart phase using a respiration compensation algorithm as presented in [73] (section 5.1) or [141] and anatomical device to vessel registration (section 5.4). This IVUS pullback to angiogram registration results in a collection of pullback paths for different respiration phases. IVUS images are registered to the points belonging to said pullback paths. The *pullback path* is reconstructed by B-spline interpolation between the reconstructed probe locations, taking into account that the pullback is performed at constant speed.

6.4.1 XIVUS algorithm for diagnosis

As a first step, an angiographic acquisition with appropriate geometry settings is selected and a set of angiograms is acquired $\{\mathcal{A}\}$. Then, the sequence phase-space is calculated and each frame is characterised by its respiration and heart phase, $\{\mathcal{A}\} = \{A(\delta_i^a, \beta_i^a)\}$ (see figure 6.2). The heart phase is quantised in N possible values and N reference angiograms $\{A(\delta_n, \beta_n^a), n = 1 \dots N\}$ are selected. These reference angiograms will serve to calculate N representative pullback paths $\{\Pi(\delta_n, \beta_n^a)\}$ at respiration phases $\{\beta_n^a, n = 1 \dots N\}$ by registering the pullback fluoroscopy frames to said reference angiograms.

Once the stenosis has been identified, a guidewire is advanced under fluoroscopy past the stenosis (see chapter 2, section 2.3.4). The IVUS probe is advanced over the guidewire and positioned at the distal end of the stenosis, the IVUS catheter is pulled slightly before starting the pullback to stretch the IVUS catheter and guarantee that the beginning of the pullback at the proximal end is translated to a movement in the distal end [58][62], and a constant speed pullback begins while IVUS images and the ECG are captured by a suitable system. Fluoroscopic frames for a complete heart cycle are necessary from at least the beginning and the end of the pullback to identify the beginning and end of the pullback path. Optionally, the fluoroscopy acquisition may be performed continuously during the pullback to improve

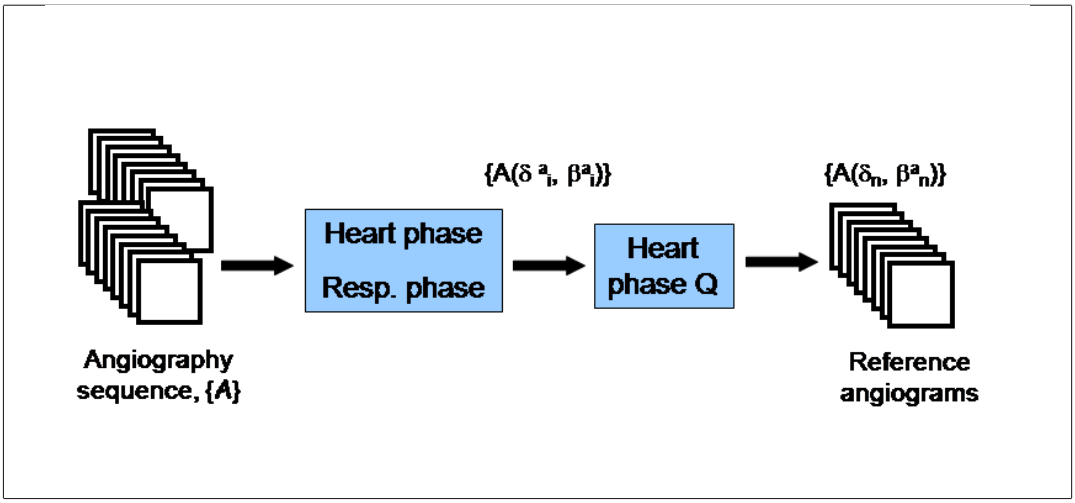

Figure 6.2: First step for XIVUS vessel reconstruction: an angiographic sequence is acquired and its phase-space calculated. The heart phase is quantised in N values and N reference angiograms are selected.

the accuracy of the pullback path reconstruction. We assume the system is capable of providing a *time stamp* that relates the acquisition instants of the fluoroscopic frames and the IVUS frames. As a result we obtain a collection of x-ray pullback frames $\{\mathcal{P}\}$ characterised by their heart and respiration phase, the probe position (x_j^p, y_j^p), and a time stamp t_j^p, $\{\mathcal{P}\} = \{P(\delta_j^p, \beta_j^p)_{(x_j^p, y_j^p, t_j^p)}\}$ (see figure 6.3). Each acquired IVUS frame I is characterised by its heart phase and a time stamp t_k^p that relates it to the fluoroscopy sequence $I(\delta_k^p)_{t_k^p}$.

Respiration compensation as explained in section 5.1 is applied to the set of angiographic frames $\{\mathcal{A}\}$ for each δ_n to estimate a heart-position function $\tilde{r}_n$ from the respiration phase of the angiograms $\{\beta_n^a, n = 1 \ldots N\}$. From this heart-position function the rigid body transformation to compensate for a difference in respiration phase $\Delta\beta$ between a reference angiogram with heart-phase δ_n^a, $A(\delta_n, \beta_n^a)$, and a pullback fluoroscopic frame with similar heart-phase but different respiration phase $P(\delta_n, \beta_n^p)_{(x_j^p, y_j^p, t_j^p)}$ is calculated . The translation vector $\vec{m}_j^n$ for respiration compensation is calculated as

$$\vec{m}_j^n = (m_{x,j}^n, m_{y,j}^n) = \tilde{r}_n(\beta_n^a) - \tilde{r}(\beta_n^p) \quad \forall n = 1 \ldots N \tag{6.1}$$

where

- $\vec{m}_j^n = (m_{x,j}^n, m_{y,j}^n)$ is the rigid-body translation vector for respiration compensation, and

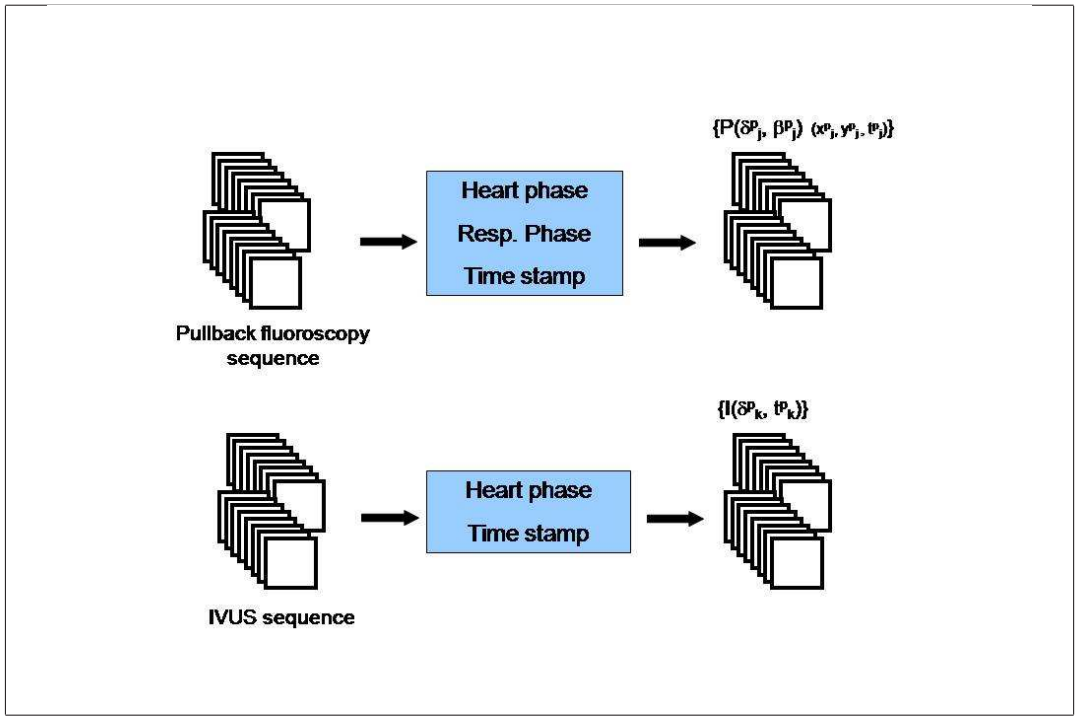

Figure 6.3: Second step for XIVUS vessel reconstruction: the pullback is imaged with fluoroscopy and its phase-space is computed.

- $\tilde{r}_n(\beta)$ is a heart-position function as in [73] (section 5.1, equation 5.3) or [141] for heart-phase δ_n.

Once the translation vector is known for each pullback frame with heart phase δ_n, the respiration-compensated positions of the IVUS probe, $(\hat{x}_j^p, \hat{y}_j^p)$, are calculated as

$$(\hat{x}_j^p, \hat{y}_j^p) = \vec{m}_{j,n} + (x_j^p, y_j^p) \tag{6.2}$$

where (x_j^p, y_j^p) is the actual position of the IVUS probe in $P(\delta_n, \beta_j^p)_{(x_j^p, y_j^p, t_j^p)}$. Anatomical device to vessel registration as explained in section 5.4 can be applied afterwards to further improve the accuracy of the reconstruction.

If for each heart phase δ_n only the respiration-compensated beginning and end positions of the pullback are available, $(\hat{x}_b^n, \hat{y}_b^n)$ and $(\hat{x}_e^n, \hat{y}_e^n)$ respectively, the pullback path $\Pi(\delta_n, \beta_n^a)$ is reconstructed from these two points and the segmented vessel (see section 5.4) visible in the angiogram by calculating its centerline [63][64], or the path of least bending energy considering the stiffness of the IVUS catheter [68] between said two points. In this case, the

pullback path can be expressed as a collection of points such as

$$\Pi(\delta_n, \beta_n^a) = \{(\hat{x}_b^n, \hat{y}_b^n), (\hat{x}_e^n, \hat{y}_e^n)\} \cup \{(x_{\text{interp}}, y_{\text{interp}})\} \tag{6.3}$$

where

- $\Pi(\delta_n, \beta_n^a)$ is the pullback path calculated for heart phase δ_n; it is expressed as a function of the respiration phase of the reference angiogram β_n^a because all the pullback frames with heart phase δ_n are respiration-compensated to match said angiogram with said respiration phase,
- $(x_{\text{interp}}, y_{\text{interp}})$ are the points resulting from interpolating between the respiration-compensated IVUS probe positions at beging and end of pullback, $(\hat{x}_b^n, \hat{y}_b^n)$ and $(\hat{x}_e^n, \hat{y}_e^n)$, with the constraint that $(x_{\text{interp}}, y_{\text{interp}})$ must belong to the vessel centerline or to the path of minimum bending energy.

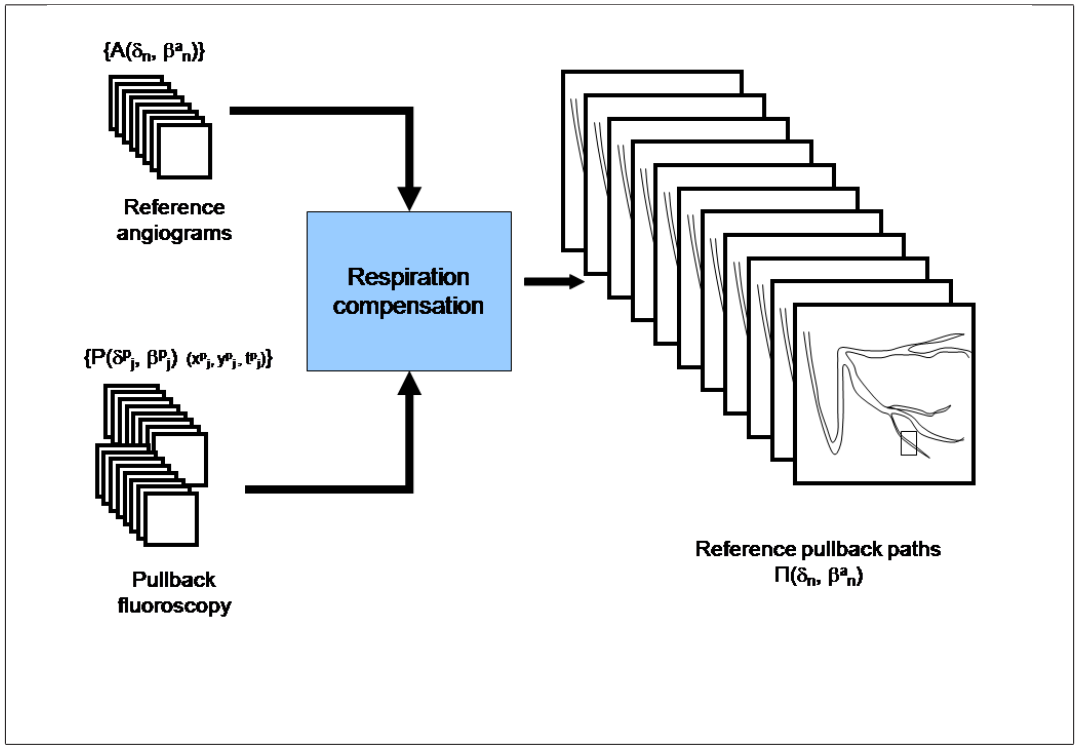

Figure 6.4: Pullback path calculation from the set of reference angiograms and the positions of the IVUS probe on the fluoroscopy pullback frames. The pullback path is highlighted with a square.

The acquisition points for each IVUS image are calculated by taking into account the constant pullback speed, v_{pull}. Let L be the length of the pullback path calculated as

$$L = \int_{t_b}^{t_e} v_{\text{pull}} dt = v_{\text{pull}}(t_e - t_b) \tag{6.4}$$

where t_b and t_e are the times when the pullback begins and ends respectively. If n_{IVUS} is the number of IVUS images acquired during the pullback for heart phase δ_n, then the IVUS frames are positioned along $\Pi(\delta_n, \beta_n^a)$ at intervals L/n_{IVUS}.

If the pullback was constantly imaged with fluoroscopy, then the pullback path for δ_n can be reconstructed by interpolation between the segmented probe positions during pullback after respiration compensation and/or device to vessel registration, $\{(\hat{x}_j^p, \hat{y}_j^p)\}$, where $j = 1 \dots k$, taking into account the known pullback speed and the constraint that the points of the pullback path should be inside the vessel. We can express this as

$$\Pi(\delta_n, \beta_n^a) = \{(\hat{x}_j^p, \hat{y}_j^p)\} \cup \{(x_{\text{interp}}, y_{\text{interp}})\} \tag{6.5}$$

where in this case, $\{(x_{\text{interp}}, y_{\text{interp}})\}$ are the points obtained by cubic-spline interpolating between $\{(\hat{x}_j^p, \hat{y}_j^p)\}$.

At the end of this procedure, N representative pullback paths have been calculated by registering the pullback fluoroscopy acquisition to N representative angiograms with heart phase δ_n and respiration phase β_n^a. Each pullback path is linked to a collection of IVUS frames with the same heart phase acquired during the IVUS probe pullback. By presenting the results to the physician (see figure 6.5), the localisation of the beginning and the end of the lesion or the site of minimum lumen diameter in the angiogram can be supported.

Our method for IVUS and 2D angiographic registration for diagnosis support bases on previously presented algorithms and conforms to the proposed image registration framework. This thesis proposes to reconstruct a 2D pullback path from a sequence of projection angiograms with various heart and respiration phases. Methods presented in the literature require biplane Cathlab systems, which are uncommon in catheterisation procedures [137]. Our methods and algorithms involved are also applicable to biplane angiography if available. By using the proposed algorithms we can reconstruct the IVUS pullback path for different heart phases while the patient is breathing freely.

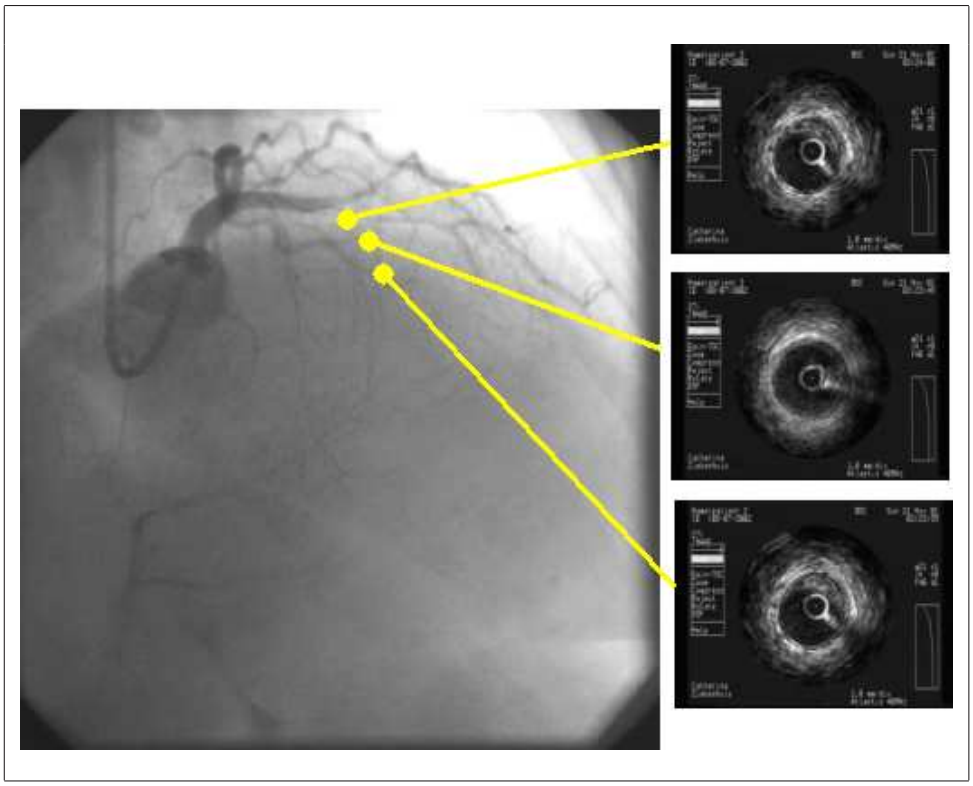

Figure 6.5: Possible user interface to aid diagnosis after IVUS and diagnostic angiography fusion. IVUS frames are registered with an angiography and can show the location of certain vessel features. The exact display of screen is in discussion with medical user interface experts.

6.5 XIVUS as positioning aid

In spite of the extensive literature on 3D vessel reconstruction fusing IVUS with angiography [56][57][60][61][62][63][64][65][66][67][68], to the author's knowledge no integrated system has yet been developed that can be applied during an intracoronary intervention for positioning support. We propose a method to register the instrument position in interventional frames with previously acquired IVUS images, in order to support accurate positioning of the interventional device.

In section 6.4 a method was introduced to reconstruct a 2D pullback path $\Pi(\delta_n, \beta_n^a)$ by registering each acquired fluoroscopic pullback frame with heart phase δ_n to a reference angiogram with the same heart phase, $A(\delta_n, \beta_n^a)$. In this section we propose a method to register each incoming live interventional

frame $F(\delta_l^f, \beta_l^f)$ to said pullback path in order to associate the current instrument position to a previously acquired IVUS frame. The overall result would be like having a (virtual) IVUS probe at the location of the interventional device. Hereby accurate instrument positioning can be achieved during the intervention.

6.5.1 XIVUS algorithm for intervention support

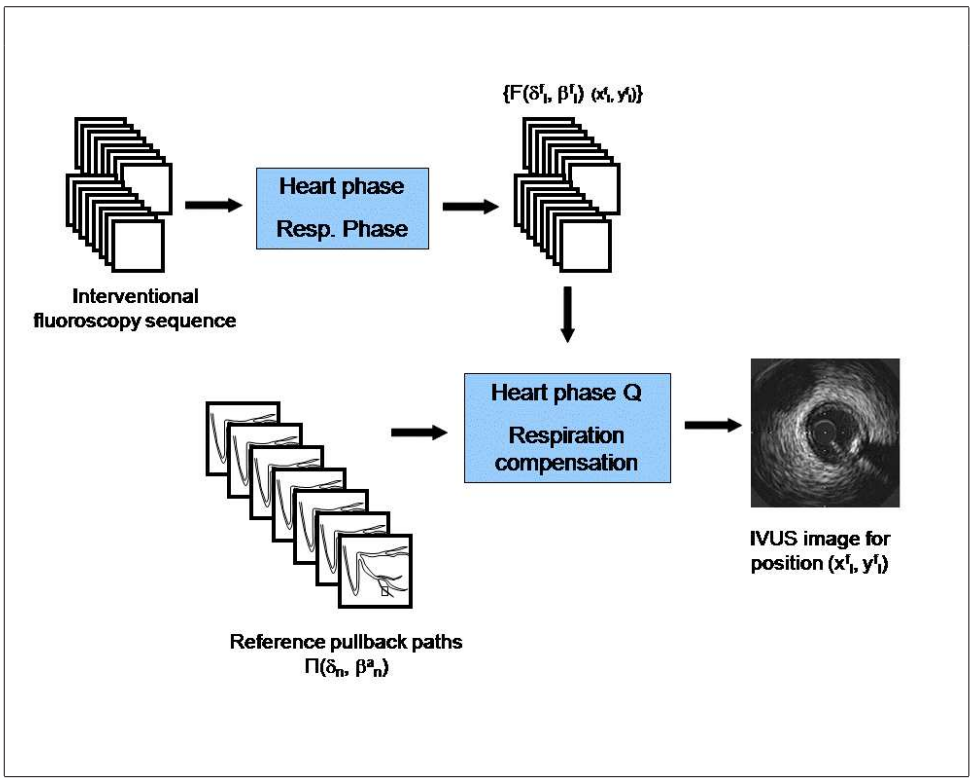

Figure 6.6: Registration of incoming live fluoroscopic frames with previously calculated pullback paths $\Pi(\delta_n, \beta_n^a)$.

A block diagram of the proposed method is shown in figure 6.6. In a previous step, N pullback paths have been calculated as explained in the previous section. For each newly acquired live fluoroscopic frame the hear and respiration phase δ_l^f and β_l^f are calculated. The heart phase is quantised into one of N bins as used for angiogram reference selection. By using respiration compensation, or anatomical device to vessel registration, or both,

the instrument is registered onto the pullback path $\Pi(\delta_n, \beta_n^a)$ for the corresponding heart phase.

The interventional device location in the fluoroscopic frame (x_l^f, y_l^f) is calculated in real-time with the methods explained in the previous chapter. The instrument position after respiration compensation and/or device to vessel registration is

$$(\hat{x}_l^f, \hat{y}_l^f) = \vec{m} + (x_l^f, y_l^f) \tag{6.6}$$

where (x_l^f, y_l^f) is the position of the interventional device in the interventional fluoroscopy frame and $\vec{m}$ is a rigid-body translation vector.

Equation 6.6 allows to register the current position of the device with a pullback path, $\Pi(\delta_n, \beta_n^a)$, and since each point along a pullback path was associated in a previous step to an IVUS frame (see section 6.4), it is possible to relate the current position of the instrument, (x_l^f, y_l^f), with an IVUS image previously acquired at the same position within the vessel.

This novel algorithm for positioning support during PTCA interventions is the first of its kind to fuse IVUS and interventional frames acquired during catheterisation interventions. If this computation is performed in real-time, then it would allow the system to provide the physician the IVUS frame that corresponds to the current instrument position during the intervention, improving instrument positioning accuracy, reducing x-ray dose for patient and physician and reducing the amount of contrast agent used.

6.6 XIVUS phantom and measurements

In vivo data acquisition for new and untested imaging protocols is in most cases not possible. Physicians are not willing to use a novel x-ray imaging protocol on a patient unless it has been clinically tested before. For the first tests of a new imaging protocol as the one presented here it is therefore common to perform the tests on a realistic phantom. Because of the lack of appropriate clinical data, a phantom was built and imaged simulating diagnostic and fluoroscopic IVUS pullback acquisitions with the same geometry settings. This phantom was designed so that independent heart and respiration motion were present.

Subsection 6.6.1 describes the different parts of the phantom that was built to tests XIVUS. Subsections 6.6.2 and 6.6.3 present the performed measurements and experiment setup respectively. Subsection 6.6.4 presents the

results from applying our XIVUS algorithm to the heart phantom.

6.6.1 XIVUS phantom

A phantom was built[1] which models respiratory and contraction movement as two independent periodic movements, and features constant speed automatic pullback. This phantom can be see seen in figure 6.7, where two different views, (a) and (b), of the phantom are shown.

A coronary vessel was simulated with a transparent plastic tube with an internal diameter of 5mm and with alternate opaque and transparent sections of known length (see figure 6.8), which were marked with partially radio-opaque metal bands. The IVUS probe and acquisition system were replaced by a fiber optic with a radio-opaque marker (see figure 6.9a) and a photodetector (see figure 6.9b) respectively. The vessel was illuminated by an external bright illumination source and the photodetector provided a signal proportional to the amount of light gathered by the fiber optic tip. As the fiber optic is pulled back, it transverses the dark and bright bands of the vessel and the signal provided by the photodetector changes accordingly.

The ECG is generated with an U-shaped infrared movement detector (see figure 6.10). When the flux of infrared light between the arms of the detector is interrupted (see figure 6.10b), there is a short pulse in the output signal.
The phantom can be dismounted in three parts:

1. the *rails* (see figure 6.11),
2. a *pullback device* and *respiration generator* (figures 6.12 and 6.13, respectively), and
3. a *moving part* (see figure 6.14).

The rails (figure 6.11a and 6.11b) provide a stable path for the backward and forward movement of the moving part. The pullback device and respiration generator can be seen in figure 6.12. One motor (right) is used to drive the respiration movement and the pullback. Both movements are coupled by two dented wheels and a chain (centre). The pullback has been implemented by a string that wraps around a turning cylinder (left) and is tied to the fiber

[1]The phantom was build with Fischer Technik®

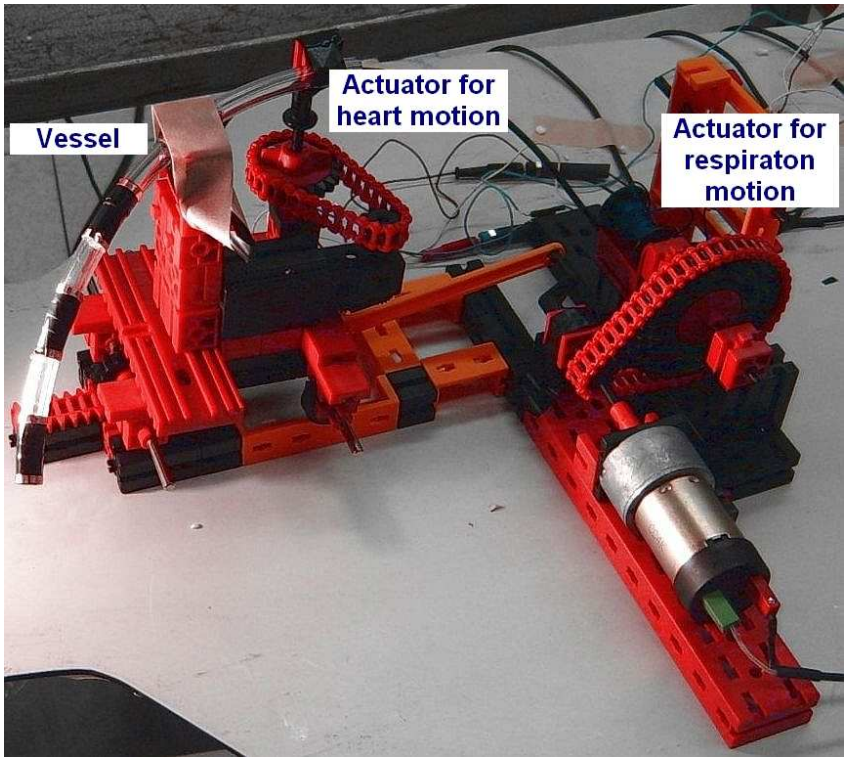

(a)

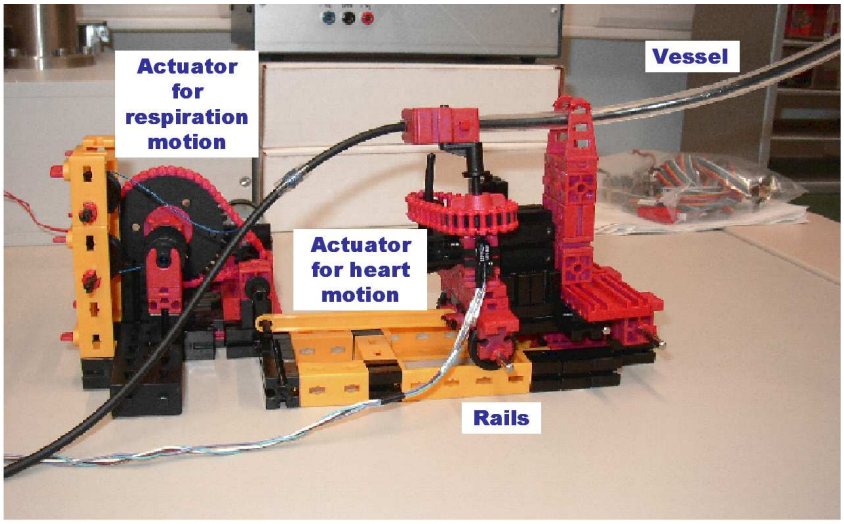

(b)

Figure 6.7: Two different views (a) and (b) of the XIVUS phantom that features independent and periodic respiration and contraction motions.

Figure 6.8: Plastic tube that mimics a vessel and that was attached to the phantom. Alternate dark bands and radio-opaque metal bands provide spatial reference.

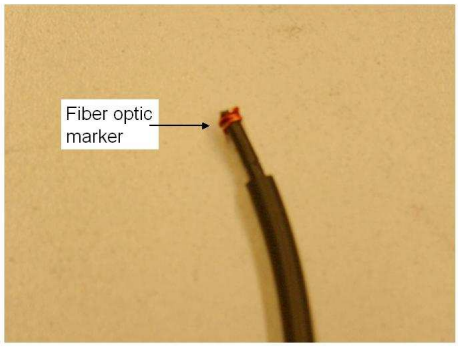

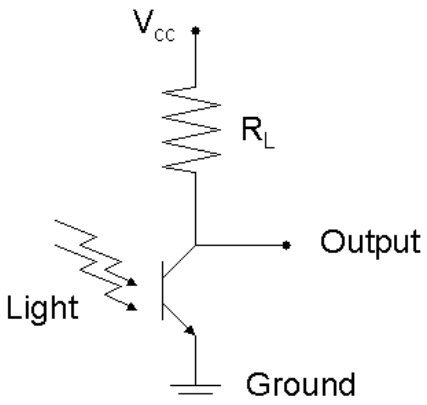

(a) (b)

Figure 6.9: (a) Fiber optic tip with a radio-opaque marker. The marker is made with a metal wire wrapped around the fiber optic. (b) Photodetector scheme attached to the end of the fiber optic. This photodetector gives a signal proportional to the amount of light gathered by the fiber optic tip.

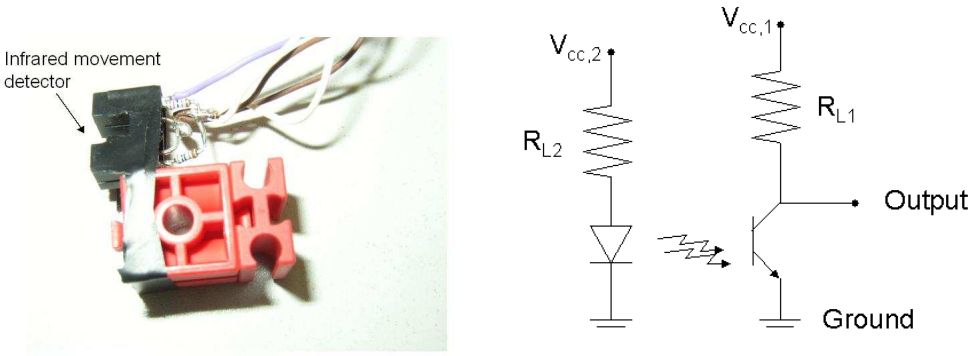

(a) (b)

Figure 6.10: (a) Infrared movement detector that generates an ECG signal. (b) Electrical scheme of the infrared movement detector.

optic. In figure 6.13 we can see an elevation 6.13a and side view 6.13b of the pullback device.

The moving part (see figure 6.14a) has an independent motor that simulates the heart contraction movement by turning a wheel (see figure 6.14b). This wheel has a jutting small screw (see figure 6.15a) that crosses the gap of the U-shaped infrared movement detector attached to the moving part (see figure 6.15b). Whenever the gap is crossed, a pulse is generated that simulates an R peak.

The moving part goes on top of the rails (see figure 6.16a) and moves backward and forward over the rails according to the respiration generator. It is connected to the respiration generator as shown in figure 6.16b. A plastic tube mimicking an artery (figure 6.17a) is attached to the heart-contraction motor (figure 6.17b) and moves according to it at the same time that the moving part moves according to respiration. To simulate the presence of a diaphragm for ulterior analysis of the x-ray images, metal bars were used as the wheels' axes (see figure 6.18).

6.6.2 XIVUS measurements on the heart phantom

The methods presented in sections 6.4 and 6.5 were tested on our heart phantom. Three angiographic acquisitions and the corresponding fluoroscopic pullback sequences were acquired with a monoplane Philips Cathlab. We

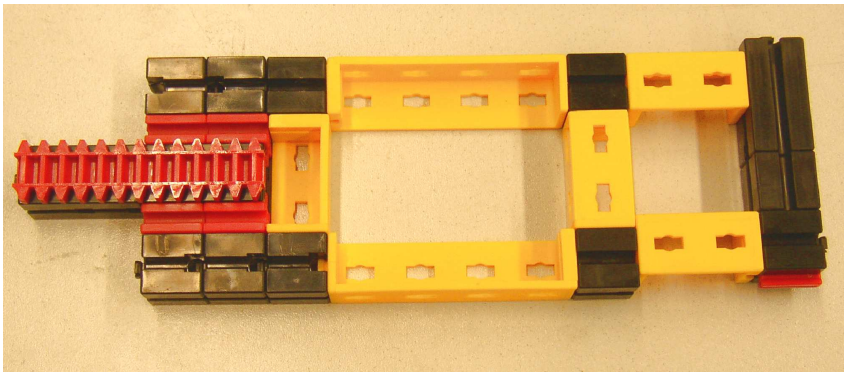

(a)

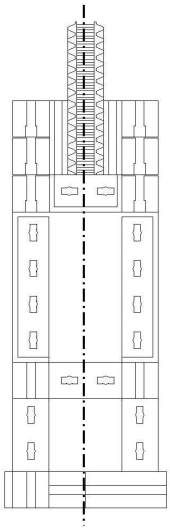

(b)

Figure 6.11: (a) The rails provide a stable path for the moving part. (b) Plan of the rails.

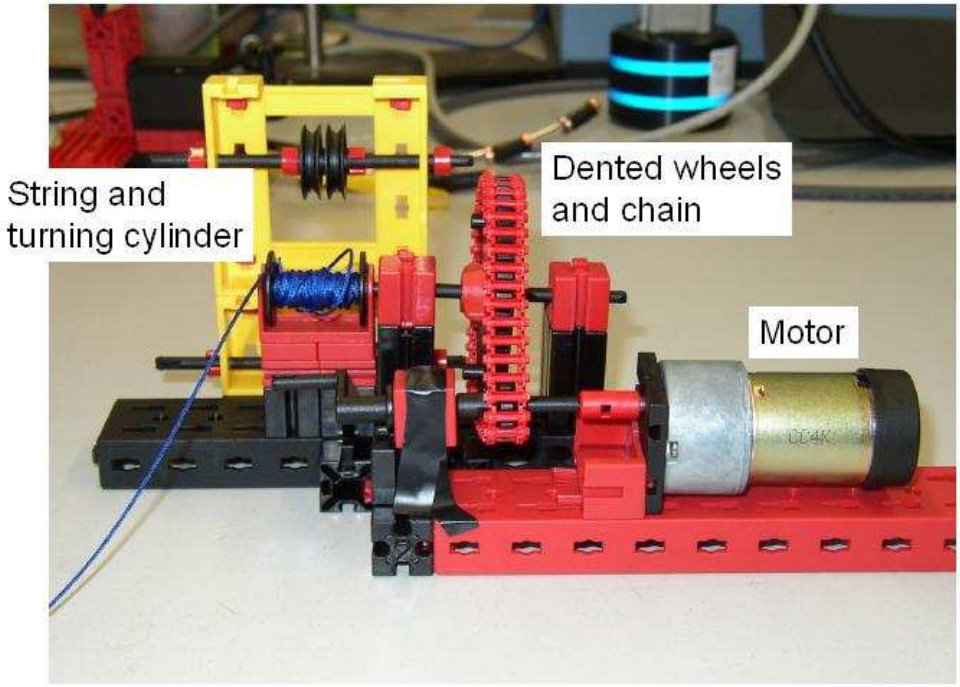

Figure 6.12: Respiration and pullback device. (a) A motor (right) drives the respiration movement and the pullback. Both movements are coupled by two dented wheels and a chain (centre).

can see in figure 6.19a and figure 6.19b examples of one acquired angiographic and one pullback fluoroscopic frames with the same heart phase δ_n but different respiration phase. We can see in figure 6.19a the plastic tube, the segmented radio-opaque markers and the metal wheel axes that provide a respiration marker. In figure 6.19b we see a pullback frame acquired with the same geometry. The fiber optic marker is visible within the plastic tube.

Automated pullback of the fiber-optic probe was provided by the phantom. No time-stamp to relate the fiber-optic output to the acquired frames was provided, so it was taken care that image acquisition, probe pullback and fiber optic acquisition were started simultaneously.

Figure 6.20a shows the difference between figures 6.19a (angiography) and 6.19b (pullback fluoroscopy) before respiration compensation, whereas figure 6.20b displays the difference between figures 6.19a and 6.19b after translat-

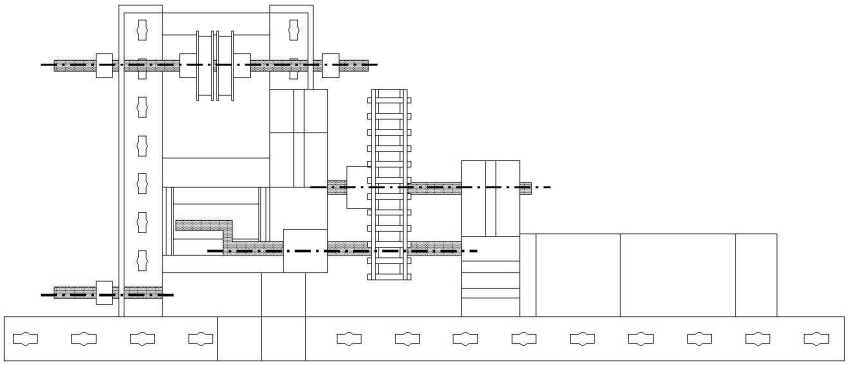

(a)

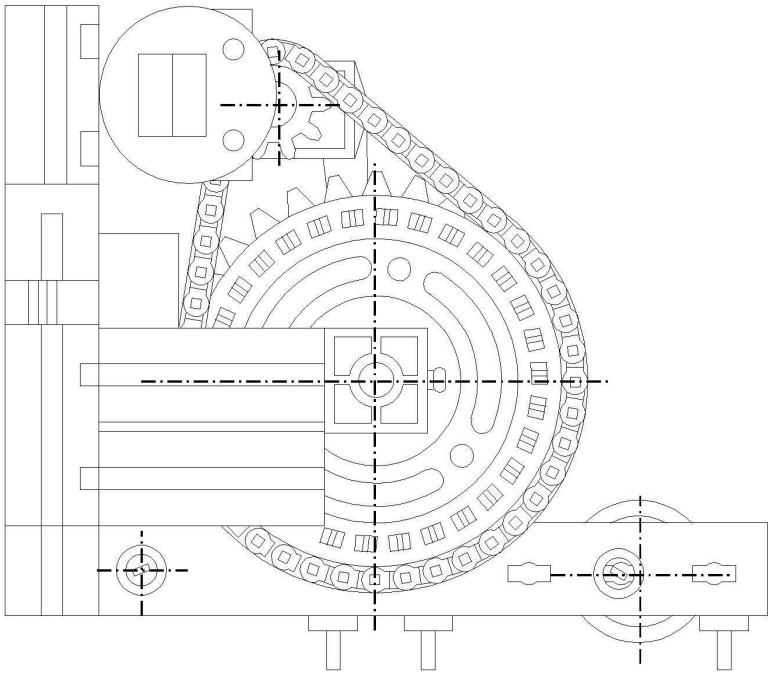

(b)

Figure 6.13: (a) Elevation of the respiration and pullback device. (b) Left view of the respiration and pullback device.

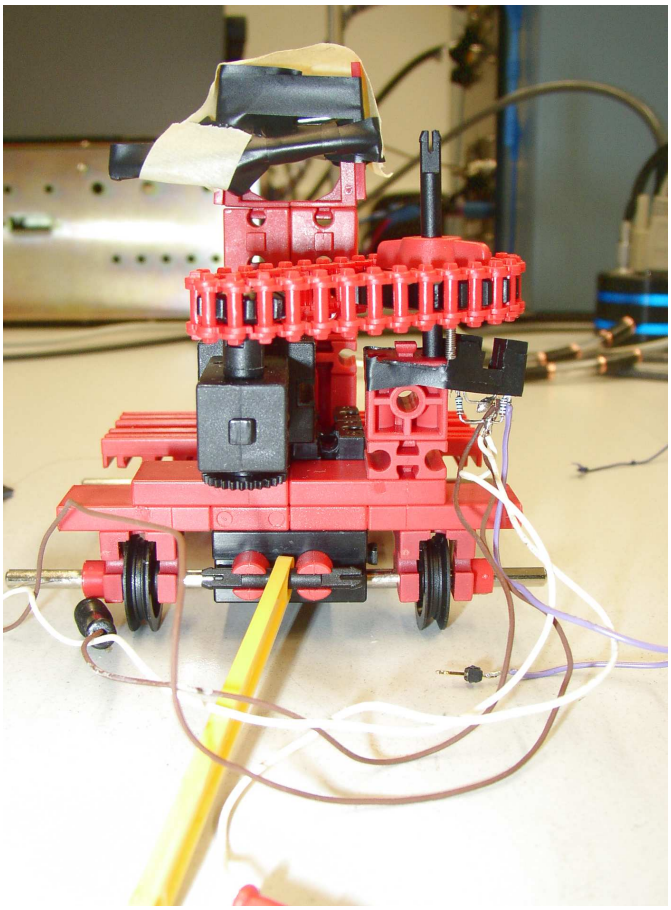

(a)

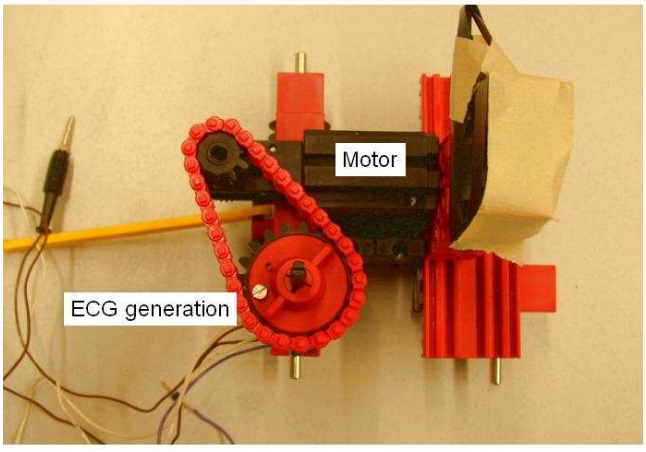

(b)

Figure 6.14: Moving part of the phantom. (a) and (b) show to different views of the moving part of the phantom. This part moves backward and forward on top of the rails (figure 6.11) with the respiration movement (see figure 6.12) and is equipped with an independent motor to simulate the heart contraction motion.

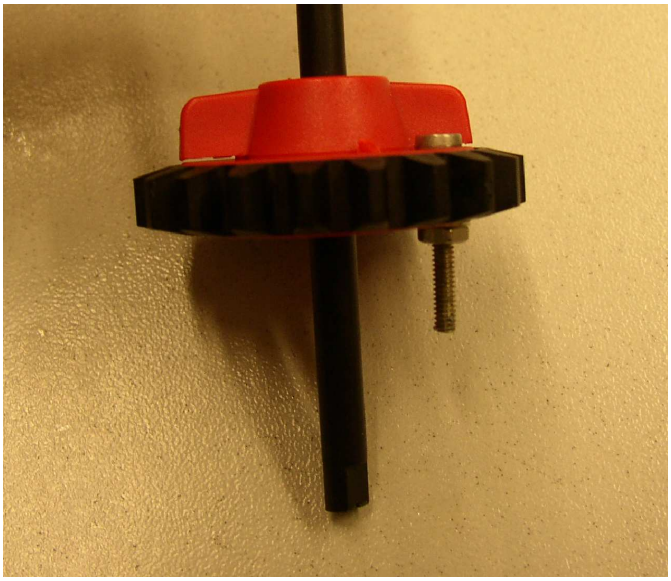

(a)

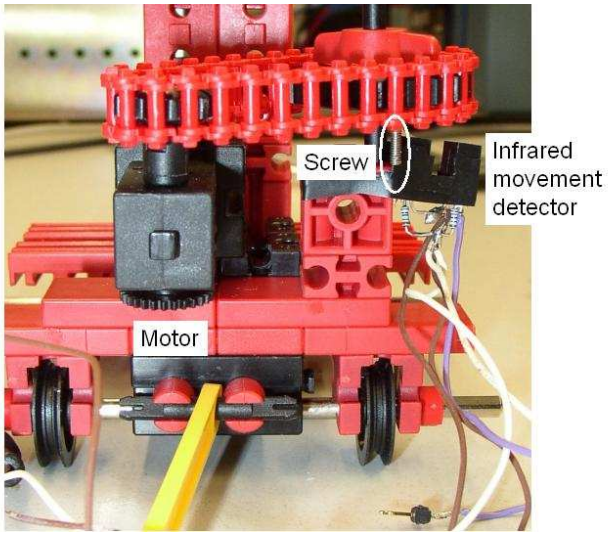

(b)

Figure 6.15: ECG generation on the moving part of the phantom. A screw (a) is attached to a wheel that turns according to the ECG generating motor, (b). When the screw crosses the gap of the U-shaped infrared detector, a pulse is generated.

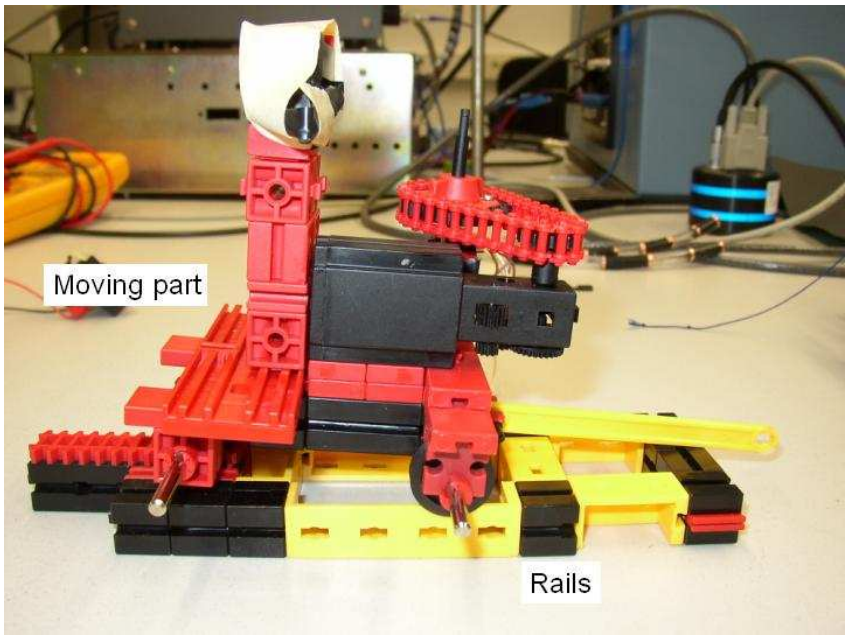

(a)

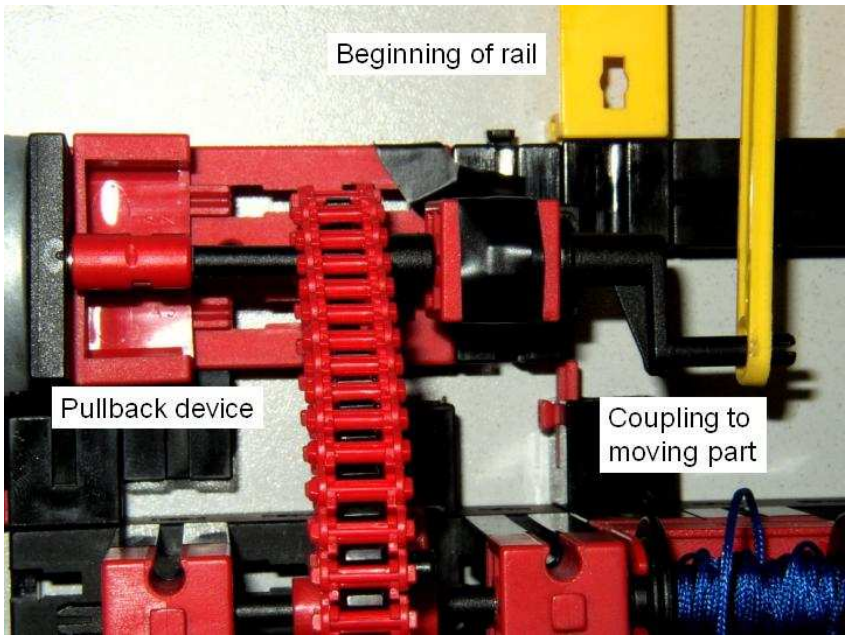

(b)

Figure 6.16: (a) The moving part of the phantom moves on top of the rails (see figure 6.11). (b) The moving part of the phantom is coupled to the pullback and respiration device (figure 6.12) as shown on the right.

(a)

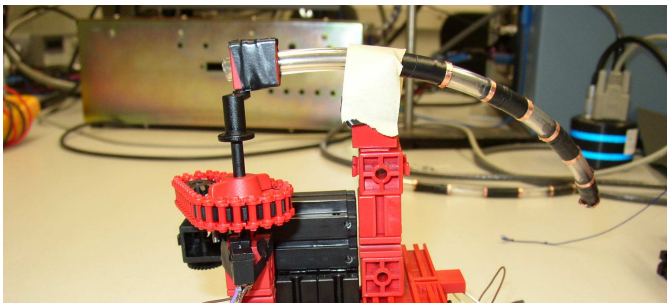

(b)

Figure 6.17: (a) Connecting part of the tube and the ECG generator. (b) The *vessel* is attached to the ECG generator wheel.

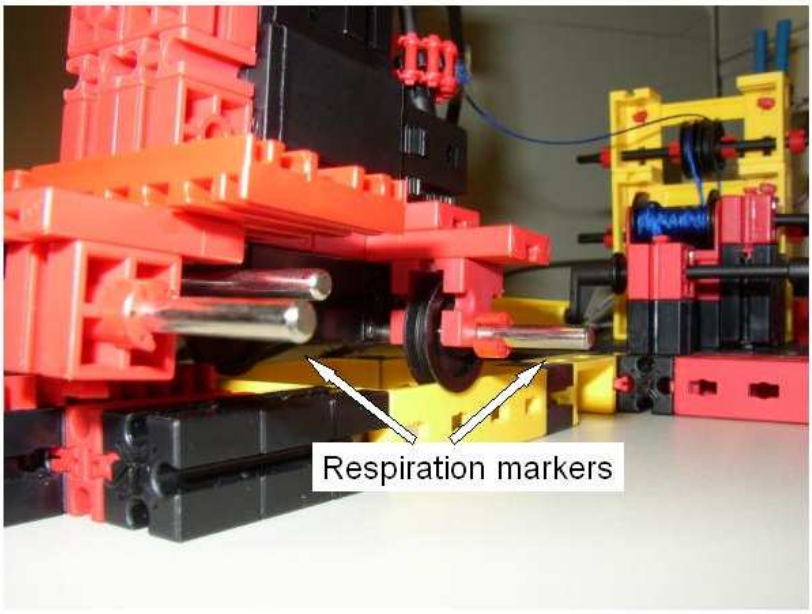

Figure 6.18: Metallic bars were used as wheels' axes to provide respiration markers in the acquired x-ray images.

ing the fluoroscopy by $\vec{m}_j^n$ calculated with respiration compensation (equation 6.1). We can see that the fiber optic marker (arrowed) now appears inside the vessel that was visible in the angiogram.

6.6.3 Experimental setup for diagnostic XIVUS

The proposed method for diagnostic pullback reconstruction and interventional positioning aid was tested in our heart phantom. The phantom was laid on a Cathlab table and illuminated with a bright light source (see figure 6.21). The phantom was imaged in angiography (see figure 6.22) and fluoroscopy mode during pullback (see figure 6.23). The pullback path was approximately $1.5cm$ long. In the angiograms (figure 6.22) it is possible to distinguish the vessel and its partially radio-opaque bands delimiting the opaque and transparent areas. The pullback was performed without changing the geometry settings or moving the phantom. We can see in figure 6.23 a small radio-opaque marker (arrowed) that indicates the position of the fiber optic aperture, that was segmented manually. The dark metal bars function as respiration markers and move according to the breathing rhythm.

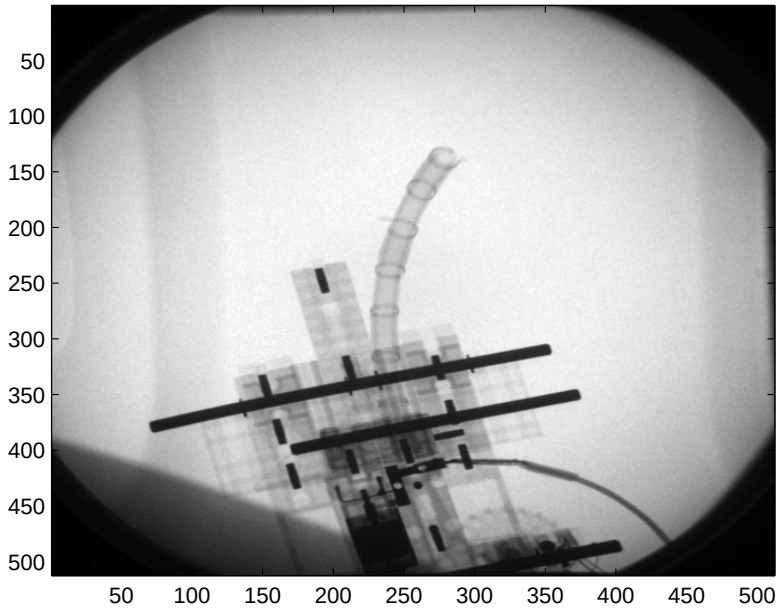

(a) Angiogram

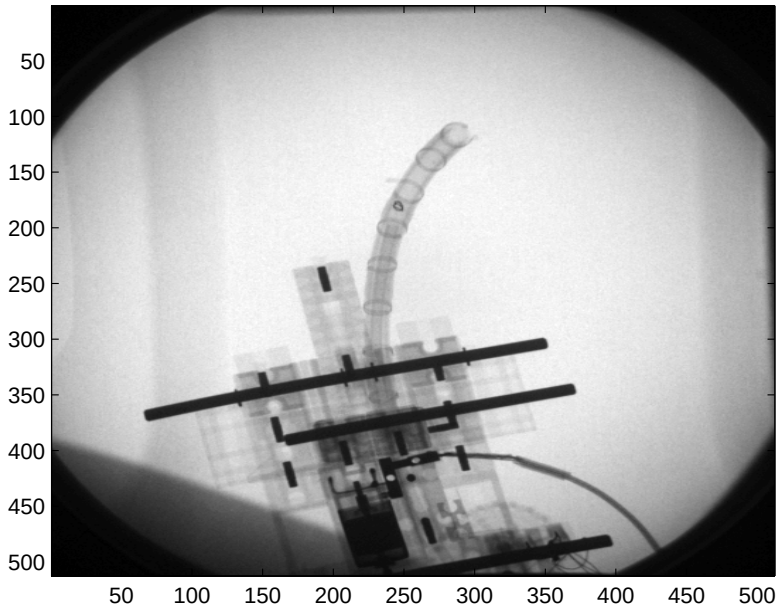

(b) Beginning of pullback

Figure 6.19: (a) Angiogram for a given heart phase δ_n. (b) Fluoroscopy frame acquired at the beginning of the pullback.

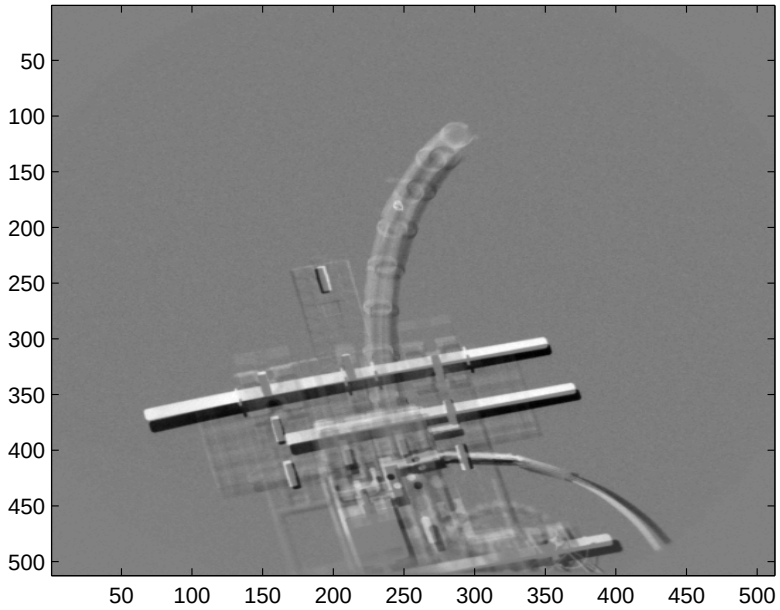

(a)

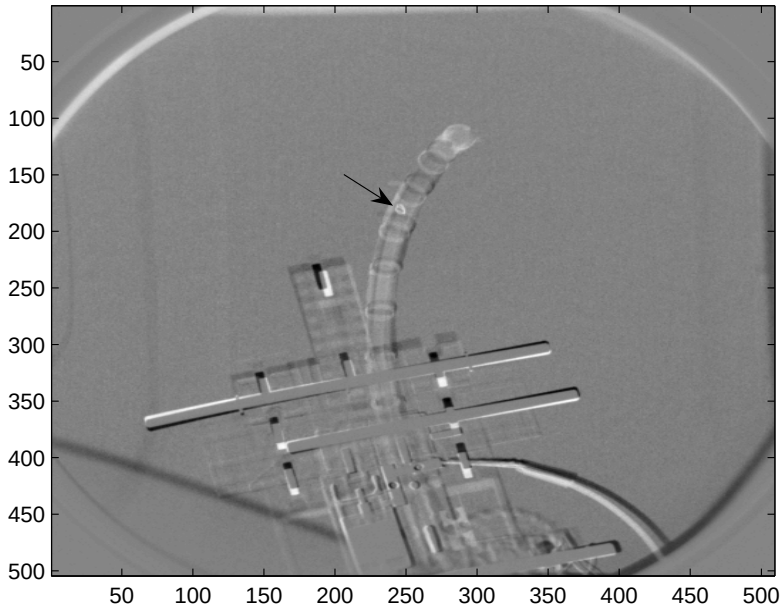

(b)

Figure 6.20: (a) Difference image between 6.19a and 6.19b before respiration compensation. (b) Difference image between 6.19a and 6.19b after respiration compensation (equation 6.1). The fiber optic marker is indicated by an arrow.

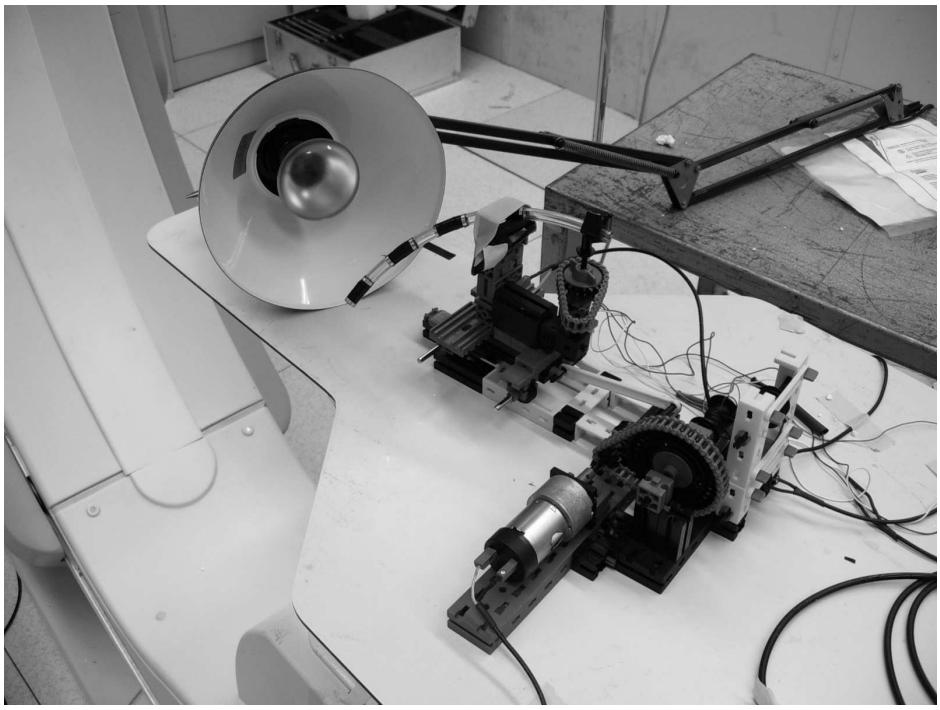

Figure 6.21: Experimental setup. The phantom lies on a Cathlab table and is next to a bright light source.

An ECG of mean frequency 1Hz was generated with a photoelectric barrier during all acquisitions (see figure 6.24a). The output of the fiber optic during pullback was detected with a phototransistor (see figure 6.24b) that provided an output voltage proportional to the amount of detected light with a dynamic range of $8V$. This signal is not a two-level digital voltage because the detected light gradually decreases or increases after an opacity change in the plastic tube. It must be taken into account that, since the fiber optic gathers light only in the forward direction and the vessel was illuminated obliquely, there will be an offset between the position of a change in vessel opacity and the position when this is reflected in the phototransistor output signal. Under our measurement conditions, this offset was determined a priori as $d_{\text{offset}} = 5\text{mm}$. We can observe that the respiration and heart contraction frequency contribute to the noise in the output signal of the phototransistor. The frequency components corresponding to respiration and heart contraction are filtered out for ulterior calculations. Given the length of the dark bands in the vessel and the time required to advance between them, it was established that the pullback speed in the phantom was 2.2mm/s.

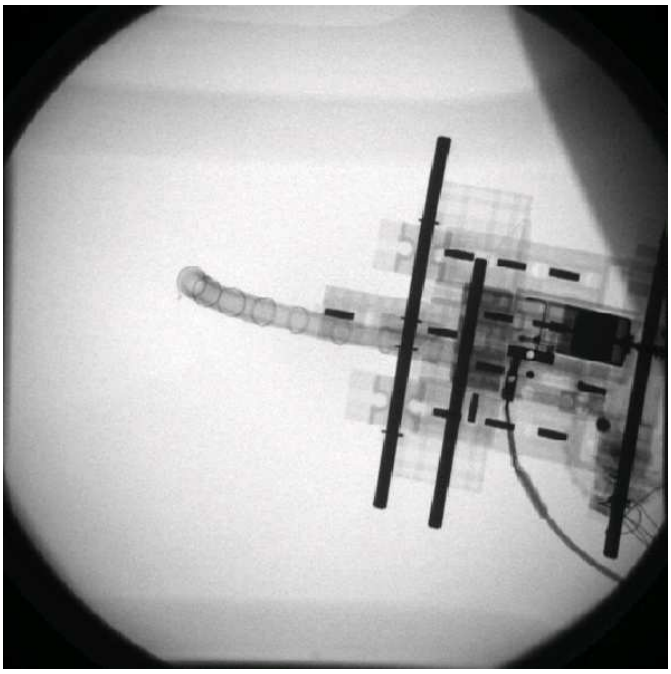

Figure 6.22: Angiographic sequences were acquired from the phantom.

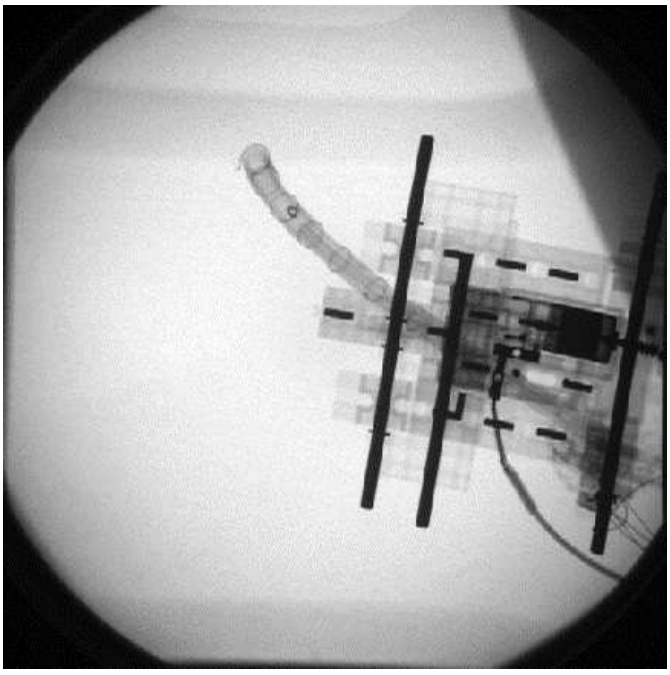

Figure 6.23: Fluoroscopic sequences were acquired from the phantom. The fiber optic marker is visible in these sequences and can be tracked during pullback.

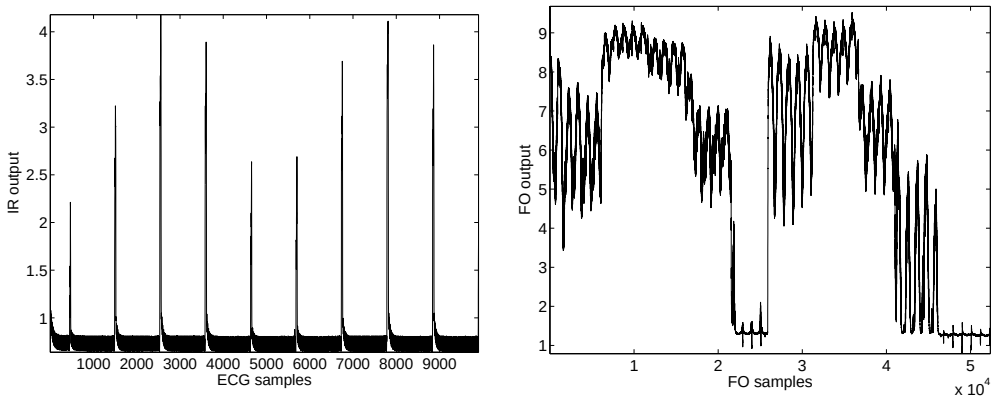

(a) (b)

Figure 6.24: (a) Artificial ECG signal generated with an IR movement detector. (b) Output of the photodetector during the fiber optic pullback.

6.6.4 Results

For the experiments, the heart phase was quantised in $N = 50$ steps and the fiber optic segmented manually to avoid the influence of detection errors in our results. For each heart phase, an angiography was selected and a pullback path $\Pi(\delta_n, \beta_n^a)$ was reconstructed from the respiration-compensated locations of the manually-segmented fiber optic marker according to equations 6.1 and 6.2. The respiration compensation algorithm presented by Eck *et al.* [73] was employed (see section 5.1). An example of respiration-compensated fiber optic positions is shown in figure 6.25 where the white dots represent the compensated positions. The pullback path was calculated by B-spline interpolating between the respiration-compensated fiber optic positions, and taking into account the known pullback speed.

To check the accuracy of the registration between the respiration-compensated positions and the pullback path, we manually checked the relative distance of the fiber optic marker to the distal band marker in the fluoroscopic pullback frames and in the calculated pullback path after respiration compensation. The absolute error for each compensated position in figure 6.25 is shown in figure 6.26. We can observe that the absolute error is below $2mm$ with a stenosis length of $12.8mm$. The accuracy of the respiration-compensated positions depends on the respiration compensation algorithm used and other respiration compensation algorithms should be further tested.

The output of the fiber optic for each position (see figure 6.27) was man-

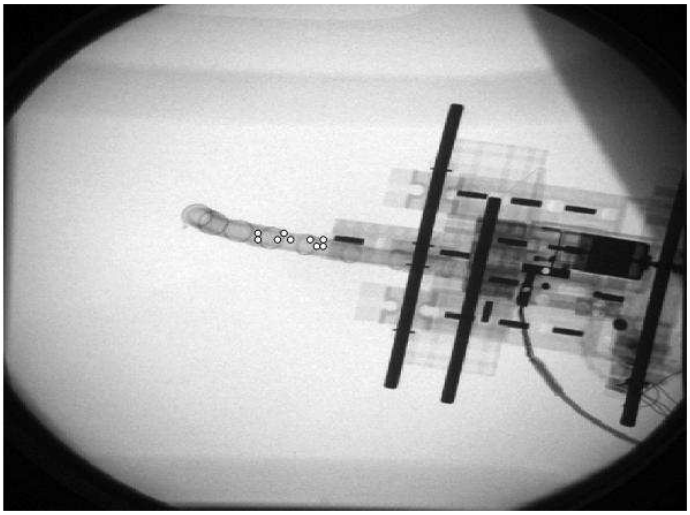

Figure 6.25: Respiration-compensated positions of the fiber optic. From these points the pullback path is reconstructed by interpolation taking into account the constant pullback speed.

ually compared to the respiration-compensated locations and we checked the measured value of $\hat{d}_{\text{offset}}$ with respect to the a priori measured d_{offset}. In figure 6.27 the shaded areas represent the frames for which the output of the phototransistor should be high because they correspond to an opaque area in the tube. Establishing a threshold of $7V$ to detect opaque (output higher than $7V$) and transparent (output lower than $7V$) areas we can observe a time offset of approximately 30 frames, which corresponds to $\hat{d}_{\text{offset}} = 5.28\text{mm}$ or an absolute error $\hat{d}_{\text{offset}} - d_{\text{offset}} = 0.28\text{mm}$.

6.7 3D vessel reconstruction

In [56][62][63][64][67][57][68] an accurate 3D pullback path is reconstructed by means of acquiring a biplane or stereo angiography of the coronary tree, and segmenting the contrasted vessels or IVUS catheter for a given heart phase and respiration phase. Then the IVUS frames are located equally spaced along this pullback path. In this section we first present the method proposed in the literature [56][62][63][64][67][57][68] to calculate the relative rotation between IVUS frames once the 3D pullback path is known. We later analyse the 3D capabilities of XIVUS as a diagnostic tool.

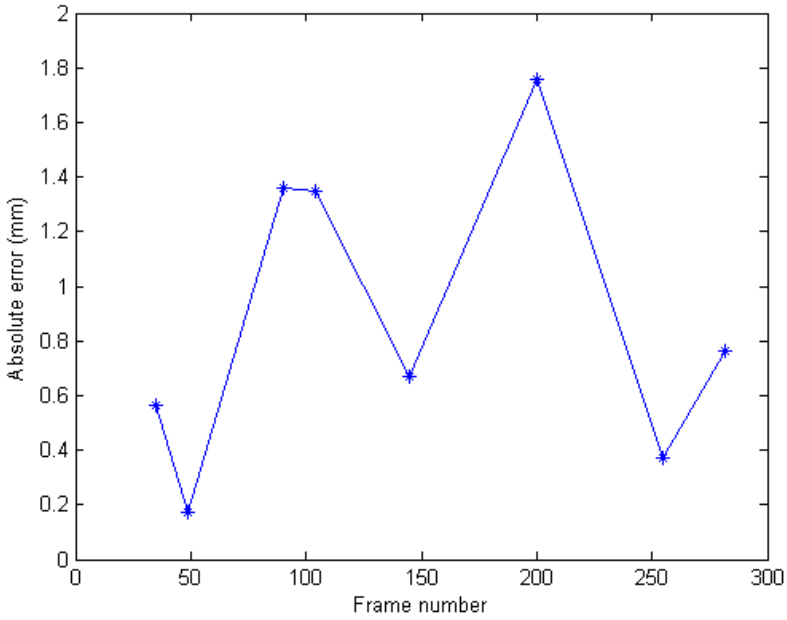

Figure 6.26: Absolute error for the respiration compensated positions of the fiber optic marker along the pullback path.

In subsection 6.7.1 we summarise the theoretical background on which previous literature bases for 3D pullback path reconstruction which relies on biplane acquisition. In subsection 6.7.2 we propose a three-dimensional ECG-gated pullback path reconstruction from the 2D projection sequence acquired during pullback, assuming constant fluoroscopic surveillance. The 3D path coordinates allow the arrangement of the IVUS cross-sections as in [54][56][66][142][139]. In subsection 6.7.3 we present the results with software simulated curves.

6.7.1 Relative orientation of IVUS frames

As a catheter is pushed into the treated coronary artery, it follows a 3D path along the vessel. To introduce the catheter into the vessel following a 3D curve, the longitudinal force applied at the distal end is translated into a torque (i.e. the actuation of two forces) whenever the path changes plane and the catheter has to bend. This results in a non-zero torsion and the catheter structure twists along its longitudinal axis [65]. As the catheter bends (or more precisely, *un*bends) as it is pulled back, it untwists, causing a relative

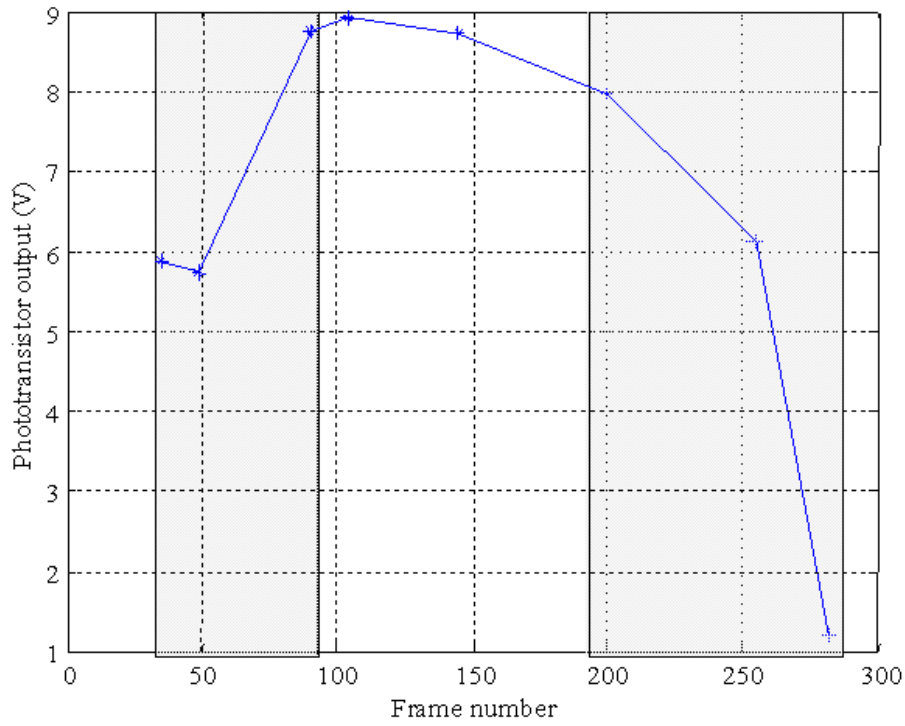

Figure 6.27: Output of the phototransistor. The shaded areas represent the frames for which the output should be high.

rotation between consecutive IVUS frames. This relative rotation can be calculated from the 3D path followed by the catheter, and has been resolved in [54][56][66][139][142] by establishing a reference frame, the Frenet-Serret frame, for each cross-section at each point along the path. The mathematical rationale of torsion and the definition of the Frenet-Serret frame will be explained in this subsection.

Let $\vec{\Pi}(t)$ be a 3D curve in space. The Frenet-Serret frame is a positively oriented orthonormal reference frame defined for each point in a 3D curve $\vec{\Pi}(t)$ and consists of a unit tangent vector $\vec{t}$, a normal vector $\vec{n}$, and a binormal vector $\vec{b}$ (see figure 6.28), defined as [143]

$$\begin{cases} \vec{t} = \frac{\vec{v}}{\|v\|} \\ \vec{n} = \frac{\vec{v}\times\vec{a}\times\vec{v}}{\|\vec{v}\times\vec{a}\times\vec{v}\|} \\ \vec{b} = \vec{t}\times\vec{n} \end{cases} \tag{6.7}$$

where $\vec{v} = \frac{d\vec{\Pi}}{dt}$ is the velocity vector, and $\vec{a}$ is the acceleration vector, defined as $\vec{a} = \frac{d\vec{v}}{dt}$. At each point, the image acquired by the IVUS probe would be on the plane defined by $\vec{n}$ and $\vec{b}$, and perpendicular to $\vec{t}$ [54][56][66][139][142].

The triplet $(\vec{t}, \vec{n}, \vec{b})$ is related to the curvature κ and torsion τ of the curve by the *Frenet-Serret* formulas,

$$\begin{cases} \frac{d\vec{t}}{dt} = \kappa\vec{n}\|\vec{v}\| \\ \frac{d\vec{n}}{dt} = -\|\vec{v}\|\kappa\vec{t} + \|\vec{v}\|\tau\vec{b} \\ \frac{d\vec{b}}{dt} = -\|\vec{v}\|\tau\vec{n} \end{cases} \tag{6.8}$$

which can also be written as

$$\frac{d}{dt}\begin{pmatrix} \vec{t} \\ \vec{n} \\ \vec{b} \end{pmatrix} = \begin{pmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{pmatrix}\begin{pmatrix} \vec{t} \\ \vec{n} \\ \vec{b} \end{pmatrix} \tag{6.9}$$

The torsion of a curve measures the rate of change of the plane defined by $\vec{t}$ and $\vec{n}$. Intuitively, we can interpret torsion as the rate at which the Frenet-Serret frame twists around the curve. Then, torsion describes the rate of change of the catheter around its longitudinal axis.

If calculated directly with equation 6.7, the Frenet-Serret frame, and more specifically $\vec{n}$, is undefined when $\vec{a} = 0$, such as inflection points $\vec{v} = 0$ or

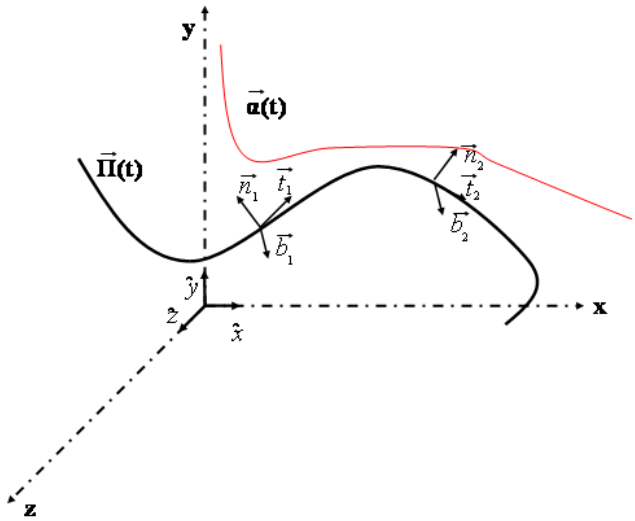

Figure 6.28: At each point of the three-dimensional curve $\Pi(\vec{t})$, the Frenet-Serret frame is composed of a tangent ($\vec{t}$), a normal ($\vec{n}$) and a binormal ($\vec{b}$) vector.

in straight segments of the curve, $\vec{v}$ = constant [143]. Furthermore, after an inflection point, $\vec{b}$ may change sign, causing a violent change of orientation in a progression of Frenet-Serret frames. To circumvent these problems, a minimum torsion calculation of the Frenet-Serret frame [54][56][143][142][139] is used. This method consists in propagating an initial reference frame along the curve, by means of rotations around an axis, so that the previous tangent vector is aligned with the next tangent vector [143]. For a discrete curve, the axis of rotation is calculated as

$$\vec{A} = \frac{\vec{t_{k-1}} \times \vec{t_k}}{\|\vec{t_{k-1}}\|\|\vec{t_k}\|} \tag{6.10}$$

and the rotation angle to align the previous tangent vector with the next one is

$$\theta = \cos^{-1} \frac{\vec{t_{k-1}} \cdot \vec{t_k}}{\|\vec{t_{k-1}}\|\|\vec{t_k}\|} \tag{6.11}$$

If $\vec{t}_{k-1} = \vec{t}_k$, then $\vec{A}$ is 0 and the frame is not rotated, just translated along the direction of $\vec{t}$.

Given a 3D pullback path segmented for a specific heart and respiration phase, the Frenet-Serret frame (see figure 6.28) allows us to arrange consecutive IVUS cross-sections with the proper relative orientation along the pullback path [54][56][66][139][142] according to the point at which they were acquired and the torsion of the 3D pullback path.

The methods presented in the literature base on vessel segmentation from biplane or stereo angiographic acquisitions. Their objective is to reconstruct the 3D vessel from fusion of angiography and IVUS for diagnosis support. One pair of biplane (or stereo) angiograms will present the heart with a specific state vector $\vec{\xi}$ from different angles. For IVUS fusion, ECG gating is used to ensure that the fused IVUS frames are acquired with the same heart phase as the biplane angiograms. These methods do not mention whether the objective of the application is to reconstruct the vessel for different heart phases.

We have proposed XIVUS as a diagnostic aid in section 6.4. XIVUS can be applied to monoplane angiograms for vessel reconstruction for different heart phases. We propose to do so for N different heart phases so that the fused information can be later utilised for accurate instrument positioning. For this later application heart and respiration phase characterisation of the acquired images is necessary as explained in chapter 4. Respiration compensation and accurate IVUS probe marker to vessel matching as explained chapter 5 are also necessary.

In order to reconstruct the 3D IVUS probe path in a similar way to [54][56][66][142][139] but based on monoplane acquisitions during IVUS probe pullback, we propose in the next subsection an algorithm to reconstruct the 3D pullback path from projection x-ray images acquired .

6.7.2 3D pullback path reconstruction

In this subsection we propose a method to reconstruct a three-dimensional path assuming that it is orthogonally projected onto the imaging plane and the pullback is performed at constant speed. We derive the 3D velocity vector from the 2D velocity of the IVUS probe as can be measured in the 2D images from the acquired sequence. The 3D velocity vector can be estimated from the 2D velocity taking into account that the speed, which is the magnitude of the velocity $\|\vec{v}\|$, is constant. Once the 3D velocity vector is known, the pullback path in 3D can be calculated.

An approach similar to the one presented here has been independently researched and published in [137]. Jourdain *et al.* reconstruct a 3D pullback path from its 2D projection considering perspective geometry. In [137], perspective geometry is considered, whereas we argue that the geometry involved in the imaging process allows for the simplification of considering orthogonal projection of the points involved in the reconstruction onto the image plane. Jourdain *et al.* present results on simulated data assuming no heart contraction nor respiration movement is present and on a static phantom.

6.7.2.1 Preliminary assumptions

We hypothesise that the dimensions of the imaged object (the coronary artery) and the displacement of the IVUS probe in the direction perpendicular to the projection plane are negligible compared to the distance between the x-ray tube and the patient, and that the geometry of the system is chosen such that the artery is almost parallel to the projection plane, to avoid foreshortening. Taking these considerations into account, we assume an orthogonal projection of the IVUS probe and the coronary arteries onto the projection plane.

Let's consider the following imaging geometry (see figure 6.29). The 3D position of the IVUS probe is represented as (x, y, z) and its projection as (X, Y). The origin of the coordinate system is placed at the x-ray source and the projection plane is perpendicular to the z axis.

In perspective projection, we have the following relation between the two-dimensional coordinates X and Y and the 3D coordinates of a point $P = (x, y, D)$

$$X = \frac{fx}{D} \qquad Y = \frac{fy}{D} \tag{6.12}$$

Let $P' = (x, y, z')$ be another point in a plane $z' = D + \Delta Z$. Its coordinates in the projection plane according to projective geometry are (X', Y')

$$\begin{aligned} X' &= \frac{fx}{D+\Delta z} \\ Y' &= \frac{fy}{D+\Delta z} \end{aligned} \tag{6.13}$$

We wish to make the simplification of $D \gg \Delta z$ so that $X \approx X'$ and $Y \approx Y'$. Typically, the distance between the x-ray source and the patient is maximised to diminish the skin dose of the patient [3], while the distance between the

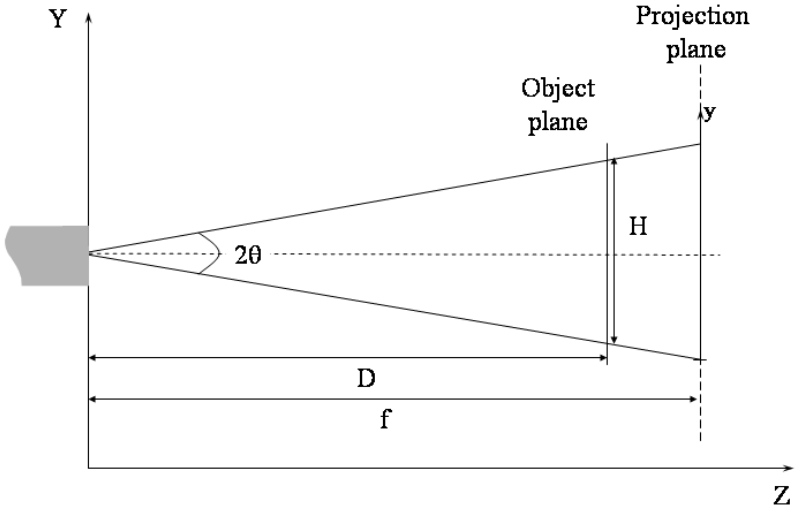

Figure 6.29: Projection geometry.

patient and the detector is minimised [3]. For optimal imaging, the heart should be positioned at the isocentre of the imaging system and the geometry is chosen such that the imaged artery is approximately parallel to the imaging plane to prevent foreshortening.

Under these assumptions, if we consider a typical source-to-patient distance of 110cm [3] and a coronary artery diameter of 5mm, the assumption $D \gg \Delta z$ holds.

If the previous assumption, $D \gg \Delta z$, does not hold, we can still consider an orthogonal projection if the angle θ in figure 6.29 is close enough to 0, $\theta \approx 0$. Considering again $D = 110cm$ and a coronary artery of length $H = 15cm$, then

$$\tan(\theta) = \frac{H/2}{D} \Rightarrow \theta \approx 6 \cdot 10^{-2} rad \tag{6.14}$$

θ as calculated above can be considered close enough to 0 to assume that the orthogonal projection assumption holds.

Based on the previous assumptions, we present a method that reconstructs the IVUS transducer 3D pullback path assuming that the projection of the path on the imaging plane is orthogonal.

6.7.2.2 3D reconstruction method

We start our analysis for a continuous curve and present our results on computer simulated curves later on.

Let the (yet unknown) parameterised 3D IVUS path be $\vec{\Pi}(t)$, where t is time. We can write with respect to a global reference system (figure 6.30)

$$\vec{\Pi}(t) = (\Pi_x(t), \Pi_y(t), \Pi_z(t)) \tag{6.15}$$

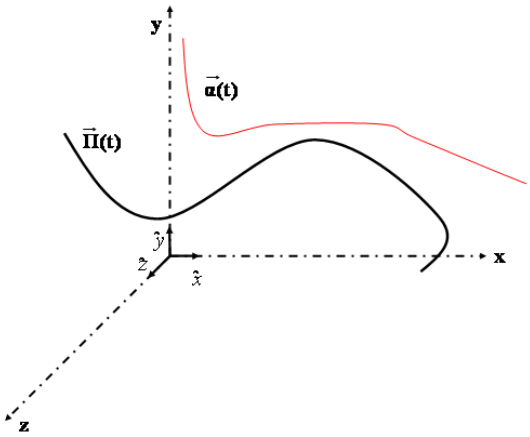

Figure 6.30: Reference system with unit vectors $\hat{x}$, $\hat{y}$ and $\hat{z}$. $\vec{\Pi}(t)$ is a three-dimensional curve and $\vec{\alpha}(t)$ its two-dimensional projection.

The IVUS probe velocity is given by

$$\vec{v}(t) = (v_x(t), v_y(t), v_z(t)) = (d\Pi_x(t)/dt, d\Pi_y(t)/dt, d\Pi_z(t)/dt) \tag{6.16}$$

The velocity is tangent to the IVUS path at each point and its norm $\|\vec{v}(t)\|$ is the constant pullback speed, v_{pull}.

In the fluoroscopic pullback sequence, an orthogonal 2D projection, $\vec{\alpha}(t)$, of the 3D path (equation 6.15) is observed (see figure 6.30). Without loss

of generality, we assume that the path is projected onto the XY plane and therefore $\vec{\alpha}(t)$ can be expressed as

$$\vec{\alpha}(t) = (\Pi_x(t), \Pi_y(t)) \tag{6.17}$$

The observed velocity in the 2D projection image will be

$$\vec{v}_{2D}(t) = (d\Pi_x(t)/dt, d\Pi_y(t)/dt) = (v_x(t), v_y(t)) \tag{6.18}$$

Knowing that the norm of the velocity is constant and equals the speed of the pullback, v_z can be calculated as

$$v_z^2(t) = \|\vec{v}(t)\|^2 - \|\vec{v}_{2D}(t)\|^2 = v_{pull}^2 - (v_x^2(t) + v_y^2(t)) \tag{6.19}$$

There is a sign indetermination in the calculation of $v_z(t)$ (equation 6.19) that will be discussed in subsection 6.7.3. Assuming that the sign indetermination is solved, the 3D path can be reconstructed using the definition of velocity

$$\vec{v}(t) = \frac{d\vec{\Pi}(t)}{dt} \Rightarrow \begin{cases} \hat{\Pi}_x(t) = x_0 + \int_{t_0}^{t} v_x dt \\ \hat{\Pi}_y(t) = y_0 + \int_{t_0}^{t} v_y dt \\ \hat{\Pi}_z(t) = z_0 + \int_{t_0}^{t} v_z dt \end{cases} \tag{6.20}$$

where

- $\hat{\Pi}_x(t)$, $\hat{\Pi}_y(t)$ and $\hat{\Pi}_z(t)$ are the estimated Π_x, Π_y and Π_z respectively,
- t_0 is the origin in time or beginning of the pullback,
- x_0 and y_0 are the coordinates of the beginning of the pullback, as measured in the corresponding pullback fluoroscopy frame, and
- z_0 is an arbitrary constant.

6.7.3 Simulations of 3D pullback reconstruction

As mentioned in previous subsections, there is a sign indetermination in the calculation of v_z (equation 6.19). At the beginning of the pullback path or whenever the curvature vanishes, the sign is undetermined. This effect is illustrated in figure 6.31. In figure 6.31a we can see two artificially reconstructed pullback paths for which the starting sign of v_z has been chosen

opposite. We can see that the dotted curve in the figure is a reflection of the other one. In figure 6.31b we illustrate the problem of the indetermination of the sign of v_z after a straight segment. As we can see, after the straight segment, the reconstructed pullback paths are mirrored versions of each other. In either case, the indetermination affects the direction of v_z and $\hat{\Pi}_z(t)$, but not the resulting volume nor the relative rotation between the Frenet-Serret frames at each point.

Regarding this indetermination problem, we should consider that usually the geometry of a monoplane Cathlab is chosen such that the stenosis is almost parallel to the projection plane. The stiffness of the catheter supports the assumption that the pullback path is a smooth and *well-behaved* curve, and that v_z does not change abruptly (it is continuous). Also knowledge about how plaque accumulates inside a vessel could be used to correct the reconstructed pullback from the analysis of the IVUS images. In [144] the dependence of plaque accumulation on vessel curvature was studied. It was found that plaque tends to accumulate at the outer wall of the curved vessel. This knowledge could be used to correct $\hat{\Pi}_z$ once the plaque has been segmented in the IVUS images (see section 6.8).

The algorithm presented above assumes continuous time t, whereas x-ray acquisition throughout the intervention is known to be done at a constant frame rate. Furthermore, because of constant heart beat and breathing, not all acquired frames can be used for 3D pullback path reconstruction. It is necessary that the frames used have the same heart and (compensated) respiration phase. A sparse sampling because of the acquisition frame rate used and heart and respiration movement introduce reconstruction errors. The frame rate for fluoroscopy is typically 8.3fr/s and that of heart-beat is typically 1Hz. We performed some simulations with a software generated vessel with a half circumference shape of radius $1mm$ contained in a plane parallel to the XZ plane in order to estimate the reconstruction error introduced by our 3D pullback path reconstruction method. We assumed a heart frequency of 1Hz and a pullback speed of 0.1 mm/s. The mean squared error for the reconstructed curve was $\varepsilon_z = 2.6.10^{-3}$ for the z coordinate.

The proposed algorithm for 3D XIVUS has not been tested due to the lack of appropriate clinical data. It would be necessary to first acquire angiographic, pullback fluoroscopic and interventional fluoroscopic according to the workflow described in section 6.3. For comparison, biplane angiography for accurate 3D vessel reconstruction would also be needed.

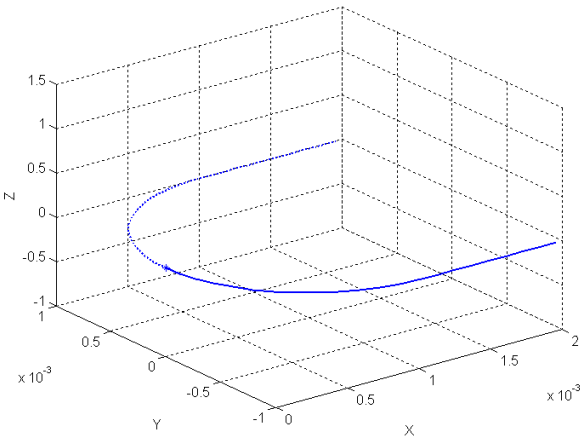

(a)

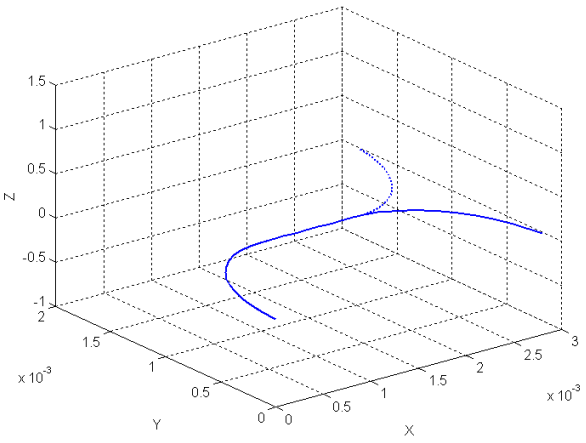

(b)

Figure 6.31: Reconstruction of a three-dimensional curve. (a) The indetermination in the starting sign of v_z allows the two possible reconstructions illustrated in the figure (solid and dotted lines), one is a mirrored version of the solid curve. (b) The indetermination in the sign of v_z causes that after a straight segment parallel to the projection plane, two possible reconstructions are possible.

6.8 Lumen segmentation

Lumen segmentation algorithms aid diagnosis by enabling the quantification of the lesion characteristics (see appendix B). Segmentation algorithms for IVUS images have been proposed in [140][145][146][147][148][149][150]. In these, deformable models driven by image features have been proposed in 2D [145][146][148][149] and 3D [150]. The advantage of using 3D deformable models for lumen segmentation is that segmentation is not driven only by image features in each IVUS plane, but profit from the vessel 3D structure and its coherence along the vessel axis.

Here we propose a 3D deformable cylinder based segmentation. Each IVUS frame is preprocessed in order to obtain salient features to guide the active cylinder. Then an active surface as defined in [151] and [152] is used. In subsection 6.8.1 we define the necessary concepts of a deformable surface and the deformation force. A much deeper insight can be found in [151]. In subsection 6.8.2 we explain the preprocessing steps on the IVUS images before applying a deformable cylinder for segmentation.

6.8.1 Deformable surfaces

According to [151], *a deformable surface is an object composed of a surface representation and an evolution law allowing the deformation of this surface to represent a different shape.* A possible representation of a surface is a *simplex mesh* $\mathcal{M}$, which is defined [152] by a set of vertices $\{\mathbf{p}_i\}$ and a specific connectivity function. Each vertex of a k-simplex mesh is connected to exactly $k+1$ neighbours. In figure 6.32 we can see a cell of a 2-simplex mesh, where

- $\mathbf{p}_i^{\perp}$ is the projection of the vertex $\mathbf{p}_i$ on the plane $\mathcal{P}_i$ defined by $\mathbf{p}_i$'s neighbours, $\{\mathbf{p}_i^1, \mathbf{p}_i^2, \mathbf{p}_i^3\}$,
- $\mathbf{c}_i$ is the centre of the circumscribed circle to the triangle $\{\mathbf{p}_i^1, \mathbf{p}_i^2, \mathbf{p}_i^3\}$, and
- r_i is the radius of the circumscribed circle to the triangle $\{\mathbf{p}_i^1, \mathbf{p}_i^2, \mathbf{p}_i^3\}$.

The *metric parameters* are defined as the centre of mass coordinates of $\mathbf{p}_i^{\perp}$ with respect to $\mathbf{p}_i$'s neighbours:

$$\begin{aligned} p_i^{\perp} &= \varepsilon_i^1 \mathbf{p}_i^1 + \varepsilon_i^2 \mathbf{p}_i^2 + \varepsilon_i^3 \mathbf{p}_i^3 \\ \varepsilon_i^1 + \varepsilon_i^2 + \varepsilon_i^3 &= 1 \end{aligned} \tag{6.21}$$

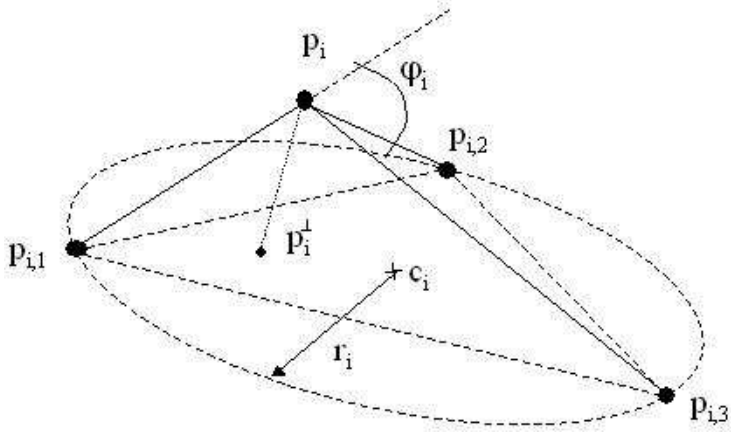

Figure 6.32: A point in a 2-simplex mesh and its three neighbours.

The metric parameters control the relative position of the vertex $\mathbf{p}_i^{\perp}$ on the plane $\mathcal{P}_i$.

The *simplex angle* $\varphi_i \in [-\pi, \pi]$ is defined by

$$\begin{cases} \sin(\varphi_i) = \frac{r_i}{R_i}\text{sign}((\mathbf{p}_i^1 - \mathbf{p}_i) \cdot \mathbf{n}_i) \\ \cos(\varphi_i) = \frac{\|\mathbf{c}_i - \mathbf{o}_i\|}{R_i}\text{sign}((\mathbf{c_i} - \mathbf{o}_i) \cdot \mathbf{n}_i) \end{cases} \tag{6.22}$$

where R_i and $\mathbf{o}_i$ are the radius and centre of the circumscribed sphere to vertices $\{\mathbf{p}_i, \mathbf{p}_i^1, \mathbf{p}_i^2, \mathbf{p}_i^3\}$ respectively . The relations in equation 6.22 are made explicit by drawing the projection of the circumscribed sphere on the plane defined by $(\mathbf{c}_i, \mathbf{o}_i, \mathbf{p}_i)$ [152] (see figure 6.33a). Because φ_i is independent of the position of $\mathbf{p}_i$ on the arch above $\mathcal{P}_i$, we can simplify its calculation by considering the case illustrated in figure 6.33b. In this figure

$$\pi = x + 2\theta \tag{6.23}$$

and $\theta = (\pi - \varphi_i)/2$. Therefore $x = \varphi_i$ and $\sin(\varphi_i) = r_i/R_i$ and $\cos(\varphi_i) = |\mathbf{c}_i - \mathbf{o}_i\|/R_i$. φ_i controls the elevation of $\mathbf{p}_i$ over $\mathcal{P}_i$. The sign(x) function in equation 6.22 determines whether the point $\mathbf{p}_i$ is located above or below the plane defined by $\{\mathbf{p}_i^1, \mathbf{p}_i^2, \mathbf{p}_i^3\}$.

On a continuous surface, the mean curvature at a point may be derived by approximating locally the surface with a sphere. On a simplex mesh, the sphere of radius R_i is the circumscribed sphere which best fits the vertices around $\mathbf{p}_i$ and therefore is analog to the concept of mean curvature of a continuous surface. The vertex *discrete mean curvature* is defined as the inverse

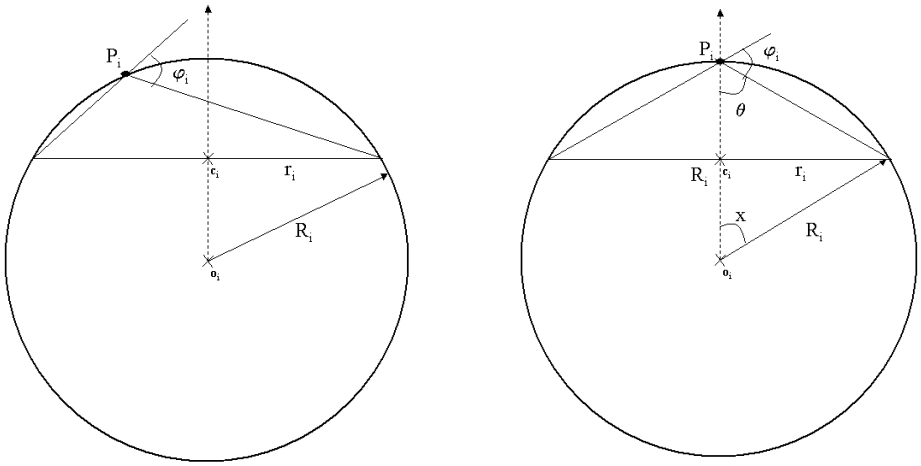

(a) (b)

Figure 6.33: (a) Result of projecting the sphere on the plane defined by $(\mathbf{c}_i, \mathbf{o}_i, \mathbf{p}_i,)$. (b) Simplified and equivalent case of (a).

of the radius of the circumscribed sphere, $H_i = \frac{1}{R_i} = \frac{\sin(\varphi_i)}{r_i}$.

From the defined parameters $(\varepsilon_i^j, \varphi_i)$, the position of the vertex can be written as

$$\mathbf{p}_i = \Big(\sum_{j=1}^{3} \varepsilon_i^j \mathbf{p}_i^j \Big) + L(\mathbf{p}_i^j, \varepsilon_i^j, \varphi_i)\mathbf{n}_i \tag{6.24}$$

where

$$\begin{aligned} L &= \frac{(r_i^2 - d_i^2)tan(\varphi_i)}{\epsilon\sqrt{r_i^2 + (r_i^2 - d_i^2)tan(\varphi_i)tan(\varphi_i)^2 + r_i}} \\ \epsilon &= \begin{cases} 1 & if\ |\varphi_i| < \pi/2 \\ -1 & if\ |\varphi_i| > \pi/2 \end{cases}, \ d_i = \|\mathbf{p}_i^{\perp} \cdot \mathbf{c}_i\| \end{aligned} \tag{6.25}$$

The first term after the equality in equation 6.24 specifies the position of $\mathbf{p}_i^{\perp}$ on $\mathcal{P}_i$, and the second term $(L(\mathbf{p}_i^j, \varepsilon_i^j, \varphi_i))$ specifies the distance of the vertex $\mathbf{p}_i$ to the plane $\mathcal{P}_i$.

The relation between L, r_i, d_i and φ_i in equation 6.25 can be seen in figure 6.34. Defining θ_u and θ_v as in the figure, then:

$$\tan(\theta_u) = \frac{(r_i - d_i)}{L} \quad \tan(\theta_v) = \frac{(r_i + d_i)}{L} \tag{6.26}$$

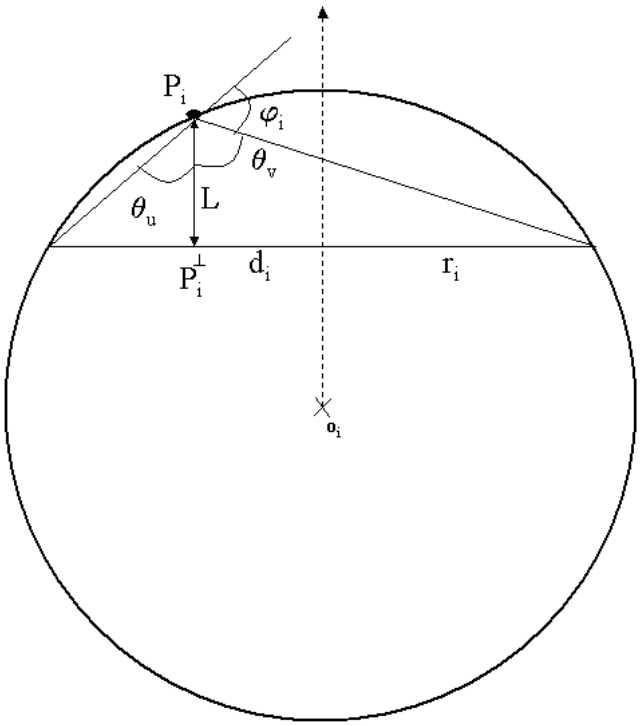

Figure 6.34: Figure that illustrates the relation between L, r_i, d_i, and ϕ_i in equation 6.25.

Since $\varphi_i = \pi - \theta_u - \theta_v$, then

$$-\tan(\varphi_i) = \tan(\theta_u + \theta_v) = \frac{\tan(\theta_u) + \tan(\theta_v)}{1 - \tan(\theta_u)\tan(\theta_v)} \tag{6.27}$$

For image segmentation, each vertex in a mesh is submitted to an internal force $\vec{f}_{\text{int}}$ and an external force $\vec{f}_{\text{ext}}$ that control and regulate its movement. The internal forces regularise the surface and prevent the vertices from dispersing and deforming into sharp bends, while the external forces are extracted from the image and carry the surface towards the image features to be segmented. The law of motion of each vertex in the mesh is expressed by the following equation:

$$\mathbf{p}_i^{t+1} = \mathbf{p}_i^t + (1-\gamma)(\mathbf{p}_i^t - \mathbf{p}_i^{t-1}) + \alpha_i \vec{f}_{\text{int}}(\mathbf{p}_i^t) + \beta_i \vec{f}_{\text{ext}}(\mathbf{p}_i^t) \tag{6.28}$$

where α_i and β_i are force weights that control the influence of the external and internal forces on the surface changes, γ is a damping factor, and t represent discrete time.

The regularisation of a simplex mesh is based on the position of each vertex with respect to its neighbours. In each basic cell, the vertex $\mathbf{p}_i$ is attracted to another position $\tilde{\mathbf{p}}_i$ in a smoother mesh. The internal force can be then decomposed in its normal and tangential component as

$$\vec{f}_{\text{int}}(\mathbf{p}_i) = \vec{f}_{\text{tg}}(\mathbf{p}_i) + \vec{f}_{\text{nr}}(\mathbf{p}_i) = (\tilde{\mathbf{p}}_i^{\perp} - \mathbf{p}_i^{\perp}) + (H(\mathbf{p}_i^j, \varepsilon_i^j, \varphi_i) - H(\mathbf{p}_i^j, \tilde{\varepsilon}_i^j, \tilde{\phi}_i))\mathbf{n}_i \tag{6.29}$$

where $\tilde{\varepsilon}_i$, $\tilde{\mathbf{p}}_i^{\perp}$ and $\tilde{\phi}_i$ are the parameters of $\tilde{\mathbf{p}}_i$. The tangential component of the internal force $f_{tg}(\mathbf{p}_i)$ controls the vertex spacing over the surface, whereas the normal component $f_n(\mathbf{p}_i)$ ensures a smooth curvature of the surface.

For lumen segmentation, the initial surface was a cylinder with N equally spaced vertices, with points located around the centre at a radius equal to that of the IVUS catheter. The connectivity between points was defined as shown in figure 6.35. We placed each vertex and its neighbours in consecutive IVUS images as shown in figure 6.35. If vertex $\mathbf{p}_i$ is in image $I(\delta, k)$, then $\mathbf{p}_{i,1}$ is located in image $I(\delta, k-1)$ and $\mathbf{p}_{i,2}$ and $\mathbf{p}_{i,3}$ are in image $I(\delta, k+1)$. We set $\{\tilde{\phi}_i\} = \phi_0$, where be ϕ_0 is the simplex angle of a cylinder (it is constant at each point of the cylinder). We also enforce uniform vertex spacing by setting $\tilde{\varepsilon}_i^j = 1/3 \quad \forall j$. The external force used to drive the deformable cylinder will be discussed in the next subsection.

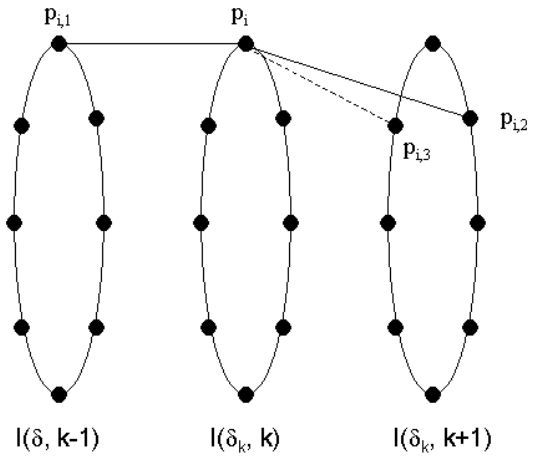

Figure 6.35: 2-simplex mesh for lumen segmentation.

6.8.2 Preprocessing of IVUS images

The result of the preprocessing will be the *external force* to drive the expansion of the deformable cylinder. We wish to have high uniform values at the lumen and low force values at the vessel wall, to stop the cylinder from growing.

The IVUS images from the database (see figure 6.36a) were transformed to polar coordinates (see figure 6.36b) which is their native format as acquired from the system. The ringdown effect can be seen in figure 6.36a as a bright ring around the centre of the image. In order to get an initialisation of the active cylinder, the ringdown effect is detected in five images and the maximum detected ringdown radius R_{ring} is used as the initial radius of the active cylinder. The ringdown effect is recognised as a bright vertical band to the left of the image and can be detected using a Hough transform after applying an edge detector operator to the image (see figure 6.37). The catheter radius R_{cath} is known, so that two parallel edges are searched for $R_{\text{cath}} < r < R_{max}$. All ulterior processing will be performed for $r > R_{\text{ring}}$.

Once the ringdown effect is detected, a 3D median filter with size 7×7 is applied to volumes centred at each pixel to reduce speckle noise and blur textures (see figure 6.38). As we can see in figure 6.38 different regions can be recognised in the image. The dark regions correspond to the lumen and to regions far away from the IVUS probe. A bright band in the middle of the figure corresponds the vessel wall. The rest of the image content either represents plaque or intima tissue.

Employing *a priori* knowledge of the image content, each IVUS frame is processed in order to classify its content in three classes corresponding to the dark lumen (and other dark tissue), the bright vessel wall, and the stenosis and other tissue. We used a k-means algorithm, that partitions the search space in k clusters, minimising over all clusters the within-cluster sums of point-to-cluster-centroid distances. The results can be seen in figure 6.39, where the plaque is separated from other tissue with similar gray values by the bright vessel wall. By identifying the cluster with the lowest centroid value, we can recognise the lumen and areas far away from the catheter. In the lumen, we are interested in uniform force values to drive the active cylinder. Therefore, we perform morphological closing and interpolation in those areas (see figure 6.40a and b).

In order to further enhance the difference between the lumen and the

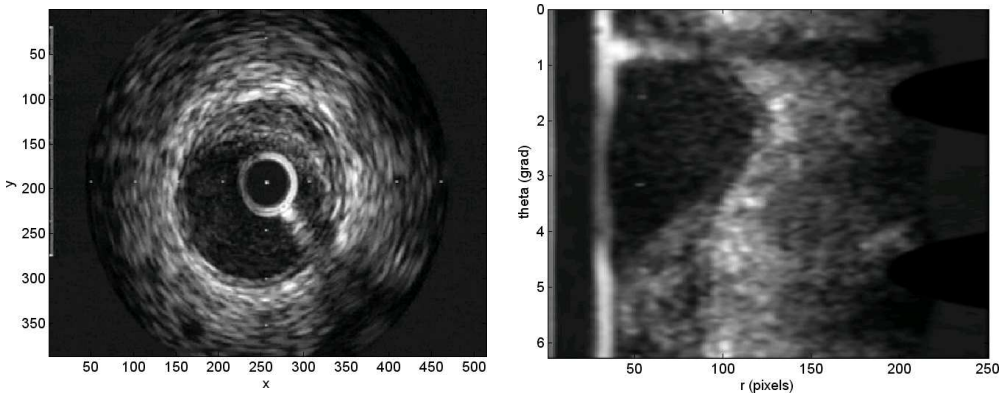

(a) (b)

Figure 6.36: IVUS frame in cartesian (a) and polar (b) form. In (b) the horizontal axis represents the radius, and the vertical axis the angle.

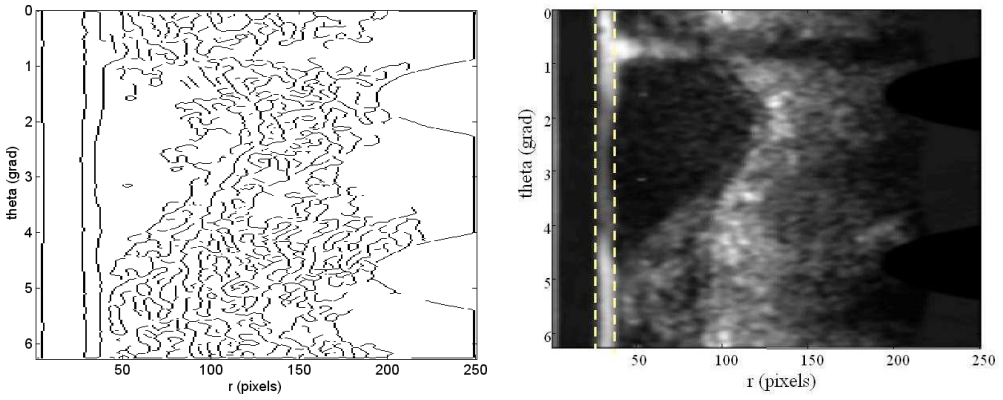

(a) (b)

Figure 6.37: (a) Result of Canny edge detection in a polar representation of an IVUS frame. (b) Result of finding the two most outstanding lines in the Hough transform.

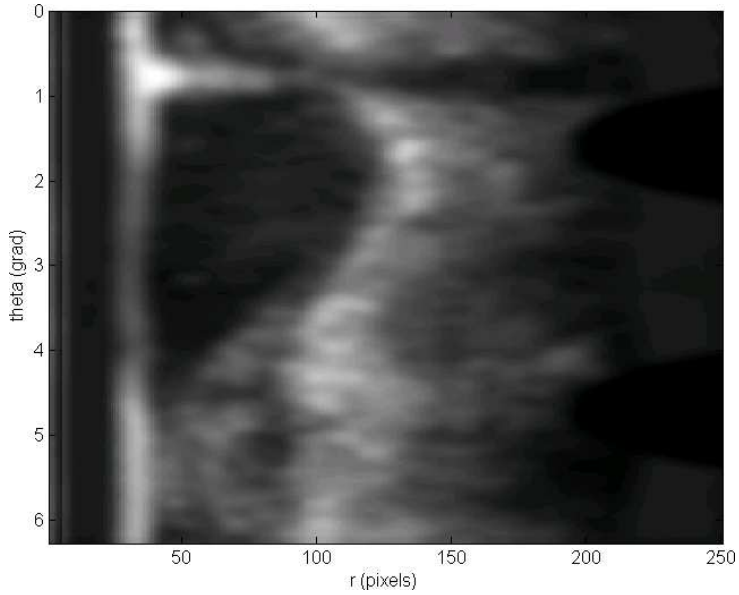

Figure 6.38: IVUS frame in figure 6.36 after being 3D median filtered.

vessel wall and other tissue, we multiply all non-lumen tissue by $k < 1$. The result can be seen in figure 6.41. This preprocessing is applied to all the acquired IVUS images for a given heart phase δ_n and used as an external force to drive the deformable cylinder. The very low values outside the lumen area will ensure that the cylinder slows down. When the active cylinder advances less than a predefined threshold, the segmentation is concluded.

The resulting external force image is converted to its cartesian representation and used to drive the deformable cylinder.

6.8.3 Segmentation results

There were nine patients from which a total of 22 IVUS pullback sequences were acquired during a constant speed pullback at 0.5mm/s [2]. X-ray sequences during diagnosis and pullback were acquired with a monoplane Philips Integris Cathlab.

In the IVUS images stenoses and stents are visible. Also common IVUS imaging artefacts appear (see appendix A), like ringdown effects, guidewire

[2] A Boston Scientific Atlantis probe at a frequency of 40 MHz and a Boston Scientific CVIS system were used

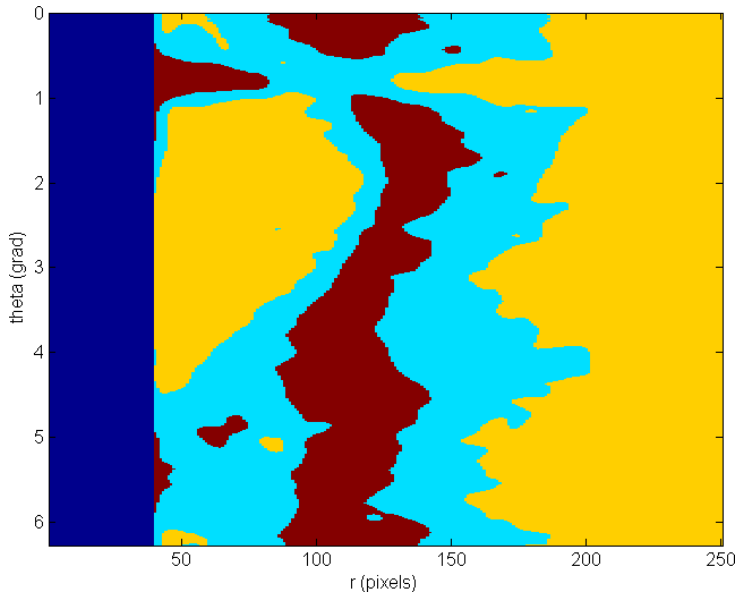

Figure 6.39: Result of applying a k-mean clustering algorithm to figure 6.38.

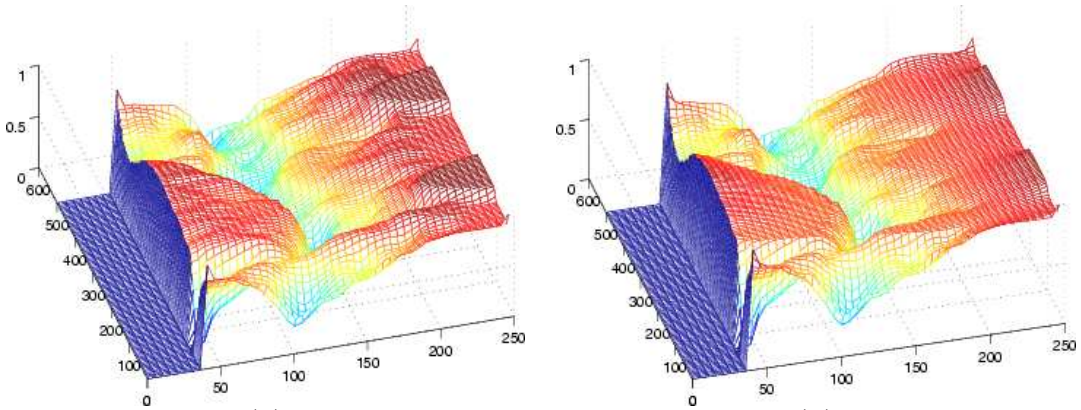

Figure 6.40: (a) IVUS image after median filtering, normalisation and inversion. (b) Result from smoothing the lumen area in (a).

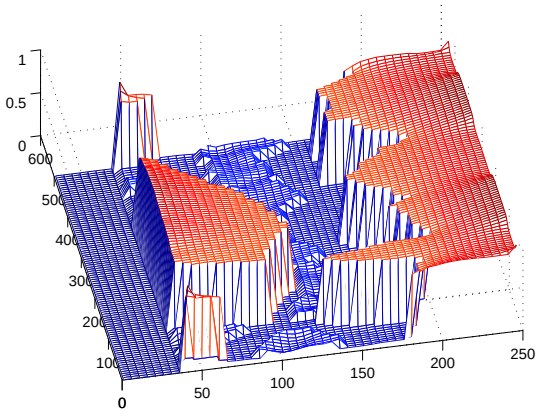

Figure 6.41: Result of multiplying the non-lumen tissue by a constant $k < 1$.

shadows and blood speckle [153].

The deformable cylinder was initialised with a radius $r = R_{ring}$ and 50 equally spaced points in each cross-section. The relative weights of the external and internal forces were empirically chosen $\alpha_i = 0.6$ and $\beta_i = 0.4$, respectively, giving more relevance to the external forces derived from the image features. If $\beta_i \gg \alpha_i$ then the resulting cylinder is very regular but does not follow the image features. On the other hand, if $\alpha_i \gg \beta_i$ then the resulting cylinder can be very irregular (it has sharp bends) and is very much affected by noise and artefacts in the image. The damping factor was chosen $\gamma = 0.3$. It controls how fast the initial cylinder takes to converge to the wished segmentation results.

Cross-sections of the segmentation results in two different images (with and without plaque) can be seen in figure 6.42.

6.9 Conclusions

A method to fuse IVUS and Cathlab information in the presence of heart beat and respiration is presented based on previously presented concepts of phase-space and device to vessel registration algorithms. Our method not only enables IVUS and x-ray fusion for diagnosis, but also supports instrument positioning.

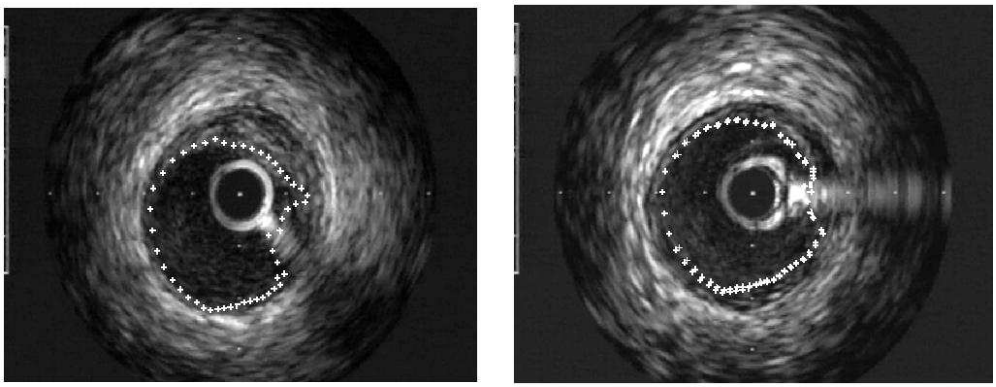

(a) (b)

Figure 6.42: Segmentation results in two IVUS images with (a) and without (b) plaque.

As a first step a set of reference pullback paths for a number of heart phases are created from IVUS pullback sequences and a number of reference angiograms. These reference paths can be used to aid diagnosis. Phase-space characterisation of the angiographic x-ray acquisitions and device to vessel registration to register each IVUS pullback frame to a position in a diagnostic angiogram are necessary. During the intervention, the position of the device is registered to the previously reconstructed pullback path such that the IVUS image corresponding to the current device location can be displayed.

Due to the lack of appropriate clinical data, our pullback reconstruction method was tested and demonstrated with a heart phantom featuring heart and respiration independent movements, and automatic constant speed pullback. Our measurements indicate that it is possible to reconstruct a two-dimensional pullback path for each quantised heart phase with an accuracy of less than $2mm$ which corresponds to ± 1 vessel radius in longitudinal direction, accurate instrument positioning. Our measurements also indicate that in the phantom case it is possible to register an instrument with a previously acquired IVUS image to an accuracy determined by the respiration compensation algorithm employed. *In vivo* clinical data acquired according to the workflow suggested in section 6.3 should be acquired in order to evaluate the presented algorithms on real clinical cases.

In this chapter we also discussed the possibilities of our method for 3D path reconstruction. When using projection x-ray for imaging, we loose depth

information. Because we work with two-dimensional projection images, an accurate three-dimensional reconstruction of the vessel is indeed impracticable and not the aim of the presented methods. We argue that the physician usually chooses the imaging plane so as to be approximately parallel to the imaged artery segment. In such a case, the information contained in the missing dimension is little, providing that the segment is not very tortuous. Given the typical length of a stenosis, the stiffness of the vessel wall and the interventional catheter, we presume that the vessel will be relatively straight. We argue that the norm of the pullback velocity vector in the direction of the missing dimension can be calculated despite a sign indetermination. According to numerical simulations, this indetermination does not affect the resulting volume nor the relative rotation between the Frenet-Serret frames at each point. *In-vivo* acquisitions and a ground truth possibly using biplane acquisitions should be made available in order to measure the accuracy of the proposed method.

As a byproduct of our investigations and unrelated to image fusion, a 3D lumen segmentation algorithm based on a deformable cylinder represented with a 2-simplex mesh is briefly presented. Further mesh configurations and driving forces should be explored. Also the computational requirements of the algorithm should be evaluated to determine its applicability in a real Cathlab setting. A ground truth (manual segmentation of the images by a physician) should also be provided to evaluate the physiological usefulness of the algorithm and further fine-tune the parameters of the deformable cylinder.

Chapter 7

Prototype

The algorithms described in the previous chapters aim at improving diagnosis and intervention outcome of PTCA interventions in a clinical setting. Therefore, steps should be taken toward a functioning prototype that meets the specifications of the application. The most critical constrains are those related to computational time and resources. The algorithms should work real-time when real-time is needed (for example, during the intervention but not necessarily during diagnosis), and should be compatible with current system resources, i.e. the end user (the hospital) should not need to upgrade the hardware of the system to an incredibly expensive processor.

We developed a prototype software architecture [76], capable of handling the image processing algorithms described in chapters 4 and 5 for navigational aid within the hard real-time constrains imposed by the application: an interventional frame *must* be displayed on time, with or without angiographic overlay. This prototype was first tested on offline clinical data and later on tested on-site. Due to the lack of appropriate clinical data and the incapability of acquiring in-vivo data, XIVUS demonstration was not contemplated. The software architecture works on a time-sharing general-purpose operating system and a specialised real-time operating system was avoided. The system resources (CPU, memory) are switched between multiple jobs and the user may interact with each of the processes running. The prototype was implemented in C++ running on a standard PC that was equipped with additional hardware for image and ECG acquisition. The prototype was tested at the Leiden University Hospital (The Netherlands) for validation of the proposed methods as well as the architecture.

This chapter is organised as follows: section 7.1 comments on implementation issues of heart and respiration phase calculation. Section 7.2 defines

the required system specifications for clinical applicability. Section 7.3 describes the different working modules in which the prototype is structured. Sections 7.4 and 7.5 present the computational structure in threads of the prototype and the communication between threads, respectively. Section 7.6 presents the computation latency results and we conclude in section 7.7

7.1 Implementation issues

For real time heart phase computation a method to estimate R peaks on a real-time ECG is necessary. According to the described workflow (see chapter 2) a diagnostic sequence is acquired first and examined by the physician. Our method, presented in subsection 7.1.1, calculates an R-peak template from the angiographic acquisition ECG while the physician examines the sequence. This template is afterwards used to detect R peaks from the ECG acquired during the intervention.

An automatic selection of a reference frame and ROI for respiration phase calculation is also needed to minimise user interaction. A method based on respiration phase curve characteristics to select an appropriate ROI and reference frame is presented in subsection 7.1.2.

7.1.1 Heart phase

According to the current Cathlab workflow, first a set of diagnostic sequences are acquired and analysed before proceeding with the intervention. Heart phase computation according to equation 4.1 involves calculating the distance between R-peaks. During fluoroscopic acquisition the heart phase has to be calculated in real-time and an estimation of the heart rate, and therefore, the R-peak distance, is necessary. We here present an algorithm that first trains on the diagnostic sequences to extract an R-peak template for ulterior R-peak detection, and estimates an initial heart rate that is later updated during the intervention as a moving average of past heart rates.

The ECG acquired for the angiographic sequence selected to subsequently guide the intervention is analysed and an R-peak template extracted. The template is calculated as a an average of the dominant peaks in the signal and it has a size of 1/4 the number of ECG data points per frame. For typical diagnostic sequences, acquired with a frame rate of 12.5 fr/s and an ECG sampling rate of 300 samples/s this corresponds to a template of

length 6 samples. The estimated R-peak template is then correlated against the ECG for the angiographic sequence and a mean and maximum values for the correlation, $\bar{m}_c$ and M_c, are calculated. These values will later be used to define positive R-peak detection. During fluoroscopic acquisition the R-peak template is correlated with the incoming data in search of an R-peak using a sliding window. The correlation results are normalised by M_c and compared with the normalised mean $\bar{m}_c/M_c$ to detect positive detection results. According on hits or misses of the template and the actual ECG, the ECG phase is computed using a moving average estimation of the pulse rate.

7.1.2 Respiration phase calculation

In order to calculate the respiration phase as explained in chapter 4 a reference frame and a ROI need to be defined automatically. The selection of the reference frame and the ROI affect the ability of the calculated respiration phase of representing the respiration state of the patient in each frame. A reference frame should be such that it corresponds to begin- or end-respiration. A suitable ROI for respiration phase calculation is such that the image content changes depending on respiration but is least affected by heart beat.

Because during normal Cathlab interventions workflow diagnostic sequences are acquired using several different geometries and later visually evaluated by the physician, the algorithms necessary for reference frame and ROI selection can be applied after each angiographic sequence has been acquired and stored, and do not need to run in real-time. A number of candidate ROIs is defined and evaluated. Reference frame selection is also applied to each of these ROIs. In the following subsections the algorithms proposed for reference frame selection (subsection 7.1.2.1) and ROI selection (subsection 7.1.2.2) are presented

7.1.2.1 Reference frame selection

A good reference frame $A_r(x, y)$ is defined as that which is acquired at begin- or end-inspiration. The criteria for reference frame selection are

- High variability of the calculated respiration phase curve $\{\beta(n),\, n = 1 \dots N\}$,
- there exists another frame $\widetilde{A}_r(x, y)$ such that the respiration phase curve calculated with such a reference $\{\widetilde{\beta}(n),\, n = 1 \dots N\}$ has a phase difference of approximately π.

In order to select a reference frame, we first calculate $\beta_t(n)$ in a given ROI, where $t = 1 \dots N$, N being the total amount of acquired diagnostic frames. Heart contraction induced noise is filtered out by low-pass filtering each $\beta_t(n)$, $\{\beta_t^{LP}(n)\}$. Frames which result in $\{\beta_r^{LP}(n)\}$ with high dynamic range are selected. Then we search for $\widetilde{A}_r(x, y)$ by calculating

$$\gamma(n) = \beta_r^{LP}(n) - \beta_j^{LP}(n), \ \forall j, \ r \neq j \tag{7.1}$$

where r corresponds to the curves/frames that show a high variability. A reference frame $A_r(x, y)$ is then selected such that $\gamma(n)$ shows minimum variability.

7.1.2.2 ROI selection

A suitable ROI for respiration phase calculation should enclose image content that changes according to breathing but is least affected by heart contraction. If a diaphragm detection algorithm like [89] is available, a ROI of size such that it encompasses the diaphragm border in several frames could be used. Assuming that the diaphragm is not contained in the field of view, then eight competing ROIs are defined as seen in figure 7.1 and evaluated as candidates for respiration phase calculation. We can see in figure 7.1 that the proposed ROIs do not span the centre of the image, since it is expected that the lesion will be centered on the image field and therefore this region will be filled with contrast agent and the resulting respiration phase would be influenced by noise caused by heart contraction movements.

For each ROI $\{\beta_t^{LP}(n)\}$ is calculated and used for residual calculation $\hat{\varepsilon}_i(n)$, where

$$\hat{\varepsilon}_i(n) = \beta_t(n) - \beta_t^{LP}(n) \tag{7.2}$$

A low variability of $\hat{\varepsilon}_i(n)$ indicates low dependency on heart beat.

7.2 System specifications

In a typical Cathlab intervention, after diagnostic sequences with different geometry settings have been acquired, the physician spends some time evaluating the lesion, deciding the therapy and choosing the system geometry for navigation and treatment. Then the physician proceeds with the intervention. This time interval between diagnosis and treatment can be used

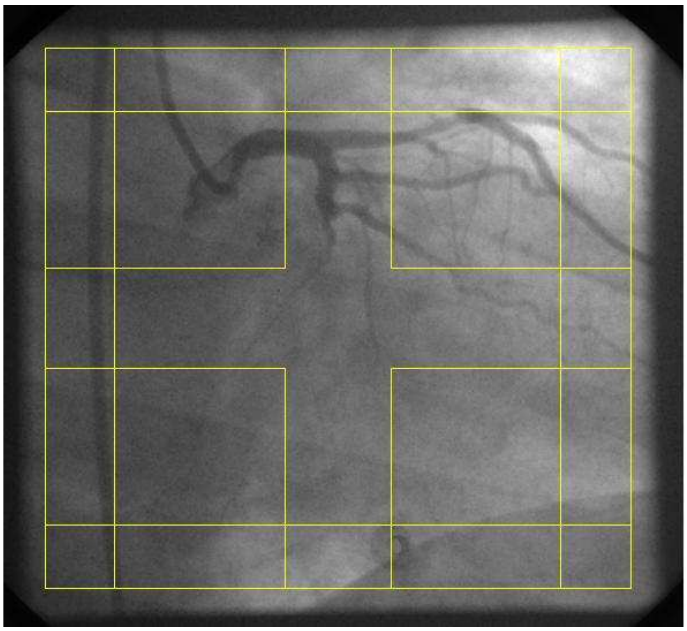

Figure 7.1: Angiogram where the defined 8 competing ROIs are depicted as rectangles along the image borders.

as computational time to evaluate the respiration and heart phase of the diagnostic frames as explained in the previous section, without interfering with the visualisation of the recently acquired angiographic sequences by the physician. During the intervention, the respiration and heart phase must be calculated in frame rate and the current interventional frame must be readily displayed on one of the monitors. Any latency between the acquisition of an interventional frame and its display on one of the monitors might disturb the physician's hand-eye coordination. Therefore, any calculation performed during the intervention must not interfere with the stream of fluoroscopic interventional images. The system must also be flexible enough to accommodate any interventional workflow. Although usually after diagnosis, the physician proceeds with the intervention, it might happen that during the procedure, contrast agent is injected again and an angiographic run is acquired. The software architecture must react accordingly to whichever workflow is enforced.

The software architecture has been designed to provide a platform for future pre-development activities and not specifically to accommodate the present respiration and heart phase calculation in x-ray acquisitions. This general approach only assumes that the acquisition hardware can deliver the

acquired frames to the main memory of the target platform. The following system specifications were defined:

1. It must be possible to implement arbitrary algorithms via a simple interface. This ensures the flexibility of the architecture as a platform for clinical validation of various methods.
2. The graphical user interface (GUI) must remain active at all times, so that the user can change algorithmic parameters when necessary.
3. The architecture must provide means to handle at least two types of image data: pre-interventional diagnostic and interventional.
4. The processing of diagnostic image sequences must be able to run in parallel to the acquisition of diagnostic and interventional sequences without detrimental effects to the live interventional image stream.
5. Whenever the acquisition mode is changed (from angiography to fluoroscopy or vice versa) the platform must react accordingly and without delay.
6. Independence of the system from the image acquisition hardware and drivers must be ensured.

7.3 System overview

The system is composed of three modules: a *data acquisition* module, an *off-line analysis* module and a *real-time* module, which perform the necessary calculations during diagnosis and the intervention respectively.

The data acquisition module is responsible for controlling the interaction with the acquisition hardware. The off-line analysis module analyses the angiographic sequences that are acquired during the diagnosis phase and constructs their phase-space. The real-time module computes in real-time the heart and respiration phase of each incoming frame, searches for a matching angiogram (in terms of heart and respiration phase) and the displays overlay result.

7.3.1 Data acquisition module

The data acquisition module acquires a synchronised video stream and ECG signal. During diagnostic angiography, the sequences are first completely

captured and subsequently stored before performing any calculation. Afterwards, the frames and the ECG will be passed to the off-line analysis module for ulterior analysis.

7.3.2 Off-line analysis module

Each angiographic sequence is acquired completely before starting any respiration or heart phase computation. Once the physician specifies the geometry settings (or selects from a list of diagnostic acquisitions which one should be used), the diagnostic frames and their corresponding ECG are input to this module, which calculates their heart and respiration phase after determines the most suitable angiogram for reference and the ROI for respiration phase calculation. Usually these calculations can be performed while the physician visually inspects a diagnostic sequence or acquires more diagnostic sequences with different geometry settings.

The off-line analysis module can be subdivided into two sub-modules (see figure 7.2): the *ECG analysis* module and the *respiration reference determination* module. The ECG analysis module calculates the heart phase as in equation 4.1. The respiration reference determination module searches for a proper angiogram reference, a ROI within the reference to compute the respiration phase and calculates the respiration phase as explained in chapter 3.6, section 4.5. To select the ROI, several possibilities are taken into account. The ROI could be selected manually by the physician. If there is a diaphragm detection algorithm implemented and the diaphragm is present in the frames, then a ROI around the diaphragm border is selected. If not, then different ROIs compete and the one that provides the most plausible respiration phase curve is selected as ROI. In this competition there are fixed candidate areas (see figure 7.3) close to corners and image borders to avoid the area centred around the coronaries, which are usually in the centre of the image.

The calculated heart and respiration phase are then stored and will be later used by the real-time module to search for a proper roadmap for each live fluoroscopic frame.

7.3.3 Real-time module

The real-time module is triggered by the acquisition of fluoroscopic frames. This module computes the respiration phase and heart phase of each incoming frame and searches for a suitable angiographic roadmap to overlay over

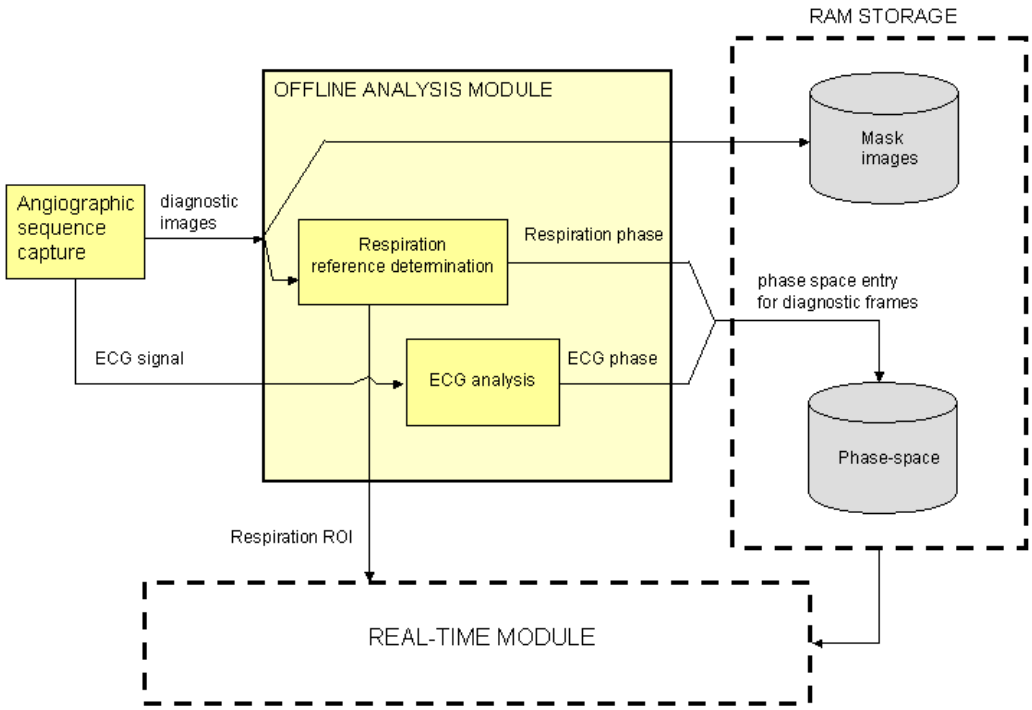

Figure 7.2: Structure of the off-line analysis module. It is composed of the respiration reference determination submodule and the ECG analysis submodule.

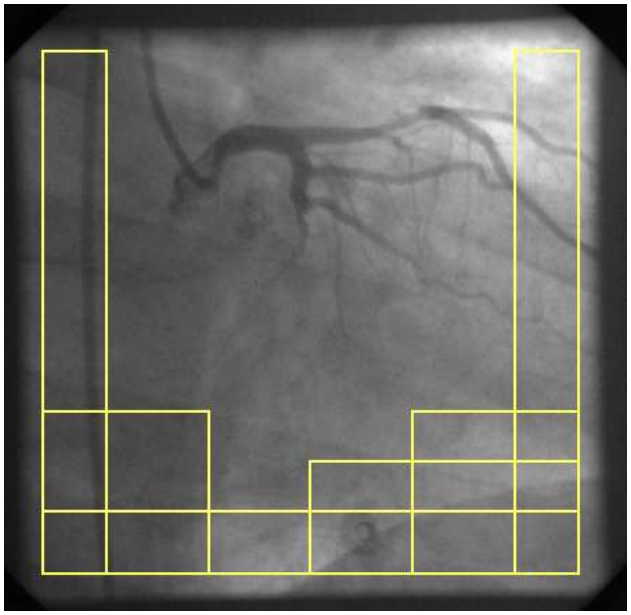

Figure 7.3: ROIs in a frame that compete for respiration phase calculation.

the current interventional frame. It receives from the offline module the ROI and the reference angiogram, and from the data acquisition module it obtains each live interventional frame and a synchronised ECG signal. Then it explores the phase-space calculated for the corresponding diagnostic acquisition to search for a roadmap to overlay to the current interventional frame. As mentioned in chapter 5, a suitable roadmap cannot always be found unless additional algorithms (respiration compensation and guidewire-to-vessel matching) are employed.

The real-time module can be further subdivided in the following submodules (see figure 7.4):

1. *Heart-phase determination module.* It computes the heart phase of the current frame given the ECG signal acquired by the acquisition module.

2. *Respiration phase determination module.* The respiration phase is computed using as reference angiogram and ROI, the ones provided by the off-line module.

3. *Phase-space matching.* Once the current interventional frame is characterised by its heart and respiration phase according to equations 4.1 and 4.7 in chapter 4, a matching (in terms of heart and respiration phase) angiogram is searched for in the stored phase-space for the diagnostic acquisition with the same geometry settings. The search does not always provide a matching angiogram as mentioned in chapter 5.

4. *Overlay module.* It fuses the interventional frame and the selected angiographic roadmap (if any). The resulting image is displayed on a monitor and shows in real-time the acquired interventional frame with an angiographic overlay.

The real-time module must work under the hard real-time constrains of our application. Only $1/f_s$ seconds of computation are available, where f_s is the frame rate during the intervention. Also, the calculations must not interfere with the visualisation of the live interventional frames.

7.4 Thread structure

The module structure proposed breaks the job into several tasks and we propose a thread structure to handle each necessary computation in parallel to speed up calculations (see figure 7.5).

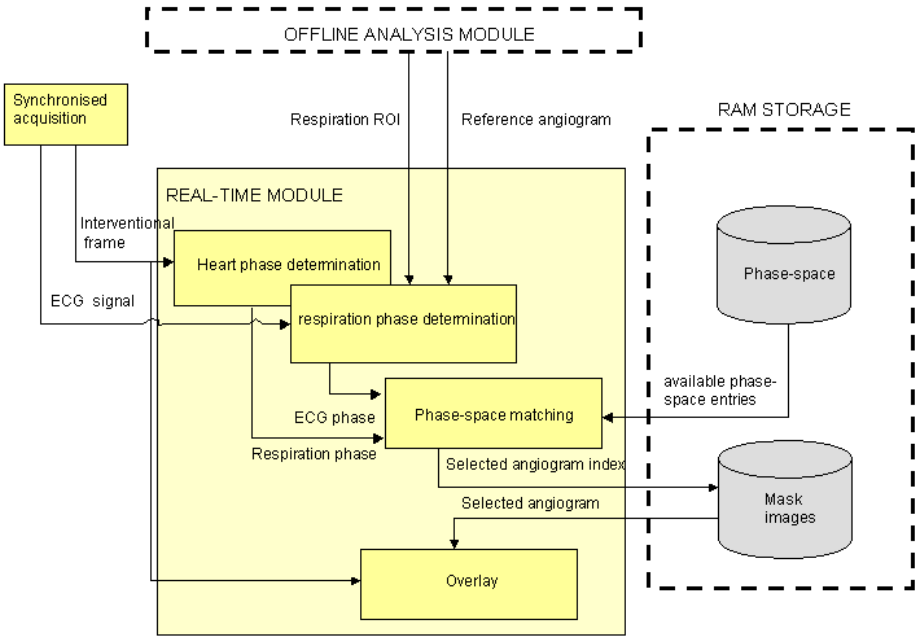

Figure 7.4: Structure of the real-time module. It is composed of the heart-phase determination, respiration reference determination, phase-space matching and overlay submodules.

> A *thread* or *light-weight process* is a basic unit of CPU utilisation and consists of a program counter, a register set and a stack space. It shares with peer threads its code section, data section and operating-system resources, collectively known as a *task* [154].

The creation of threads and CPU switching among peer threads is inexpensive, compared with context switches among heavyweight processes because no memory management-related work needs to be done, hereby allowing parallelism [154]. Therefore, implementing each task in a thread allows that at any time any thread can relinquish the CPU to the real-time thread to avoid a delay in the visualisation of the live interventional frames.

The proposed software architecture is the following [76] (see figure 7.5):

GUI and administration/control thread. It is the initial thread and handles all user interaction and administration/control tasks. This thread creates the other threads when needed and distributes the data acquired by the data acquisition module to the off-line or real-time analysis thread. It is also responsible of storage activities when CPU processing time is available (see subsection 7.4.5).

Data acquisition thread. The communication with the hardware grabber and its drivers is controlled by this thread. When data is acquired, it informs the control thread and proceeds to hand the acquired data to the GUI and administration thread. Because the acquisition of frames must always be active to avoid ignoring a frame during the intervention, this thread has high priority.

Off-line preparation of diagnostic data. This thread implements the tasks of the off-line module (subsection 7.3.2). It calculates the heart and respiration phase of the acquired diagnostic frames. While this thread is active, it must not interfere with the visualisation of the diagnostic data by the physician, which has high priority.

Real-time image enhancement. This thread performs the real-time calculations that allow for heart and respiration phase calculation, respiration compensation, vessel to instrument matching or any image processing for the current frame. These computations should run in frame rate.

Research and documentation thread. The data from the procedure is stored on hard disk for subsequent validation and analysis. This task

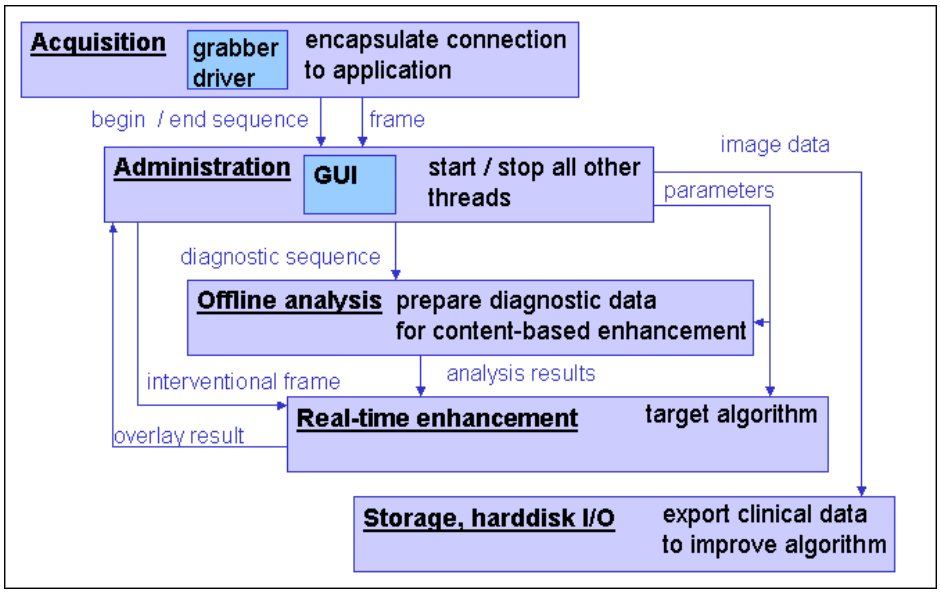

Figure 7.5: Thread structure of the proposed software architecture. The threads are in charge of the following tasks: administration, data acquisition, off-line analysis, real-time analysis and storage of the acquired data. The arrows represent data passed between threads.

has low priority because it only has a documentation purpose and is not crucial for the intervention outcome.

In figure 7.5 the arrows represent data that is handed from one thread to another. The GUI thread passes the algorithm parameters introduced by the user and the corresponding acquired frames to the off-line and real-time threads. It also receives the notifications of the beginning or end of an acquisition and the acquired frames. Thread communication will be further explained in section 7.5.

7.4.1 GUI and administration/control thread

The GUI and administration/control thread is created when the application is started, and is responsible for creating the other threads. It manages all operating system and user activity and communicates the data acquired by the acquisition to the corresponding threads. It is also responsible for displaying in the live interventional frames.

This thread receives the user input from the GUI and acts accordingly. In the current prototype implementation, the user manually identifies the

acquired sequences as diagnostic angiography or interventional fluoroscopy. This information determines the creation of an off-line or a real-time analysis thread. Also the angiographic sequence that corresponds to the interventional system geometry as well as the algorithmic parameters are selected manually. Because the user must be able at all times to interact with the system (e. g. to change parameters) with minimum latency, it is necessary to ensure that the message queue for this thread is kept open at all times. Therefore, no data analysis nor hardware communication activities are allowed in this thread.

After the administration thread creates an acquisition thread, it enters reactive mode, i.e. it waits until it recognises that data is available and starts either an off-line analysis thread if it is angiographic data, or a real-time thread if it is fluoroscopic data and a corresponding diagnostic sequence is available. When a diagnostic sequence is acquired, it is added to a list of all diagnostic acquisitions that are waiting to be processed by the off-line thread. During the intervention, the availability of an angiographic sequence acquired with the same geometry settings is checked by the user in order to start a real-time thread.

7.4.2 Data acquisition thread

The data acquisition thread is responsible for handling all driver-related calls that control the interaction of the application and the acquisition hardware. It also sends all acquired frames (diagnostic or interventional) to the administration thread for further handling. Depending on whether the acquisition hardware provides a callback informing of an acquired frame or not, there are two different execution modes.

If the system provides a callback that announces the acquisition of a new frame, the thread sleeps until the notification of a new frame is received. Whenever a new frame is acquired, it is passed to the administration thread. If such hardware is available, the thread does not have to poll for new frames and can *sleep* while a notification of a new frame is not received. This has the advantage that the acquisition thread does not consume CPU cycles while waiting for a new frame.

If the acquisition hardware does not provide a callback for each incoming frame, then the acquisition thread has to poll the hardware for a frame. Small latency between the availability of a frame and the moment it is read would

require a highly prioritised thread and frequent polling, but this would result in an intensive use of resources, which would slow down the calculations performed by other threads and waste CPU resources. A polling scheme has been designed which minimises CPU use and avoids jitter artefacts due to the time-lag between the availability of the frame and its detection. The mechanism is based on the knowledge that a frame is not to be expected at all times. The thread polls the hardware for a frame with a frequency which is adapted to the actual frame rate and relinquishes the time slice whenever a frame is not expected.

Figure 7.6 shows the time line of the acquisition thread during the acquisition of a sequence when no callbacks are provided by the grabber driver. On initialisation, the expected acquisition frame rate f_s is given as a parameter by the user and therefore we know that the time between frames is $1/f_s$. An estimation of the sleep time t_{sleep} for the acquisition is set to one fourth of the acquisition period $t^0_{sleep} = 1/(4f_s)$. When the acquisition thread is started, it polls the hardware for a new frame. If none is available, then the thread sleeps for another t^0_{sleep}. When the first frame is detected, a notification about the start of an acquisition and the acquired frame are sent to the administration thread. Then t_{sleep} is set to $t^1_{sleep} = 0.9f_s$ and another poll is made after t^1_{sleep} has elapsed. This process is illustrated in figure 7.6a. If a frame is discovered, then t_{sleep} must be updated and reduced because a frame was already waiting when the poll was made. The sleep time for the next frame is set to $t^2_{sleep} = 0.9t^1_{sleep} = 0.81f_s$. If a frame is not available after t^2_{sleep}, the acquisition thread returns to sleeping, but since a new frame should be available shortly, we update the sleep time to $t^3_{sleep} = 0.1f_s$. This sleep time is not changed until a new frame is detected. This case is illustrated figure 7.6b. Then, the sleep time is again set to $t_{sleep} = t^0_{sleep}$. If the amount of time slept continuously exceeds a threshold $t_{thres} = 2/f_s$ the end of the sequence is inferred and the respective notification is sent to the administration thread (see figure 7.6c).

7.4.3 Off-line analysis thread

The off-line analysis thread receives the complete diagnostic sequences and the algorithmic parameters from the administration thread. It is accountable for the calculation of the respiration and heart phase of the diagnostic acquisitions and creating their phase-space. It also processes the angiograms for visualisation enhancement. These computations are usually performed either while the physician acquires more diagnostic sequences or while it inspects

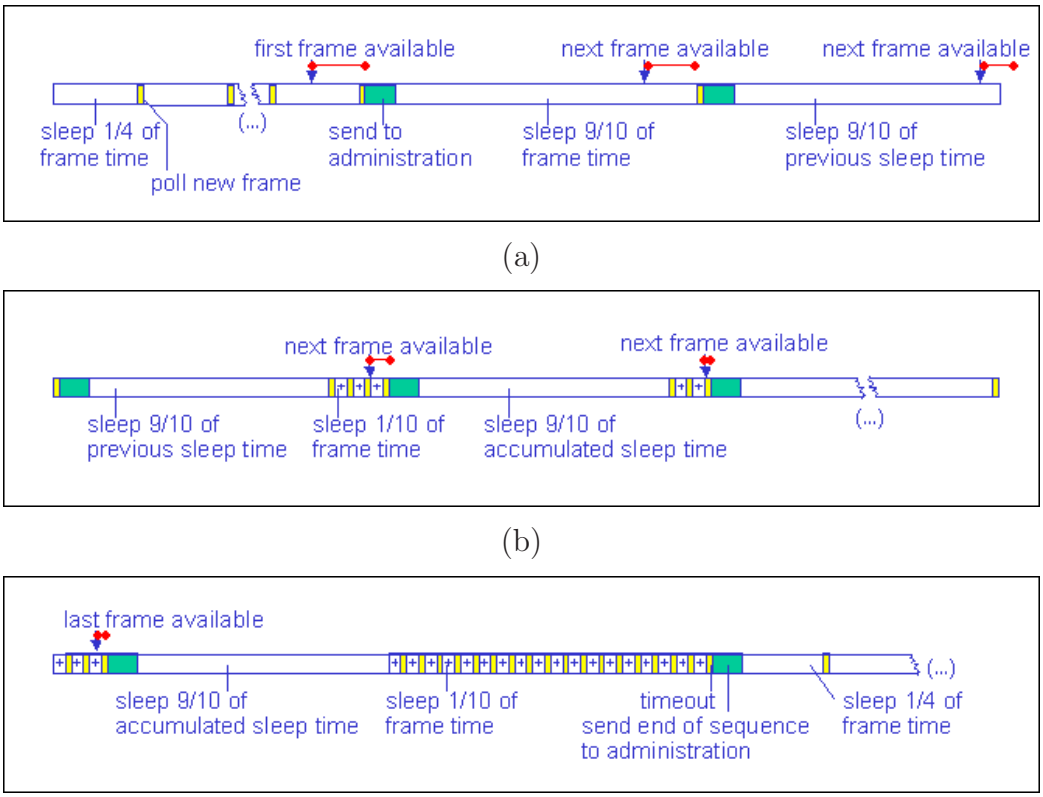

Figure 7.6: Synchronisation diagram when the hardware does not provide a call-back whenever a frame is acquired. In the first frames of a sequence, the mechanism synchronises to the acquisition (a). The sleep time changes depending on whether a frame was read or not (b) and (c).

in the work station the acquired sequences to evaluate the lesion. Therefore, there are not any real-time constrains. Nevertheless, we must ensure that the calculations do not interfere with the visualisation of the diagnostic data.

In the proposed architecture, all diagnostic data to be processed is arranged in a list. When a diagnostic acquisition is completed, an entry is created at the end of this list. Only when the list contains data that needs to be processed, does the administration thread create an off-line thread to handle this data. Whenever an off-line analysis finishes, the list is once more tested for entries that require processing. If the list is empty, the off-line thread ends and will be restarted only when more diagnostic sequences are acquired.

7.4.4 Real-time processing thread

This thread implements the tasks of the real-time module (see subsection 7.3.3): heart- and respiration-phase calculation for each incoming live frame, phase-space search for an adequate angiographic roadmap, and overlay of said roadmap (if found) with the current interventional frame.

Because an overlay can be found only if the imaging geometry of the diagnostic and interventional acquisitions is the same, this thread is started if and only if a diagnostic sequence with the same geometry settings was previously acquired and was chosen by the user (see section 7.4.1). The phase-space of the diagnostic sequence and the algorithm parameters are provided by the administration thread.

All the above tasks must be completed in frame rate to be of any use to the system. If the processing has not finished when the next interventional frame is acquired, the thread quits trying to find a suitable roadmap overlay and the administration thread will display just the interventional frame.

7.4.5 Storage and memory preservation

For documentation and further algorithm development activities, it would be useful to store the acquired data during an intervention. This is not a core activity of the system and is assigned a low-priority thread. Therefore, system resources for immediate storage on a hard-disc or other storage media are not available. Because the amount of acquired data is usually very high, the main memory of the system has not sufficient capacity. The proposed strategy for storage of the acquired sequences and the preservation of main memory is based on the observation that interventions are performed in distinct steps with pauses in between (i. e. at the end of an acquisition or when interventional devices are exchanged). These pauses, during which no high-priority tasks (like visualisation of interventional frames) are performed, can be employed to store the acquired data.

We propose to store the acquired image stream in the system's main memory as long as there is space. When the memory's limit is reached, any further acquired image will not be stored. At the end of each acquisition, a low-prioritised task, which writes the sequence from main memory to hard disc whenever computational resources are available, is started. The previously allocated main memory becomes again available to store more recently

acquired interventional sequences. Just like in off-line processing, a queue is used to gather all acquired data that requires storage on hard disc.

7.5 Thread communication

If we take another look at figure 7.5, we can see that there is certain data that is handed from one thread to another. The acquisition thread notifies the GUI thread about the beginning or end of a sequence, and transfers it the acquired frame. The GUI thread communicates the algorithm parameters introduces by the user, and the corresponding acquired frames to the off-line and real-time threads.

The communication between the different threads is message-based and uses the *interprocess communication* (IC) facility provided by the operating system, that allows processes to communicate and to synchronies their actions [154]. By using IC we avoid resorting to shared memory, which might result in data inconsistencies when several threads are working in parallel.

For simplification of IC, the message handler was implemented *re-entrant*, i.e. the code does not rely on singleton data and works on private local variables and can therefore be executed by two or more threads at the same time [154]. This mechanism speeds thread communication and reduces latency. It is used when the processing of the message does not require a reaction of the GUI and can therefore be directly executed by the thread which issued the message.

7.6 Results

A Dell Precision WAS 530 $2 \times 2.4GHz$ workstation was selected as target platform. The prototype was implemented on a standard PC equipped with image and ECG acquisition hardware [76]. The PC grabber card interfaces to the x-ray detector and writes the image stream to the PC's main memory using *direct memory access* (DMA). The frame grabber does not provide a callback when frames are acquired and therefore has to be polled (see subsection 7.4.2). The buffered analog output of an ECG device is read via an isolated amplifier and a National Instruments data acquisition card, which writes the ECG samples into main memory using DMA. The system was configured to acquire diagnostic and interventional 512^2 frames at $f_s = 12.5$fr/s.

The ECG system provided in parallel 300samples/s.

Pre-recorded clinical sequences with 2^{10} gray levels were analysed by the prototype using an acquisition emulator that reads clinical stored data and presents it to the application imitating a Cathlab system. The latency (t_L) and the performance of the synchronization mechanisms of the acquisition thread were tested. The latency was measured in three cases: when only diagnostic data was acquired (figure 7.7a), diagnostic data was acquired during off-line processing of previously acquired diagnostic data (figure 7.7b), and while the real-time thread was active (figure 7.7c). In the figures, the continuous line represents total latency, the dashed line represents the acquisition latency, and the dotted line represents the time for computation and display.

We can see on all three figures that, at the start of the intervention, a high amount of computation time is required. This time exceeds the frame period of $T_s = 1/f_s = 80$ms and, therefore the second and third frames are not processed and are displayed directly (this effect is visible in the step in the computation time per frame line in figure 7.7c). The image acquisition introduces a maximum latency of 5.7ms with an average latency of 3.1ms. The frame grabber is polled a maximum of 4 times per frame, and an average of 2.4 times.

The time to pass a frame to the administration thread and display it was 7.1ms if only the acquisition was active, 9.1ms when the off-line analysis was simultaneously running and 15.6ms when the real-time thread was active.

7.7 Conclusions

A flexible software architecture platform has been proposed to implement the algorithms described in chapter 3.6. The architecture is capable of handling off-line as well as real-time data and of administering computational resources such that the execution of the necessary operations does not interfere with the clinical workflow. The tasks of acquiring the images and calculating the heart and respiration phase for diagnostic and interventional images are distributed in modules, which are implemented with a thread structure. Each thread is a light-weight process that allows for reduced latency during computation. The results of the computations performed in each thread are communicated to other threads using IC.

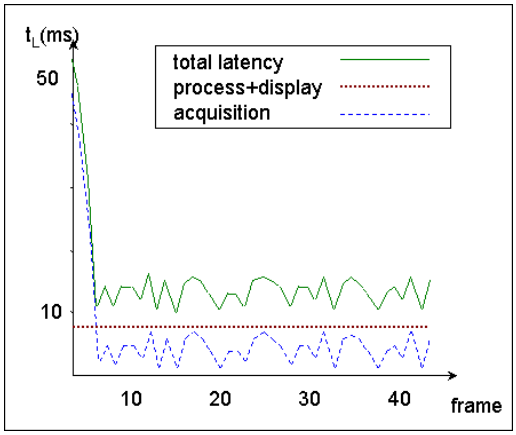

(a)

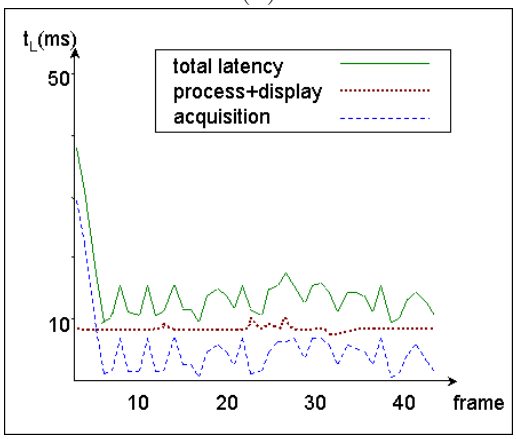

(b)

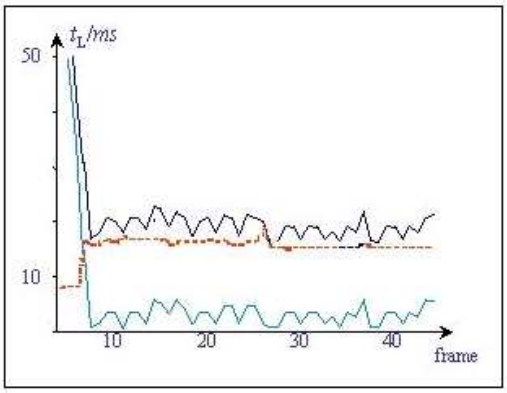

(c)

Figure 7.7: Latency measurements when (a) acquisition only, (b) off-line analysis, and (c) real-time processing are running.

In the suggested software architecture, diagnostic data analysis is performed by an off-line thread with relaxed time constrains. This thread runs in parallel to the visualisation of the diagnostic data by the physician. Whenever an interventional acquisition starts, a real-time thread performs the necessary calculations for navigational aid. An administration and data acquisition threads control the visualisation and acquisition of the frames respectively. The administration thread is also responsible for providing each active thread with the necessary information to perform its calculations.

Measurements are required that quantify the trade-off between available computational resources and the jitter of the input frames. In the current configuration, a worst-case-jitter of $6ms$ ($\ll T_s$) can occur, and the acquisition thread polls the frame grabber an average 2.4 for each frame. Jitter as well as CPU cycles consumed by the acquisition thread depend directly on the parameters that determine the sleeping time of the algorithm. Means to reduce the total latency by parameter optimization is work in progress.

Chapter 8

Conclusions

During a catheterisation procedure with a state-of-the-art Cathlab, the physician manipulates the instruments inside the patient's body by looking at one monitor displaying a static angiogram or *roadmap*, which shows the coronary arteries filled with contrast dye, while on another monitor the movement of the instrument can be seen in real-time. The role of the roadmap is to help the physician *estimate* the instrument's location. Because the static roadmap and the live images are displayed in different monitors, there is not a direct relation between their image content. Furthermore, since the heart is constantly moving inside the thorax, the roadmap does not always display the heart in the same state (i.e. shape and position) as in the interventional frames and a direct overlay will not improve the images presented to the physician. We propose to aid navigation and positioning of the devices by fusing diagnostic angiograms onto interventional frames in real-time, and by fusing projection x-ray images with IVUS by a novel multi-modality image fusion technique based on anatomical features.

Multi-modality image fusion registers images acquired with different imaging modalities to provide complementary information about anatomical structures. Multi-modality registration is usually performed on images from static body structures with the same image content. X-ray guided PTCA interventions can benefit from multi-modality registration by overlaying angiographic roadmaps to fluoroscopic interventional frames, in order to provide the physician with a fused image displaying vessels and interventional instruments simultaneously on the same monitor to aid navigation and positioning of devices. The difficulty lies in that images do not present similar image features (opaque vessels or interventional instruments) and the body structures being imaged are constantly moving. Therefore image feature-based fusion is not feasible. This thesis presents a novel framework for anatomical-based reg-

istration. An image is characterised by the state of the imaged anatomical structure. In PTCA interventions this is the heart state $\vec{\xi}$ that describes its position and shape, which are influenced by heart contraction, respiration, patient and table movement, and imaging geometry. By employing this framework, angiogram to interventional fluoroscopy registration is effectively achieved and presented to the physician for improved navigation and positioning during catheterisation interventions.

Novel algorithms to calculate $\vec{\xi}$ for PTCA applications are presented and evaluated on clinical data. These include heart state characterisation by heart and respiration phases, and phase-space representation for fast search of a suitable overlay for live fluoroscopic frames. The influence of $\Delta\vec{\xi}$ in the resulting fused image is also studied. New algorithms to compensate for such difference are proposed and evaluated on clinical data. These include a novel vessel detection, a real-time capable guidewire detection, device marker segmentation algorithm and vessel to device registration.

In all imaging geometries independent of the presence of a diaphragm or other assumed image content, the algorithm overlays onto each interventional frame an angiogram that shows the heart in the same contraction state and position within the image, hereby providing a suitable updated roadmap for navigation and positioning. As we have shown, each acquired image can be characterised (in the absence of patient or table movement) by the heart contraction and respiration phase. Each image in a diagnostic or interventional sequence can thus be represented as a point in a two-dimensional graph or *phase-space* [73][70][71][72], where nearby points correspond to images that display the heart in a similar state. To accomplish this task, the contraction or *heart phase* is determined from the ECG signal, whereas the *respiration phase* is calculated by comparing the image content to that of an angiographic reference. The similarity measure employed is *mutual information* because of its independence from x-ray dose and its sensitivity towards soft tissue changes [74]. It also allows to compare interventional frames, where only soft tissue and the interventional device is visible, with a reference angiogram where the vessels are filled with a radio-opaque contrast dye. The similarity measure is affected by breathing and heart contraction, which is removed by frequency analysis. Unlike previous approaches [19][20][21][22][23][71][72], the calculation of the respiration phase does not depend on the presence of the diaphragm in the image. The performance of MI in different areas of the image was compared with a bronze standard (section 4.5.2). It was shown that, although an area near the diaphragm gives a good respiration phase curve, other areas of the image are also affected by breathing and the

proposed similarity measure is able to measure the respiration state of the patient. We also concluded that, given a certain heart phase, breathing causes an affine transformation on the heart as suggested by [20]. The performance of mutual information was compared to other local correlation [95] and normalised generalised mutual information measures [78]. MI proved to have better performance than competing methods like LC, $nGMI_h^\alpha$ and $NGMI_h^\alpha$ (section 4.6).

The method is further extended to achieve compensation of $\Delta\vec{\xi}$ by presenting novel algorithms for accurate device to vessel registration (chapter 5). Once a proper roadmap is found for the current interventional frame, the location of the instrument with respect to the contrasted vessel might not coincide. This might be caused by small differences in the respiration or heart phase, by slight involuntary movements of the patient or by displacement of soft tissue. Two methods to refine the registration of the vessels visible in the roadmap and the instrument in the interventional frames are commented. The method in [73] (section 5.1) overcomes the sparse sampling of the phase-space by estimating a heart-position function parameterised by the respiration phase. The reported registration accuracy of this method is 3.3 pixels. The method in [32][121] (section 5.4) proposes a device to vessel match algorithm capable of running in frame rate. The method proposes a vessel detection filter (section 5.2.1), *directional stamping*, to segment vessel-like structures from the diagnostic roadmap and the guidewire from the interventional frames. An instrument marker segmentation algorithm based on *pdf* estimation is presented in section 5.3. A multiplicative distance transform with an hyperbolic kernel (section 5.4) ensures the registration of the segmented device with the nearest vessel.

A novel multimodality application that fuses x-ray and IVUS was presented in chapter 6. This method differs from those in [56][60][61][62][63][64][65][66][67][57][68] in that it fuses *projection* x-ray images with IVUS for diagnosis *and positioning* purposes. To perform this task, the phase-space representation and respiration compensation are employed to reconstruct a two-dimensional pullback path for all possible (quantised) heart phases and a given respiration phase. The reconstructed pullback path is fused with the acquired IVUS images and presented to the physician for diagnostic purposes (section 6.4). During the intervention, the fluoroscopic frames are registered to the pullback path for the same heart phase and the current position of the instrument can be related to that of an IVUS frame to aid accurate positioning of the instrument with respect to the lesion (section 6.5). The three-dimensional capabilities of our method are also discussed in subsection

6.7 and its limitations discussed in section 6.7.3. A lumen segmentation algorithm based on a deformable cylinder and a k-mean clustering driving force is presented in section 6.8. The XIVUS algorithm for diagnosis was tested on a heart phantom featuring independent respiration and contraction movements. Our results indicate that the two-dimensional pullback path can be reconstructed with an accuracy of $2mm$. This value depends on the respiration compensation method employed.

Some of the above mentioned algorithms have been implemented in a real-time prototype architecture [76]. The software architecture was developed to ease transfer from research activities to a clinical setting, is flexible and real-time capable. Computation is divided into off-line preparation of angiograms and real-time analysis of interventional frames. The software architecture is organised in threads that take care of the following tasks: administration and user interaction, acquisition of data (frames and ECG), off-line calculations, and real-time calculations. The proposed software architecture runs in frame rate and is therefore suitable for a clinical setting.

8.1 Further work

There are several areas in which this work can be continued. Diagnostic sequences are characterised by inflow, stationary state and outflow of contrast agent. Only the frames in which the coronary vessels are filled with contrast dye are useful for navigational purposes. A method to detect such angiograms would be useful and save computational time, since the respiration and heart phase would not be calculated for frames where the arteries are not visible. Also arbitrary injections of contrast dye during the intervention to improve the visualisation of the area might be processed as any other roadmap and reduce the sparseness of the phase-space. Such a method is proposed in [155].

The field of multimodality registration is an active one and new or improved measures are being proposed. Their dependence on contrast agent inflow/outflow or the presence of other moving objects (like the interventional devices) should be explored and evaluated. As for respiration compensation, the iterative solution proposed by Eck *et al.* in [73] assumes that a difference in respiration phase corresponds to a rigid body translation, and that respiration is not a hysteretic process. These assumptions contradict the results in [52][71] and section 4.5.2. A respiration compensation algorithm should take both effects into account to provide a more accurate compensation.

An instrument marker segmentation algorithm has also been proposed. There is extensive literature on object segmentation or guidewire segmentation in medical images. Nevertheless, the area of automatic guidewire or device marker segmentation has not been extensively explored [122][124][125][126]. An algorithm to segment device markers (or any dark moving object) based on *pdf* estimation was proposed. This algorithm needs initialisation time in order to collect enough samples to estimate a meaningful *pdf*. The real-time capabilities of the algorithm after initialisation should be tested. Further testing on clinical data should also be performed.

The x-ray and IVUS fusion proposed in chapter 6 differs from the method proposed in the literature [56][60][61][62][63][64][65][66][67][57][68]. Our algorithm does not require a biplane acquisition system and aids positioning of the instrument during the procedure. This method was (partially) evaluated on a heart phantom. Measurements as to registration accuracy for two-dimensional pullback reconstruction, and instrument to IVUS registration should be performed on clinical data. A method to reconstruct the three-dimensional pullback path was proposed. Phantom measurements comparing the results of reconstructing the pullback path using biplane angiography and our method should be performed to asses the validity of the proposed algorithm. Further testing of our lumen segmentation algorithm should be done to adjust the parameters of the deformable cylinder and explore other driving forces.

A lumen segmentation algorithm based on deformable surfaces is presented but not compared to a gold standard. Other mesh configurations and driving forces should be explored. Also the computational requirements of the algorithm should be evaluated to determine its applicability in a real Cathlab setting.

These methods have been presented on phantom data and, where available, clinical data. However, every new method includes changes in the Cathlab workflow that initially require additional x-ray and contrast agent and longer interventional time. Our XIVUS algorithms have been technically validated but a strategy to introduce these methods into the medical workflow in order to acquire the necessary data for clinical validation is required. The workflow and data acquisition protocols explained in this thesis should be followed for such purposes.

Furthermore, an appropriate user interface that suits the physicians needs

and wishes should be designed for the proposed applications. In the course of this work, several physicians were asked as to which type of image overlay they preferred, the answer being in some cases a stroboscopic update, in other cases a fadeout effect of the overlaid angiogram was preferred. Additionally, vessel enhancement can be applied, but always avoiding to diminish the diagnostic value of the images as they are acquired.

Appendix A

Artefacts in IVUS images

IVUS images suffer from artefacts, some of which are inherent to the imaging modality, while others are caused by the catheter movement and position. In the following sections we give a brief description of the more important artefacts [153].

A.1 Near-field artefacts

Near-field artefacts appear in the vicinity of the catheter. They are related to the catheter itself or to the blood flow.

A.1.1 Ringdown

Ringdown artefacts are produced by acoustic oscillations in the transducer. These oscillations result in high amplitude signals, that can be seen in an image as a bright halo around the central position of the catheter (figure A.1). They are present in all IVUS images and create an area of uncertainty around the catheter, which might complicate the segmentation of the inner wall of the artery from the lumen.

A.1.2 Blood speckle

Blood speckle is caused by the echogenicity of blood cells, which increases with the applied ultrasound frequency and as the blood velocity decreases (figure A.2). Blood speckle can make it difficult to distinguish the lumen

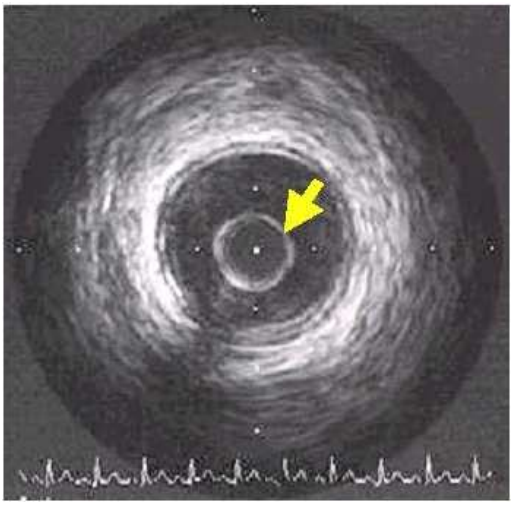

Figure A.1: The ringdown artefact (arrowed) appears around the catheter (centre of the image). (Image courtesy of TCT [2]).

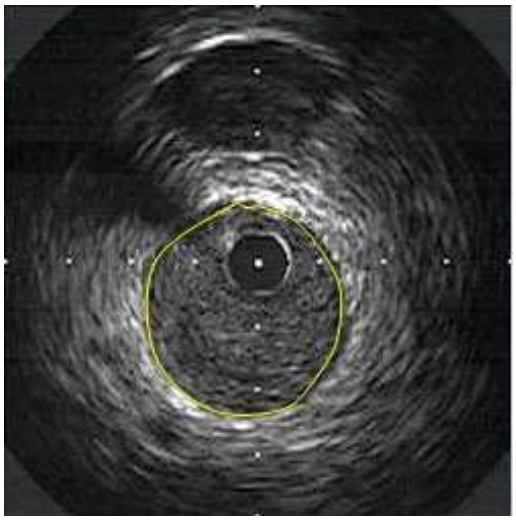

Figure A.2: Blood speckle artefact: blood cells become clearly visible inside the lumen (line). (Image courtesy of TCT [2]).

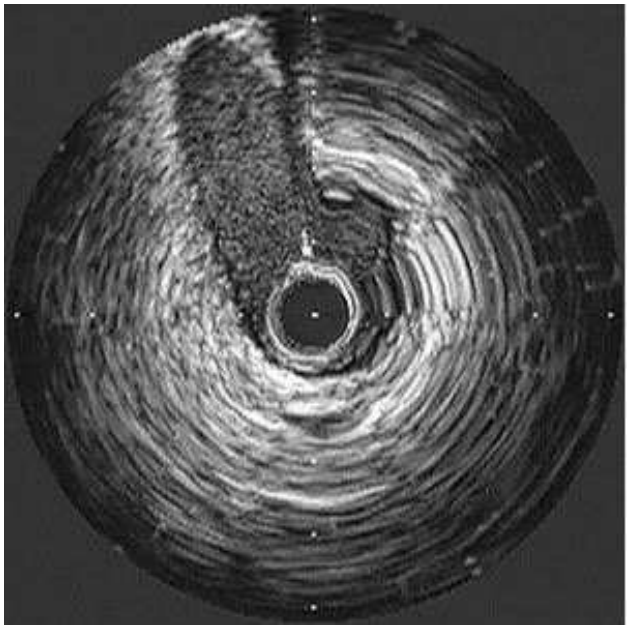

Figure A.3: NURD distortion [6] caused by the non-constant rotation speed of the IVUS transducer in mechanically driven catheters. (Image courtesy of TCT [2])

from the intima and makes necessary a flush of saline or contrast agent to clear the lumen.

A.2 Motion artefacts

These artefacts are caused by the non-uniform movement of the transducer at the catheter tip either during the motion along the vessel, or during the rotation of the transducer in mechanically driven catheters.

A.2.1 Non-Uniform Rotational Distortion

Non-Uniform Rotational Distortion (NURD) artefacts appear only in mechanically driven catheters. They are caused by a variation of the rotation speed of the transducer, which is caused by friction along the drive cable (figure A.3). The friction can originate from the sheath, a bending of the catheter, etc.

A.2.2 Cyclic distortion

The position of the catheter inside the vessel is not stable due to heart contraction and breathing, which also induces vessel pulsation. If a scanning is not completed before either of these occur, a cyclic deformation of the image appears.

A.2.3 Heart and breathing motion

The shape of the vessel changes during heart contraction and breathing [54][156][157]. To eliminate the influence of heart phase, ECG-gated devices are used. Unfortunately, the breathing motion, which was mentioned in early articles [157], has not been investigated yet.

A.3 Position artefacts

The course of the transducer does not follow the vessel centreline [67]. If this is not taken into account when reconstructing the 3D vessel tree it will cause a substantial error. Approaches that image the complete IVUS pullback [60][61][62], that identify the IVUS catheter before pullback as the pullback path [56][65][66][67][57] and that calculate a path of minimum bending energy between the starting and ending point of the pullback path [68] have been proposed (see chapter 6). Other position artefacts that should be taken into account are mentioned in the next subsections.

A.3.1 Transducer obliquity

Transducer obliquity happens when the imaging plane is not perpendicular to the vessel centreline. It causes geometric distortion and can result in vessel dimension overestimation. Furthermore, the reflection of ultrasound waves depends on the angle and so transducer obliquity can cause a decrease in image quality.

A.3.2 Vessel curvature

When reconstructing a 3D volume from a stack of IVUS images, the vessel curvature has to be taken into account. In the past, it has been common to consider IVUS images as parallel in space but lately, efforts have been made

in order to correct this with the use of coronary angiography [56][62][66][68].

A.3.3 Spatial orientation

Throughout the ECG cycle, and during the transducer pullback, the position of the IVUS probe within the vessel and its orientation changes. Also a flush of saline or contrast agent will cause the axial orientation of the catheter to change [56]. Direct stacking of consecutive IVUS frames [156][157] is unfeasible since the same features will appear at different locations in consecutive images.

It is possible to find the relative orientation between consecutive IVUS frames from the Frenet-Serret formulas, but still the overall orientation is unknown unless the artery shape in 3D is known. This information can be obtained from biplane angiography.

Appendix B

Measurements from IVUS images

IVUS images show the physician the morphology of the inner wall of the coronary arteries. A lesion is identified not only by low lumen area compared with that of a reference, but it is possible to identify the type of lesion, and thus provide the means to select the appropriate therapy and instruments. A number of measurements can be performed on IVUS images. These can be qualitative (tissue and lesion identification) or quantitative (lumen area, degree of stenosis,...). Measurement related to areas and diameters of lumen and/or other structures can be found already incorporated in existing products. With the help of angiography and a proper 3D reconstruction, it is also possible to measure volumes [56][62][66][68].

B.1 Morphology with IVUS

With IVUS imaging it is possible to view the structure of the atherosclerotic plaque. This information can help to select the correct treatment or stent/balloon size and improve the probability of success and decrease the probability of re-stenosis. IVUS also provides the means to check stent deployment.

B.1.1 Fibrotic/hard plaque

The majority of coronary lesions are caused by fibrotic plaque. It is highly reflective [56], and appears brighter than the adventitia. Its echogenicity in-

creases with its fibrousness. Sometimes it can cause acoustic shadowing like calcium (see section B.1.4).

B.1.2 Soft Plaque

Soft plaque is texturally homogeneous and produces a reduction in the lumen area. It has low echogenicity and is not as bright as the adventitia.

B.1.3 Vulnerable plaque

Vulnerable plaque is recognised to be the main cause of acute myocardial infarction [158][159]. Vulnerable plaque is composed of a lipid pool covered by a thick fibrous cap. The cap appears bright in the IVUS images followed by an area of low echogenicity [158][159]. While fibrotic or soft plaque reduces the lumen area, vulnerable plaque, in its early stages produces an outward growth (positive remodelling), so the lumen area is maintained and the diseased spot is invisible in angiography [159].

B.1.4 Calcium

Coronary arteries become calcified only in those areas that are atherosclerotic, following a process similar to bone mineralisation. The amount of calcification is directly related to the response of the lesion to catheter treatment [57]. While the amount of overall calcification in a vessel may predict the overall degree of vulnerable plaque, the calcification of an individual lesion does not indicate the vulnerability of the plaque [158].

Calcium is very hard to detect in angiographic images, while IVUS is very sensitive to it. Calcificated areas are highly reflective and produce acoustic shadowing (figure B.1) [56][158]. In an US image it will appear like a white bright area followed by a dark area. Depending on how near the lumen the calcification is, it is labelled as superficial or deep.

B.1.5 Stent placement

Depending on the design, stents appear in an IVUS frame as highly echogenic points or arcs (figure B.2). Comparing their position with the lumen-intima

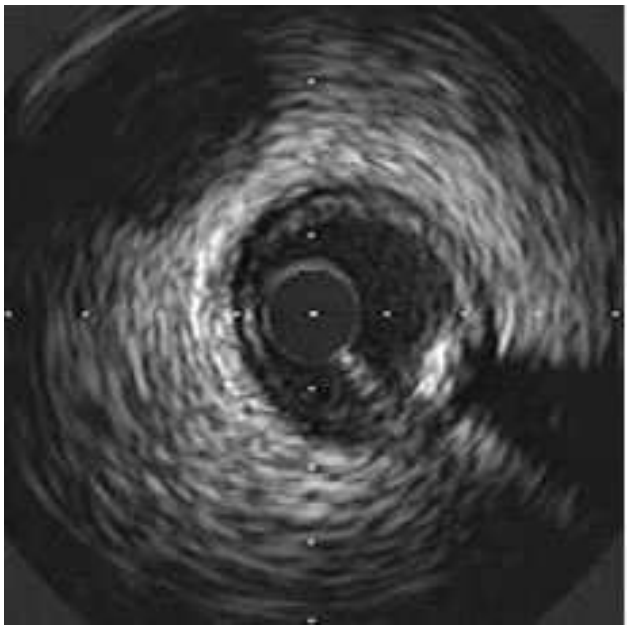

Figure B.1: We can see a calcium deposit at 4 o'clock marked by a very bright area followed by a shadow. (Image courtesy of TCT [2]).

boundary it is possible to ascertain the correct deployment of the stent. Common stent problems like incomplete apposition, incomplete expansion or edge tear can be easily seen with the help of IVUS.

B.2 Quantitative measurements

Currently IVUS images are used to carry out a number of quantitative measurements that help diagnose the stenosis. In the following sections we present these measures as defined in [160].

B.2.1 Lumen measurements

Lumen area and eccentricity are indicative of the gravity of the stenosis. In order to do comparative measurements, a reference segment is chosen (proximal or distal) and the lumen area at one spot is compared to the reference to detect a change of the lumen area.

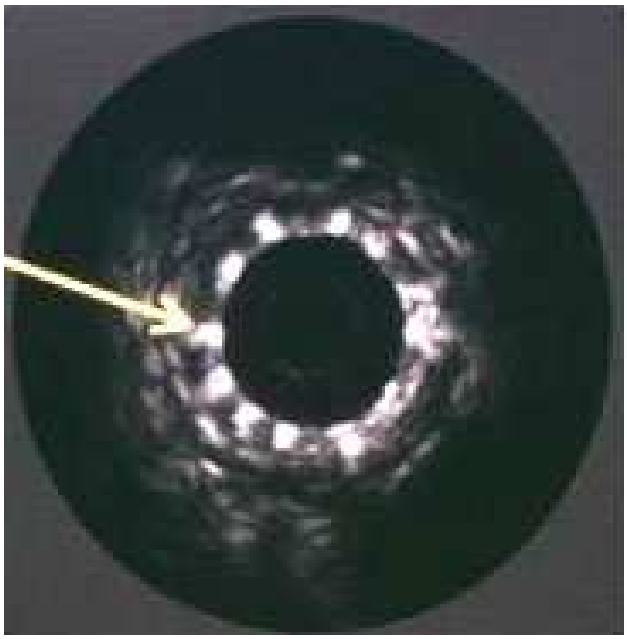

Figure B.2: Stent as seen in an IVUS image. The bright dots represent the struts of the stent (Image courtesy of TCT [2]).

B.2.1.1 Lumen Cross-Sectional Area

The Lumen Cross-Sectional Area (CSA) is the area bounded by the luminal border (figure B.3).

B.2.1.2 Maximum and minimum lumen diameters

The maximum and minimum lumen diameters (MLD and mld) are the longest and shortest diameter through the centre of mass of the lumen.

B.2.1.3 Lumen eccentricity

Lumen eccentricity (LE) is defined as

$$LE = \frac{MLD - mld}{MLD} \tag{B.1}$$

where

- MLD is the maximum lumen diameter, and

Figure B.3: Lumen CSA measurement. (Image courtesy of G. Dangas [7]).

- *mld* is the minimum lumen diameter.

B.2.1.4 Lumen area stenosis

The lumen area stenosis (LAS) represents the area within the EEM occupied by the stenosis, and is defined as

$$LAS = \frac{CSA_{RL} - CSA_{ML}}{CSA_{RL}} \tag{B.2}$$

where

- CSA_{RL} is the reference lumen cross-sectional area, and
- CSA_{ML} is the minimum lumen cross-sectional area.

B.2.2 External elastic membrane

The *external elastic membrane* (EEM) separates de adventitia or outer vessel layer from the media layer. The EEM is very difficult to identify when large side branches originate or when the shadow caused by calcium deposits subtends more than 90°. In healthy vessels the EEM has a circular shape, while in diseased vessels it takes an eccentric shape.

B.2.2.1 EEM CSA

EEM CSA is the area enclosed by the EEM border (figure B.4). It is representative of the total vessel area.

Figure B.4: EEM CSA measurement. (Image courtesy of G. Dangas [7]).

B.2.2.2 Minimum and maximum diameters

The minimum and maximum diameters of the EEM are the maximum and minimum diameters of the EEM border measured from the vessel centre.

B.2.3 Atheroma measurements

The atheroma measurements quantify the extent of the plaque and indicate the importance of the atheroma in the vessel.

B.2.3.1 Plaque plus media CSA

Since the media cannot be easily segmented this measure provides an estimation of the extent of the plaque (figure B.5). It is defined as

$$CSA_{PM} = CSA_{EEM} - CSA_L \tag{B.3}$$

where

- CSA_{PM} is the plaque plus media cross-sectional area,
- CSA_{EEM} is the EEM cross-sectional area, and
- CSA_L is the lumen cross-sectional area.

Figure B.5: Plaque plus media CSA in an IVUS image. (Image courtesy of G. Dangas [7]).

If there is a stent, the stent CSA (see section B.2.4) is used instead of the CSA_{EEM}.

B.2.3.2 Maximum and minimum plaque plus media thickness

The maximum and minimum plaque plus media thickness (MPM and mpm) are the largest and shortest distance from the intima edge to the EEM along any line passing through the centre of the lumen.

B.2.3.3 Plaque plus media eccentricity

The plaque plus media eccentricity (E_{PM}) is defined by the following equation as

$$E_{PM} = \frac{MPM - mpm}{MPM} \tag{B.4}$$

where

- E_{PM} is the plaque plus media eccentricity,
- MPM is the maximum plaque plus media thickness, and
- mpm is the minimum plaque plus media thickness.

B.2.3.4 Plaque burden

The plaque burden (PB) is defined as

$$PB = \frac{CSA_{PM}}{CSA_{EEM}} \tag{B.5}$$

where

- PB is the plaque burden,
- CSA_{PM} is the plaque plus media cross-sectional area, and
- CSA_{EEM} is the EEM cross-sectional area.

B.2.4 Stent measurements

The stent struts are highly echogenic, and appear as very bright dots or arcs in the image. Stent measurements can be a good indicator of the intervention outcome, because they provide information about the stent struts apposition and its proximity to the artery walls.

B.2.4.1 Stent CSA

It is defined as the area bounded by the stent border (figure B.6).

Figure B.6: Stent CSA in an IVUS image. (Image courtesy of G. Dangas [7]).

B.2.4.2 Maximum and minimum stent diameter

The maximum and minimum stent diameters (MSD and msd) are the longest and shortest diameters through the centre of mass of the stent.

B.2.4.3 Stent symmetry

The stent symmetry (SS) is defined as

$$SS = \frac{MSD - msd}{MSD} \tag{B.6}$$

where

- SS is the stent symmetry,
- MSD is the maximum stent diameter, and
- msd is the minimum stent diameter.

B.2.4.4 Stent expansion

Stent expansion is the minimum stent CSA compared to the reference CSA.

B.2.5 Calcium measurements

Calcium appears as bright echoes that obstruct the penetration of ultrasound and create a trailing shadow. Because of this *acoustic shadowing* only the leading edge of the calcium deposit appears, preventing the measurement of the deposit depth.

B.2.5.1 Arc of calcium

The arc of calcium is measured in degrees by using an electronic protractor centred on the lumen as shown on figure B.7.

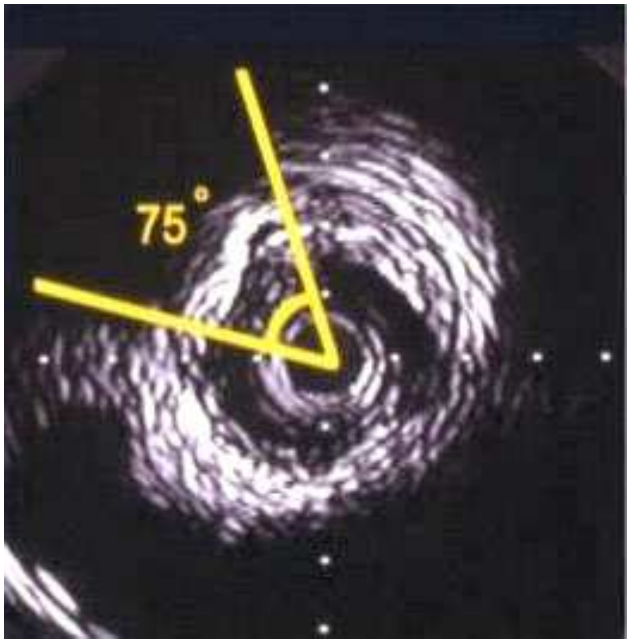

Figure B.7: Measurement of the arc of calcium. (Image courtesy TCT [2]).

B.2.6 Length measurements

A length measurement of the lesion can be performed with the help of a constant speed pullback. Knowing the time elapsed and the constant speed of the pullback, it is possible to calculate the length of any feature. If the information about the shape and curvature of the vessel is available, it is also possible to measure volumes.

B.3 Relevant measurements

According to the American College of Cardiology (ACC) [160] the basic quantitative measurements that should be performed with IVUS are:

- Minimum Lumen Diameter, which is directly related with the stent area, and the pressure of the balloon to be used,
- Minimum Stent Area, which compared to the lumen area can indicate the correct deployment of the stent,
- Plaque Burden, which is related to the severity of the lesion.

We also consider important the stenosis length in order to select a proper instrument size.

As to morphology, the ACC considers important to ascertain calcium, thrombus or dissection. The exact characterisation of the plaque (vulnerable, hard, ...) is not considered crucial. The information about vulnerable plaque, if encountered at all can be beneficial for prevention.

Bibliography

[1] Agency for Healthcare Research and Quality, "HCUPnet, Healthcare Cost and Utilization Project." http://hcupnet.ahrq.gov/.

[2] "G. Vince, Transcatheter Cardiovascular Therapeutics (online image)." http://www.tctmd.com/expert-presentations/multi-slide.html?product_id=80&large_image_p=0&start_idx=1, Slide2.jpg, Accessed September 2004.

[3] J.W. Hirshfeld Jr. *et al.*, "ACCF/AHA/HRS/SCAI Clinical competence statement on physician knowledge to optimize patient safety and image quality in fluoroscopically guided invasive cardiovascular procedures," *Journal of the American College of Cardiology*, vol. 44, pp. 2259–2282, December 2004.

[4] "Mid-atlantic surgical associates." http://www.heartsurgeons.com/ed2.html, angios.jpg, Accessed September 2004.

[5] "Jomed In Vision Gold." http://www.jomed.com/products/invisiongold/productinfo/in−vision−gold.html, invision.tif, Accessed June 2003.

[6] "Transcatheter Cardiovascular Therapeutics (online image), Non-uniform Rotation Distortion (NURD)." http://www.tctmd.com/ display/case/image/1447/, Non-uniform%20Rotation%20Distortion%20(NURD).jpg, Accessed August 2004.

[7] G. Dangas, "Mechanisms of Restenosis: Intravascular Ultrasound Insights (online image)." http://www.tctmd.com/display/expert/slides/4708/Slide5.jpg, Slide5.jpg, Accessed August 2004.

[8] American Heart Association, "Heart disease and stroke statistics - 2004 update." http://www.americanheart.org/downloadable/heart/1079736729696HDSStats2004UpdateREV3-19-04.pdf.

[9] "Stanford hospital and clinics (online image)." http://www. stanford-hospital.com/healthLib/atoz/cardiac/arteries.html, arteries.gif, Accessed January 2005.

[10] P. Kearney, "All you want to know about coronary arteries and angiography," *Heartwise*, vol. 2, no. 3, pp. 10–12, 1999.

[11] P. Libby, "Atherosclerosis: The new view," *Scientific American*, pp. 46–55, May 2002.

[12] S. Waxman, F.Ishibashi, and J. Muller, "Detection and Treatment of Vulnerable Plaques and Vulnerable Patients: Novel Approaches to Prevention of Coronary Events," *Journal of the American Heart Association. Circulation*, vol. 114, pp. 2390–2411, April 2006.

[13] L. Spagnoli, E. Bonanno, G. Sangiorgi, and A. Mauriello, "Role of inflamation in atherosclerosis," *Journal of Nuclear Medicine*, vol. 48, pp. 1800–1815, October 2007.

[14] "G. Vince, transcatheter Cardiovascular Therapeutics (online image)." http://www.tctmd.com/expert-presentations/multi-slide.html?product_id=11&large_image_p=0&start_idx=10, Slide25.jpg, Accessed September 2004.

[15] "Cimit." http://www. cimit.org/vplaque.html, vp2.jpg, Accessed January 2005.

[16] A. Farb, A. Burke, A. Tang, Y. Liang, P. Mannan, J. Smialek, and R. Virmani, "Coronary plaque erosion without rupture into a lipid core: a frequent cause of coronary thrombosis in sudden coronary death," *Journal of the American Heart Association. Circulation*, vol. 93, p. 1354–1364, 1996.

[17] A. Schmermund and R. Erbel, "Unstable coronary plaque and its relation to coronary calcium," *Journal of the American Heart Association. Circulation*, vol. 104, pp. 1682–1687, 2001.

[18] M. Sonka and J. Fitzpatrick, eds., *Handbook of medical imaging. Volume 2: Medical image processing and analysis*. ISBN 0-8194-3622-4.

[19] T. Sachs, C. Meyer, J. Pauly, B. Hu, D. Nishimura, and A. Macovski, "The real-time interactive 3-D-DVA for robust coronary MRA," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 73–79, February 2000.

[20] D. Manke, P. Rösch, K. Nehrke, P. Börnert, and O. Dössel, "Model evaluation and calibration for prospective respiratory motion correction in coronary MR angiography based on 3-D image registration," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1132–1141, September 2002.

[21] R. Ehman and J. Felmlee, "Adaptive technique for high-definition MR imaging of moving structures," *Radiology*, vol. 173, pp. 255–263, 1989.

[22] K. Nehrke, P. Bornert, J. Goren, J. Smink, and J. Bock, "On the performance and accuracy of 2D navigator pulses," *Magen. Reson. Imag.*, vol. 17, no. 8, pp. 1173–1181, 1999.

[23] R. Ehman, M. McNamara, M. Pallack, H. Hricak, and C. Higgins, "Magnetic resonance imaging with respiratory gating: Techniques and advantages," *American Journal Roentegenol.*, vol. 143, pp. 1175–1182, December 1984.

[24] S. Bonnet, A. Koenig, S. Rouy, P. Hugonnard, R. Guillemaud, and P. Grangeat, "Dynamic x-ray computed tomography," in *Proceedings of the IEEE*, vol. 91, pp. 1574–1587, October 2003.

[25] I. Buliev, C. Badea, Z. Kolitsi, and N. Pallikarakis, "Estimation of the heart respiratory motion with applicatiosn for cone beam computed tomography imaging: a simulation study," *IEEE Transactiosn on Information Techonology in Biomedicine*, vol. 7, pp. 404–411, December 2003.

[26] D. Parker, J. Wu, C. Tang, F. Gang, P. Frederick, and H. Marshall, "3D reconstruction of coronary arteries: correlation of CT calcium to vessel location," in *Proceedings Computers in Cardiology*, pp. 149–152, 1989.

[27] M. Kachelriess, S. Ulzheimer, and W. Kalender, "ECG-correlated imaging of the heart with subsecond multislice spiral CT," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 888–901, September 2000.

[28] M.J. Budoff *et al.*, "Assessment of coronary artery disease by cardiac computed tomography," *Journal of the American Heart Association. Circulation*, vol. 5, p. 35, 2006.

[29] C. Hong, C. Becker, R. Bruening, and M. Reise, "Multislice spiral computed tomography in cardiac imaging," *Journal of Hong Kong College of Radiologist*, vol. 4, pp. 252–258, December 2001.

[30] "Philips medical systems monoplane cathlab (online image)." http://www.medical.philips.com/main/products/cardiovascular/products/allura_flat_detector/index.html, flat_detector_ceiling.jpg, Accessed May 2004.

[31] "Philips medical systems biplane cathlab (online image)." http://www.medical.philips.com/main/products/cardiovascular/products/allura1215biplane/index.html, allura1215biplane.jpg, Accessed August 2004.

[32] J. Bredno, B. Martín Leung, and K. Eck, "Algorithmic solutions for live device-to-vessel match," in *Proceedings SPIE Medical Imaging* (J. Fitzpatrick and M. Sonka, eds.), vol. 5370 of *Image Processing*, pp. 1486–1497, 2004.

[33] "Philips medical systems cathlab procedure (online image)." http://www.medical.philips.com/main/products/cardiovascular/products/allura_flat_detector/features/, integris.jpg, Accessed May 2004.

[34] L. Verhoeven and E. Mast, "Coronary x-ray angiography: 40 years of experience." http://www.medical.philips.com/main/news/assets/docs/medicamundi/mm_vol43_no2/mm_vol43_no2_article_coronary_xray_angiography.pdf, Accessed December 2004.

[35] H. Amols, "Intravascular brachytherapy physics: introduction," in *Proceedings* 22^{nd} *annual EMBS International Conference*, vol. 2, pp. 1573–1576, 2000.

[36] J. Haller and M. Wholey, "The current status of laser angioplasty: coronary and peripheral results," in *Proceedings of the IEEE*, vol. 80, pp. 861–868, 1992.

[37] A. Wahle, J. Lopez, E. Pennington, S. Meeks, K. Braddy, J. Fox, T. Brennan, J. Buatti, J. Rossen, and M. Sonka, "Effects of vessel geometry and catheter position on dose delivery in intracoronary brachytherapy," *IEEE Transactions on Biomedical Engineering*, vol. 50, pp. 1286–1295, November 2003.

[38] M. Sabate, G. Pimentel, C. Prieta, J. Corral, C. Banuelos, D. Angiolillo, F. Alfonso, R. Hernandez-Antolin, J. Escaned, P. Fantidis, C. Fernandez, A. Fernandez-Ortiz, R. Moreno, and C. Macaya, "Intracoronary brachytherapy after stenting de novo lesions in diabetic patients:

results of a randomized intravascular ultrasound study," *Journal of the American College of Cardiology*, vol. 44, pp. 520–527, August 2004.

[39] "Cardio-vascular Center Hokkaido Ohno Hospital." http://www.cvc-ohno.or.jp/web/cardio/pci/pci2.html, Accessed August 2004.

[40] M. Prokop and M. Galanski, *Spiral and Multislice Computed Tomography of the Body*. Thieme, 2003.

[41] M. F. J. R. A. Lutz, K. Klingenbeck-Regn, "Spiral scanning computed tomography apparatus, and method for operating same, for cardiac imaging, United States Patent, No.: US 151588," 2001.

[42] N. Dendukuri, K. Chiu, and J. Brophy, "Validity of electron beam computed tomography for coronary artery disease: asystematic review and meta-analysis," *BMC Medicine*, vol. 5, p. 35, 2006.

[43] "Electron computed beam tomography." http://www.ncbi.nlm.nih.gov/books/bv.fcgi?indexed=google&rid=hstat6.section.622.

[44] S. U. C. C. A. Laboratory, "Can Routine Ultrasound Influence Stent Expansion (CRUISE)." http://www-med.stanford.edu/school/ccal/trials/, Accessed December 2004.

[45] S. U. C. C. A. Laboratory, "Angiography *vs.* Intravascular Ultrasound Directed Coronary Stent Placement (AVID)." http://www-med.stanford.edu/ school/ccal/trials/, Accessed December 2004.

[46] P. Oemrawsingh, M. Schalij, and E. van der Wall, "Guidance of long stent implantation by intravascular ultrasound has a significantly better outcome compared to implantation guided by angiography alone: final results of a randomized study (TULIP)," *Journal of the American College of Cardiology*, vol. 37, p. 20A (abstract), March 2001.

[47] E. Nix, A. Som, M. Rothman, and A.R.Pacey, "Ultrasonic visualization and catheters therefor, United States Patent, No.: US 6283921 B1, Date of Patent: September 4, 2001."

[48] Y. Jang, S. Salmon, D. White, and R. Garcia, "Catheter system having imaging, balloon angioplasty, and stent deployment capabilities, and methods of use for guided stent deployment, United States Patent, No.: US 6074362," 1997.

[49] Y. Jang, S. Salmon, D. White, and R. Garcia, "Catheter system having imaging, balloon angioplasty, and stent deployment capabilities, and methods of use for guided stent deployment, United States Patent, No.: US 5749848," 1995.

[50] P. Yock, "Method and apparatus for intravascular ultrasonography, United States Patent, No.: US 5676151," 1993.

[51] P. Feyter, "In-vivo imaging of acute coronary syndromes," *Journal of Clinical and Basic Cardiology*, vol. 3, no. 2, pp. 95–98, 2000.

[52] K. McLeish, D. Hill, D. Atkinson, J. Blackall, and R. Razavi, "A study of the motion and deformation of the heart due to respiration," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1142–1150, September 2002.

[53] A. Wahle, H. Oswald, and E. Fleck, "3D heart-vessel reconstruction from biplane angiograms," *IEEE Computer Graphics and Applications*, vol. 16, pp. 65–73, January 1996.

[54] A. Wahle, G. Prause, S. DeJong, and M. Sonka, "Determination of the absolute axial orientation of intracoronary ultrasound images in fusion with biplane angiography," *Computers in Cardiology*, pp. 153–156, September 1998.

[55] A. Wahle, G. Prause, S. DeJong, and M.Sonka, "Limitations of the manual pullback in intracoronary ultrasound imaging," in *Proceedings 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS), Hong Kong*, vol. 1, (Piscataway NJ), pp. 506–509, IEEE Press, 1998.

[56] A. Wahle, G. Prause, C. von Birgelen, R. Erbel, and M. Sonka, "Fusion of angiography and intravascular ultrasound in-vivo: establishing the absolute 3-D frame orientation," *IEEE Transaction on Biomedical Engineering - Biomedical Data Fusion*, vol. 46, pp. 1176–1180, October 1999.

[57] D. Rotger, P. Radeva, C. Cañero, J. Villanueva, J. Mauri, E. Fernandez, A. Tovar, and t. . V. Valle

[58] D. Rotger, P. Radeva, E. Fernandez-Nofrerias, and J. Mauri, "Multimodal registration of intravascular ultrasound images and angiography," in *Proceedings XX Congreso Anual de la Sociedad Española de Ingenieria Biomedica*, vol. 1, pp. 137–140, November 2002.

[59] D. Rotger, P. Radeva, and O. Rodriguez, "Vessel tortuosity extraction from ivus images," in *Computers in Cardiology*, vol. 33, pp. 689–692, September 2006.

[60] L. Kock, P. Kearney, R. Erbel, T. Roth, R. Brennecke, and J. Meyer, "Three dimensional reconstruction of intracoronary ultrasound images: roadmapping with simultaneously digitised coronary angiograms," in *Proceedings Computers in Cardiology*, pp. 89–91, 1993.

[61] J. Evans, K. Ng, S. Wiet, M. Vonesh, W. Burns, M. Radvany, B. Kane, C. Davidson, S. Roth, B. Kramer, S. Meyers, and D. McPherson, "Accurate three-dimensional reconstruction of intravascular ultrasound data, spatially correct three-dimensional reconstructions," *Circulation*, vol. 93, pp. 567–576, February 1996.

[62] R. Shekhar, R. Cothren, D. Vince, and J. Cornhill, "Fusion of intravascular ultrasound and biplane angiography for three-dimensional reconstruction of coronary arteries," in *Proceedings Computers in Cardiology*, pp. 5–8, 1996.

[63] O. Pujol, C. Cañero, P. Radeva, R. Toledo, J. Saludes, D. Gil, J. Villanueva, J. Mauri, B. Garcia, J. Gomez, A. Cequier, and E. Esplugas, "Three-dimensional reconstruction of coronary tree using intravasuclar ultrasound images," in *Proceedings Comput< ers in Cardiology*, vol. 26, pp. 265–268, 1999.

[64] J. Lengyel, D. Greenberg, and R. Popp, "Time-dependent three-dimensional intravascular ultrasound," in *Proceedings SIGGRAPH*, Computer Graphics and Interactive Techniques, pp. 457–464, 1995.

[65] M. Laban, J. Oomen, C. Slager, J. Wentzel, R. Krams, J. Schuurbiers, A. den Beer, C. von Birgelen, P. Serruys, and P. de Feijter, "ANGUS: a new approach to three-dimensional reconstruction of coronary vessels by combined use of angiography and intravascular ultrasound," *Computers in Cardiology*, pp. 325–328, 1995.

[66] A. Wahle, G. Prause, S. DeJong, and M. Sonka, "Geometrically correct 3-D reconstruction of intravascular ultrasound images by fusion with biplane angiography-methods and validation," *IEEE Transactions on Medical Imaging*, vol. 18, pp. 686–699, August 1999.

[67] A. Wahle, S. Mitchell, M. Olszewski, R. Long, and M. Sonka, "Accurate visualization and quantification of coronary vasculature by 3-D/4-D fusion from biplane angiography and intravascular ultrasound," in

European Biomedical Optics Week (EBiOS 2000): Biomonitoring and Endoscopy Technologies (I. Gannot, Y. V. Gulyaev, T. G. Papazoglou, and C. F. P. van Swol, eds.), vol. 4158, (Bellingham WA), pp. 144–155, SPIE Europto, 2000/2001.

[68] C. Molina, G. Prause, P. Radeva, and M. Sonka, "3-D catheter path reconstruction from biplane angiography using 3D snakes," in *Proceedings SPIE Medical Imaging*, Imaging Processing, pp. 504–512, 1998.

[69] Y. Wang, S. Riederer, and R. Ehman, "Respiratory motion of the heart: kinematics and the implications for the spatial resolution in coronary imaging," *Magn. Reson. Med.*, vol. 33, pp. 713–719, May 1995.

[70] D. Ro, L. Axel, and G. Herman, "Computed masks in coronary subtraction imaging," *IEEE Transactions on Medical Imaging*, vol. 6, pp. 292–296, December 1987.

[71] G. Shechter, *Respiratory motion of the heart: implications for magnetic resonance coronary angiography.* PhD thesis, John Hopkins University, 2004.

[72] G. Shechter, C. Ozturk, J. Resar, and E. McVeigh, "Free breathing respiratory motion of the heart measured from x-ray coronary angigrams," in *Proceedings SPIE Medical Imaging* (A. Amini and A. Manduca, eds.), vol. 5369 of *Physiology, Function and Structure from Medical Images*, pp. 34–43, 2004.

[73] K. Eck, I. Waechter, and J. Bredno, "Synthesis of angiographic images using iterative approximation," in *Proceedings SPIE Medical Imaging* (J. Fitzpatrick and M. Sonka, eds.), vol. 5370 of *Imaging Processing*, pp. 163–171, 2004.

[74] B. Martín Leung, K. Eck, I. Stuke, J. Bredno, and T. Aach, "Mutual information based respiration detection," *Computer Assisted Radiology and Surgery*, pp. 1085–1092, 2003.

[75] R. Close, J. Whiting, X. Da, and N. Eigler, "Stabilized display of coronary x-ray image sequences," in *Proceedings SPIE: Medical Imaging* (J. R.L. Galloway, ed.), vol. 5367 of *Visualization, Image-Guided Procedures and Display*, pp. 294–301, 2004.

[76] J. Bredno, B. Martín Leung, and K. Eck, "Software architecture for live enhancement of medical images," in *Proceedings SPIE Electronic*

Imaging (N. Kehtarnavaz and P. Laplante, eds.), vol. 5297 of *Real-Time Imaging VIII*, pp. 122–123, 2004.

[77] P. Viola, *Alignment by Maximization of Mutual Information*. PhD thesis, MIT Artificial Intelligence Lab, 1995.

[78] A. Bardera, M. Reiyas, and I. Boada, "Normalized simlarity measures for medical image registration," in *Proceedings SPIE Medical Imaging* (M. S. J.M. Fitzpatrick, ed.), vol. 5370 of *Image Processing*, pp. 108–118, 2004.

[79] I. Stuke, S. Mollus, A. Condurache, K. Eck, and T. Aach, "Cardio dynamic subtraction angiography (CDSA)," in *Proceedings of the Second Joint EMBS/BMES Conference*, vol. 2, pp. 915–916, 2002.

[80] H. Timinger, S. Krueger, J. Borgert, and R. Grewer, "Motion compensation for interventional navigation on 3D static roadmaps based on an affine model and gating," *Physics in Medicine and Biology*, vol. 49, pp. 719–732, March 2004.

[81] R. Botnar, M. Stuber, P. Danias, K. Kissinger, and W. Manning, "Improved coronary artery definition with T2-weighted, free breathing, three-dimensional coronary MRA," *Circulation*, vol. 99, pp. 3139–3148, June 1999.

[82] "Standford hospital and clinics (online image)." http://www. stanfordhospital.com/healthLib/atoz/cardiac/ecg.html, ecg.gif, Accessed June 2004.

[83] K. Eck, S. Mollus, I. Stuke, A. Condurache, and T. Aach, "Interventional positioning and navigation aids for the x-ray Cathlab. Part 1." Internal report, Philips Research Germany, 1581/2002, 2002.

[84] B. Sturm, K. Powerll, S. Halliburton, and R. D. White, "Image fusion of 4D cardiac CTA and MR images," *Applied Imagery Pattern Recognition Workshop*, pp. 21–24, 2001.

[85] G. Moody, R. Mark, A. Zoccola, and S. Mantero, "Derivation of respiratory signals from multi-lead ECGs," *Computers in Cardiology*, vol. 12, pp. 113–116, 1985.

[86] G. Moody, R. Mark, M. Bump, J. Weinstein, A. Berman, J. Mietus, and A. Goldberger, "Clinical validation of the ECG-Derived Respiration (EDR) technique," *Computers in Cardiology*, vol. 13, pp. 507–510, 1986.

[87] D. Caggiano and S. Reisman, "Respiration derived from the electrocardiogram: a quantitative comparison of three different methods," in *Proceedings of the 1996 IEEE Twenty-Second Annual Northeast Bioengineering Conference*, pp. 103–104, 1996.

[88] P. Zhang, W. Tapp, S. Reisman, and B. Natelson, "Respiration response curve analysis of heart rate variability," *IEEE Transactions on Biomedical Engineering*, vol. 44, pp. 321–325, April 1997.

[89] A. Condurache, T. Aach, K. Eck, J. Bredno, and T. Stehle, "Fast and robust diaphragm detection and tracking in cardiac x-ray projection images," in *Medical Imaging 2005: Image Processing* (J. M. Fitzpatrick and J. M. Reinhardt, eds.), (San Diego, USA), pp. 1766–1775, SPIE Vol. 5747, February 2005.

[90] J. Klingensmith and D. Vince, "Respiratory motion correction for geometrically correct three-dimensional reconstructions of coronary arteries," *IEEE Computers in Cardiology*, vol. 29, pp. 545–548, September 2002.

[91] H. Nakai, K. Ishihara, Y. Miyake, and M. Watanabe, "Non-restrictive visual respiration monitoring," in *Proceedings 15th International Conference on Pattern Recognition*, vol. 4, pp. 647–651, 2000.

[92] T. Buzug, J. Weese, C. Fassnacht, and C. Lorenz, "Elastic matching based on motion vector fields obtained with a histogram-based similarity measure for DSA-image correction," in *Proceedings of CARS 1997*, pp. 139–144, 1997.

[93] T. Buzug, J. Weese, C. Fassnacht, and C. Lorenz, "Elastic matching based on motion vector fields obtained with a histogram-based similarity measure for DSA-image correction," in *Computer Assisted Radiogology and Surgery*, pp. 139–144, 1997.

[94] T. Buzug and J. Weese, "Image registration for DSA quality enhancement," *Journal of Computerized Medical Imaging and Graphics*, vol. 22, no. 2, pp. 103–112, 1998.

[95] T. Netsch, P. Rösch, A. Muiswinkel, and J. Weese, "Towards real-time multi-modality 3-D medical image registration," in *Eighth IEEE International Conference on Computer Vision*, vol. 1, pp. 718–725, 2001.

[96] P. Viola and W. M. Wells, "Alignment by maximization of mutual information," in *Proceedings Fifth International Conference on Computer Vision*, pp. 16–23, 1995.

[97] J. Weese, G. Penney, P. Desmedt, T. Buzug, D. Hill, and D. Hawkes, "Voxel-based 2-D/3-D registration of fluoroscopy images and CT scans for image-guided surgery," *IEEE Transactions on Information Technology in Biomedicine*, vol. 1, pp. 284–293, December 1997.

[98] Y. Zhu and S. Cochoff, "Likelihood maximization approach to image registration," *IEEE Transactions on Image Processing*, vol. 11, pp. 1417–1426, December 2002.

[99] P. Cachier and X. Pennec, "3D non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions," in *Proceedings IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 182–189, 2000.

[100] L. Brown, "A survey of image registration techniques," *ACM Computing Surveys*, vol. 24, pp. 325–376, December 1992.

[101] P. Rogelj and S. Kovacic, "Local similarity measures for multimodal image matching," *Image and Signal Processing and Analysis*, pp. 81–86, 2000.

[102] J. Maintz, E. Meijering, and M. Viergever, "General multimodal elastic registration based on mutual information," in *Proceedings SPIE Medical Imaging* (K. Hanson, ed.), vol. 3338 of *Image Processing*, pp. 144–154, 1998.

[103] R. Shekhar and V. Zagrodsky, "Mutual information-based rigid and nonrigid registration of ultrasound volumes," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 9–22, January 2002.

[104] J. Pluim, J. Maintz, and M. Viergever, "Mutual-information-based registration of medical images: a survey," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 986–1004, August 2003.

[105] T. Brückner, R. Lucht, and G. Brix, "Comparison of rigid and elastic matching of dynamic magnetic resonance mammographic images by mutual information," *Medical Physics*, vol. 10, pp. 2456–2451, October 2000.

[106] Y. He, A. B. Hamza, and H. Krim, "A generalized divergence measure for robust image registration," *IEEE Transactions on Signal Processing*, vol. 51, pp. 1211–1220, May 2003.

[107] L. Freire and F. Godinho, "Registration by maximization of mutual information, a cross validation study," in *Proceedings IEEE International Symposium on BioInformatics and Biomedical Engineering*, pp. 322–329, 2000.

[108] C. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, July and October 1948.

[109] P. Rösch, T. Blaffert, and J. Weese, "Multi-modality image registration using local correlation," in *Computer Assisted Radiology and Surgery (CARS)*, pp. 228–232, November 1999.

[110] J. Weese, P. Rösch, T. Netsch, T. Blaffert, and M. Quist, "Gray-value based registration of ct and mr images by maximation of local correlation," in *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pp. 656–663, Springer-Verlag, 1999.

[111] M. Wachowiak, R. Smolikova, J. Zurada, and A. S. Elmaghraby, "A supervised learning approach to landmark-based elastic biomedical image registration and interpolation," in *Proceedings of the International Joint Conference on Neural Networks*, vol. 2, pp. 1625–1630, 2002.

[112] M. Wachowiak, R. Smolíková, G. Tourassi, and A. Elmaghraby, "Similarity metrics based on nonadditive entropies for 2D-3D multimodal biomedical image registration," in *Proceedings SPIE Medical Imaging* (M. Sonka and J. Fitzpatrick, eds.), vol. 5032 of *Image Processing*, pp. 1090–1100, 2003.

[113] P. Jizba and T. Arimitsu, "The world according to Rényi: Thermodynamics of multifractal systems." eprint arXiv:cond-mat/0207707, http://oai:arXiv.org:cond-mat/0301343, 2003.

[114] C. Tsallis, "Nonextensive statistics: Theoretical, experimental and computational evidences and connections." eprint arXiv:cond-mat/9903356, http://oai:arXiv.org:cond-mat/9903356, 1999.

[115] A. Hamza and H. Krim, "Jensen-Reńyi divergence measure: theoretical and computational perspectives," in *Proceedings IEEE International Symposium on Information Theory (ISIT)*, p. 257, June 2003.

[116] A. Plastino and A. Plastino, "Tsallis entropy and Jayne's information theory formalism," *Brazilian Journal of Physics*, vol. 29, pp. 50–60, March 1999.

[117] J. Burbea and C. Rao, "On the convexity of some divergence measures based on entropy functions," vol. 28, pp. 489–495, May 1982.

[118] J. G. Lopera, J. M. Aroza, A. R. Pérez, and R. R. Roldán, "An analysis of edge detection by using the jensen-shannon divergence," *IEEE Transactions on Medical Imaging*, vol. 13, pp. 35–56, December 2000.

[119] A. Condurache, K. Eck, J. Bredno, and T. Aach, "Statistical-model based identification of complete vessel-tree frames in coronary angiograms," in *Proceedings SPIE Electronic Imaging: Computational Imaging*, 2004.

[120] M. Schrijver, *Angiographic image analysis to assess the severity of coronary stenoses*. PhD thesis, Twente University Press, 2002.

[121] J. Bredno, K. Eck, B. Martín Leung, and P. Rongen, "Method to match interventional devices and vessels, European Patent Application," 2003.

[122] S. Baert and W. Niessen, "Guide wire tracking during endovascular interventions," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 965–972, August 2003.

[123] S. Baert, E. van de Kraats, T. van Walsum, M. Viergever, and W. Niessen, "Biplane image sequences for integrated display in 3-D vasculature," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 1252–1258, October 2003.

[124] G. Koning, J. Tuinenburg, E. Hekking, J. Peelen, A. van Weert, D. Bergkamp, B. Goedhart, and J. Reiber, "A novel measurement technique to asses the effects of coronary brachytherapy in clinical trials," *IEEE Transactions on Medical Imaging*, vol. 21, pp. 1254–1263, October 2002.

[125] M. de Bruijne, W. Niessen, J. Maintz, and M. Viergever, "Localization and segmentation of aortic endografts using marker detection," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 473–482, April 2003.

[126] R. van der Weide, C. Bakker, and M. Viergever, "Localization of intravascular devices with paramagnetic markers in MR images," *IEEE Transactions on Medical Imaging*, vol. 20, pp. 1061–1071, October 2001.

[127] G. Babich and O. Camps, "Weighted Parzen windows for pattern classification," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, pp. 567–570, May 1996.

[128] N. Kwak and C.-H. Choi, "Input feature selection by mutual information based on Parzen window," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 1667–1671, December 2002.

[129] K. Fukunaga and R. Hayes, "The reduced Parzen classifier," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 423–425, April 1989.

[130] B. Jeon and D. Landgrebe, "Fast Parzen density estimation using clustering-based branch and bound," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16, pp. 950–954, September 1994.

[131] Y. Dit-Yan and C. Chow, "Parzen-window network intrusion detectors," in *Proceedings 16th International Conference on Pattern Recognition*, vol. 4, pp. 385–388, 2002.

[132] R. P. R. U. Glassboro, NJ, "Lecture 8: Non-parametric density estimation." http://users.rowan.edu/∼polikar/CLASSES/ECE504_04/Lecture8.pdf, 2003.

[133] D. Sherknies, J. Meunier, R. Mongrain, and J. Tardif, "Three-dimensional trajectory assessment of an ivus trnsducer from single-plane cienanagiograms: A phantom," *IEEE Transactions on Biomedical Engineering*, vol. 52, pp. 543–549, March 2005.

[134] K. Eck, B. Martín Leung, J. Bredno, and J. Weese, "Device and method for displaying ultrasound images of a vessel, European Patent EP 03102691," June 2006.

[135] B. Martín Leung, K. Eck, J. Bredno, and T. Aach, "X-IVUS: Integrated x-ray and IVUS system for the Cathlab," in *Proceedings SPIE Medical Imaging* (J. R.L. Galloway and K. Cleary, eds.), vol. 5744 of *Visualization, Image-guided procedures, and Display*, pp. 378–387, April 2005.

[136] J. Gravensen, "Differential geometry and design of shape and motion." http://www.tdu.dk/, Accessed March 2004.

[137] M. Jourdain, J. Meunier, R. Mongrain, D. Sherknies, J. Weng, and J.-C. Tardif, "3D reconstruction of an IVUS transducer trajectory with a

single view in cineangiography," in *Proceedings SPIE Medical Imaging* (J. Fitzpatrick and J. Reinhardt, eds.), vol. 5747, pp. 1430–1437, April 2005.

[138] L. Mauri, A. O'Malley, D. Cutlip, K. Ho, J. Popma, M. Chauhan, D. Baim, D. Cohen, and R. Kuntz, "Effects of stent length and lesion length on coronary restenosis," *American Journal of Cardiology*, vol. 93, pp. 1340–1346, June 2004.

[139] A. Wahle, G. 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), pp. 1094–1104, SPIE Proceedings, 1999.

[140] M. Sonka, W. Liang, X. Zhang, S. DeJong, S. Collins, and C. McKay, "Three-dimensional automated segmentation of coronary wall and plaque from intravascular ultrasound pullback sequences," in *Proceedings Computers in Cardiology*, pp. 637–640, 1995.

[141] K. Eck, J. Bredno, and T. Stehle, "Absolute alignment and extrapolation of breathing states using image similarity," in *Proceedings SPIE Medical Imaging 2005* (J. R.L. Galloway and K. Cleary, eds.), vol. 5744 of *Visualization, Image-guided procedures, and Display*, pp. 79–86, April 2005.

[142] G. Prause, S. DeJong, C. McKay, and M. Sonka, "Towards a geometrically correct 3-D reconstruction of tortuous coronary arteries based on biplane angiography and intravascular ultrasound," *International Journal of Cardiac Imaging*, vol. 13, pp. 451–462, December 1997.

[143] J. Bloomenthal, *Graphic Gems*, ch. Calculation of reference frames along a space curve. Academic Press, 1990.

[144] A. Wahle, R. Medina, K. Braddy, J. Fox, T. Brennan, J. Lopez, J. Rossen, and M. Sonka, "Impact of local vessel curvature on the circumferential plaque distribution in coronary arteries," in *Proceedings SPIE Medical Imaging* (A. Clough and A. Amini, eds.), vol. 5031 of *Physiology and Function*, pp. 204–213, 2003.

[145] C. Cañero, O. Pujol, R. Toledo, J. Saludes, D. Gil, J. Villanueva, J. Mauri, B. Garcia, J. Gomez, A. Cequier, and E. Esplugas, "Optimal

stent implantation: three-dimensional evaluation of the mutual position of stent and vessel via intracoronary echocadiography," *Computers in Cardiology*, vol. 26, pp. 261–264, September 1999.

[146] D. Gil, P. Radeva, J. Saludes, and J. Mauri, "Automatic segmentation of artery wall in coronary IVUS images: a probabilistic approach," *Computers in Cardiology*, vol. 27, pp. 687–690, September 2000.

[147] P. Brathwaite, K. Chandran, D. McPherson, and E. Dove, "Lumen detection in human IVUS images using region-growing," *Computers in Cardiology*, pp. 37–40, September 1996.

[148] O. Pujol, D. Rotger, P. Radeva, O. Rodriguez, and J. Mauri, "Near real-time plaque segmentation of ivus," *Computers in Cardiology*, pp. 69–72, September 2003.

[149] M. Plissit, D. Fotiadis, L. Michalis, and G. Bozios, "An automated method for lumen and media-adventitia border detection in a sequence of IVUS frames," *IEEE Transactions on Information Technology in Biomedicine*, vol. 8, pp. 131–140, June 2004.

[150] J. Klingensmith, R. Shekhar, and D. Vince, "Evaluation of three-dimensional segmentation algorithms for the identification of luminal and medial-adventitial borders in intravascular ultrasound images," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 996–1011, December 2000.

[151] J.Montagnat, H. Delingette, N. Scapel, and N. Ayache, *Representation, shape, topology and evolution of deformable surfaces. Application to 3D medical image segmentation.* INRIA, May 2000.

[152] H. Delingette, *Simplex meshes: a general representation for 3D shape reconstruction.* INRIA, March 1994.

[153] B. Martín Leung and J. Weese, "Intravascular ultrasound: Overview." Internal report, Philips Research Germany.

[154] A. Silberschatz and P. Galvin, *Operating System Concepts.* Addison-Wesley Publishing Company, 1998.

[155] A. Condurache, T. Aach, K. Eck, and J. Bredno, "Fast detection and processing of arbitrary contrast agent injections in coronary angiography and fluoroscopy," in *Bildverarbeitung fuer die Medizin*, Algorithmen, Systeme, Anwendungen, pp. 5–9, 2004.

[156] N. Bruining, C. von Birgelen, M. Mallus, P. de Feyter, E. de Vrey, W. Li, F. Prati, P. Serruys, and J. Roelandt, "ECG-gated ICUS image acquisition combined with a semi-automated contour detection provides accurate analysis of vessel dimensions," *Computers in Cardiology*, pp. 53–56, September 1996.

[157] N. Bruining, C. von Birgelen, F. Prati, M. Mallus, W. Li, W. D. Hoed, M. Patijn, P. D. Feyter, P. Serruys, and J. Roelandt, "Dynamic three-dimensional reconstruction of ICUS images based on an ECG-gated pull-back device," in *Engineering in Medicine and Biology Society, 1996. Bridging Disciplines for Biomedicine. Proceedings of the 18th Annual International Conference of the IEEE*, vol. 2, pp. 664–665, 1997.

[158] P. Schoenhagen, S. Halliburton, R. White, D. Vince, S. Nissen, and E. Tuzcu, "Characterization of coronary atherosclerotic plaques and the significance of vessel calcification," *Applied Radiology*, vol. 30, no. 11, pp. 46–53, 2001.

[159] A. El-Shafei and M. Kern, "New techniques for the evaluation of the vulnerable plaque," *Journal Invansive Cardiology*, vol. 14, pp. 129–137, March 2002.

[160] G. Mintz and S. N. *et al.*, "Clinical expert consensus document on standards for acquisition, measurement and reporting of intravascular ultrasound studies (IVUS)," *Journal of the American College of Cardiology*, vol. 37, pp. 1478–1492, April 2001.

Curriculum Vitae

Name Bárbara Martín Leung
Contact b_leung@mail.com
Born in Madrid, Spain 12.09.1977

Education

Sep. 1995 - Sep. 2000 **Telecommunication Engineering** at *Escuela Técnica Superior de Ingeniería de Telecomunicación*, *Universidad Politécnica de Madrid* (Spain).

Oct. 2000 - June 2001 **Master thesis** at the Royal Institute of Technology (Sweden), Speech, Music and Hearing Department. Title: **Waveform Interpolation speech coding for the Internet**.

Sep. 2001 - Dec. 2004 **Research Associate** at the University of Lübeck (Germany), Institute for Signal Processing. Medical image processing for cardiovascular interventions.

Sep. 2005 - July 2006 **Chinese** at the Beijing Foreign Studies University (China). Scholarship.

Work Experience

June 2000 - Aug. 2000 **Airtel S.A** (nowadays Vodafone), Internship, Interconnection and Roaming Group. Interconnection with other Spanish operators and roaming enabling.

Feb. 2005 - May 2005 **Philips Medical Systems**, Software Engineer, X-ray Predevelopment group. Adaptation of medical image processing algorithms to Philips Cathlab systems. Main tools: C++ Visual Studio.

May 2007 - Feb. 2009 **AEQ**, R&D Engineer. Professional audio broadcasting. Main tools: C, Nucleus OS, Netdisturb.

Publications

Martin Leung, B.; Eck, K.; Bredno, J.; Aach, T.: *XIVUS: Integrated X-ray and IVUS system for the Cathlab*, Proc. SPIE Medical Imaging 2005: Visualization, Image-Guided Procedures, and Display, Vol. 5744, Page(s): 378-387.

Bredno, J.; Martin Leung, B.; Eck, K.: *Software architecture for live enhancement of medical images*, Proc. SPIE Electronic Imaging 2004: Real-Time Imaging, Vol. 5297, Page(s): 122-133.

Bredno, J.; Martin Leung, B.; Eck, K.: *Algorithmic solutions for live device-to-vessel match*, Proc. SPIE Medical Imaging 2004, Vol. 5370, Page(s): 1486-1497.

Bredno, J. Martin Leung, B.; Eck, K.: *X-IVUS Options for integrating local information into x-ray controlled interventions*, Internal technical note TN 32/2003.

Bredno, J.; Eck, K.; Martin Leung, B.; Rongen, P.M.J.: *Device and method for combined display of angiograms and current X-ray images*, US20060257006.

Martin Leung, B.; Eck, K.; Bredno, J.; Weese, J. : *Device and method for displaying ultrasound images of a vessel*, European Patent EP 03102691, 2003.

Martin Leung, B.; Eck, K.; Bredno, J.; Aach, T. : *Mutual Information Based Respiration Detection*, Computer Assisted Radiology and Surgery 2002, Page(s): 1085-1092.

Martin Leung, B.; Eck, K. : *Enhanced Smart Mask Navigation support for intracoronary applications*, Internal report TR 1612/2002.